DOMAIN RESTRICTION: THE PROBLEM OF THE VARIABLE LOCATION REVISITED

Diego Feinmann

Institut Jean Nicod, Département d'Etudes Cognitives, ENS, EHESS, CNRS, PSL Research University

Abstract

Two theories of implicit domain restriction have gained considerable prominence over the last two decades. According to von Fintel (1994), quantifiers come with covert restrictors and, as a result of this, induce domain restriction; according to Stanley (2002; Stanley and Szabó 2000), by contrast, nouns, as opposed to quantifiers, come with covert restrictors. In this article, I do three things. First, I assess previous arguments for and against these two accounts and show that none of them is conclusive. Second, I advance a novel empirical argument based on the observed pragmatic behaviour of bare nouns, an argument that falsifies Stanley's theory while providing clear evidence in support of von Fintel's (1994). Finally, I discuss the relevance of the bare noun data in the context of another important debate—namely, whether domain restriction is a local mechanism only, or whether it can also be achieved by global means.[*]

1   INTRODUCTION

It is a well-known fact that quantified NPs are context-dependent. For example, *every school*, in the examples below, means something different in each case: in (1)a, *every* quantifies over the set of UK state schools—(1)a has the reading in (1)b—whereas in (2)a, it does over the set of French state schools—(2)a has the reading in (2)b. This phenomenon goes by the name of (quantifier) domain restriction.

---

(1)   Context: The UK Government has decided to reduce the budget for education.

        a. Every school is going to suffer.
        b. Every *UK state school* is going to suffer.

(2)   Context: The French Government has decided to reduce the budget for education.

        a. Every school is going to suffer.
        b. Every *French state school* is going to suffer.

To handle this sort of data, a number of theoretical approaches are, *prima facie*, available—for an overview, see, for example, von Fintel (1994: §2.2.2.) or Stanley and Szabó (2000: §4). In this article, I will focus on one of such approaches, namely, the *domain variable approach*, which, over the last two decades, has gained considerable prominence. On this approach, implicit restrictions are represented in the syntax by means of a variable (in the case of DPs, a variable ranging over predicates) whose value is given by context. Following von Fintel (1994), I will notate this variable with an italic capital C; also following this author, I will sometimes refer to $g(C)$, the denotation of C under some assignment g, as the *resource domain*.[1]

The domain variable approach gave rise to two influential theories. One of these theories has it that quantifiers are context-sensitive; according to it, quantifiers come with a domain variable or covert restrictor C (i.e. C appears next to the quantifier at LF) and $g(C)$, or the resource domain, is intersected with the denotation of the NP.[2] This version of the domain variable approach is associated with von Fintel (1994)'s seminal work, and I shall refer to it as *Quantifier Restriction Theory* (QRT, for short).[3] The other theory, due to Stanley and Szabó (2000) and Stanley (2002), has it that nouns are context-sensitive; according to it, nouns come with a covert restrictor C (i.e. C appears next to the head noun at LF) and $g(C)$ is intersected with the denotation of the head noun (because of this, Stanley and Szabó's account goes by the name of *Nominal Restriction Theory*, NRT henceforth). Table 1 illustrates these two proposals.

---

[1] To allow for binding, C is often assumed to have a more complex structure: the standard move is to say that C is composed of a functional variable f (of type $\langle e,\langle e,t\rangle\rangle$) and an argumental variable x (of type e); see von Fintel (1994) and Stanley and Szabó (2000).

[2] $g(C)$ is a one-place predicate meaning, that is, the characteristic function of a set of entities. For ease of exposition, and unless it is necessary to do otherwise, I shall identify one-place predicates with sets of entities, rather than with the characteristic functions thereof.

[3] See also Westerståhl (1984) and Martí (2003).

Table 1. (1)a/(2)a through the lens of QRT and NRT.

QRT	$_S[_{DP}[_D [\text{Every}_C] _{NP}[_N[\text{school}]]] _{VP}[\ldots]]$
	Interpretation: $[\![\text{Every}_C \text{ NP VP}]\!]^g = 1$ iff $[\![\text{NP}]\!]^g \cap g(C) \subseteq [\![\text{VP}]\!]^g$
NRT	$_S[_{DP}[_D [\text{Every}] _{NP}[_N[\text{school}_C]]] _{VP}[\ldots]]$
	Interpretation: $[\![\text{school}_C]\!]^g = [\![\text{school}]\!]^g \cap g(C)$

It should be noted that whether $g(C)$ is intersected with the head noun (as in NRT) or the NP (as in QRT) is not an arbitrary choice: it follows from where at LF C is stipulated to be located and standard assumptions about compositionality (cf. Larson and Segal 1995, Heim and Kratzer 1998). Indeed, if C is located next to the determiner-quantifier, then $g(C)$ cannot be intersected with $[\![N]\!]^g$ (i.e. if it was intersected with $[\![N]\!]^g$, then the meaning of the NP would not be exclusively determined by the meaning of its immediate daughters); and, if C is located next to the head noun, then $g(C)$ cannot be intersected with the $[\![NP]\!]^g$ (i.e. if C and the head noun occur next to each other at LF, then $g(C)$ must be composed with the denotation of the head noun).

In the present article, I shall be concerned only with the phenomenon of domain restriction in DPs; it is worth noting, however, that QRT, as conceived in von Fintel (1994), is a general account of quantifier domain restriction, one which applies to VPs as much as to DPs. Take (3), for example.

(3)   John always orders café crème.

(3) may be true if uttered by the waiter of Café de Flore (where John goes on Tuesdays and Thursdays) but false if uttered by the waiter of Café Oberkampf (where John goes on Mondays and Fridays): indeed, John may always order a café crème when he goes to Café de Flore, but never when he goes to Café Oberkampf. To deal with the context-dependence of (3), von Fintel (1994) posits that adverbial quantifiers such as *always*, in the same way as determiner-quantifiers, come with a covert restrictor (in the case of adverbial quantifiers, this restrictor is a variable over sets of situations).

NRT, as opposed to QRT, is just an account of domain restriction in DPs and it is not clear how it could be extended to deal with (3). The fact that NRT cannot account for the context-dependence of (3) is not a strong argument against NRT, I believe: NRT advocates, for example, could argue that the sort of context-dependence that a sentence such as (3) exhibits is of a different nature from the context-dependence of (1) and (2)—and thus refrain from providing a uniform treatment of these data. I simply want to note that QRT is, or at least aspires to be, a general account of domain restriction in natural language, the guiding idea being that quantifiers (and not just determiner-quantifiers) come with covert restrictors.

In the early 2000s, Stanley (2002; Stanley and Szabó 2000) offered three arguments in support of NRT (and against QRT), arguments which turned out to be very influential; Stanley's arguments involve considerations about cross-sentential anaphora, superlative modifiers and the pragmatic behaviour of relative adjectives. Soon afterwards, Kratzer (2004) presented the *Fake Philosopher* argument, later echoed in Gillon (2006), Schwarz (2009), and von Fintel (2014). This argument, which draws on data reported in Breheny (2003), has been taken to provide evidence against NRT. In view of the conflicting evidence, the debate came to a deadlock, and some theorists have even suggested that the domain variable approach may have to be abandoned.[4]

In this article, I will revisit the problem of the location of *C*. To begin with, I will show that the arguments put forward in Stanley (2002) and Stanley and Szabó (2000) in support of NRT are unsuccessful; though voices have been raised against NRT (e.g. Giannakidou 2004, von Fintel 2014), a thorough refutation of these arguments has, to my knowledge, not yet been provided. Second, I will examine the *Fake Philosopher* argument. This argument, although it indicates that NP-level restrictions exist, as stipulated in QRT, does not rule out NRT, nor can it be taken to provide conclusive evidence for QRT (or so will I argue). Third, I will put forward a novel empirical argument based on the observed pragmatic behaviour of bare nouns, an argument that disproves NRT while providing support for QRT. To conclude, I will discuss the relevance of the bare noun data in the context of another important debate—namely, whether a general account of

---

[4] 'We currently have a number of good arguments supporting conflicting conclusions about where in the structure domain restriction variables are introduced. Unless we can debunk one set of these arguments, the outlook for this type of approach is not very promising.' (Schwarz 2009: 108-9)

domain restriction should also permit 'global' restrictions, as in Kratzer (2007) and Schwarz (2009, 2012).

2   REASSESSING THE ARGUMENTS
2.1   Stanley's arguments in support of NRT
2.1.1   The argument from anaphora

Stanley and Szabó (2000)—S&S henceforth—observe that (4) can have two possible readings, depending on how cross-sentential anaphora is resolved.

(4)   <u>Context:</u> Talking about a certain village…

   Most people regularly scream. They are crazy.

   Reading A: The people in the village are crazy.
   Reading B: The people in the village who regularly scream are crazy.

On the basis of this data point, S&S put forward two arguments in support of their account of domain restriction. The first one is tied up to Reading A and goes as follows:

> Consider the first reading [Reading A], that everyone in the village is crazy. Ideally, one would wish to say that cross-sentential anaphora of this sort requires antecedents that are constituents (nodes) of a preceding logical form. However, if the domain variable co-habits a terminal node with 'most', there is no single node in the logical form of the first sentence of [(4)] whose associated semantic value is the set of people in the village. In our favoured approach, however, there is such a node: the one labelled '⟨people, *i*⟩'. So, our favoured approach provides a far more natural account of the first reading of the second sentence of [(4)]. (Stanley and Szabó 2000: 257).

It is not clear to me in which sense NRT provides 'a far more natural account' of Reading A. S&S appear to be working under the assumption that *they,* in the second sentence of (4), refers to the same referent as *people$_C$* (in their notation, '⟨people, *i*⟩'). If this were true, then, indeed, Reading A could be derived in a straightforward manner (if NRT were to be assumed): *they* could simply co-refer with *people$_C$*. The issue is that *people$_C$* is of type ⟨*e,t*⟩, whereas *they* is of type *e* (i.e. *they* denotes a plural individual (Link 1983)): hence, *they* cannot co-refer with *people$_C$*. To derive

Reading A via co-reference, the node that S&S need isn't *people_C* but, rather, *the people_C*, which is neither present in the NRT LF nor in the QRT LF (see Fig. 1 below).[5, 6]

Figure 1. The LFs of (4) according to NRT and QRT, respectively.
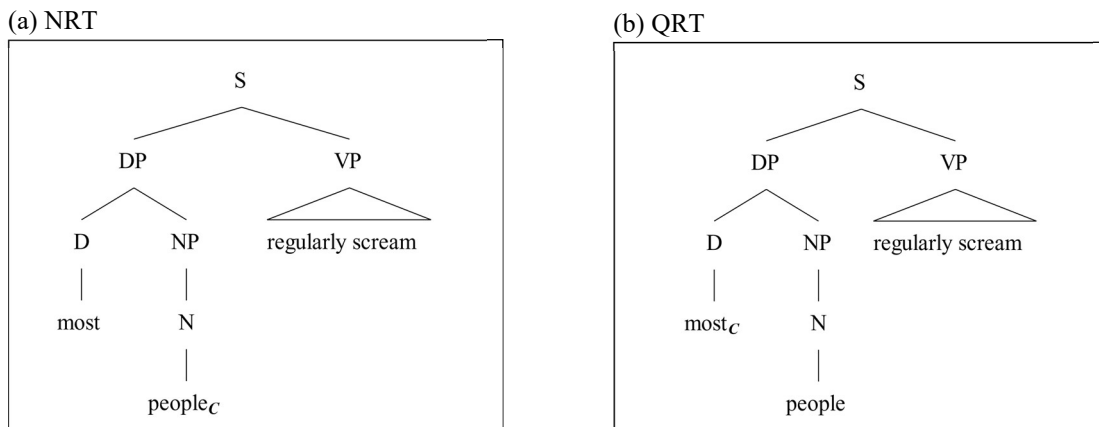
The second argument is tied up to Reading B. This reading, just like Reading A, cannot be derived via co-reference, no matter whether NRT or QRT is assumed: indeed, on either analysis, there just isn't a linguistic constituent whose denotation is *the people in the village who regularly scream* (cf. Fig. 1). S&S note, however, that, if the syntax in Fig. 1(a) were to be assumed, then Neale's (1990) account of e-type pronoun interpretation could be invoked to derive Reading B.[7] On this account, *they* in (4) would be interpreted according to the following rule:

---

[5] In order to derive Reading A via co-reference, S&S would need to make non-trivial semantic assumptions, e.g.: claiming that nouns are of type *e*—see Lasersohn (2020) for a recent proposal—and/or positing some type-shifting mechanism; but, if such a route were to be taken, it is not clear to me that the resulting account could be called 'natural' (I guess that would depend on what the benchmark for naturalness is, as well as on the precise details of the account).

[6] An anonymous reviewer points out that S&S's argument would make more sense if it was interpreted against the backdrop of Elbourne's (2001, 2005) description-theoretic approach to e-type anaphora. This is only partly true. On Elbourne's approach, pronouns are definite articles followed by an NP which is deleted under identity with a preceding NP: for example, in (4), if NRT were to be assumed, the LF of *they* would be *the people_C*, and the NP would be deleted under identity with *people_C* in the first sentence. Notice, however, that NRT has no clear advantage over QRT on this approach, at least not in connection with the derivation of Reading A; indeed, if QRT were to be assumed, the LF of *they* would be *the_C people* and the NP would be deleted under identity with *people* in the first sentence.

[7] Following Nouwen (2020), by 'e-type pronoun' I mean a pronoun that (i) has a linguistic antecedent, (ii) is not bound by its antecedent in the sense that the pronoun occurs in the syntactic scope of its antecedent and its reference co-varies with the antecedent, and (iii) doesn't refer to the same referent its antecedent refers to.

(5)   If x is a pronoun that is anaphoric on, but not c-commanded by a non-maximal quantifier '[Dx:Fx]' that occurs in an antecedent clause '[Dx:Fx](Gx)', then x is interpreted as '[the x: Fx&Gx]'.[8]

S&S reason as follows:

> [N]ote what happens if we apply [(5)] to [(4)], on the assumption that the domain variable co-habits the node of the quantifier 'most'. In constructing the definite description which gives the interpretation of the pronoun in the second sentence, we drop the quantifier and lose the domain variable with it. 'They' is then interpreted as [the x: person(x) & regularly-scream(x)], which results in an unrestricted reading of the second sentence, according to which everyone in the universe who regularly screams is crazy. By contrast, if, as on our favored approach, the variable co-habits a node with 'person', 'they' is interpreted as [the x: ⟨person, i⟩(x) & regularly-scream(x)], we obtain the desired reading of the second sentence, according to which it expresses the proposition that every person in the village who regularly screams is crazy. (Stanley and Szabó 2000: 257-8)

S&S's argument, it seems to me, relies on a far too literal understanding of (5). Neale's rule has not been formulated to deal with quantified sentences in which restrictions to quantifier domains are represented in the syntax. If (5) were to be re-formulated as in (6), then Reading B could be derived using the LF in Fig. 1(b) (the QRT analysis).

(6)   If x is a pronoun that is anaphoric on, but not c-commanded by a non-maximal quantifier '[D$_C x$:Fx]' that occurs in an antecedent clause '[D$_C x$:Fx](Gx)', then x is interpreted as '[the x: **C**x&Fx&Gx]'.

It is also worth noting that S&S's (second) argument is conditional on Neale's (1990) account being correct (or, at least, on the right track). This account, however, faces several difficulties; for recent discussions, see, for example, King and Lewis (2018) and Nouwen (2020).

To conclude, S&S's arguments based on example (4), as far as I can tell, do not constitute evidence in support of NRT (and/or against QRT).

---

[8] Neale's (1990: 266) rule (P5b). A quantifier '[Dx:Fx]' is non-maximal iff there is some G for which '[Dx:Fx](Gx)' is true but '[every x: Fx] (Gx)' is false. Hence, *most* is non-maximal.

2.1.2   The argument from superlatives

The second argument, reported by Stanley (2002, ascribed to Delia Graff Fara, p.c.), centres on the pragmatic behaviour of superlative noun phrases. Consider (7)a, which, in the context stipulated, has the reading in (7)b.

(7)   Context: Talking about Cornell students…

            a. The tallest student is nice.
            b. The tallest student from Cornell is nice.

Stanley (2002) notes that, under the assumption that *tallest* takes the head noun as its argument and returns a set consisting of the tallest individual in the set denoted by the head noun, attaching *C* to the determiner yields an odd result: namely, that (7) can only be truthfully uttered if the tallest student in the world happens to be a student at Cornell. If, on the other hand, *C* were to be attached to the noun, as stipulated in NRT, the correct interpretation would be obtained, as *student* would be restricted to *Cornell student* prior to being composed with *tallest*.

This is not a strong argument in support of NRT: any modern semantic treatment of *-est* can generate the correct truth-conditions for (7)a, and it can do so without stipulating that domain restriction is computed on the noun.[9] Let's consider, for example, Heim's (1999) semantics for the superlative morpheme, which, incidentally, builds upon von Fintel's (1994) account of domain restriction.

Heim (1999) assumes that a gradable adjective denotes a function from degrees to $\langle e,t \rangle$ functions, following Seuren (1973), Cresswell (1976), and others. The semantics of *tall*, for example, is as follows:

(8)   For any degree d and individual x, $[\![tall]\!](d)(x) = 1$ iff x is tall to degree d (i.e. iff x's height is equal or greater than d).[10]

---

[9] See, for instance, Heim (1999), Farkas and Kiss (2000), or Sharvit and Stateva (2002).

[10] According to (8), every gradable adjective meaning R is *downward monotonic*, in the sense that the following holds of R: $\forall x,d,d'[[R(d)(x) = 1 \land d > d'] \rightarrow R(d')(x) = 1]$. Hence, on this account, *John is four feet tall* entails *John is three feet tall*.

On Heim's analysis, modified nominal expressions such as *tall student*, where the modifier is a gradable adjective, also denote functions from degrees to $\langle e,t \rangle$ functions.[11]

As is standard in semantics, Heim (1999) treats the superlative morpheme *-est* as a degree quantifier and, following von Fintel (1994), stipulates that this morpheme, by virtue of being a quantifier, introduces a domain variable.[12] The upshot of this is that the meaning of *-est* ends up taking three arguments: a domain argument (i.e. the value of C under some assignment g), a gradable adjective meaning, and an individual. Heim's (1999) semantics for *-est* is as follows:[13]

(9)   Let x be an individual, R a gradable adjective meaning, and K the relevant domain:

$[\![\text{-est}]\!](K)(R)(x)$ is only defined if $x \in K \wedge \forall y[y \in K \rightarrow \exists d[R(d)(y) = 1]]$; whenever defined, $[\![\text{-est}]\!](K)(R)(x) = 1$ iff $\exists d[R(d)(x) = 1 \wedge \forall y[[y \neq x \wedge y \in K] \rightarrow R(d)(y) = 0]]$

Let's now return to *tallest student*—on Heim's (1999) account, *tallest_C student*—, which she analyses as having the LF-constituency [*-est_C* [*tall student*]] (rather than [[*-est_C tall*] *student*]); according to the definition in (9), $[\![\text{-est}]\!]$ takes K, the relevant domain or comparison class, as argument (in the case at hand, the set of Cornell students), $[\![\text{-est}]\!](K)$ takes R, a gradable adjective meaning, as argument (in the case at hand, the denotation of the modified nominal *tall student*), while $[\![\text{-est}]\!](K)(R)$ takes x, an individual, as argument: $[\![\text{-est}]\!](K)(R)(x) = 1$ iff, for some d, x is R to degree d and no other individual in K is R to degree d.

(9), as shown in (10), predicts the intuitively correct reading of the sentence *the tallest student is nice*.[14]

---

[11] This move, of course, raises a number of technical issues, which Heim (1999) is aware of (for example, she considers the possibility that the semantics of gradable adjectives and modified nominal expressions may be related by a type-shifting rule). These issues, however, go beyond the scope of this article.

[12] Heim (1999: 3) writes: 'Following von Fintel (1994), we may localize the context-dependency of quantifiers like *every* in an extra argument, a phonetically unrealized predicate variable that appears next to the determiner at LF and receives a value from the context of utterance. Adapted to the case at hand, this suggests that *-est* likewise takes an additional argument'.

[13] In (9), I am providing the curried version of Heim's entry.

[14] I follow the notational convention from Heim and Kratzer (1998), where $\lambda x : \psi \, . \, \alpha$ is written for the function that maps x to α if ψ, and otherwise incurs a presupposition failure.

(10)   a. $[\![\text{tallest}_C \text{ student}]\!]^g = [\![\text{-est}]\!]^g \,([\![C]\!]^g)\,([\![\text{tall student}]\!]^g) = \lambda x_e : x \in g(C) \land \forall y \,[y \in g(C) \to$
$\exists d[\text{tall-student}'(d)(y) = 1]] \,.\, \exists d[\text{tall-student}'(d)(x) = 1 \land \forall y[[y \neq x \land y \in g(C)] \to \text{tall-}$
$\text{student}'(d)(y) = 0]]$

b. $[\![\text{the}]\!]^g = \lambda f_{\langle e,t\rangle} : \exists!x[f(x) = 1] \,.\, \iota x[f(x) = 1]$

c. $[\![\text{the}]\!]^g \,([\![\text{tallest}_C \text{ student}]\!]^g)$ is only defined if $\exists!x[[\![\text{tallest}_C \text{ student}]\!]^g(x) = 1]$—for that to
be the case, any x that satisfies the condition just written, has also to satisfy the
presupposition of 'tallest$_C$ student'; whenever defined,
$[\![\text{the}]\!]^g \,([\![\text{tallest}_C \text{ student}]\!]^g) = \iota x[\exists d[\text{tall-student}'(d)(x) = 1 \land \forall y[[y \neq x \land y \in g(C)] \to \text{tall-}$
$\text{student}'(d)(y) = 0]]]$

d. $[\![\text{is nice}]\!]^g \,([\![\text{the tallest}_C \text{ student}]\!]^g)$ inherits the presupposition of 'the tallest$_C$ student';
whenever defined, $[\![\text{is nice}]\!]^g \,([\![\text{the tallest}_C \text{ student}]\!]^g) = 1$ iff nice$'(\iota x[\exists d \,[\text{tall-student}'(d)(x)$
$= 1 \land \forall y \,[[y \neq x \land y \in g(C)] \to \text{tall-student}'(d)(y) = 0]]]) = 1$[15]

According to (10)d, the sentence *the tallest student is nice*, in the context stipulated in (7) and
whenever defined, is true iff the unique individual x such that, for some d, x is a tall student to
degree d and no other individual in the set of Cornell students is a tall student to degree d is nice.

It should be noted that Heim's (1999) proposed treatment of *-est* is compatible with a QRT
treatment of the definite article, as shown below.

(11) $[\![\text{the}_{C'}]\!]^g = \lambda f_{\langle e,t\rangle} : \exists!x[x \in g(C') \land f(x) = 1] \,.\, \iota x[x \in g(C') \land f(x) = 1]$

Indeed, if (11) was used (instead of (10)b), there would be an additional failure condition—namely,
that whoever is the tallest man in $g(C)$ is also in $g(C')$. Under the (reasonable) assumption that
interpreters do not domain-restrict to provoke presupposition failure, the extra domain-restricting
device introduced by (11) will not have any tangible impact on the meaning of (7)a. It could also
be proposed that the surface article is in fact part of the discontinuous item *the -est* (e.g. Szabolcsi
1986); on this account, there would just be one covert restrictor. I will not develop the details of
such an account here.

---

[15] For ease of explanation, I am here assuming that *nice* isn't a gradable adjective (though, of course, it is).

To conclude, it is worth noting that Heim's (1999) semantics has far greater empirical coverage than NRT. Indeed, NRT has no resources to deal, for example, with (12)a, which has the reading in (12)b.

(12)   Context: Talking about three brothers…
         a. The tallest is nice.
         b. The tallest (among the three brothers) is nice.

In (12)a, there is no set-denoting noun that can be restricted and, as a result, there is no way to derive (12)a's reading in the stipulated context via NRT.[16] Heim's (1999) semantics, by contrast, can handle (12)a in a straightforward manner.

2.1.3   The argument from context-sensitive adjectives

The third argument, put forward in Stanley (2002), is as follows. Stanley observes that (13) may be true if talking about Smith's piano-playing at a dinner party, but not true if talking about Smith's piano-playing at a formal concert setting.

(13)   Smith is a remarkable pianist.

Next, he argues that the perceived context-dependence of (13) can be captured if one assumes that nouns introduce covert restrictors, as stipulated in NRT: in one case the denotation of *pianist* would be intersected with the set of dinner party musicians and, in the other, with the set of professional musicians.

This is not an argument in support of NRT though. (14), in the same way as (13), may be true if talking about Smith's piano-playing at a dinner party, but not true if talking about Smith's piano-playing at a formal concert setting: in (14), however, there is no set-denoting noun to be restricted.

---

[16] Any attempt to argue that, in (12)a, there is an unpronounced noun at LF will face a non-trivial issue: indeed, such an account appears to predict that (i)a can have the reading in (i)b (and it clearly cannot).

(i)   Context: Talking about three brothers…
         a. # The tall is nice.
         b. The tall (brother) is nice.

(14)   Smith is remarkable.

The standard assumption that relative adjectives access a comparison class (e.g. Bartsch and Vennemann 1972, Klein 1980) is enough to account for the context-dependence of (13) and (14). No appeal to domain restriction is needed.

The fact that none of Stanley's (and Stanley and Szabó's) arguments goes through does not disprove NRT, however. To rule out NRT, one needs to provide direct or indirect evidence that nouns cannot induce domain restriction.

2.2   The Fake Philosopher argument

Consider (15)a below; QRT predicts (15)a to have the reading in (15)b, while NRT predicts it to have the reading in (15)c. The observation is that the reading that NRT predicts isn't available (Kratzer 2004). (In this example, and in examples that follow, the rightwards arrow with tail signals that the sentence being pointed at is a reading of the sentence under scrutiny, while the rightwards arrow with tail and stroke signals that the sentence being pointed at isn't a reading of the sentence under scrutiny).

(15)   Context: Talking about American people…
         a. Every fake philosopher is from Idaho. (Kratzer 2004; cf. Breheny 2003)[17]

         b. ↣ Every [**American** [fake philosopher]] is from Idaho. (QRT reading)
         c. ↣ Every [fake [**American** philosopher]] is from Idaho. (NRT reading)

This seems correct; consider, for example, the existence of a genuine philosopher who is from Paris (France) but pretends to be American: such a person would not be a counterexample to (15)a

---

[17] Example (15)a is reported in Kratzer (2004); however, it makes the same point as Breheny's (2003) example (21)b.

in the stipulated context, nor would it be a counterexample to (15)b, but it would count as a counterexample to (15)c.[18]

It should be noted that Breheny (2003), who discusses examples that behave just like (15)a, also draws attention to an example, given in (16)a below, which he takes to suggest that domain restriction can, in some cases, be computed in the scope of intensional adjectives.

(16) a. Our boss sends every former employee a Christmas card.
      b. ↝ Our boss sends every [former [employee of our company]] a Christmas card.

(16)a, quite clearly, can have the reading in (16)b; can it be concluded, on the basis of this example, that domain restriction is being computed on *employee*? I do not think that it can; *employee*, just like *foreigner* or *enemy*, is a relational (context-sensitive) noun: the determination of its actual denotation requires access to contextual information (Mitchell 1986, Partee 1989, Kennedy and McNally 2005, among others). Indeed, one isn't an employee *simpliciter*, but an employee *of some organisation or other*. There is thus no need to appeal to domain restriction to account for (16)a: presumably, (16)a has the reading that it has because the denotation of *employee* takes an implicit argument.

If the analysis that I am advocating is along the right lines, then the *seeming* domain restriction effect observed in (16)a should disappear as soon as *employee* is replaced by another noun (a noun whose denotation isn't context-dependent). This prediction is born out; consider, for example, (17).

(17) <u>Context</u>: In John's company, there is always a doctor on site, just in case one of the company's workers needs medical assistance. This isn't a very appealing job and, therefore, the on-site doctor tends to change quite frequently. John tells Jenny, a newly hired employee…

      a. ? Our boss sends every former doctor a Christmas card.
      b. ↝ Our boss sends every [former [company-doctor]] a Christmas card.

---

[18] For any world w, $[\![$[fake [American philosopher]]$]\!]^{w,g}$ = {$x : x$ is either a fake American and real philosopher in w or an American and fake philosopher in w or a fake American and a fake philosopher in w}. Evidence that this the correct denotation of [*fake* [*American philosopher*]] comes from the fact that one can truly say 'she is a fake American philosopher' if/when *she* refers to (a) a real philosopher who pretends to from the US, (b) a US citizen who pretends to be a philosopher, or (c) a person who pretends to be from the US and also pretends to be a philosopher.

(17)a, as far as I can tell, doesn't have (17)b as a possible *reading*.[19] Indeed, even in the context stipulated, (17)a entails that the people who are sent Christmas cards are no longer doctors—hence its pragmatic oddness: why would the boss be sending Christmas cards to people who, for some reason or other, have quit the medical profession? (17)b, by contrast, does not entail this and feels apposite: it makes sense that the boss may want to send Christmas cards to doctors who no longer work for the company but did at some point. The important empirical point that (15)a makes (i.e. contextual domain restriction below an intensional adjective isn't possible) stands: as shown above, (16)a is confounded.[20]

The *Fake Philosopher* argument, at least *prima facie*, poses a double challenge to NRT: (i) NRT fails to predict (15)b, a reading that exists; (ii) NRT predicts (15)c, a reading that does not exist. (i) shows that the resource domain can be intersected with the NP denotation, as stipulated in QRT;

---

[19] Given the context stipulated, (17)a could be 'made sense of': that is, it could plausibly be repaired into something like (17)b. However, the fact that (17)a strikes one as an odd sentence indicates that *doctor* is not restricted to *company doctor*: if it was, then no oddness would be perceived, because (17)a would have (17)b as a reading, and (17)b isn't odd in the stipulated context—cf. (15)a, for example, which reads naturally as 'every [American [fake philosopher]] is from Idaho'.

[20] Breheny (2003), it should be noted, also discusses two examples—given as (ii) and (iii) below—that are claimed to involve multiple bound dependencies. On the basis of these examples, he argues that domain restriction below an intensional adjective is possible.

(ii)   Context: Though art dealers are entirely scrupulous, many famous artists have become paranoid due to the long-term effects of absinthe; in particular, they each imagine that the otherwise scrupulous dealers are out to ruin the market value of just their paintings.

a. Every paranoid artist thinks no dealer will stop at selling every forged painting.
b. Every$_x$ paranoid artist thinks no$_y$ dealer will stop at selling every$_z$ [[forged [painting by x]] coming into y's possession].

(iii)   a. Every government makes allowance for the fact that no type of watermarking will prevent the circulation of every counterfeit note.
b. Every$_x$ government makes allowance for the fact that no$_y$ type of watermarking will prevent the circulation of every$_z$ [counterfeit [note of currency of x marked by y]].

The claim is that (ii)a and (iii)a have (ii)b and (iii)b as readings, respectively. I absolutely fail to get these readings and would need to see experimental evidence to be persuaded that such readings exist. (ii)a, to me, even in the context stipulated, means something like 'every$_x$ paranoid artist thinks no$_y$ dealer will stop at selling every$_z$ [[forged painting] coming into y's possession]'; indeed, (ii)a seems false in a world in which every paranoid artist thinks that the relevant dealers would willingly sell every fake copy of their works while refusing to sell fake copies of other artists' works (if (ii)a had (ii)b as a reading, then (ii)a would be true in such a world). Likewise, (iii)a, as far as I can tell, reads as 'every$_x$ government makes allowance for the fact that no$_y$ type of watermarking will prevent the circulation of every$_z$ [[counterfeit note] present in government x's territory]]'; indeed, (iii)a seems false if, say, the Polish Government, while making allowance for the fact that watermarking on Złoty banknotes will fail to prevent the circulation of counterfeit Złotys, makes no allowance for the fact that watermarking on Euro banknotes will fail to prevent the circulation of counterfeit Euros within Poland (if (iii)a had the reading in (iii)b, then it would not be possible, on the basis of this information, to establish whether (iii)a is true or false).

however, does (i) refute NRT? It depends. There are two possible construals of NRT: one can construe NRT as entailing that all domain restriction effects observed in DPs are, in fact, *nominal* restrictions; one could also construe NRT as being agnostic as to whether restrictions *à la* von Fintel are possible: after all, both determiner-quantifiers and nouns could come with covert restrictors. Let's call the former construal 'Strong NRT' and, the latter, 'Weak NRT'. (i), as far as I can tell, rules out Strong NRT—the version of NRT that Stanley appears to advocate[21]—, but it doesn't rule out Weak NRT. Let's now move to (ii); does (ii) rule out Weak NRT? In what follows, I will argue that it doesn't: (ii) presupposes that, if nouns came with covert restrictors, nominal restriction would be computed below an intensional adjective. This, as will be shown, cannot be presupposed.

Consider (18):

(18)   Every school in the UK is free.

Such a sentence, if uttered in a context that does not make salient any particular subset of schools, would strike anyone as false; however, if the interpreter of (18) could domain-restrict *school in the UK* to *state school in the UK*, (18) wouldn't be judged to be false. This goes to show that one cannot domain-restrict to an arbitrary set of entities, not even in cases in which doing so would save the sentence from falsehood—and hence keep afloat the presumption that the speaker complies with the maxim of Quality (Grice 1975).

The standard assumption, which I think is in-keeping with the domain restriction data, is that g(C), the resource domain, is (the characteristic function of) the set of entities that is being talked about (or is relevant) in a given context (cf. von Fintel 1994, Gillon 2006, Schwarz 2009). For example, if (18) were to be uttered in a context that made it clear that British state schools (and only these schools) are being talked about, then the set of entities that *every* quantifies over would be restricted to the set of UK state schools. Consider (19) below, for example.

---

[21] Consider, for example, the following passage from Stanley (2002: 372): 'According to NRT, the intuitive restriction on quantificational determiners such as "every", "some", and "most" is not due, as may seem obvious, to a restriction on the quantificational expressions themselves, but rather to a restriction on the nominal complements of these determiners.'

(19)   <u>Context</u>: Two British civil servants and an education specialist from the World Bank are having a discussion around specific problems facing the British state schooling system. The education specialist points out that, according to her sources, several schools in the North of the country have started to charge fees, which imposes a heavy burden on many families. She advises that Government funds should be allocated to help the affected families. One of the British civil servants, visibly confused, replies:

But… is this even possible? Every school in the UK is free.

What the contrast between (18) and (19) suggests is that contextual domain restriction is a relevance-guided mechanism, one that connects what has been uttered to what is being talked about.

Thus, for the resource domain to be identified with the set of American citizens (in a given *world*; see ftn. 23), (15)a should be uttered in a context in which American citizens (and only American citizens) are being talked about. Such a context could be, for example, one in which the Question Under Discussion (or QUD; Roberts 1996/2012) is *where are various kinds of Americans from?*, a QUD that just cares about American citizens.[22] Now, if *C* were to be attached to the head noun, as stipulated in NRT, the resulting NP would be [*fake* [*American philosopher*]], whose denotation includes, for example, German philosophers who pretend to be Americans. The upshot of this is that, in (15)a, *every* would end up quantifying over the set of fake American philosophers—a set that has among its members individuals who are not being talked about.[23]

Let's assume that the following constraint is in place: contextual domain restriction can be performed only if the denotation of the NP (after the restriction has been performed) is either the set of entities that is being talked about or a strict subset of this set. Once this constraint is in place, (non-vacuous) contextual domain restriction, provided that the domain of discourse is rich enough,

---

[22] That is, to provide an answer to such a QUD, it's enough to restrict one's attention to those people who are American citizens: for example, Spanish, Swedish, or French citizens, insofar as they are not also American citizens, can be safely ignored.

[23] A technical note: to model the meaning of a 'fake'-containing DP, an intensional semantics is needed: indeed, 'fake' is an intensional adjective: it takes the intension of its complement as argument (the type of 'fake' is $\langle\langle s,\langle e,t\rangle\rangle,\langle e,t\rangle\rangle$). In the intensional semantics laid out in von Fintel and Heim (2011), the composition of the NRT LF would proceed as follows. $[\![\text{philosopher}_C]\!]^{w,g}$ would be interpreted by intersecting $[\![\text{philosopher}]\!]^{w,g}$ (the extension of 'philosopher' in w) with $[\![C]\!]^{w,g}$ (the extension of C under g in w). Then, $[\![\text{philosopher}_C]\!]^{w,g}$ would be composed with $[\![\text{fake}]\!]^{w,g}$ via *Intensional Functional Application* (i.e. If α is a branching node and $\{\beta, \gamma\}$ the set of its daughters, then, for any world w and assignment g: if $[\![\beta]\!]^{w,g}$ is a function whose domain contains $\lambda w.[\![\gamma]\!]^{w,g}$, then $[\![\alpha]\!]^{w,g} = [\![\beta]\!]^{w,g}(\lambda w.[\![\gamma]\!]^{w,g}))$. Hence, $[\![\text{fake philosopher}_C]\!]^{w,g} = [\![\text{fake}]\!]^{w,g}(\lambda w.[\![\text{philosopher}_C]\!]^{w,g})$.

will never be computed below an intensional adjective such as *fake* (even if the noun came with a *C* variable). The reason for this is straightforward: as illustrated in (20), the set of entities that results from composing the restricted noun with *fake* (via *Intensional Functional Application*; see ftn. 23) will fail to be a subset of the set of entities that is being talked about—provided that the domain of discourse has among its elements at least one fake American who is either a real or fake philosopher (see ftn. 18).

(20)   If $\{x \in \mathcal{D} : x$ is a fake American and either a real or fake philosopher in $w\} \neq \emptyset$, where $\mathcal{D}$ is the set of all possible individuals, then…[24]

$$\underbrace{[\![\text{fake [American philosopher]}]\!]^{w,g}}_{\textit{NP denotation}} \nsubseteq \underbrace{\{x \in \mathcal{D} : x \text{ is American in } w\}}_{\textit{what is being talked about}}$$

I am, of course, stipulating that there is such a constraint. Notice, however, that it is not unreasonable to think that a constraint of this sort exists: if it did not, domain restriction could have the effect of introducing entities that are not being talked about (and are therefore irrelevant) in the discourse, while the function of this mechanism appears to be restricting quantifier domains to those entities that are being talked about.

In sum, the *Fake Philosopher* argument shows that the resource domain can be intersected with the NP denotation. This is important because it indicates that there is something right about QRT. This argument, however, doesn't disprove Weak NRT. First, showing that NP-level restrictions do exist does not rule out the possibility that N-level restrictions also exist. In addition, as argued above, there are reasons to suspect that, if nouns were to come with covert restrictors, considerations about relevance may prevent N-level restrictions from being performed in sentences such as (15)a.

There is an additional consideration: though the *Fake Philosopher* argument poses no challenge to QRT, neither does it pose a challenge to an alternative account (which I do not know whether anyone has ever explicitly put forward): *C* variables are NP sisters. Such an account would predict NP-level restrictions but, unlike QRT, it would not make quantifiers context-sensitive. Thus, on

---

[24] I am here assuming that $\mathcal{D}$ is fixed across worlds.

the basis of the pragmatic behaviour (15)a, it cannot be concluded that quantifiers come with domain variables, as stipulated in QRT: it can only be concluded that there is such a thing as NP-level restrictions.

In the next section, I shall present a novel empirical argument which does two things: on the one hand, it provides clear evidence that domain restriction in DPs is induced, as stipulated in QRT, by the determiner-quantifier; on the other hand, it shows that, if nouns were to come with covert restrictors (as stipulated in NRT), then a host of unattested readings would (incorrectly) be predicted to be possible.

3   Existential bare nouns and domain restriction

If one were able to show that, after removing a determiner, domain restriction ceases to be possible, then one could reject NRT: this account, indeed, predicts domain restriction *to be independent* from there being a determiner (overt or covert) in the structure. With this in mind, I shall test whether domain restriction is possible with plural and mass bare nouns (in non-generic, existential uses[25]): to my knowledge, this has not been done before.[26]

3.1   The argument from bare nouns

Let's consider the following example, which comes from Stanley (2002):

(21)   Context: Pastor Hannah is concerned about the fact that someone has been drinking the holy water in her church on warm summer days. In a discussion with John, John confesses:

I drank a little water last week.

---

[25] The issue of whether bare nouns in *generic* uses support domain restriction remains unsettled. Examples have been reported showing that generic DPs cannot be domain-restricted (e.g. Krifka 1987, Teichman 2015), but there are also cases of (what looks like) generic DPs where domain restriction appears to be at work (e.g. Condoravdi 1994). There is an additional problem: judgments about whether generic DPs can be domain-restricted are confounded by several factors, such as the well-known exception tolerance property of generics as well as the fact that these DPs appear to resist explicit restrictions to a time or place (Carlson 1977, 1982).

[26] I am only aware of some examples (*not* minimal pairs) reported in Arregui (2008), which suggest that indefinites without an overt determiner may fail to take restricted domains in the scope of negation. The minimal pairs presented here reveal that determinerless indefinites do not support domain restriction, irrespective of whether negation is present.

What John expresses, as Stanley (2002: 19) remarks, is the proposition that John drank a little of the church's holy water. Stanley appears to take (21) to indicate that the mass noun, as opposed to the determiner-quantifier, induces domain restriction (it is not clear to me why). The data that I will now discuss provides evidence of the exact opposite.

Let's consider (22), against the context stipulated in (21): if *water* could be restricted to *holy water from the church*, then (22)b should have the reading in (22)b'. (22)b, however, does not have (22)b' as a reading: indeed, (22)b is irrelevant in the context stipulated, whereas (22)b' isn't.

(22)   a. I drank a little water last week.
        a'. ↣ I drank a little holy water from the church last week.
        b. ?? I drank water last week.
        b'. ↣ I drank holy water from the church last week.

On the other hand, (22)a has the reading in (22)a': indeed, in the context stipulated, (22)a and (22)a' can be used interchangeably to the same effect. The contrast between (22)b and (22)a—(22)b is odd, (22)a isn't—shows that *a little*, the determiner-quantifier, is responsible for inducing domain restriction (notice that (22)a works too if *a little* is replaced by *some*).

The pattern in (22) can be reproduced using a bare plural noun like *apples*, as illustrated in (23).

(23)   Context: Pastor Hannah is concerned about the fact that someone has been eating the apples from the churchyard's apple tree on warm summer days. In a discussion with John, John confesses:

        a. I ate some apples last week. (Please forgive me.)
        a'. ↣ I ate some apples from the churchyard's apple tree last week. (Please forgive me.)
        b. ?? I ate apples last week. (Please forgive me.)
        b'. ↣ I ate apples from the churchyard's apple tree last week. (Please forgive me.)

If *apples* could be restricted to *apples from the churchyard's apple tree*, then (23)b should have the reading in (23)b'. (23)b, however, does not have (23)b' as a reading: if it did, uttering (23)b should be as acceptable as uttering (23)b' and, quite clearly, it is not. Conversely, (23)a does have the reading in (23)a': John could confess his misdeed by uttering either (23)a or (23)a'. As in the

previous example, the contrast between (23)a and (23)b shows that *some*, the determiner-quantifier, is responsible for inducing domain restriction.

More or less the same results are obtained if the target sentence is negated, as illustrated in (24).

(24)   Context: Pastor Hannah is concerned about the fact that someone has been eating the apples from the churchyard's apple tree on warm summer days. In a discussion with John, John defends himself by saying:

       a. I didn't eat any apples last week.
       a'. ↣ I didn't eat any apples from the churchyard's apple tree last week.
       b. ? I didn't eat apples last week.
       b'. ↛ I didn't eat apples from the churchyard's apple tree last week.

The contrasts in (24) are a bit less sharp than the ones reported in (23), the reason for this being two-fold. First, *any* is typically associated with domain-widening effects (Kadmon and Landman 1993); as a result, its presence is expected to interfere with the mechanism of domain restriction. In addition, (24)b is not pragmatically infelicitous—or, at least, not as infelicitous as (23)b. Indeed, on the (reasonable) assumption that the (implicit) question that John is addressing is *have you eaten apples from the churchyard's apple tree last week?*, then (24)b entails a negative answer to this question (i.e. *I didn't eat apples last week* entails *I didn't eat apples from the churchyard's apple tree last week*), whereas (23)b is intuitively irrelevant (that is, it does not settle the question). However, there's a still a clear contrast between (24)a and (24)b: John could truly utter (24)a even if he had eaten some apples from the market on the week in question; an utterance of (24)b, by contrast, appears to commit John to the stronger proposition that he didn't eat apples *tout court*, neither from the churchyard apple tree nor from anywhere else.

Given these data, it should be clear that NRT cannot be right; if such a mechanism existed, then (22)b, (23)b and (24)b should have the readings in (22)b', (23)b' and (24)b', respectively, and they clearly don't. As stipulated in von Fintel (1994), it must be the determiner-quantifier, as opposed to the head noun, that induces domain restriction in DPs.[27]

---

[27] The issue of how the existential force of (22)b, (23)b and (24)b should be derived is a highly contentious one—see Carlson (1999) for a survey of different approaches. If D is taken to be occupied by a silent existential determiner then, in the light of the data reported here, that determiner would have to be stipulated not to introduce a domain variable. There are, however, a host of

3.2 Challenges to the bare noun data

I will here discuss three data points, which were brought to my attention by anonymous reviewers. These data, at least on first inspection, appear to contradict the generalisation made in the previous section, i.e.: bare nouns in existential uses do not support domain restriction.

3.2.1 Potential counterexample #1

To begin with, consider (25).

(25) a. At the embassy reception, Alice ate vegan snacks.
      b. At the embassy reception, Alice ate vegan snacks at the embassy reception.

It could be claimed that *snacks* in (25)a is restricted to *snacks at the embassy reception*. Such a claim entails that (25)a has the reading in (25)b.

If one were to hear (25)a, one would infer that the snacks that Alice ate were at the embassy reception. This is not controversial; what is controversial, however, is the claim that (25)a has the *reading* in (25)b. It is a fact about the world (and it is common ground) that, in order for x to eat y, x has to be in the same place as y; as a result of this, the *contextual* meanings of (25)a and (25)b are bound to be identical (irrespective of whether *snacks* in (25)a is restricted to *snacks at the embassy reception*): indeed, $[\![(25)a]\!] \cap C = [\![(25)b]\!] \cap C$, where C stands for the Stalnakerian Context Set (the set of worlds compatible with all the propositions in the common ground). Since the interpretation of (25)a can be accounted for without the need of positing domain restriction, it surely cannot be taken to be evidence for it.

3.2.2 Potential counterexample #2

Here's another example:

---

proposals that derive the existential reading of bare nouns without positing something like a silent *some*; for instance, existential closure (e.g. Diesing 1992), type-shifting (e.g. Chierchia 1998), quantification over stages (Carlson 1977), or default interpretation of an empty D (Longobardi 1994).

(26)   Context: The US army is taking part in a nighttime military battle. Two US army units, units 202 and 203, have established a signaling system: if unit 202 flashes green lights, unit 203 is to advance; if, instead, 202 flashes red lights, 203 is to retreat. Officer cadet Smith, from unit 203, is in charge of scanning the battlefield with his binoculars looking for unit 202's signals (he does not know unit 202's precise location). He suddenly exclaims:

I see red lights!

In the context stipulated, *I see red lights!* appears to mean something like *I see red lights… and these lights are (probably) being flashed from unit 202*. The question is the following: is this an instance of nominal restriction (of *lights* being restricted at the compositional stage to *lights being flashed from unit 202*) or, rather, an instance of an implicature, with (26) implicating that Smith believes that the red lights that he sees are being flashed from unit 202? To address this question, the dissent test—a test which relies on the propositional anaphor *that*—can be applied (Tonhauser 2012).

(27)   Context: The one in (26).

Lieutenant Johns:	a. Any signal?
Cadet officer Smith:	b₁. **I see red lights!**
Cadet officer Yoo:	c. # *That* can't be true: I'm also watching and can only see red lights being flashed from enemy airplanes.

Lieutenant Johns:	a. Any signal?
Cadet officer Smith:	b₂. **I see red lights being flashed from unit 202!**
Cadet officer Yoo:	c. *That* can't be true: I'm also watching and can only see red lights being flashed from enemy airplanes.

Lieutenant Johns:	a. Any signal?
Cadet officer Smith:	b₃. **I see red lights (being flashed from somewhere)!**
Cadet officer Yoo:	c. # *That* can't be true: I'm also watching and can only see red lights being flashed from enemy airplanes.

If (27)b₁ were to entail that the lights that speaker sees are being flashed from unit 202 (due *to* lights being restricted to *lights being flashed from unit 202*), then the pair (27)b₁-c would be as felicitous as (27)b₂-c. It clearly isn't, however; in fact, (27)b₁-c feels as deviant as (27)b₃-c.

The above can be contrasted with an actual case of domain restriction, in which contextual information does get into the truth-conditions of the sentence. Let's apply the dissent test to (23), for example.

(28)   Context: Pastor Hannah is concerned about the fact that someone has been eating the apples from the churchyard's apple tree on warm summer days. In a discussion with John, John confesses…

John:	a₁. OK. I'll confess… **I ate some apples last week.**
John's sister:	b. But *that*'s not true John! Last week you only ate apples from the market. You didn't even come to the church!
John:	a₂. OK. I'll confess… **I ate some apples from the churchyard's apple tree last week.**
John's sister:	b. But *that*'s not true John! Last week you only ate apples from the market. You didn't even come to the church!
John:	a₃. OK. I'll confess… **I ate some apples (from somewhere) last week.**
John's sister:	b. # But *that*'s not true John! Last week you only ate apples from the market. You didn't even come to the church!

The data, I think, are clear: (28)b can be used to object to both (28)a₁ and (28)a₂, but it cannot be used to object to (28)a₃. This indicates that (28)a₁ must have the reading in (28)a₂: that is, in (28)a₁, *apples* must be implicitly strengthened to *apples from the churchyard's apple tree* (exactly what one expects if domain restriction is in operation).

Though (28)b works well as an objection to (28)a₁, I am willing to concede that it works even better if used to object to (28)a₂. But there is nothing surprising about this: (28)a₁ entails that John ate some apples from the churchyard's apple tree, but only if contextual domain restriction is computed; (28)a₂, by contrast, entails that John ate some apples from the churchyard's apple tree irrespective of context. Hence, (28)a₁, unlike (28)a₂, is ambiguous between an unrestricted reading (*I ate some apples from somewhere last week*) and a restricted one (*I ate some apples from the churchyard's apple tree last week*). Though the latter reading is the most salient (for obvious

pragmatic reasons), the former, which lurks behind the scenes, can occasionally blur the judgment.[28]

It is worth stressing that, insofar as (28)a_1's restricted reading is accessed, both (28)a_1-b and (28)a_2-b are acceptable: it is just (28)a_3-b that is deviant. In (27), the situation is different: neither (27)b_1-c nor (27)b_3-c work (both are deviant), the only felicitous exchange being (27)b_2-c. On the basis of these observations, it can be concluded that, in (26), *lights* is not restricted to *lights being flashed from unit 202* (if it were, then (27)b_1-c would be an acceptable exchange, just like (28)a_1-b is): rather, (26) triggers a pragmatic inference (roughly paraphrasable as 'Smith believes that the red lights that he sees are being flashed from unit 202') which creates the illusion that domain restriction is in operation.

An anonymous reviewer raises an interesting question: if (26) gives rise to an implicature (a pragmatic inference) that creates the illusion that the domain of quantification is being restricted, then, why is it that (23)b, for example, does not create the same illusion? In what follows, I will first argue that (26)'s truth-conditional content, in the stipulated context, constitutes relevant information, whereas (23)b's truth-conditional content doesn't; subsequently, I will show two things: first, that (26) gives rise to an implicature that creates the illusion of domain restriction *only if* (26)'s truth-conditional content is relevant and, second, that (23)b does behave just like (26)—it *also* creates the illusion of domain restriction if, in the context in which it is uttered, its truth-

---

[28] By 'lurks behind the scenes', I just mean that the probability of (28)a_1's unrestricted reading being the reading intended by the speaker is very low—given pragmatic considerations—but, quite possibly, not 0. One can in fact increase the pragmatic pressure against the unrestricted reading and, by so doing, improve the quality of the judgment. Consider (iv) below, for example. Here, (iv)a_1's unrestricted reading is irrelevant, just like (28)a_1's, but, in addition, is (contextually) tautological; hence, in (iv)a_1, the probability of the unrestricted reading being the intended reading is bound to be even lower than in (28)a_1. Unsurprisingly, (iv)a_1-b and (iv)a_2-b exhibit no contrast whatsoever (in terms of acceptability).

(iv)   Context: Every apple from the churchyard's apple tree has been eaten. Pastor Hannah allows John to eat one or two apples a day, but he isn't allowed to eat them all at once. John's sister tells Pastor Hannah that John broke the rules and, in a fit of hunger, ate every single apple from the churchyard's apple tree. John contests this accusation by exclaiming:

John: a_1. I didn't eat every apple!
John's sister: b. That's not true John! You did eat every apple from the churchyard's apple tree.

John: a_2. I didn't eat every apple from the churchyard's apple tree!
John's sister: b. That's not true John! You did eat every apple from the churchyard's apple tree.

John: a_3. I didn't eat every apple (in the world)!
John's sister: b. # That's not true John! You did eat every apple from the churchyard's apple tree.

conditional content constitutes relevant information. To do this, a brief excursus into the notion of relevance will be necessary.

3.2.2.1 Relevance: a brief excursus

In 'classical' formal pragmatics, the notion of relevance is QUD-dependent; a QUD (or Question Under Discussion) is a partition—typically, of the Stalnakerian Context Set (Roberts 1996/2012). In a nutshell, a proposition is relevant iff it is a partial answer: that is, iff it eliminates some (and only some) cells of the QUD that it addresses.[29] This approach to modelling relevance, which has been popularised by Roberts (1996/2012), originates in Groenendijk and Stokhof's (1984) seminal work on questions and answers.

The classical notion of relevance has be questioned (including by Roberts herself) for being rather restrictive; Simons, Tonhauser, Beaver, and Roberts (2010: 316-7), for example, write:

> [Roberts's (1996/2012)] definition of Relevance is overly restrictive and should be weakened at least to allow for discourse moves which merely raise or lower the probability of some answer to the QUD being correct. Consider for example the sequence: Q: "Is it going to rain?" A: "It's cloudy." A's utterance does not contextually entail an answer to the QUD (at least not in Pittsburgh, PA). Intuitively, it is relevant because it somewhat raises the probability of an affirmative answer to the QUD.

The idea of characterising relevance (relative to a QUD) in probabilistic terms can be traced back to at least Büring (2003); the basic idea is the following: given a proposition p, a QUD Q, and a probability measure P that assigns a probability between 0 and 1 to each subset of Logical Space, p is relevant if updating P with p shifts the probabilistic weights among Q's cells. In this framework, P is the common ground and QUDs are defined as partitions of Logical Space.

This probabilistic account of relevance, however, does not quite explain the contrast between (29)b and (30)b: while the former feels relevant, the latter hardly does.

(29) a. Is Mary coming to the party?
      b. She's tired.

---

[29] More formally, p is a partial answer to Q (where p is a proposition and Q a partition of C, the Context Set) iff $p \cap C \neq \emptyset$ and, furthermore, there is an $X \subset Q$ such that $p \cap C \subseteq \bigcup X$ (such that p contextually entails $\bigcup X$) (cf. Groenendijk and Stokhof 1990).

(30)   a. Is Mary coming to the party?
         b. She isn't tired.

But why? (30)b, just like (29)b, does shift the probabilistic weights among the QUD cells: indeed, given a reasonable P, one would expect P(YES-cell │ ⟦(30)b⟧) to be greater than P(YES-cell): after all, not being tired, is a reason to go (as opposed to not go) to a party. However, (30)b, though probabilistically relevant, hardly feels relevant. It seems that, in order for relevance to kick in, updating P with *p* should *significantly* shift the probabilistic weights among the QUD cells. (I won't discuss the technical details of probabilistic relevance; for an explicit threshold-sensitive measure of relevance based on Kullback–Leibler divergence, see Feinmann (2020)). With this extra assumption on board, the contrast between (29)b and (30)b can be accounted for: given a reasonable P, P(YES-cell │ ⟦(30)b⟧) is bound to be greater than P(YES-cell), but only marginally greater: there are far too many reasons for not going to a party; indeed, if one learns that Mary isn't tired, one doesn't learn much (intuitively): only perhaps that, if she ends up not coming, it won't be because she is tired. By contrast, and given a reasonable P, P(YES-cell │ ⟦(29)b⟧) is bound to be a lot smaller than P(YES-cell): being tired, after all, is the kind of thing that would make someone refrain from going to a party.

3.2.2.2   The illusion of domain restriction

Let's now come back to (26), repeated as below (31):

(31)   <u>Context:</u> The US army is taking part in a nighttime military battle. Two US army units, units 202 and 203, have established a signaling system: if unit 202 flashes green lights, unit 203 is to advance; if, instead, 202 flashes red lights, 203 is to retreat. Officer cadet Smith, from unit 203, is in charge of scanning the battlefield with his binoculars looking for unit 202's signals (he does not know unit 202's precise location). He suddenly exclaims:

         I see red lights!
         ⤳ Smith believes that the lights that he sees are being flashed from unit 202.

The first thing to notice is that for a signaling system of the sort envisaged in (31) to be successful, it has to be the case that, in the battlefield in question, red and green lights are a rare sight: if red

and green lights were constantly being flashed from everywhere, then the signaling system established between units 202 and 203 would be hopeless. Thus, on the assumption that the soldiers from units 202 and 203 established a signaling system which they themselves think has at least some chances to succeed, it must be common ground—prior to (31) being uttered—that, if green or red lights are flashed from somewhere in the battlefield, the probability of these lights being flashed from unit 202 is much greater than the probability of these lights being flashed from somewhere else. Relative to such a common ground and the 2-cell QUD *any signal from unit 202?* (the most salient QUD given the context stipulated), (31) is expected to be relevant: this is because P(YES-cell │ ⟦(31)⟧), the conditional probability of the YES-cell given ⟦(31)⟧, will typically be a lot greater than P(YES-cell).

If, however, the common ground was such that the sentence *I see red lights*, relative to the QUD *any signal from unit 202?*, failed to secure relevance, then the said sentence would not be interpreted as in (31): that is, it would not give rise to the implicature that Smith believes that the lights in question are being flashed from unit 202 (and, consequently, the illusion of domain restriction would not be generated). To illustrate this, consider (32) below:

(32)   Context: The US army is taking part in a nighttime military battle. Two US army units, units 202 and 203, have established a signaling system: if unit 202 flashes green lights, unit 203 is to advance; if, instead, 202 flashes red lights, 203 is to retreat. Just before the battle starts, unit 203 members are notified that units 204-5, 206-7 and 208-9 are also using the green/red light signaling system (i.e. if units 204/206/208 flash green lights, units 205/207/209 are to advance, respectively; if units 204/206/208 flash red lights, units 205/207/209 are to retreat, respectively). This is unfortunate: in the event of green/red lights, there is no way for unit 203 to know whether such lights are being flashed from unit 202 or, alternatively, from units 204, 206, or 208. Officer cadet Smith, from unit 203, who is very much aware of this issue (like everyone else in unit 203) suddenly exclaims:

    a. I see red lights!
    ↝ Smith believes that the lights that he sees are being flashed from unit 202.
    b. I see red lights being flashed from unit 202!

(32)a, quite clearly, is not relevant relative to the QUD *any signal from unit 202?* (nor does it implicate that Smith believes that the lights that he sees are being flashed from unit 202); in the context stipulated, as far as I can tell, (32)a means *I see red lights simpliciter*. (32)b, by contrast,

is relevant, though it raises the question of how come Smith was able to establish, in the middle of the night, that the red lights that he sees are being flashed from unit 202 and not from one of the other units.

Let's now consider (23)b, for example, repeated as (33)b below.

(33)   Context: Pastor Hannah is concerned about the fact that someone has been eating the apples from the churchyard's apple tree on warm summer days. In a discussion with John, John confesses:

a. I ate some apples last week. (Please forgive me.)
a'. ↦ I ate some apples from the churchyard's apple tree last week. (Please forgive me.)
b. ?? I ate apples last week. (Please forgive me.)
b'. ↦↦ I ate apples from the churchyard's apple tree last week. (Please forgive me.)

(33)b does not have the reading in (33)b' and, unlike (31), is pragmatically deviant. There is a straightforward reason for this, I think: (33)b, unlike (31), is *irrelevant*. Indeed, given the 2-cell QUD *have you eaten any of the apples from the churchyard's apple tree?* and a reasonable common ground, learning that (33)b is true (that is, learning that John ate some apples from somewhere) only marginally increases the probability of the QUD's YES-cell: there are so many apples that John could have eaten that, learning that he ate some, doesn't do much in terms of helping us to resolve the issue at hand—namely, whether John ate apples from the churchyard apple tree.

If the common ground in (33) were to be tweaked so that (33)b became relevant, then, an utterance of (33)b, just like an utterance of (31), would create the illusion of domain restriction. Consider, for example, (34):

(34)   Context: Pastor Hannah's church is located in a small village of which three things are true: (i) there are no shops; (ii) its dwellers eat only local produce; (iii) there are only two apples trees, one in the churchyard and the other right next to Sheriff's station. Pastor Hannah is concerned about the fact that someone has been eating the apples from the churchyard's apple tree without paying for the fruit. He suspects that John may be responsible for the misdeed: John is the sort of person who would be scared of stealing apples from the Sheriff's station (the village sheriff is known to be violent and unforgiving) but not from the churchyard (John knows Pastor Hannah is too nice to punish him). In a discussion with John, John confesses:

a. I ate apples last week. (Please forgive me.)
b. I ate some apples last week. (Please forgive me.)

(34)a, according to my judgments, is a lot better than (33)b; this isn't surprising: *I ate apples last week*, given the context stipulated in (34), and given the 2-cell QUD *have you eaten any of the apples from the churchyard's apple tree?*, is relevant: technically, it (significantly) raises the probability of the QUD's YES-cell. There is of course a contrast between (34)a and (34)b: whereas (34)b can entail (via domain restriction) that the apples that John ate come from the churchyard's apple tree, (34)a cannot (because (34)a does not support domain restriction). This contrast, however, is much milder than the one between (33)a and (33)b, the reason for this being self-evident: in (34), one is contrasting two sentences, neither of which constitutes a relevance violation; in (33), one is contrasting two sentences, one of which, namely (33)b, constitutes a relevance violation.

To sum up, (31) and (34)a are both relevant; furthermore, while (31) implicates that Smith believes that the red lights that he sees are being flashed from unit 202, (34)a implicates that John believes that the apples that he ate come from the churchyard apple tree. The question of how these implicatures are derived is beyond the scope of this paper; however, it is clear that, in order for these implicatures to get off the ground, relevance must be satisfied in the first place, as exposed by the contrasts between (31) and (32)a, and between (33)b and (34)a.[30]

3.2.3   Potential counterexample #3

To conclude, I will look at yet another example; in (35)a, at least on first inspection, *glasses* appears to be restricted to *glasses from vitrine #55*.

---

[30] Without going into much detail, it is not hard to see where these implicatures are coming from. Take (31), for example; when officer cadet Smith utters *I see red lights*, he knows that his utterance will have the effect of (massively) increasing the probability mass of the QUD's YES-cell; if Smith didn't believe that, in fact, the red lights that he sees are being flashed from unit 202, then, by uttering *I see red lights*, he would be misleading his audience (i.e. though *I see red lights* would be true, it would have the effect of making Smith's audience believe something that he does not himself believe). On the assumption that speakers are cooperative (in the broadly Gricean sense) and, among other things, do not mislead their audiences, it must be the case that Smith believes (or at least thinks more likely than not) that the red lights that he sees are being flashed from unit 202. See Feinmann (2020) for a discussion on these inferences.

(35)   Context: Paula works as a cleaner in The Drinking Glass Museum. The establishment has five floors packed with drinking glasses from the Tudor period onwards. Paula, vividly alarmed, tells Carol, the museum manager…[31]

> *Paula*:   Some stemware is missing from vitrine #55!
> *Carol*:   a. (I know.) I borrowed Victorian sherry glasses for a small exhibit that is taking place on the 3rd floor.
> b. ↣ I borrowed Victorian sherry glasses from vitrine #55 for a small exhibit that is taking place on the 3rd floor.

Is then (35)a a counterexample to the generalisation made in the previous section? I will show that it isn't. Note that the verb *borrow* is a three-place predicate, i.e.: [[borrow]]$(x)(y)(z) = 1$ iff x borrows y from z; although in (35)a the latter of these arguments has no phonological realisation, that does not mean that the argument isn't there. Thus, there are two competing explanations for this data point: either (35)a has the reading in (35)b because *glasses* introduces a covert restrictor (as stipulated in NRT) or, alternatively, it has the reading that it has because *borrow* takes an implicit argument (in this case, *from vitrine #55*).

In (36), in order to shed light on this issue, I have replaced *borrow* for (transitive) *break*.

(36)   Context: The same as in (35).

> *Paula*:   Some stemware is missing from vitrine #55!
> *Carol*:   a. ? (I know.) I broke Victorian sherry glasses last week. My fault.
> b. ↣ I broke Victorian sherry glasses from vitrine #55 last week.

If *glasses* could be restricted to *glasses from vitrine #55*, as predicted by NRT, and if this was indeed what is going on in (35)a, then it would be a complete mystery why (36)a does not have the reading in (36)b. It can thus be concluded that, in (35)a, it is *borrow*, and not *glasses*, that induces contextual enrichment.[32] It should also be noted that, if an overt determiner-quantifier is incorporated, as in (37) below, domain restriction becomes possible, exactly as predicted by QRT.

---

[31] I thank an anonymous SuB24 reviewer for this example.
[32] See Williams (2015: ch. 5) for a recent discussion on verbs and implicit arguments.

(37)   Context: The same as in (35).

Paula: Some stemware is missing from vitrine #55!
Carol: a. (I know.) I broke **some** Victorian sherry glasses last week. My fault.
             a'. ↣ I broke some Victorian sherry glasses from vitrine #55 last week.
             b. (I know.) I broke **a few** Victorian sherry glasses last week.
             b'. ↣ I broke a few Victorian sherry glasses from vitrine #55 last week.
             c. (I know.) I broke **three** Victorian sherry glasses last week.
             c'. ↣ I broke three Victorian sherry glasses from vitrine #55 last week.

These data, like the data discussed in the previous section, indicate that NRT cannot be right: if nouns came with covert restrictors, then (36)a would have the reading in (36)b, and it clearly does not. In turn, the fact that (37)a-b-c have the readings in (37)a'-b'-c' (respectively) provides support for QRT.

3.3   A note on global restrictions

To conclude, I would like to note that the bare noun data discussed in § 3.1 is not just evidence against NRT; it is also evidence against a wholly different approach to domain restriction, which Westerståhl (1984) calls the *flexible universe strategy* and I will simply refer to as the *global approach*.[33] Unlike in QRT or NRT, on this account, restrictions are not locally intertwined with the recursive semantics: instead, they are accomplished by restricting the domain of evaluation for the whole sentence.[34, 35]

Restricting the domain (or universe) of discourse can handle simple cases—such as (38)a below; however, as noted in Westerståhl (1984) and later in Soames (1986), it has no resources to deal with cases such as (39)a, in which two different DP in the same sentence are evaluated relative to different (contextually provided) sets of entities.

---

[33] This approach is also known as the *implicit strategy* (Neale 1990) and the *model theoretic approach* (Stanley and Szabó 2000).
[34] As von Fintel (1994: 29, ftn. 15) remarks, this is an approach with a venerable pedigree, which includes luminaries such as John Wallis, August de Morgan, and George Boole.
[35] It is worth noting that, on the global approach, restrictions, although they are not locally intertwined with the recursive semantics, are nonetheless predicted to affect truth-conditional content. Thus, examples such as (26) or (34)a cannot be regarded as involving global domain restriction.

(38)   Context: The UK Government has decided to reduce the budget for education.

        a. Every school is going to suffer.
        b. ↦ Every *UK state school* is going to suffer.

(39)   Context: The MIT Linguistics & Philosophy department consists of linguists and philosophers, naturally. On one particular year, three academics are shortlisted for a teaching position, a linguist and two philosophers, but the department is only allowed to recommend one of them.[36]

        a. Every linguist voted for the linguist.
        b. ↦ Every linguist *from the MIT Linguistics & Philosophy department* voted for the linguist *from the pool of job candidates.*

Indeed, in (39)a, *every linguist* and *the linguist* are clearly not being evaluated relative to the same (restricted) universe of discourse; if they were, the uniqueness presupposition triggered by the definite article would not be satisfied. To derive the correct truth-conditions for (39)a in the stipulated context, each quantified DP needs to be able to access its own resource domain (as in QRT or NRT).[37]

Examples such as (39)a, it is worth noting, are not evidence against the existence of global restrictions: rather, this sort of examples indicate that, if there were to be something like a global restriction mechanism, then such a mechanism could not account for all the quantifier domain restriction data; as Kratzer (2007) remarks, domain restriction effects could, at least *prima facie*, be the byproducts of a number of mechanisms, including global restrictions.

Now, unlike (39)a and other examples in the same vein, the data I that have presented in § 3.1 does call into question the idea that a global domain restriction mechanism exists. Let's take a look at (23) again, repeated as (40) below: if the universe of discourse could be restricted to exclude those apples that aren't relevant in the context (most apples), then it should be possible for (40)b to have

---

[36] Adapted from Cooper (1996).
[37] Soames's (1986) classic example (i.e. *everyone is asleep and being monitored by a research assistant*) makes essentially the same point as (39); see also Lewis (1973) and McCawley (1979).

the reading in (40)b'. But this isn't the case: whereas (40)a has the reading in (40)a', (40)b doesn't have (40)b' as a reading.

(40)   Context: Pastor Hannah is concerned about the fact that someone has been eating the apples from the churchyard's apple tree on warm summer days. In a discussion with John, John confesses:

a. I ate some apples last week. (Please forgive me.)
a'. ↝ I ate some apples from the churchyard's apple tree last week. (Please forgive me.)
b. ?? I ate apples last week. (Please forgive me.)
b'. ↝̸ I ate apples from the churchyard's apple tree last week. (Please forgive me.)

Contrasts such as this, it seems to me, put pressure on 'hybrid' accounts of domain restriction, that is, accounts which allow each determiner-quantifier to access its own resource domain but, *in addition*, permit global restrictions, such as Kratzer (2007) and Schwarz (2009, 2012).[38] Further investigations would be needed but, at least on the basis of the data reported in § 3.1, it appears to be the case that quantifier domain restriction is, exclusively, a local mechanism.

4   CONCLUSION

In this article, I have considered Stanley's (2002; Stanley and Szabó 2000) arguments in favour of NRT, as well as the *Fake Philosopher* argument, an argument that has been given against NRT. Stanley's arguments have been shown to be not conclusive; the *Fake Philosopher* argument, in turn, has been called into question on the basis that it is does not rule out Weak NRT. In addition, I have presented a novel empirical argument based on the pragmatic behaviour of bare nouns, an argument that constitutes evidence against Weak NRT and in support of QRT.

---

[38] The proposals in Kratzer (2007) and Schwarz (2009, 2012) are couched in the framework of situation semantics (Barwise and Perry 1983); in this framework, global restrictions come about due to the fact that, to paraphrase Kratzer (2007), 'assertions are about particular actual situations, and that those situations can be smaller or bigger parts of the actual world'. For the point that I'm making here, it does not matter whether global restrictions are cashed out in terms of relativisation to a (contextually restricted) universe of discourse or in terms of relativisation to a (contextually provided) situation: either approach predicts that (40)b should have a reading that, as a matter of fact, doesn't have.

REFERENCES

Arregui, Ana. 2008. "Some Remarks on Domain Widening." In *Proceedings of the 27th West Coast Conference on Formal Linguistics*, 45–53.

Bartsch, Renate, and Theo Vennemann. 1972. "Relative Adjectives and Comparison." *UCLA Papers in Syntax* 2: 80–107.

Barwise, Jon, and John Perry. 1983. *Situations and Attitudes*. MIT Press.

Breheny, Richard. 2003. "A Lexical Account of Implicit (Bound) Contextual Dependence." In *Semantics and Linguistic Theory*, 13:55–72.

Büring, Daniel. 2003. "On D-Trees, Beans, And B-Accents." *Linguistics and Philosophy* 26 (5): 511–45. https://doi.org/10.1023/A:1025887707652.

Carlson, Greg. 1999. "No Lack of Determination." *Glot International*, 3–8.

Carlson, Greg N. 1977. "Reference to Kinds in English." PhD Thesis, University of Massachusetts, Amherst.

———. 1982. "Generic Terms and Generic Sentences." *Journal of Philosophical Logic* 11 (2): 145–81.

Chierchia, Gennaro. 1998. "Reference to Kinds across Language." *Natural Language Semantics* 6 (4): 339–405. https://doi.org/10.1023/A:1008324218506.

Condoravdi, Cleo A. 1994. "Descriptions in Context." PhD Thesis, Yale University.

Cooper, Robin. 1996. "The Role of Situations in Generalized Quantifiers." In *The Handbook of Contemporary Semantic Theory*, edited by Shalom Lappin, 65–86. Oxford: Blackwell.

Cresswell, Max J. 1976. "The Semantics of Degree." In *Montague Grammar*, edited by Barbara Partee, 261–92. New York: Academic Press.

Diesing, Molly. 1992. *Indefinites*. MIT press.

Farkas, Donka F., and Katalin E. Kiss. 2000. "On the Comparative and Absolute Readings of Superlatives." *Natural Language & Linguistic Theory* 18 (3): 417–55.

Feinmann, Diego. 2020. "Lexical Illusions, (Non-)Maximality, and Invisible Gaps." PhD Thesis (revised version), University of Sheffield.

Fintel, Kai von. 1994. "Restrictions on Quantifier Domains." PhD Thesis, University of Massachusetts at Amherst.

———. 2014. "Quantifier Domain Restriction." Presented at the Semantics Workshop, Rutgers University.

Gillon, Carrie. 2006. "The Semantics of Determiners: Domain Restriction in Skwxwú7mesh." PhD Thesis, University of British Columbia.

Groenendijk, Jeroen Antonius Gerardus, and Martin Johan Bastiaan Stokhof. 1984. "Studies on the Semantics of Questions and the Pragmatics of Answers." PhD Thesis, University of Amsterdam.

———. 1990. "Partitioning Logical Space." Annotated handout.

Heim, Irene. 1999. "Notes on Superlatives." Manuscript. MIT.

Heim, Irene, and Angelika Kratzer. 1998. *Semantics in Generative Grammar*. Oxford: Blackwell.

Kadmon, Nirit, and Fred Landman. 1993. "Any." *Linguistics and Philosophy* 16 (4): 353–422.

King, Jeffrey C., and Karen S. Lewis. 2018. "Anaphora." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2018. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/fall2018/entries/anaphora/.

Klein, Ewan. 1980. "A Semantics for Positive and Comparative Adjectives." *Linguistics and Philosophy* 4 (1): 1–45.

Kratzer, Angelika. 2004. "Covert Quantifier Restrictions in Natural Languages." Palazzo Feltrenelli in Gargnano, Italy.

Krifka, Manfred. 1987. "An Outline of Genericity." Presented at the SNS-Bericht 87-25, Seminar fur Natiirlich-Sprachliche Systeme, Universität Tübingen.

Larson, Richard K., and Gabriel Segal. 1995. *Knowledge of Meaning: An Introduction to Semantic Theory*. Bradford Books.

Lasersohn, Peter. 2020. "Common Nouns as Modally Non-Rigid Restricted Variables." *Linguistics and Philosophy*, March. https://doi.org/10.1007/s10988-019-09293-4.

Lewis, David. 1973. *Counterfactuals*. Oxford: Blackwell.

Link, Godehard. 1983. "The Logical Analysis of Plurals and Mass Terms: A Lattice-Theoretic Approach." In *Meaning, Use and Interpretation of Language*, edited by Arnim von Stechow, Christoph Schwarze, and Rainer Bäuerle, 302–23. Berlin: De Gruyter.

Longobardi, Giuseppe. 1994. "Reference and Proper Names: A Theory of N-Movement in Syntax and Logical Form." *Linguistic Inquiry* 25 (4): 609–65.

Martí, Luisa. 2003. "Contextual Variables." PhD Thesis, University of Connecticut.

McCawley, James D. 1979. "Presupposition and Discourse Structure." *Syntax and Semantics* 11: 371–88.

Neale, Stephen. 1990. *Descriptions*. MIT Press.

Nouwen, Rick. 2020. "E-Type Pronouns: Congressmen, Sheep and Paychecks." In *The Wiley Blackwell Companion to Semantics*, edited by Lisa Matthewson, Cécile Meier, Hotze Rullman, Ede Zimmermann, and Daniel Gutzmann. Oxford: Wiley-Blackwell.

Roberts, Craige. (1996) 2012. "Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics." *Semantics and Pragmatics* 5. https://doi.org/10.3765/sp.5.6.

Schwarz, Florian. 2009. "Two Types of Definites in Natural Language." PhD Thesis, University of Massachusetts at Amherst.

———. 2012. "Situation Pronouns in Determiner Phrases." *Natural Language Semantics* 20 (4): 431–75. https://doi.org/10.1007/s11050-012-9086-1.

Seuren, Pieter A. M. 1973. "The Comparative." In *Generative Grammar in Europe*, edited by F. Kiefer and N. Ruwet, 528–64. Dordrecht: Reidel.

Sharvit, Yael, and Penka Stateva. 2002. "Superlative Expressions, Context, and Focus." *Linguistics and Philosophy* 25 (4): 453–504.

Simons, Mandy, Judith Tonhauser, David Beaver, and Craige Roberts. 2010. "What Projects and Why." *Semantics and Linguistic Theory* 20 (0): 309–27. https://doi.org/10.3765/salt.v20i0.2584.

Soames, Scott. 1986. "Incomplete Definite Descriptions." *Notre Dame Journal of Formal Logic* 27 (3): 349–75.

Stanley, Jason. 2002. "Nominal Restriction." In *Logical Form and Language*, edited by Preyer Gerhard and Georg Peter, 365–388. Oxford: Clarendon Press.

Stanley, Jason, and Zoltán Gendler Szabó. 2000. "On Quantifier Domain Restriction." *Mind & Language* 15 (2–3): 219–61.

Teichman, Matt. 2015. "Characterizing Kinds." PhD Thesis, University of Chicago.

Tonhauser, Judith. 2012. "Diagnosing (Not-) at-Issue Content." In *Proceedings of Semantics of Under-Represented Languages of the Americas (SULA)*, 6:239–54.

Westerståhl, Dag. 1984. "Determiners and Context Sets." In *Generalized Quantifiers in Natural Language*, Johan van Benthem and Alice ter Meulen, 46–71. Berlin: De Gruyter.

Williams, Alexander. 2015. *Arguments in Syntax and Semantics*. Cambridge University Press.