Complex demonstratives, hidden arguments, and presupposition*

Ethan Nowak

May 12, 2019

Abstract

Standard semantic theories predict that non-deictic readings for complex demonstratives should be much more widely available than they in fact are. If such readings are the result of a lexical ambiguity, as Kaplan (1977) and others suggest, we should expect them to be available wherever a definite description can be used. The same prediction follows from 'hidden argument' theories like the ones described by King (2001) and Elbourne (2005). Wolter (2006), however, has shown that complex demonstratives admit non-deictic interpretations only when a precise set of structural constrains are met. In this paper, I argue that Wolter's results, properly understood, upend the philosophical status quo. They fatally undermine the ambiguity theory and demand a fundamental rethinking of the hidden argument approach.

1 Introduction

Philosophers' interest in complex demonstratives has traditionally been focused on deictic instances; i.e., instances on which a demonstrative is used to pick out an object from the context of utterance, as in:

(1) **That river** looks treacherous (while pointing at a river).

Most philosophical work on the topic has revolved around the question of how to represent the semantic contribution of the pointing gesture (or directed gaze,

*Forthcoming in *Synthese*. Thanks to Peter Epstein, Peter Jenks, Nathan Klinedinst, Petr Kusliy, Brian Leahy, John MacFarlane, Eliot Michaelson, Line Mikkelsen, Barbara Partee, Michael Rieppel, Daniel Rothschild, Ekaterina Vostrikova, Seth Yalcin, and two anonymous referees for helpful comments and discussion.

or scrutable referential intention, etc.), as well as the question of how to represent the contribution of the predicate from which the complex demonstrative is formed.[1]

As King (2001) made clear, however, it is a striking fact about complex demonstratives that not all uses conform to the familiar deictic paradigm. The demonstratives from sentences like the following, for example, are interpreted like definite descriptions:[2]

(2) Every king$_i$ cherishes **that cleric who crowned him$_i$.**

(3) **That candidate who receives the most certified votes** will become mayor.

Examples like (2) and (3), which involve what have been called 'non-deictic' or 'non-referential' demonstratives, and 'demonstrative descriptions', have been the subject of controversy in the philosophical literature. King (1999, 2001, 2008), Roberts (2002), Swanson (2005), and Nowak (2014, 2018, 2019) argue that non-deictic data undermine the classic direct reference treatment of demonstratives. Kaplan (1977), Davies (1982), Neale (1993), Dever (2001), Salmon (2002, 2006, 2008), Corrazza (2003), Comorovski (2007), Braun (2008), and Georgi (2012), on the other hand, maintain that directly referential semantic theories should only be required to explain deictic data.

Those philosophers who take non-deictic data seriously typically offer theories that treat complex demonstratives as a special kind of definite description. On the most prominent version of this line, which is defended by King (2001) and Elbourne (2005), complex demonstratives are formed from two arguments, an overt one that is supplied by the predicate from which the demonstrative expression is formed, and a covert one that takes its value from the context. When a demonstrative is used deictically, the hidden argument place is saturated by a substantive property, and when a demonstrative is used non-deictically, the hidden argument place is saturated by a trivial argument, which results in a semantic structure that is equivalent to a definite description.

[1] On the question of how deictic effects should be modeled in a semantic theory for demonstratives, compare Kaplan (1977, 1989), Devitt (1981), Wettstein (1984), and Reimer (1991). On the question of the semantic status of the predicate from which a complex demonstrative is formed, compare Kaplan (1977), Larson and Segal (1995), Lepore and Ludwig (2000), and Glanzberg and Siegel (2006).

[2] In the terminology introduced by King (2001), (2) is an example of a 'quantifying-in' or 'QI' demonstrative, while (3) involves a 'no demonstration, no speaker reference' or 'NDNS' instance. For present purposes, the differences between the various sorts of non-deictic uses that King identifies is less important than noticing that unlike (1), both (2) and (3) admit readings on which they are truth-conditionally equivalent to the variations that would be formed if the word 'the' were substituted for 'that'. Although it is not a perfect diagnostic tool, Wolter (2006) notes that non-deictic demonstratives tend to take parenthetical disclaimers like 'if such there be' or 'whoever he is' more naturally than (singular) deictic demonstratives, which by their nature involve reference to a particular individual.

By far the more common response to non-deictic demonstratives, however, has been to preserve the shape of traditional semantic theories by claiming that non-deictic demonstratives are not really demonstratives at all. Philosophers inclined to this line typically explain the appearance of non-deictic demonstratives by saying 'that' is ambiguous as between a genuine demonstrative expression and an expression that means what 'the' does.

Although a variety of theoretical arguments have been made against the ambiguity treatment of 'that', philosophers attracted to the view have largely remained unpersuaded. One of the reasons for this may be that there is a level of description at which it can be hard to see what exactly the choice between the two views amounts to. On the ambiguity treatment, the lexicon is complicated by the addition of two distinct demonstrative determiners, each with a comparatively simple semantics. On the hidden argument theory, the lexicon is simple, in that it features just one instance of the determiner 'that', but the semantics for that item is more complicated. Fundamentally, however, both views make the same empirical prediction: that complex demonstratives should always admit of a description-type reading.

One of the goals of the present paper is to break this stalemate by calling attention to a point the parties on both sides of the debate have overlooked. As Wolter (2006, 2007, 2009) makes clear, the range of demonstrative constructions that support description-type interpretations is in fact strictly limited.[3] Despite the fact that the demonstrative from sentence (5) involves what we would expect to be a semantically insignificant variation on the one from sentence (4), the latter can be interpreted as though it were equivalent to a definite description, while the former cannot.[4]

(4) **That guy who wrote *Waverley*** also wrote *Ivanhoe*.

(5) #**That author of *Waverley*** also wrote *Ivanhoe*.

Definite descriptions, however, are felicitous in each of the following configurations:

(6) **The guy who wrote *Waverley*** also wrote *Ivanhoe*.

(7) **The author of *Waverley*** also wrote *Ivanhoe*.

[3] Wolter (2006) credits Maclaran (1982) with an early mention of the phenomenon in question.

[4] I use the hash sign to flag sentences that are infelicitous under the relevant interpretation. It is important to note that there are certain 'emotive' uses of demonstratives, which Partee (p.c.) claims are often used to create or attempt to create a kind of evaluative solidarity between speaker and listener, that would make 'that author of *Waverley*' felicitous even when interpreted non-deictically. Compare: *That (darn) author of 'Waverley'—he turns up everywhere in philosophy, doesn't he?!* Since emotive uses raise issues we cannot settle here, I have tried to structure my examples in a way that defeats emotional shading. For discussion, see Lakoff (1974), Wolter (2006), and Potts and Schwarz (2010).

Both the ambiguity treatment of 'that' and the hidden argument approach predict that non-deictic interpretations should be available for demonstratives wherever a definite description is. (6) and (7) reveal that this prediction is wrong.

In this paper, I follow Wolter in taking this pattern in the data involving non-deictic demonstratives to reveal that the internal semantic structure of complex demonstratives is significantly more complex than philosophers have appreciated.[5] I will argue that while this fact amounts to a fatal flaw for the ambiguity theory, hidden argument theories in the style of King (2001) and Elbourne (2005) can be modified in an intuitively-plausible way to account for the observed subtleties.

The modification I propose is based on the idea from Reimer (1991) that 'when a speaker uses an expression of the form "that *F*" to refer to a particular *F*, there is an implication to the effect that the intended *F* is somehow "discriminated" with respect to all other *F*s' (pg 178). In order for an *F* to be 'discriminated' with regard to all other *F*s, of course, there must *be* other *F*s. I will claim that this is the key to understanding Wolter's demonstratives—to put the point roughly, what licenses (4) is the fact that there are many guys, only one of whom wrote *Waverley*. What rules out (5), on the other hand, is the fact that there is just one author of *Waverley*.[6]

To implement the idea that saying 'that *F*' essentially amounts to saying which of the *F*s, I claim that the English demonstrative determiner 'that' is like the ordinary definite article in that it carries a presupposition of uniqueness, but different in that it carries an anti-uniqueness presupposition as well. Instead of taking a single property-type argument and returning the unique object that satisfies it, 'that' takes two property-type arguments and returns the unique object that satisfies both *on the condition that there be more than one satisfier of the first argument*.[7]

The radical part of this story is that making it work compositionally requires adopting another innovation from Wolter (2006). We must claim, that is, that at least some relative clauses occur in a syntactic and semantic configuration different from the one most theorists expect.[8] On that alternative configuration, which is generally

[5]I focus my discussion here on examples (4) and (5) for the sake of clarity and concreteness. Crucially, I do not intend to claim that the phenomenon involved is unique to relative clauses or relational nouns, much less to my particular example sentences. Whether there are other constructions that would meet the semantic and structural constraints I describe, or which constructions those might be, will have to remain a question for future work. Wolter (2006: 141-146) provides a more complete survey of the possibilities, some of which I discuss on pages 22 and 27.

[6]In a context in which it is mutually known that only Scott wrote *Waverley*, the question 'which guy are you talking about?' is natural, and 'the one who wrote *Waverley*' is an answer. The question 'which author of *Waverley* are you talking about?' on the other hand, is defective.

[7]As I hope will be clear from the discussion, I do not mean to endorse a Fregean treatment of the definite article. I will often speak as though 'the referent' or 'the extension' of a demonstrative in a context is an individual, but this is for the sake of convenience; the view I describe could perfectly well be implemented in a Russellian framework, for example.

[8]On Wolter's view, the crucial work is done not by looking at the cardinality of the predicate

attributed to Bach and Cooper (1978), instead of the noun and the relative clause combining to form a constituent that picks out a single property as in (8), the determiner takes the noun as one argument and the relative clause as another, as in (9):

(8)

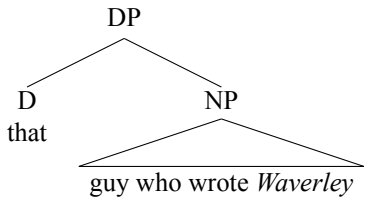

(9)

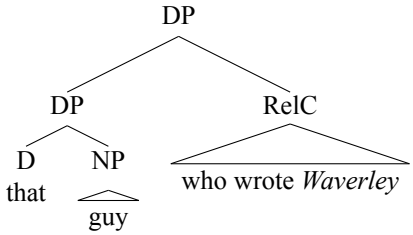

Many philosophers would understandably balk at the prospect of endorsing a revisionary syntactic and semantic structure for relative clause on the basis of some problematic English data. As I will show, however, Wolter's demonstrative data are not the only piece of evidence that is relevant here. In fact, Lin (2003) and del Gobbo (2003) argue that the availability of these two structures is required to explain the pattern of deictic and non-deictic interpretations associated with demonstratives in Mandarin. When Mandarin demonstratives occur in structure (8), they are interpreted deictically, and when they occur in structure (9), they are interpreted non-deictically. Since the differences in the two Mandarin structures are marked explicitly, they provide a direct source of evidence for the kind of semantic theory advanced here.

The plan for the paper is as follows. In the next section, I show how the two most prominent philosophical treatments of non-deictic demonstratives overgenerate, by predicting that a non-deictic demonstrative should be acceptable wherever a definite description is. In section 3, I describe a version of the hidden argument semantics based on the idea that one of a demonstrative determiner's two arguments must serve to restrict the other. In section 4, I show how that idea allows us to derive the right results with regard to a wide range of otherwise problematic extensional data. In section 5, I explain the predictions the proposal makes with regard to modal sentences of a sort that philosophers have traditionally been concerned with, and in section 6, I acknowledge some outstanding issues the proposal raises.

---

associated with the demonstrative, but by the machinery of situations, as developed after the fashion of Barwise and Perry (1983). According to Wolter, the difference between the demonstrative determiner and the ordinary definite article is that the demonstrative determiner requires that its argument be evaluated with regard to a non-default situation; in the case of deictic demonstratives, this requirement is met by 'zooming in' or 'out' of the situation given by the context, and in the case of non-deictic demonstratives, by the fact that the non-standard syntactic position of certain relative clauses allows them to be interpreted with regard to a different situation from the one used to interpret the determiner's first argument. A side-by-side comparison of the two proposals will have to wait for another day, but I hope the present work will at the very least establish how her basic approach might be rendered in terms of the simpler semantic framework that is standard in the philosophical literature on demonstratives.

2 The overgeneration problem

As we have just seen, deictic and non-deictic demonstratives appear to be formed from the same basic components (examples repeated):

(1) **That river** looks treacherous (while pointing at a river).

(2) Every king$_i$ cherishes **that cleric who crowned him$_i$.**

This raises a challenge for traditional philosophical thinking about demonstratives. On the standard picture, which descends from Kaplan (1977), the extension of a demonstrative in a context is the object ostended or intended by the speaker of the context. While this kind of treatment seems reasonable enough when applied to examples like (1), as King (2001) shows, it breaks down when applied to examples like (2). Regardless of how exactly we characterize the contribution of the predicative material from which the demonstrative is formed, it just is not plausible to think that the extension of the demonstrative from (2) must be an object from the context, much less one that is determined by the speaker's referential intentions.

2.1 Ambiguity theories

The most common response to the problem posed by non-deictic demonstratives is to claim that they are not really demonstratives at all. Many philosophers have attempted to preserve a directly referential semantics for deictic demonstratives by doing just this and saying that the determiner 'that' is lexically ambiguous. On this treatment, one member of a pair of homonyms, 'that$_1$', is said to be used to form bona fide demonstratives, while the other, 'that$_2$' is interpreted just like the definite article.[9] If the ambiguity theory were right, some apparent demonstratives, like the one from (4, repeated), would turn out to be definite descriptions in disguise:

(4) **That guy who wrote *Waverley*** also wrote *Ivanhoe*.

Ambiguity theorists explain the non-deictic readings of sentences like (4) by offering disambiguations like the following:

(10) **That$_2$** (= the) **guy who wrote *Waverley*** also wrote *Ivanhoe*.

There are two major problems for the ambiguity theory. The first is that deictic and non-deictic interpretations are available for demonstratives formed from what

[9]Dever (2001) and Georgi (2012) explicitly defend this treatment, which is also implied by remarks from Kaplan (1977), Corrazza (2003), and Braun (2008).

appear to be the same lexical items not just in English but in a wide variety of other languages, as well.[10] To explain this fact, the ambiguity theorist would have to claim that parallel ambiguities recur in a significant range of unrelated and distantly-related languages.

If the ambiguity theory made the right predictions about all of the relevant data, such an awkward consequence might be one that some theorists would be willing to tolerate. But in fact, no plausible ambiguity treatment will result in an empirically adequate theory.

The contrast between examples (4) and (5) (repeated below) reveals that non-deictic demonstratives are not interchangeable with definite descriptions. Since (5) would be felicitous if 'that' were replaced by 'the', there can be no question of 'that_2' being equivalent to 'the':

(4) **That guy who wrote *Waverley*** also wrote *Ivanhoe*.

(5) #**That author of *Waverley*** also wrote *Ivanhoe*.

A determined ambiguity theorist might respond to data like these by offering a more sophisticated version of the theory. Instead of saying that 'that_2' is semantically equivalent to 'the', she might claim that it is equivalent to a hitherto unremarked determiner that appears to be semantically like 'the', but which is only licensed in certain environments.

As we will see, however, in order to make the right predictions about the distribution of those environments, the ambiguity theorist would have to offer a treatment on the order of complexity of the one that will eventually be described here. But developing such a sophisticated account of structural licensing, only to attach it as a footnote to the familiar direct reference semantics for deictic demonstratives, would take a view that was uncomfortably ad hoc to begin with—as far as I can tell, the only argument in favor of the ambiguity treatment is that it would insulate a popular semantic view from counter-example—and make it significantly less plausible. I will assume going forward that the only serious contenders are theories that can generate both deictic and non-deictic interpretations using the same basic semantic machinery.

2.2 Familiar hidden argument theories

Apart from the ambiguity treatment, the most popular alternative approach to the non-deictic demonstrative data involves what we might call a 'hidden argument'

[10] I discuss the cross-linguistic facts in detail in Nowak (2018). For additional historical considerations that recommend against an ambiguity theory, see also Roberts (2002).

analysis.[11] On the hidden argument theory, all complex demonstratives have the same basic semantic structure:

(11) that F = the x: [$F(x)$ & $G(x)$]

The two most prominent versions of the hidden argument theory are due to King (2001) and Elbourne (2005). King defends a Russellian version of the view, on which 'that' combines with two arguments to make a generalized quantifier, while Elbourne develops a Fregean version, on which 'that' takes two arguments and returns the unique individual that satisfies both.

When a demonstrative is used deictically, both authors say that the G argument place is saturated by a hidden argument that corresponds to an identificational property. For King, that property is determined by the speaker's referential intentions, while for Elbourne, the hidden argument place is occupied by a variable over identificational properties, the value of which is set by a contextually-determined assignment function. On either version of the theory, if someone standing in the Desolation Wilderness points up at Mount Agassiz and says:

(12) **That mountain** is lovely!

her demonstrative will be represented along the following lines:

(13) the x: [mountain(x) & identical-to-Mount-Agassiz(x)]

For King, when a demonstrative is used non-deictically, the G argument place is saturated by a trivial property, like the property of being self-identical. Elbourne does not explicitly address the case of non-deictic demonstratives. One of the aims of his semantics, however, is to unify the treatment of demonstratives and definite descriptions. To derive attributive readings for typical definite descriptions, he proposes saturating the determiner's second argument position with a trivial property. So, it is not much of a stretch to think that on either version of the hidden argument theory, the demonstrative from (4, repeated):

(4) **That guy who wrote *Waverley*** also wrote *Ivanhoe*.

[11] There are semantic theories that can generate non-deictic interpretations for complex demonstratives that neither count as ambiguity theories nor hidden argument theories. Roberts (2002) offers a dynamic semantics for demonstratives that can explain non-deictic readings of a certain limited range of uses. In Nowak (2018), however, I explain in detail why her semantics cannot discriminate between the acceptable and unacceptable data presented here. Elbourne (2008) offers a semantics designed to explain cases of deferred ostension. While the machinery Elbourne describes might be adapted to produce non-deictic readings for demonstratives, none of that machinery offers a way to discriminate between our data, either. But discrimination is precisely the key to having an empirically adequate account.

would be represented roughly as per:

(14) the x: [guy-who-wrote-*Waverley*(x) & self-identical(x)]

Since every object is self-identical, the second argument essentially drops out of the derivation, and the demonstrative is interpreted as though it were equivalent to the definite description:

(15) the guy who wrote *Waverley*

Although the hidden argument theory easily generates the desired reading on which the demonstrative from (4) is truth-conditionally equivalent to a definite description, the theory offers no way of ruling out such a reading for the demonstrative from (5, repeated):

(5) #**That author of *Waverley*** also wrote *Ivanhoe*.

If the hidden argument theory were right, we would expect to be able to produce such an interpretation by representing the demonstrative as follows:

(16) the x: [author-of-*Waverley*(x) & self-identical(x)]

The infelicity of (5) shows that the hidden argument theory as described so far must be wrong; non-deictic interpretations for demonstratives are not the result of merely replacing a substantive argument with a trivial one.

3 A candidate solution

3.1 Some intuitive background

The hidden argument theory is built on attractive premises. Instead of explaining the behavior of demonstratives by deploying (for example) sui generis operators that take a definite description and make it directly referential, the hidden argument theorist claims that demonstratives involve the same semantic operations as other determiners. She also promises to explain a wider range of data than her competitor; her analysis is designed to cover both familiar deictic uses of demonstratives as well as non-deictic uses.

As we have just seen, however, the hidden argument theory as described so far is not empirically adequate. We can fix this problem without giving up the advantages

of the two-argument architecture if we modify some of the details. In order to appreciate the way the required modifications work, we will have to look more closely at the possible syntactic and semantic structures that might be associated with the demonstratives from (4) and (5). Before doing that, however, it will be helpful to establish some context, by taking a broad look at the role demonstratives typically play in our communicative practices.

In many familiar cases, demonstratives supply answers to the question 'which one?' Consider a butcher-shop vignette:

(17) A: Number 49?
B: Yes, I'd like a rib-eye steak, please.
A: Which one?
B: **That one**, please. (pointing)

This dialogue illustrates an extremely common pattern of use of deictic complex demonstratives. By saying 'rib-eye steak', B calls attention to a particular class of individuals, and by pointing, selects one from among them. Very frequently, when we use complex demonstratives deictically, a set of candidate referents is provided by the predicate from which the demonstrative is formed, and we pick out one by way of a gesture, or by its salience, or whatever.

The idea that complex demonstratives serve to pick one object out from a set of candidates has long been prominent in the philosophical literature. Witness, for example:

> A not implausible view about the reference of expressions of the form 'that *F*', is that such expressions refer to *F*'s which have somehow been 'discriminated' from all other *F*'s. After all, when a speaker uses an expression of the form 'that *F*' to refer to a particular *F*, there is an implication to the effect that the intended *F* is somehow 'discriminated' with respect to all other *F*'s. (Reimer 1991: 178)

The foregoing suggests a way of distinguishing the felicitous (4, repeated) from the infelicitous (5, repeated):

(4) **That guy who wrote *Waverley*** also wrote *Ivanhoe*.

(5) #**That author of *Waverley*** also wrote *Ivanhoe*.

At first glance, it seems like (4) should admit an interpretation on which the relative clause is used to perform the kind of restricting work that would ordinarily be done by a pointing gesture; there are many individuals that satisfy the predicate 'guy', and the relative clause provides a way of selecting just one from among them. On the other hand, there does not appear to be any way of saying the same thing about

(5), since there is no way of restricting the predicate 'author of Waverley', which already picks out a single individual. Indeed, a common reaction to the contrast between examples (4) and (5) is to point out this difference.[12]

Standard versions of the hidden argument theory are not well-positioned to implement an explanation along these lines, however. Neither King nor Elbourne offer a specific analysis of the structure of restrictive relative clauses, but it is clear that both authors subscribe to some version of the familiar picture on which a relative clause and the noun that it modifies combine to form a single property-type constituent.[13]

As we have said, for King (2001), the English determiner 'that' is a quantifier expression, the behavior of which can be modeled using the following fragment of a formal language:[14]

> The propositional frame expressed by [[Thatξ Σ]Ψ] in *c* is [[THAT$_{f(c),h(c)}$ ξ Σ']Ψ'], where *f* is a function from contexts to propositional frames and *h* is a function that maps each context $\langle i, w, t\rangle$ to either J, the property of being jointly instantiated, or J$_{wt}$ the property of being jointly instantiated in *w,t*, where if $f(c) = [\xi = * o]$ or $[o = * \xi]$, for some individual *o*, then $h(c)$ = J$_{w,t}$. Otherwise, $h(c)$ = J; and THAT$_{f(c),h(c)}$ is the result of saturating the second and third argument places in the 4-place relation expressed by 'that' (i.e., THAT: __ and __ are uniquely __ in an object and it is __) with $f(c)$ and $h(c)$ respectively (i.e., __ and $f(c)$ are uniquely jointly instantiated/jointly instantiated in *w,t* in an object and it is __); and with Σ', Ψ' as above. (King 2001:165)

On this treatment, the syntactically-realized arguments Σ and Ψ saturate two of the four places associated with the determiner 'that'. Adapted to fit English, and applied to the case of (18), for example, King's approach would see the determiner's two syntactic arguments providing the properties of being a hominid who discovered fire and the property of being a genius:

(18) **That hominid who discovered fire** was a genius.

[12] Wolter (2003) describes, but does not endorse, a view of the sort I will develop here. There, she alleges a contrast between *The first hero who kills the dragon* and *#That first hero who kills the dragon*, which she explains by claiming that the demonstrative version is unacceptable because "it does not contain an NP complement that denotes a set with more than one element" (pg. 23). She takes the effect, however, to be due to a conversational implicature, as opposed to "an entailment or a presupposition".

[13] Heim and Kratzer (1998) credit Quine (1960) with an early description of this idea, which, in various syntactic guises, is the contemporary standard. For three prominent contemporary developments, see Kayne (1994), Heim and Kratzer (1998), and Sauerland (2003).

[14] Here, ξ is a variable, and Σ and Ψ are formulas, as expected.

When a demonstrative sentence is used in a context in which the speaker does not intend to refer to a particular object, the *h* function returns the property of being jointly instantiated, and the *f* function a trivial property, like the property of being self-identical. Those properties saturate the remaining two argument places associated with 'that', and the upshot is that (18) expresses truth conditions that we might paraphrase as follows:

(19) The property of being a hominid who discovered fire and the property of being self-identical are uniquely jointly satisfied by an object *x* and *x* is a genius.

As noted in the previous section, this ingenious treatment makes exactly the right predictions with regard to the class of data that motivated it, i.e., non-deictic uses of demonstratives like the one from (18).[15] Crucially, however, by treating the relative clause 'who discovered fire' and the noun 'hominid' as together providing the determiner with a single property, the treatment leaves us no way to implement the idea that the relative clause serve as a restrictor on the set of hominids. So, where the subtler pattern evidenced by Wolter's data is concerned, we are left with no way of making the required discrimination.

Elbourne's version of the hidden argument theory runs into the same difficulty. Elbourne says that the hidden argument associated with a demonstrative expression is supplied by an index on the determiner:

(20) [[that$_i$] F]

With regard to an assignment that maps i to the property G, (20) is interpreted as though it were equivalent to:

(21) the *x*: [*F*(*x*) & *G*(*x*)]

This treatment suggests that the demonstrative from (18) would be represented as per:

(22) [[that$_i$] hominid-who-discovered-fire]

To derive a non-deictic interpretation, (22) would have to be evaluated with regard to a variable assignment that maps the demonstrative index to a property like the property of being self-identical. Such an assignment would result in something equivalent to:

[15] Of course, the treatment also makes the right predictions about *deictic* demonstratives; the fact that it provides a fundamentally unified treatment is one of its primary selling points. In contexts in which the speaker has a determinate object in mind, the *f* function will return an identificational property involving that object, and the *h* function the property of being jointly instantiated in w, t. This guarantees that the demonstrative expression will 'refer' rigidly to the object in question.

(23) the x: [hominid-who-discovered-fire(x) & self-identical(x)]

Regardless of whether we take a King-style or an Elbourne-style approach to the hidden argument theory, as long as we think that the relative clause and the noun it modifies combine to form a single argument, we will have no way of distinguishing the felicitous (4, repeated) from the infelicitous (5, repeated):

(4) **That guy who wrote *Waverley*** also wrote *Ivanhoe*.

(5) #**That author of *Waverley*** also wrote *Ivanhoe*.

Regardless of how exactly we say the determiner works, that is, there will be no way to avoid the fact that the property contributed by 'guy who wrote *Waverley*' is the same as the property contributed by 'author of *Waverley*' (*modulo* an irrelevant gender presupposition introduced by 'guy'). In order to avoid this problem, we need some way of separating a noun from the relative clause that modifies it; with regard to our example (4), that is, we need some way of treating 'guy' and 'who wrote *Waverley*' as two separate arguments for the determiner.

3.2 The structure of relative clause

As it turns out, a syntactic configuration that would allow us to do exactly what is needed has long been the topic of discussion among linguists. Ross (1967) described a structure for restrictive relative clauses that according to Stockwell et al. (1973) was the standard for the time. On that structure, a determiner combines with a noun to form a constituent, which in turn combines with the relative clause:

(24)

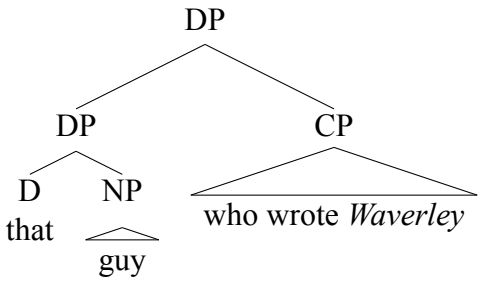

Partee (1975) argued that this configuration—which, because of the syntactic category names of the era, came to be known as the 'NP-S' configuration—would violate compositionality, and could thus be ruled out on semantic grounds. The problem, as she saw it, was precisely that 'guy' and 'who wrote *Waverley*' do not form a constituent. If 'the' is understood along familiar lines, and if 'guy' and 'who wrote *Waverley*' simply pick out the properties of being a guy and having written *Waverley*, respectively, this means there will be no way of deriving the expected extension:

(25)

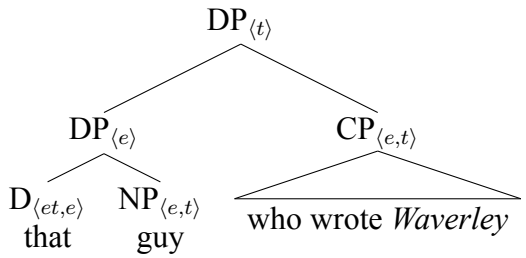

If there is only one guy, the higher DP from (24) will have a truth-value as its extension, instead of picking out the unique author of *Waverley*. If there is more than one guy, the extension of (24) will be undefined. Neither of these results is acceptable.

Bach and Cooper (1978), however, showed that this objection could be avoided, by describing a simple way of making the NP-S structure produce the expected compositional outcome. Their solution was to effectively raise the type of the determiner, by inserting a variable over properties into its semantic representation. When a relative clause occurs in the NP-S configuration, they say, instead of:

(26) $[\![\textbf{the}]\!] = \lambda f.\iota x : f(x) = 1$

the determiner is interpreted as though it introduced a resource variable, *R*:

(27) $[\![\textbf{the}]\!] = \lambda f.\iota x : f(x) = 1$ and $R(x) = 1$

A construction-specific composition principle allows the property picked out by the relative clause to provide the value of the *R* variable, so that the semantic value for (24) turns out to be:

(28) $\iota x : x$ is a man and x wrote *Waverley*

Bach and Cooper argued that the NP-S structure was required to explain the composition of relative clauses in Hittite.[16] Where English is concerned, however, they saw their work as a proof of concept, since the data involving definite descriptions that they considered are data that could just as easily be handled using the structure that is standard today, on which a noun and a relative clause combine to form a constituent.

Even if English data were the only data we had access to, it would be reasonable to take the contrast between (4) and (5) to provide precisely the kind of argument Bach and Cooper would have needed to vindicate their proposal—their structure

[16] Larson (1982) showed how Bach and Cooper's analysis might be extended to deal with a certain variety of relative clause construction in Walpiri. An anonymous referee points out that the Hittite and Walpiri data might nowadays be explained as involving extraposition.

for relative clause fills a theoretical role the standard structure cannot.[17] After all, regardless of how exactly we explain the difference between the demonstratives that license non-deictic interpretations and the ones that do not, it is hard to see how the puzzle could possibly be solved if there were no way of separating the NP from the relative clause. If 'guy' and 'who wrote *Waverley*' form a constituent, that constituent will pick out a property which, for all the determiner cares, is the same as the property picked out by 'author of *Waverley*'.

Importantly, however, the contrast between the English examples (4) and (5) is not the only reason for thinking that relative clauses sometimes occur in the familiar low configuration, and sometimes high, in the NP-S configuration. In fact, Lin (2003) and del Gobbo (2003) argue that the NP-S analysis is required to make sense of a parallel contrast in the interpretations associated with certain demonstrative constructions in Mandarin, a contrast that Huang (1982) credits Chao (1968) with first remarking on.[18] Consider:

(29) neiben wo zuotian mai de shu
that I yesterday buy DE book
'That book, which I bought yesterday'

(30) wo zuotian mai de neiben shu
I yesterday buy DE that book
'The book that I bought yesterday'

(31) na-yi-ge [chouyan de] ren
that-one-CL smoke DE person
'That person that smokes'

(32) [chouyan de] na-yi-ge ren
smoke DE that-one-CL person
'The person that smokes'

Huang (1982) maintained that the relative clause from (29) is used non-restrictively, to add parenthetical information about the extension of a a deictic demonstrative, while the relative clause from (30) plays a restrictive role, helping to determine the extension of the demonstrative phrase. Lin (2003) offers a variety of syntactic and

[17] Although I became convinced of the need to separate at least some relative clauses from the nouns they modify before encountering Wolter's excellent work, she deserves the credit for the idea of applying the NP-S structure to the problematic data involving non-deictic demonstratives.

[18] I am indebted to Peter Jenks. for bringing the literature on Mandarin demonstratives to my attention, and to an anonymous referee for further helpful discussion. Examples (29) and (30) are from Huang (1982). Examples (31) and (32) are from del Gobbo (2003). I transcribe both authors' glosses without modification. Del Gobbo uses the abbreviation 'CL' to indicate a classifier, and 'DE' the word 'de', which is a modification marker.

semantic arguments against the idea that the relative clause in (29) is non-restrictive, but he describes a broad agreement among authors about the fact that the first syntactic configuration systematically produces deictic interpretations, while the interpretations associated with the second are non-deictic.

This is especially interesting for our purposes because the word order in Mandarin provides a non-speculative way of determining how high the attachment site of a relative clause is. Unlike in English, that is, the claimed syntactic difference between the NP-S relative clause and the standard structure is manifest on the surface. Lin and del Gobbo analyze Huang's demonstrative-initial constructions using the following structure:

(33)

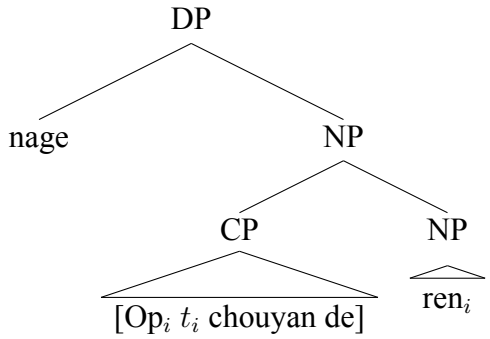

In the relative clause-initial constructions, on the other hand, they claim that the demonstrative combines first with an unmodified NP and later with the relative clause:

(34)

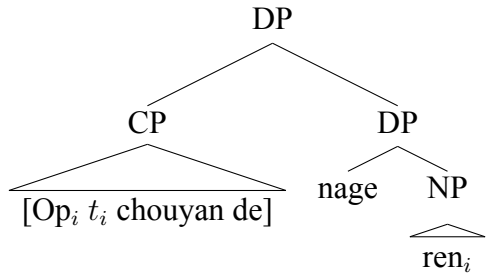

How can the semantic derivation for (34) proceed compositionally? And why should a structure like (33) produce a deictic interpretation, while structure (34) is interpreted non-deictically? Lin and del Gobbo answer both questions at once by employing the same machinery Bach and Cooper used to avoid Partee's challenge about compositionality. Lin writes:

> Context-dependency can be nicely captured by introducing to the usual translations of determiners an extra property variable whose value is filled by a variable assignment. However, if there is overt linguistic

> material denoting a property around the property variable, the variable can be filled by that property… (Lin 2003: 230)

In other words, Lin and del Gobbo treat the demonstrative determiner in roughly the same way as King and Elbourne, claiming that it performs a semantic operation not on one property, but two.[19] If (33) provides the right structure for (31), then when the noun and the relative clause form a constituent, they jointly occupy only one of the two argument places introduced by the demonstrative determiner, making the second available for contextual saturation. On the other hand, when the structure of the demonstrative expression allows 'overt linguistic material' (here in the form of a high relative clause) to supply a distinct property-type argument, there is no work left for the context to do, and the expected interpretation turns out like a definite description.[20]

The upshot for us is that Mandarin demonstrative constructions involving relative clauses provide a model for understanding their English counterparts. As we will see, if we combine the idea that English relative clauses can occur in either of the two structures described here with the idea that the demonstrative determiner introduces presuppositions involving both uniqueness and anti-uniqueness, we can explain the puzzling data we began with.[21]

3.3 A job for presupposition

We turned our attention to questions about the structure of relative clause because we wanted a way to explain the contrast between the English sentences (4) and (5). We thought we might be able to make progress on that contrast by saying that the predicative material from the former sentence is structured in a way that allows one of the determiner's arguments to perform a kind of restriction operation on the other, while the material from the latter sentence is not. We needed something like the NP-S structure to make that explanation compositionally possible. Now that we see that there are independent reasons for thinking that at least some relative clauses indeed occur in that configuration, we are in a position to pull the rest of the pieces together.

[19] Lin (2003) endorses a Fregean semantics on which the determiner takes two arguments and returns their unique joint satisfier, while del Gobbo (2003) treats the determiner in quantificational terms, in roughly the way King (2001) does.

[20] An anonymous referee points out that this treatment echoes the idea from Vendler (1967) and Cooper (1975) that restrictive relative clauses can sometimes occupy the position of a covert restrictor variable.

[21] Yang (2005) describes data involving Mandarin possessive constructions that can be used to mount an argument similar to the one sketched here. Possessor phrases in Mandarin can attach either high or low with regard to a demonstrative element, with the attachment site marked by word order. Partee (2006) analyzes the high-attached possessors as non-deictic demonstratives.

As we have seen, King and Elbourne invite us to think of demonstratives as definite descriptions that sometimes take an identificational property as the value of a hidden argument, and sometimes a trivial property. This flexibility allows them to offer a unified analysis of both deictic and non-deictic data, but as we saw, it causes their theories to overgenerate. If we add a certain presuppositional restriction to the basic architecture of the hidden argument analysis, we can retain its breadth of application while avoiding the overgeneration problem.

Like King and Elbourne, I propose that we treat demonstratives that appear to have the form:

(35) that *F*

as though they really involved the determiner's taking two arguments, *F* and *G*, and performing a description-type operation on them, so that (35) is interpreted:

(36) the *x*: [*F*(*x*) & *G*(*x*)]

As on the hidden argument theory, I take the property that occupies the first argument place in the schema, F, to be supplied by the predicative material from which the complex demonstrative is formed.

Instead of following the hidden argument theorist in saying that the property that occupies the second argument place, G, is always covert, I claim that it is covert in deictic cases, but overt in non-deictic cases. The fact that certain syntactic and semantic environments make an explicit second argument available, while others do not, will play an important role explaining the curious pattern in the data concerning non-deictic demonstratives.

The key difference between my proposal and the hidden argument theory and between my proposal and the proposals described by Lin and del Gobbo, concerns the relationship between the two arguments taken by the determiner. Instead of letting 'that' return the singleton intersection of any two properties, I propose limiting its application along the following lines:

(37) that $F = \begin{cases} \text{[the } x\text{: } [F(x) \text{ \& } G(x)]] \text{ iff } (F \cap G) \subset F \\ \text{otherwise undefined} \end{cases}$

(37) is meant to capture the intuition that when someone utters a (singular) complex demonstrative, the predicate she uses introduces a set of candidates from which a single individual is to be picked. Our formulation works by adding a new presupposition to the presupposition of uniqueness that is standardly supposed to be a part of the semantics of 'the'. This presupposition requires that the demonstrative's second argument restrict its first argument in the following sense:

Definition A property G is a **restrictor** on another property F just in case the intersection of $\{x : F(x)\}$ and $\{x : G(x)\}$ is a proper subset of $\{x : F(x)\}$.

The imprecise formulations from (36) and (37) are meant to underscore the fact that I would prefer to remain agnostic about details that I take to be irrelevant where the primary point of this paper is concerned. So, for example, while it would be natural to think that when demonstratives are used deicitically, an identificational property occupies the syntactic position that can be occupied by a relative clause in a non-deicitc construction, I see no reason to stake a claim with regard to the question (in my discussion to follow, I will assume that structure, but nothing important turns on it).

By the same token, since my aim here is not to settle the question of Frege vs. Russell, I characterize the demonstrative determiner in terms of 'the', without saying how exactly 'the' should be understood. In the discussion to follow, I assume a Fregean approach to definite descriptions so that I can talk simply about 'the referent' of a certain demonstrative instead of about a function from properties to truth values. For concreteness' sake, then, I will assume that *modulo* a host of irrelevant details concerning the distal/proximal distinction, animacy/inanimacy requirements, and similar:

(38) $[\![\textbf{that}]\!] = \lambda f \lambda g$: the intersection of $\{x : f(x) = 1\}$ and $\{x : g(x) = 1\}$ is a proper subset of $\{x : f(x) = 1\}$. $\iota x : f(x) = g(x) = 1$

As far as I can tell, however, none of the relevant features of my proposal depend on this assumption. The restriction presupposition that is at the heart of my proposal could easily be built into either of Elbourne's or King's version of the hidden argument theory, as well as theories like the one described by Roberts (2002).

4 Extensional results

4.1 Paradigmatic deictic data

If we apply our analysis to a typical deictic demonstrative, we can quickly verify that it makes intuitively accurate extensional predictions. Suppose, to take a standard sort of example, that someone utters (39) while pointing towards Maryam Mirzakhani:

(39) **That woman** won a Fields medal.

On our view, the property of being a woman saturates the first argument place introduced by the determiner, and the property of being identical to Mirzakhani saturates

the second. Since there are women other than her, the property of being identical to Mirzakhani is a restrictor on the property of being a woman, according to our definition; in other words, $(\{woman\} \cap \{Mirzakhani\}) \subset \{woman\}$. Once we verify that the restriction presupposition is met, we apply our schema and end up with the following representation for the demonstrative:

(40) the *x*: [woman(*x*) & identical-to-Mirzakhani(*x*)]

Mirzakhani is the unique individual that is a woman and that is identical to Mirzakhani, so our treatment predicts that she herself will be the extension of 'that woman', when the expression is uttered in the context described.

How exactly does the property of Mirzakhani come to saturate the second argument place associated with the demonstrative determiner? Although this question is interesting and important, it raises issues that far outrun the scope of the present work.

On one compositionally plausible story that fits well with a familiar treatment of other referring expressions, we might claim that the semantic representations for simple deictic demonstratives feature variables over individuals. To extend that story to cover complex demonstratives, too, we could claim that those variables are type-shifted by something like the *ident* functor described by Partee (1986). On such a view, deictic demonstratives would not be sensitive to the context of utterance, per se, but to the pragmatically-determined value of an assignment function. For present purposes, however, I cannot see any reason for thinking that the details matter. Let the value of the hidden argument be set by a speaker's referential intentions, by her gestures, or by whatever mechanism is described in your favorite theory. The upshot, as far as anything I hope to establish here is concerned, is an assignment-sensitive representation that looks something like this:

(41)

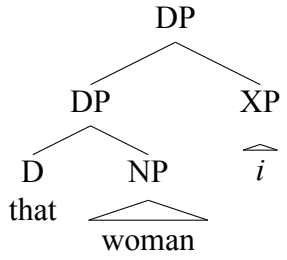

With regard to an assignment that maps *i* to Mirzakhani, (41) amounts to:[22]

[22] I use strikeout notation to indicate unpronounced material.

(42)

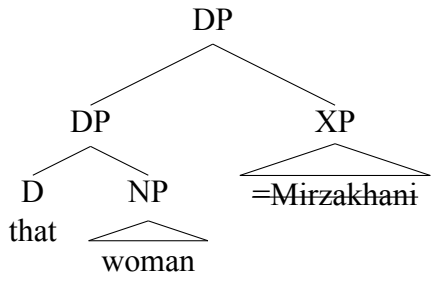

4.2 Acceptable non-deictic data

The major advantage of our view over the hidden argument theory becomes evident when we consider non-deictic data; adding the restriction presupposition to our semantics gives us a way to explain the pattern that obtains in those data. On our view, English non-deictic demonstratives are derived in the same way as their Mandarin analogues, by means of a high-attached relative clause (example repeated from page 13):

(24)

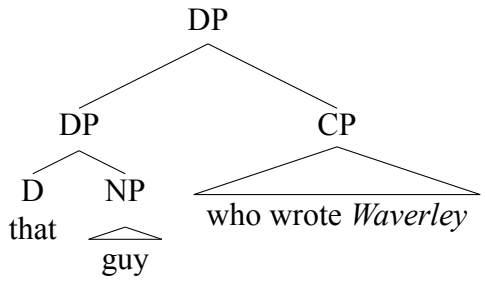

Standard assumptions about local semantic composition allow us to use structure (24) to derive the result we expect for non-deictic demonstratives. In this structure, 'that' finds both of the arguments it requires in the syntax; 'guy' saturates one argument place, and 'who wrote *Waverley*' saturates the second.

To compute the extension of the string, we start by checking to see that it satisfies our restriction presupposition. Since there are guys who did not write *Waverley*, the expression 'who wrote *Waverley*' is a restrictor on 'guy', according to our definition, which means the derivation can proceed. The next step is to apply the schema from (37), which yields the following representation:

(43) the x: [guy(x) & wrote-*Waverley*(x)]

The unique individual that satisfies the predicates 'guy' and 'wrote *Waverley*' is Scott himself, which means that on our theory, 'that guy who wrote *Waverley* picks out just what we expect it to.

4.3 Unacceptable non-deictic data

In addition to making the right predictions about demonstratives that allow non-deictic interpretations, if we combine our presupposition requirement with the idea that relative clauses can occur in two different positions, we open up a way to make the right predictions about those demonstratives which do *not* allow such interpretations. Consider our example (5, repeated):

(5) #**That author of *Waverley*** also wrote *Ivanhoe*.

In a nutshell, the problem with (5) is that it only makes a single argument—'author of *Waverley*'—available to the demonstrative determiner. Since overgeneration errors rule out the idea of using trivial properties to fill the second argument place introduced by 'that', this means that the relevant constructions end up with semantic representations that are incomplete.

On standard thinking, the matrix of the demonstrative from (5) involves two semantically significant constituents. The first—the word 'author'—is what is sometimes called a 'relational' noun; it picks out the two-place relation that obtains between authors and the things they write. The second—the word '*Waverley*'—is a proper name for a book that is probably discussed more than it is read. The entire expression 'author of *Waverley*' is formed when 'author' takes '*Waverley*' as an argument (there are good reasons to think 'of' is merely a phonetic marker of the argument relation in this construction).[23] In the demonstrative from (5), in other words, the *x*-wrote-*y* relation is partially saturated by the book *Waverley*, and the result is an expression that picks out the property of having written *Waverley*.

If we follow standard practice and treat 'author of *Waverley*' as though its extension is a property, we can use that property to saturate one of the argument places introduced by the demonstrative determiner. The result, however, is the semantically incomplete:

(44) the *x*: [author-of-*Waverley*(x) & $G(x)$]

If we split the expression into its basic constituents, on the other hand, we end up with a two-place relation and an individual, not the two properties we need to fill out our template for demonstratives.

The significance of this problem can be made vivid by contrast with the relative clause case. While it is important that availability of high attached relative clauses makes it syntactically plausible for us to separate nouns from the restrictive relatives

[23] Although there is broad support for the standard treatment of relational nouns among both linguists and philosophers, an anonymous referee points out that there are dissenters, too. See Adger (2012), chapter 4 for discussion.

that modify them, the real key to our explanation of the viable non-deictic examples is the semantic fact that the guy who wrote *Waverley* is both a guy and is a thing that wrote *Waverley*. Even if the syntax allowed it—which it does not—a parallel treatment of relational genitives would result in nonsense: the author of *Waverley* is not the unique thing that is both an author-of-*y* and which is identical to the novel *Waverley*.

Importantly, this explanation of the difference in acceptability between (4) and (5) (repeated):

(4) **That guy who wrote *Waverley*** also wrote *Ivanhoe*.

(5) #**That author of *Waverley*** also wrote *Ivanhoe*.'that author of *Waverley*'

is not intended to apply *only* to relational genitives and relative clauses. A survey of all the possible syntactically well-formed complex demonstrative configurations is beyond the scope of the present work, but if our theory is correct, we should expect any given configuration to license a non-deictic interpretation only if there is a way to extract two arguments of the appropriate type from it and if those arguments are compatible with the presupposition described here.[24]

For the sake of illustration, consider how the strategy might be applied in the case of superlatives:

(45) **The fastest rider** will take a hefty purse.

(46) **#That fastest rider** will take a hefty purse.

(47) **That rider who rides faster than all the rest** will take a hefty purse.

(45) is a perfectly ordinary string that involves a definite description. If we replace the definite article from that description with the demonstrative determiner, the result—contrary to what would be predicted by the standard approaches to non-deictic data—is the degraded (46). The felicity of (47) shows that it is not the property of being faster than everyone else, per se, that causes this problem.

As we did in the case of the construction involving a relational noun—and again, holding in abeyance worries raised by the syntactic implausibility of treating 'fastest' and 'rider' as two separate arguments—we can explain the markedness of (46) by pointing to the fact that the fastest rider is not the unique individual that is both fastest and a rider. In order for the word 'fastest' to work the way we expect, it has to modify 'rider', which means that the demonstrative determiner takes the con-

[24]For a discussion of the range of constructions that license non-deictic interpretations—all of which I take to be compatible with the account offered here—again, see Wolter (2006: 141-146).

stituent ‘that rider’ as a single argument, leaving the second argument position unsaturated.[25]

An anonymous referee wonders whether pre-nominal adjectives should make a non-deictic interpretation available. I imagine the answer will depend on the syntactic possibilities offered by the language in question, as well as on the choice of adjective; if the pre-nominal adjective is not a restrictor on the first argument to the determiner, our semantics will leave the extension of the demonstratives undefined, and even if it is a restrictor, the two arguments must be jointly satisfied by a unique individual. In English, however, as Wolter (2006) observes, pre-nominal adjectives do not admit non-deictic interpretations:

(48) #**That unhelpful person** will be fired.

(49) #**Those friendly applicants** will be hired.

(50) #**Those legal immigrants** were granted citizenship. (adapted from Wolter 2006: 143)

If the account developed here is right, this is exactly the result we should expect. The prenominal adjective forms a constituent with the noun it modifies and returns a single property-type argument for the determiner. In order for the complex demonstrative to have the expected semantic type, this means a second argument must be supplied. So, we should expect the demonstrative from (48), for example, to be felicitous only when used deictically. In fact, this is what we find:

(51) **That unhelpful person** will be fired. (pointing at a certain person)

4.4 More complicated cases

Two-author scenarios

So far, we have been primarily concerned to explain why the expression ‘that author of *Waverley*’ does not admit the same non-deictic readings as the semantically similar ‘that guy who wrote *Waverley*’. Our discussion raises a few issues that deserve clarification.

One of those issues concerns the possibility of using ‘author of *Waverley*’ to form a *deictic* demonstrative. As far as the syntax we have relied on is concerned, a structure like the following should be permissible:

[25] Similar reasoning will rule out any construction that would rely on a non-intersective adjective to saturate one of the required argument places.

(52)

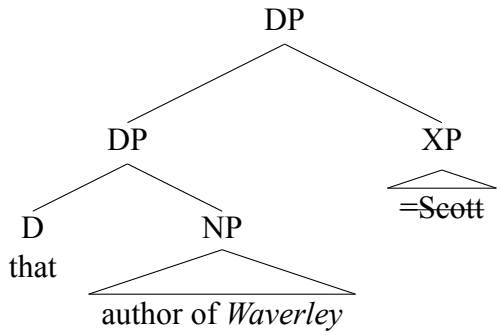

If the theory we have described so far is correct, however, we should expect (5) to be marked even in a context in which someone utters it while pointing towards Sir Walter Scott. After all, since there is only one author of *Waverley*, there is no way to restrict the predicate 'author of *Waverley*', and thus no way of satisfying the presupposition demonstratives introduce. In fact, this prediction appears to be borne out by the data; 'that author of *Waverley*' is just as bad taken deictically as it is taken non-deictically.[26]

Another point that deserves emphasis is that the considerations advanced above do not rule out strings involving relational nouns, full stop. If the account we have offered is right, 'that author of *N*' should be perfectly acceptable in a context that would support a restriction on the extension of 'author of *N*'.

Since we know that only Scott wrote *Waverley*, example (5) will not help in making this clear. If we take up the case of a book we know had two authors, however, the situation changes; this provides an additional source of evidence that our presupposition requirement is on the right track.[27]

Imagine that we show up at a book signing hosted by Russell and Whitehead. In such a scenario, you might say to me:

(53) **That author of *Principia*** (gesturing at one) looks friendly, but I wouldn't try to get an autograph from **that one** (gesturing at the other).

On our theory, the first instance of 'that author of *Principia*' from (53) would be interpreted:

[26] A violation of the restriction presupposition is not the only thing that makes this demonstrative bad; I suggest in note 35 on page 33 that the predicate from which a demonstrative is formed is supposed to help listeners figure out which object a speaker has in mind. Since the property of having written *Waverley* is not one that Scott wears on his sleeve, so to speak, 'that author of *Waverley*' fails in this regard, too.

[27] Examples like these reflect the fact that the evaluation of presuppositions proceeds with regard not simply to the facts that characterize the speech situation, but with regard to the conversational background. People who believe that *Waverley* was a collaboration, for example, would presumably have different intuitions about the felicity of the deictic demonstrative from people who know that it was Scott's work.

(54) the *x*: [author-of-*Principia*(*x*) & identical-to-Whitehead(x)]

by way of the following structure:

(55)

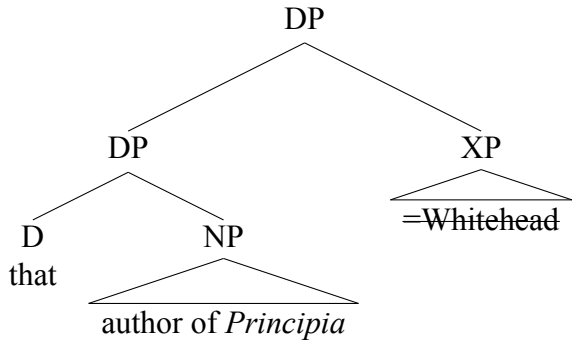

Since there are two individuals in the scenario described that each satisfy the property of having written *Principia*, the property of being identical to Whitehead counts as a restrictor in the sense of our (37), which means the demonstrative presupposition is met. The unique individual that wrote *Principia* and is identical to Whitehead is Whitehead, so we predict that the demonstrative will be felicitous and refer to Whitehead, just as our intuitions demand.

Similar considerations would allow the expression 'that author of *Principia*' to be used non-deictically in a linguistic environment that provided a way of picking out just one of the authors. Consider:

(56) **That author of *Principia*** who spent time in jail was famous for his political views.

(57)

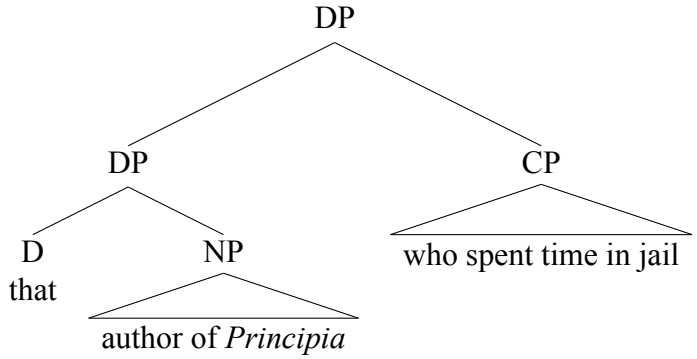

As before, the first step in our derivation is to check that the predicate from which the demonstrative is formed is multiply-satisfied. Then we check to see whether there is an appropriate restrictor; since only Russell spent time in jail, the relative clause can serve in that role. The only thing that wrote *Principia* and served time in prison was Russell, so we predict that the demonstrative from (56) picks him out.

Deixis with relative clauses

The analysis of demonstratives offered here depends on the fact that relative clauses are at least sometimes found in a syntactically high position, in the NP-S configuration. Making the analysis work, however, does not require claiming that *all* relative clauses take that structure. In fact, to make the right predictions about deictic demonstratives that involve relative clauses, we need the clauses to be available in the familiar lower position, too. Consider the demonstrative from the following sentence:

(58) See **that guy who just topped out with no rope?** (pointing towards Alex Honnold)

In order to derive a deictic interpretation for (58), we rely on a structure that makes 'guy who just topped out with no rope' a single argument, so that the second argument position can be saturated by an identificational property, i.e., the property of being identical to Alex Honnold:

(59)

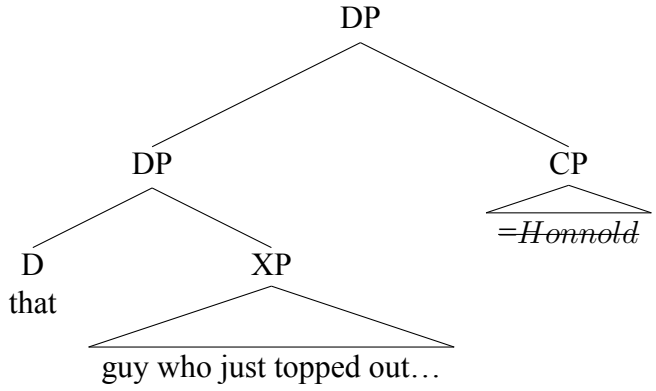

Non-deictic demonstratives without relative clauses

An anonymous referee reports having no difficulty hearing a non-deictic reading for demonstratives like the ones from:

(60) Every academic cherishes that first paper of theirs.

(61) Every university professor cherishes that first paper of hers. (King 2001: 40)

The referee worries that sentences like these reveal that the central phenomenon presented here, and the analysis I offer of it, are less general than they at first appear. (60) and (61) do not involve a relative clause, and they might thus seem to be unlikely candidates for the restrictor treatment I have described.

Given the controversial status of the syntax and semantics of possessive genitive constructions in English, a proper discussion of this sort of example will require a paper of its own.[28] My inclination, however, would in fact be to treat the possessive genitive that appears in these constructions in the same way I treat relative clause, by attaching it higher than is typically expected. If it is plausible to treat the possessive genitive not as an argument but as a modifier, as many have claimed, then detaching it from the noun it modifies will not result in problems of the sort I claim undermine the relational genitive constructions we looked at in section 4.3.

5 Intensional results

5.1 Non-deictic data

The apparent truth conditions of sentences involving modal operators and non-deictic demonstratives favor the hidden argument theory over traditional direct reference semantics for demonstratives.[29] Our semantics offers a similar advantage.

Consider the following example:

(62) If Åsa$_i$ had won the election, she$_i$ would definitely have embraced **that elector who cast the deciding vote.**

Understood naturally, (62) is true just in case, in the nearest world in which Åsa wins the election, she hugs whichever individual from that world cast the deciding vote. In other words, if the nearest world in which Åsa wins the election is one in which Elizabeth Warren cast the deciding vote, the sentence will be true just in case Åsa hugs Elizabeth Warren at that world.

This interpretation is easily derived if we treat the demonstrative from (62) in the way we have suggested here. We say the relative clause 'who cast the deciding vote' serves as a restrictor on the set of electors, and we predict that the truth-conditional contribution of demonstrative will be the same as the contribution made by the definite description from:

(63) If Åsa$_i$ had won the election, she$_i$ would definitely have embraced **the elector who cast the deciding vote.**

If we were to endorse a theory that treated all demonstratives as rigid designators, on the other hand, we would have no way of generating the required interpretation

[28] Compare for example Lyons (1986), Barker (1995, 2011), Partee and Borschev (2003), Vikner and Anker (2002), and Peters and Westerståhl (2013).

[29] For a detailed exploration of the issues involved here, compare Nowak (2014).

for (62). Suppose, for example, that Antonin Scalia in fact cast the deciding vote, and that the conservative candidate therefore won the election instead of progressive Åsa. If we treat *that elector who cast the deciding vote* as though its extension at every world were the same as its extension at the actual world, we would end up having to say that (62) is true just in case Åsa hugs Scalia at the nearest world in which she wins the election. While there might be circumstances in which someone would want to express this idea—maybe the speaker intends to communicate that a victory would be so significant that Åsa would even reconcile with Scalia—the most natural reading of (62) is the reading on which the claim made is the claim that Åsa would have hugged whoever it turned out to be that handed her the victory.

5.2 Deictic data

One of the primary motivations for direct reference is the intuition that deictic demonstratives are rigid designators. If someone points out Semyon, who is wearing a poncho, and says:

(64) **That guy in the poncho** might have been late.

most people will agree that the proposition expressed is true just in case there is an accessible world in which Semyon is late.[30]

It is easy to make the intuitive prediction using our semantics. We say that the demonstrative from (64) is essentially equivalent to:

(65) the x: [guy-in-a-poncho(x) & identical-to-Semyon(x)]

Because this representation involves the property of being identical to Semyon, there is no chance that our demonstrative will pick out some other individual at some other world; if the demonstrative picks out anything anywhere, it picks out Semyon.[31] If we say that the identificational property is supplied by means of a variable over individuals—perhaps one that is type-raised along the lines suggested here in section 4.1—our formalism will make clear why this would be: individual variables are not sensitive to the permutations of the world of evaluation that are wrought by modal operators.

The fact that our analysis involves the idea that demonstratives are a special kind of definite description, however, will likely make some philosophers uneasy. Definite descriptions are commonly supposed to give rise to what are known as 'scope ambiguities'. Consider the following example:

[30] As far as I can tell, none of the details concerning the semantics of counterfactuals affect any of the claims I make here.

[31] Compare the notion of weak rigidity from Devitt (2004).

(66) I could have had lunch with the president.

(66) appears to admit both of the following two paraphrases:

(67) The president is such that there is an accessible world in which I have lunch with him.

(68) There is an accessible world in which I am having lunch with whichever person is the president at that world.

It is standard practice to explain these two readings by saying that sentences like (66) are ambiguous at the level of semantic representation; (67) is the result of treating the definite description as though it has scope over the modal operator, while (68) is the result of treating the definite description as thought it scopes under the operator.[32]

If demonstratives are semantically similar to descriptions in the way we have proposed here, as long as other things are equal, we should expect them to occur in both scope configurations. In fact, we have already looked at a case that suggests that they *do* occur in both positions. In the previous section, we saw how non-deictic demonstratives are most naturally interpreted non-rigidly. As with definite descriptions, we allow the extension of a non-deictic demonstrative to vary across possible worlds by embedding the expression under a modal operator. If the non-deictic demonstratives takes wide scope with regard to the operator, on the other hand, the result is a rigid reading. For non-deictic demonstratives, the narrow scope readings are the most natural, but it is not hard to hear that both are available.

In the case of deictic demonstratives, tracking the possibilities is less straightforward. As we have seen, our intuitions suggest that deictic demonstratives admit only rigid readings. But, on the view we have developed here, that is what we would expect regardless of the scope options, since deictic demonstratives are formed from identificational properties. Consider (64) and (65) again:

(64) **That guy in the poncho** might have been late.

(65) the x: [guy-in-a-poncho(x) & identical-to-Semyon(x)]

As long as we use (65) to interpret the demonstrative from (64), the two scope paraphrases will come out truth-conditionally indiscernible in most cases:

(69) The guy in the poncho who is identical to Semyon is such that there is an accessible world in which he is late.

[32] The classic source on the subject is Russell (1905/2005). For more recent alternatives, compare Percus (2000) and Keshet (2010).

(70) There is an accessible world in which the guy who is wearing a poncho and who is identical to Semyon is late.

With regard to worlds in which Semyon is wearing a poncho, (69) and (70) amount to the same thing. The question about scope possibilities is still an important one, though, because unless Semyon's relationship with his poncho is a basic fact of metaphysics, we should not expect him to be wearing it at every world.

Consider the following example:

(71) **That guy in the poncho** could have worn a jacket instead.

Intuitively, this sentence should be true in a context like the one described above just in case there is an accessible world in which Semyon is wearing a jacket instead of a poncho. We have no trouble generating these truth-conditions using (65); we simply say that the demonstrative takes wide scope with regard to the modal. But if both scope positions are structurally available, we might expect there to be another reading of the sentence, too. We might expect there to be a reading on which the sentence is either false or 'gappy', depending on whether the description is understood along Russellian or Fregean lines. Other things being equal, that is, we might expect (71) to admit both of the following paraphrases:

(72) The guy in the poncho who is identical to Semyon is such that there is an accessible world in which he is wearing a jacket. (OK)

(73) There is an accessible world in which the guy in the striped poncho who is identical to Semyon is wearing not a poncho but a jacket. (Contradiction!)

It is very difficult, however, to hear (71) as expressing anything but the straightforwardly contingent proposition expressed by (72). Is this a problem for our semantics? Are we committed to generating a class of 'missing' defective readings for sentences like (71)?

The answer to both questions is 'no'. First of all, it is important to remember that standard practice in the industry is to guarantee the rigidity of demonstrative expressions by fiat. The operator that is at the heart of Kaplan's formal system is precisely a rigidifying operator, and both Elbourne and King stipulate that when a complex demonstrative is used deictically, the predicate from which it is formed is evaluated with regard to the world of the context. If we treat the semantic proposal described here as a modification of existing versions of the hidden argument theory, then, and if we follow those proposals with regard to the question of which worlds should be used to evaluate the predicative material associated with 'that', we can guarantee that our proposal will fare just as well as those do.

Even without building such a stipulation into the semantics, however, I think we can make good sense of the intuitive intensional data.[33] In general, when a particular construction admits a felicitous reading, it is notoriously difficult to determine whether it admits defective alternatives, too, and arguments based on the existence of such alternatives must be taken with a significant dose of salt. Even if we set this point of method aside for the sake of argument, though, there is a convincing explanation of the absent narrow scope readings. In fact, the scope possibilities licensed by deictic and non-deictic demonstratives parallel the possibilities licensed by definite descriptions.

Rothschild (2007) employs the following data, among others, to show that definite descriptions do not uniformly admit two scope possibilities with regard to modal operators:[34]

(74) Mary-Sue could have been married to the president.

(75) Hans might have been the person I talked to the whole time.

(74) clearly admits both scope options. If Grover Cleveland were president, someone could use the sentence to say that there is an accessible world in which Mary-Sue marries Cleveland. Alternatively, (74) could be used to make a claim about how well Mary-Sue does in high-stakes social events; it can easily be understood to mean that there is an accessible world in which she is married to whoever happens to be the president at that world.

(75), on the other hand, seems to admit only a reading on which the description takes wide scope with regard to the modal. For example, it is easy to imagine relying on the availability of the wide-scope reading to express an epistemic claim. If someone notices that I spent the entire costume party talking to a particular individual and asks who it was, I might answer with (75). I might say use the sentence, in other words, to say that:

(76) The person I talked to is such that for all I know, he might have been Hans.

As Rothschild points out, however, if it was in fact the case that I divided my attention at the party equally among the guests, (75) sounds bizarre. There is no reading of the sentence on which it expresses the idea that there is a world accessible from the context in which I spent all my time talking to a single person, Hans. This is a surprising result, however. If a narrow-scope reading of the description ‘the person

[33] It is important to note that if the strategy I describe here works, it will work for existing versions of the hidden argument theory, as well

[34] These examples are numbered (1) and (5) in Rothschild (2007), and are found on pages 71 and 78.

I talked to the whole time' were available, we would expect (75) to serve as a vehicle for exactly that proposition. On standard assumptions about the way parties go, such a world should be accessible in most contexts. So why can we not hear (75) this way?

Rothschild's answer is that a difference in the presuppositions associated with the two descriptions is responsible for the different scope possibilities.[35] The description from (74), 'the president', is an example of what he calls a 'role-type' description; in normal conversational situations, it will be part of the common ground that at any world of evaluation, a unique individual will satisfy this description. The description from (75), on the other hand, is 'particularized'; unless the conversation has provided a specific reason to think otherwise, it will not typically be assumed that the description 'the person I talked to the whole time will pick out anyone at all. As Rothschild puts things, "we naturally think that over a relevant set of possibilities" (2007: 93) there will be a unique president in each, while we have no reason to expect there to be a unique person I spent the whole party talking to.

When we go to evaluate the claim made by (74), it is easy to see that the uniqueness presupposition introduced by 'the' is met, whether we interpret the description with wide or narrow scope.[36] When we go to evaluate (75), however, we are left to confront a stark contrast. If we take the description to have wide scope, we can accommodate the presupposition introduced by 'the' relatively straightforwardly, by assuming that there must in fact have been a unique individual the speaker spent the party with.[37] The accommodation that would be required to license the narrow-scope reading, on the other hand, is (usually) a bridge too far. Unless it is somehow clear—maybe because of the history of the conversation, maybe because of mutual knowledge of the way I typically allocate my time at parties—that the counterfactual possibilities we are countenancing are all possibilities in which there is just one person I talk to, the uniqueness presupposition will not be met across the range of worlds that are live options, and the description will result in infelicity.

Our data involving deictic and non-deictic demonstratives exhibit just the pattern of readings Rothschild's data do, and they appear to admit of the same explanation. Non-deictic demonstratives, like the following (repeated), license both wide-scope

[35]In an earlier version of this paper, I sketched a pragmatic explanation broadly similar to the one described by Hawthorne and Manley (2012). I took as my point of departure something Lepore and Ludwig (2000) point out: that the predicate from which a complex demonstrative is formed is generally assumed to be one that listeners can use to identify the object in question. Properties objects have in counterfactual scenarios are not properties that help people identify objects in the actual world; imagine someone saying *Quick, pull on the lever that could have been red!* In future work, I hope to further develop this line.

[36]Rothschild treats 'the' as though it introduces a presupposition about salience, but this does not materially affect the applicability of his idea in the present connection, as far as I can tell.

[37]That accommodation can occur here is highlighted by the fact that it is natural to imagine someone responding to (75) by saying *I didn't realize you'd talked to just one person*.

and narrow-scope interpretations:

(62) If Åsa$_i$ had won the election, she$_i$ would definitely have embraced **that elector who cast the deciding vote.**

The demonstrative 'that elector who cast the deciding vote' is a paradigmatic role-type description: in the context of a conversation about elections, we naturally assume that in any given counterfactual scenario, there will be a unique person who cast the deciding vote.[38]

Deictic demonstratives, on the other hand, are the limit case of particularized descriptions. Since deictic demonstratives involve an identificational property, the only way to coerce a role reading is to change the semantic type of the demonstrative, as arguably takes place when you point at a person doing something no one would want to do and say:

(77) Don't be **that guy**.

If we take a typical deictic demonstrative, like (71, repeated):

(71) **That guy in the poncho** could have worn a jacket instead.

we can see how a presupposition failure comes into play to explain the 'missing' narrow scope reading. The presuppositions introduced by 'the guy who is wearing a poncho and who is identical to Semyon'—a rough paraphrase of the sort of thing we use to analyze the demonstrative—are only met at worlds where there is a unique poncho-wearing Semyon. But those worlds hardly form a natural class, much less the class the typical conversational participants would have in mind when encountering (71).

It is instructive to compare (71) with the following variation involving a description:

(78) The guy in the poncho could have worn a jacket instead.

(78) admits exactly the scope possibilities (71) does. If it is uttered in the context we described earlier, it can be used to say that Semyon could have worn a jacket. But there is no defective reading of the sentence, i.e., no reading on which it means that there is an accessible world in which there is a unique guy in a poncho who is wearing a jacket instead of a poncho. No one will claim that this 'missing' reading

[38] Of course, there may be metaphysical or epistemological questions about what it takes to be the person who cast the deciding vote. I do not mean for the discussion to turn on them, and ask people to substitute a charitable alternative demonstrative if they find that this one raises such worries.

impugns the status of 'the guy in the poncho' as a description, and no one should believe that the absence of a defective reading tells against the idea that 'that guy in the poncho' is a kind of description, either.

6 Outstanding issues

I hope that the ground covered so far will be enough to show that the challenge of explaining the distribution of non-deictic interpretations for demonstratives is a challenge that must be taken seriously, and that the approach sketched here deserves to be counted as a live option. Turning that approach into a fully-developed theory, of course, will involve a significant amount of further work.[39] To be upfront about the extent to which the present material constitutes a first, rather than a last installment, I close with a sketch of some worries the project raises.

6.1 'Only': problem 1

An anonymous referee points out that sentences like the following appear to be a problem for our theory:

(79) Every candidate that receives exactly one vote$_i$ will later contact **that unique voter who voted for her$_i$**.

(80) The resume is **that lone document that can make or break our chances of getting a job**.

(81) Fear is **that sole thing that keeps us from taking the very risks that make us who we are**.[40]

The demonstratives from these sentences can clearly be interpreted non-deictically. If the analysis developed here is correct, the relative clauses 'who voted for her', 'that can make or break…', and 'that keeps us from…' are attached high, and thus do not form constituents with the nouns they modify. If 'unique', 'sole', 'lone', and similar adjectives require that their complements have a cardinality of one, there is a difficulty here. Without the restriction provided by the relative clause, none of

[39] In addition to the worries mentioned below, questions about plurals, the distal/proximal distinction, and construction-specific constraints on anaphora would have to be settled. A full treatment would presumably also have to speak to the question of whether high-adjoined relative clauses are available in general, or whether the construction is one specific to demonstratives. I hope to be able to address these and other questions going forward.

[40] Examples (80) and (81) are attested; compare http://www.nwnit.com/nwn-blog/dont-fear-the-resume/ and http://www.blogher.com/little-white-dress. Thanks to the referee for these data.

these nouns should pick out a singleton set, and we should expect a presupposition failure.

I agree that these data constitute a puzzle that needs to be solved. I do not think, however, that they amount to a clearly fatal flaw for the view described so far.

For one thing, the challenge posed by sentences like these is really best presented as a dilemma. If the arguments I have made here are successful, regardless of how exactly the details of the compositional semantics go, something very like the NP-S configuration for relative clause will be required to explain the distribution of non-deictic interpretations for demonstratives. As long as the relative clause and the expression it modifies form a constituent, the determiner will take the two as a single argument, with nothing to distinguish the result from the illicit counterpart construction formed from the relational genitive. Any structure that separates the relative clause and the constituent it modifies, however, will separate the material we would expect to be required to hang together for 'sole', etc. to apply to.

The fact that adjectives from the target class are somewhat mysterious to begin with provides another reason to hesitate before deciding to reject the current proposal on the basis of data like (80)–(81). As the following examples from Rothschild (2006a,b) show, adjectival 'only', which on the face of things would seem to pattern with 'sole', 'unique', 'lone', and similar, licenses negative polarity items:

(82) a. *I saw the man who'd ever gone to Universal Studios.
b. I saw the only man who'd ever gone to Universal Studios.

We get the same result if we replace the definite description with a non-deictic demonstrative:

(83) a. *I saw that man who'd ever gone to Universal Studios.
b. I saw that only man who'd ever gone to Universal Studios.

Although 'unique', 'lone', and 'sole' sound stranger than 'only' in the environment in question, there is still a clear contrast in the following set of pairs:

(84) a. *I saw the/that man who'd ever gone to Universal Studios.
b. ??I saw the/that unique man who'd ever gone to Universal Studios.

(85) a. *I saw the man who'd ever gone to Universal Studios.
b. ?I saw the/that lone man who'd ever gone to Universal Studios.

(86) a. *I saw the man who'd ever gone to Universal Studios.
b. ?I saw the/that sole man who'd ever gone to Universal Studios.

The fact that any of these items should license NPIs is surprising. If we stick with the orthodox view of 'the' (or an analogous treatment of 'that'), on which the determiner introduces some form of commitment to uniqueness, none of the obvious candidate semantic values for adjectives like 'unique' should make any difference to the NPI possibilities.[41] If, on the other hand, we give up the idea that the determiners themselves introduce restrictions on the cardinality of their complements, we might open up some space to explain the contrast in the pairs given above. But many philosophers will balk at the cost of breaking with a long-established and generally successful research tradition.

There is no room here to consider all the data that might tell in favor of one or another possible analysis of 'unique', 'lone', 'sole', and so on, or to consider treatments of definites that abandon the idea of uniqueness.[42] For our purposes, what is important to notice is that there are significant open questions in the vicinity. While examples like (79)–(81) should certainly be taken as a constraint on a theory of demonstratives, data involving demonstratives might just as well be taken to constrain future work on the adjectives in question. Once all the considerations are in, it might turn out that those adjectives have their cardinality requirements met nonlocally, or those requirements might be different from they way they appear at first inspection.[43]

6.2 'Only': problem 2

Nathan Klinedinst points out a different problem involving cardinality. The following discourse is perfectly natural:

(87) I love cats. All cats. If there were only one cat left on Earth, I would find **that cat** and adopt it.

The demonstrative from (87) is used anaphorically, so we might assume that the hidden argument place is occupied by a variable over individuals or identificational properties. Ordinarily, that would put the restrictor in a position to perform a nontrivial operation on the predicate 'cat'. This example, however, is structured so that $\{x : x \text{ is a cat}\}$ is already a singleton, which means it can admit of no further restriction. On our analysis, this should result in a presupposition failure.

[41] Suppose we say, for example, that these adjectives perform a test on their complements, returning the semantic value of the complement if the complement is one cardinal, and crashing the derivation otherwise. That seems like a reasonable first hypothesis about the way 'only' or 'unique' might work. But this semantics offers no obvious explanation of the NPI facts.

[42] Compare Heim (1982), Szabó (2000, 2005), and Roberts (2003).

[43] Of course, it might also turn out that the NPI data can be used to mount an argument against the kind of view developed here. If the relative clause from (83) is outside of the scope of 'only', as it would appear to be, we may have to face a new mystery about how the NPI could be licensed.

If we want to employ the kind of anti-uniqueness presupposition we have relied on here to explain the pattern in the data involving non-deictic demonstratives, we will have to come up with a compelling story about why demonstratives from strings like (87) do not cause a crash. Although I am not in a position to offer such a story now, I hope to be able to do so in the future. For now, let me simply note that there are enough subtleties in the relevant data to make a closer look seem warranted.

Consider, for example, the following variation on (87):

(88) *I love cats. All cats. If there were only one cat left on Earth, I would find **the cat** and adopt it.

If 'the cat' were evaluated with regard to the local context of the antecedent, then other things being equal, we would expect there to be no problem with this description. After all, in the situation described, there is just one cat. 'The cat', however, is infelicitous in (88). One prominent and plausible explanation of this fact invokes the idea of competition between 'the' and 'it'. Contrast the following variation on the example:

(89) I love cats. All cats. If there were only one cat left on Earth, I would find **it** and adopt it.

The idea that the problem with (88) is due to competition with the pronoun is reinforced by the acceptability of the following:

(90) Imagine that there were only two domestic animals left on Earth, a cat and a dog. I would find **the cat** and adopt it.

When there are alternative candidate antecedents available for the pronoun, the definite description is permitted, presumably because it has a job to do in helping to distinguish between them.[44]

If the definite description from (89) is licensed by the fact that the predicate from which it is formed has a substantial role to play in distinguishing one from among the candidate interpretations, and if 'that' does not involve an anti-uniqueness presupposition, we should expect it to be interchangeable with 'the' in this example. That, however, is not what we find:

(91) #Imagine that there were only two domestic animals left on Earth, a cat and a dog. I would find **that cat** and adopt it.

[44]Compare Schlenker (2005), Johnson (2013).

The markedness of (91), and indeed, the contrast between (87), Klinedinst's perfectly natural cat sentence, and (88), the marked variation formed from a definite description, suggests that there is work to be done here—standard accounts should lead us to expect definite descriptions and demonstratives to stand or fall together. While I do not mean to suggest that these data point towards a clear explanation of how (87) might escape the anti-uniqueness presupposition that has been our focus here, I hope the complexities in the data will make the search for an answer appear worthwhile.

Along similar lines, an anonymous referee observes that a class of putative counterexamples can be developed that do not rely on anaphora:

(92) That rhino is the only one still alive.

(93) That atomic bomb is the last one remaining.

These sentences are unimpeachable, but they appear to violate the presupposition I have claimed is introduced by 'that', since in the context as specified by (92), there is only one rhino left, and in the context as specified by (93), there is only one atomic bomb left. As in the case of the other sorts of problematic data canvassed thus far, I am not sure exactly what to say about these sentences. I take them to raise an important question about where exactly the presuppositions introduced by the demonstrative determiner are checked. For now, I take that question to be an open one, and a question that should not prevent us from moving forward in developing the restrictor theory of 'that' further.

Although the point has received less theoretical attention than it deserves, it has long been recognized that restrictive relative clauses and close appositives cannot be used to modify expressions that already pick out a singleton. So, for example:

(94) #The author of *Waverley* who wrote *Ivanhoe* was Scott.

(95) #Frege who wrote the *Begriffschrift* worked for Zeiss.

Now, if the mechanism is responsible for this effect were a general one, as it appears to be, and if the cardinality of 'the rhino' and 'the atomic bomb' were checked with regard to the contexts specified by (92) and (93), we would expect the following pair to be marked in the way (94) and (95) are:

(96) The rhino that you see before you is the only one still alive.

(97) The atomic bomb that you see before you is the last one remaining.

The fact that (96) and (97) are perfectly felicitous suggests that the predicates 'rhino' and 'atomic bomb' are treated as though they did not pick out singletons—if they did, the restrictive relative clause 'that you see before you' would be forbidden.

Moving from this observation to a defense of the restrictor view developed here will no doubt involve substantial further work. At first glance, however, I take data like these to provide a set of constraints that we will have to draw on in realizing that work, as opposed to reasons to think that it will not be possible.

7 The upshot

We began this paper with two empirical questions:

a. What licenses a non-deictic reading for the demonstrative from (4)?

(4) **That guy who wrote *Waverley*** also wrote *Ivanhoe*.

b. What rules out such a reading for the demonstrative from (5)?

(5) #**That author of *Waverley*** also wrote *Ivanhoe*.

So far, we have seen how leading semantic proposals fail to provide a satisfactory answer to those questions, and we have developed the outlines of an alternative that does better. If we supplement the classical direct reference story about demonstratives with the idea that some instances of 'that' mean what 'the' means, we end up with no way of ruling out (5). The same problem obtains if we think that 'that' takes any two property-type arguments and returns their unique satisfier. If we claim, on the other hand, that 'that' takes two property-type arguments and requires that the second non-trivially restrict the first, we can make the right predictions about our problematic data.

In addition to the fact that our alternative builds on the intuitively plausible idea that saying 'that *F*' amounts to saying something about one from among the *F*s, it shows that the characteristic behavior of demonstratives can be explained using what we might think of as 'off-the-shelf' components, instead of the special operators that have typically been employed. The idea that determiners check the cardinality of their complements, for example, is familiar from standard treatments of 'the', 'many', 'both', and so on, while the crucial structural innovation we rely on—the NP-S analysis of relative clause—was developed for independent reasons. While there is clearly more work to be done, I hope the present will be taken to give at least a sense of one of the forms that work might take.

References

Adger, D. (2012). *A Syntax of Substance*. MIT Press.

Bach, E. and Cooper, R. (1978). The NP-S analysis of relative clauses and compositional semantics. *Linguistics and Philosophy*, 2(1):145–150.

Barker, C. (1995). *Possessive Descriptions*. Stanford: CSLI.

Barker, C. (2011). Possessives and relational nouns. In Maienborn, C., von Heusinger, K., and Portner, P., editors, *Semantics: An International Handbook of Natural Language Meaning*, volume 2, pages 1109–30. Mouton de Gruyter.

Barwise, J. and Perry, J. (1983). *Situations and Attitudes*. MIT Press.

Braun, D. (2008). Complex demonstratives and their singular contents. *Linguistics and Philosophy*, 31:57–99.

Chao, Y. R. (1968). *A Grammar of Spoken Chinese*. University of California Press.

Comorovski, I. (2007). Constituent questions and the copula of specification. In Comorovski, I. and von Heusinger, K., editors, *Existence: Semantics and Syntax*, pages 49–78. Springer.

Cooper, R. (1975). *Montague's Semantic Theory and Transformational Syntax*. PhD thesis, University of Massachussetts.

Corrazza, E. (2003). Complex demonstratives *qua* singular terms. *Erkenntnis*, 59:263–283.

Davies, M. (1982). Individuation and the semantics of demonstratives. *Journal of Philosophical Logic*, 11(287-310).

del Gobbo, F. (2003). *Appositives at the Interface*. PhD thesis, University of California at Irvine.

Dever, J. (2001). Complex demonstratives. *Linguistics and Philosophy*, 24(3):271–330.

Devitt, M. (1981). *Designation*. Columbia University Press.

Devitt, M. (2004). The case for referential descriptions. In Reimer, M. and Bezuidenhout, A., editors, *Descriptions and Beyond*, pages 234–260. Clarendon Press.

Elbourne, P. (2005). *Situations and Individuals*. MIT Press.

Elbourne, P. (2008). Demonstratives as individual concepts. *Linguistics and Philosophy*, 31(4):409–466.

Georgi, G. (2012). Reference and ambiguity in complex demonstratives. In Kabasenche, W. P., O'Rourke, M., and Slater, M. H., editors, *Reference and Referring*, pages 357–384. MIT Press.

Glanzberg, M. and Siegel, S. (2006). Presupposition and policing in complex demonstratives. *Noûs*, 40(1):1–42.

Hawthorne, J. and Manley, D. (2012). *The Reference Book*. Oxford University Press.

Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases*. PhD thesis, University of Massachussetts, Amherst.

Heim, I. and Kratzer, A. (1998). *Semantics in Generative Grammar*. Blackwell.

Huang, C.-T. J. (1982). *Logical Relations in Chinese and the Theory of Grammar*. PhD thesis, MIT.

Johnson, K. (2013). Pronouns vs. definite descriptions. In Becker, M., Grinstead, J., and Rothman, J., editors, *Generative Linguistics and Acquisition: Studies in honor of Nina M. Hyams*, number 54 in Language Acquisition and Language Disorders, pages 157–184. John Benjamins.

Kaplan, D. (1977). Demonstratives. In Almog, J., Perry, J., and Wettstein, H., editors, *Themes from Kaplan*, pages 481–563. Oxford University Press.

Kaplan, D. (1989). Afterthoughts. In Almog, J., Perry, J., and Wettstein, H., editors, *Themes from Kaplan*, pages 565–614. Oxford University Press.

Kayne, R. S. (1994). *The Antisymmetry of Syntax*. MIT Press.

Keshet, E. (2010). Split intensionality. *Linguistics and Philosophy*, 33(4):251–283.

King, J. C. (1999). Are complex 'that' phrases devices of direct reference? *Noûs*, 33(2):155–182.

King, J. C. (2001). *Complex Demonstratives: A Quantificational Account*. MIT Press.

King, J. C. (2008). Complex demonstratives, QI uses, and direct reference. *Philosophical Review*, 117(1):99–117.

Lakoff, R. (1974). Remarks on 'this' and 'that'. *Proceedings of the Chicago Linguistics Society*, 10:345–356.

Larson, R. (1982). A note on the interpretation of adjoined relative clauses. *Linguistics and Philosophy*, 5:473–482.

Larson, R. and Segal, G. (1995). *Knowledge of Meaning*. MIT Press.

Lepore, E. and Ludwig, K. (2000). The semantics and pragmatics of complex demonstratives. *Mind*, 109:199–240.

Lin, J.-W. (2003). On restrictive and non-restrictive relative clauses in Mandarin Chinese. *The Tsing Hua Journal of Chinese Studies*, 33(1):199–240.

Lyons, C. (1986). The syntax of english genitive constructions. *Journal of Linguistics*, 22(1):123–143.

Maclaran, R. (1982). *The Semantics and Pragmatics of English Demonstratives*. PhD thesis, Cornell Univesity.

Neale, S. (1993). Term limits. *Philosophical Perspectives*, 7:89–123.

Nowak, E. (2014). Demonstratives without rigidity or ambiguity. *Linguistics and Philosophy*, 37(5):409–436.

Nowak, E. (2018). Really complex demonstratives.

Nowak, E. (2019). No context, no content, no problem. *Mind and Language*.

Partee, B. (1975). Montague grammar and transformational grammar. *Linguistic Inquiry*, 6(2):203–300.

Partee, B. (1986). Noun phrase interpretation and type-shifting. In Groenendijk, J., de Jongh, D., and Stokhof, M., editors, *Studies in discourse representation theory and the theory of generalized quantifiers*, chapter 5, pages 115–143. Walter de Gruyter.

Partee, B. (2006). A note on Mandarin possessives, demonstratives, and definiteness. In Birner, B. J. and Ward, G., editors, *Drawing the Boundaries of Meaning: Neo-Gricean Studies in pragmatics and semantics in honor of Laurence R. Horn*, volume 80, pages 263–280. John Benjamins.

Partee, B. and Borschev, V. (2003). Genitives, relational nouns, and argument-modifier ambiguity. In Lang, E., Maienborn, C., and Fabricius-Hansen, C., editors, *Modifying Adjuncts*, pages 67–112. Mouton de Gruyter.

Percus, O. (2000). Constraints on some other variables in syntax. *Natural Language Semantics*, 8(3):173–229.

Peters, S. and Westerståhl, D. (2013). The semantics of possessives. *Language*, 89(4):713–759.

Potts, C. and Schwarz, F. (2010). Affective 'this'. *Linguistic Issues in Language Technology-LiLT*, 4(5):1–30.

Quine, W. (1960). *Word and Object*. MIT Press.

Reimer, M. (1991). Demonstratives, demonstrations, and demonstrata. *Philosophical Studies*, 63:187–202.

Roberts, C. (2002). Demonstratives as definites. In van Deemter, K. and Kibble, R., editors, *Information Sharing*. Stanford: CSLI.

Roberts, C. (2003). Uniqueness in definite noun phrases. *Linguistics and Philosophy*, 26:287–350.

Ross, J. R. (1967). *Constraints on Variables in Syntax*. PhD thesis, MIT.

Rothschild, D. (2006a). Definite descriptions and negative polarity. *Manuscript*.

Rothschild, D. (2006b). Non-monotonic NPI-licensing, definite descriptions, and grammaticalized implicatures. In Gibson, M. and Howell, J., editors, *Proceedings of SALT*, volume XVI, pages 228–240. Cornell University Press.

Rothschild, D. (2007). Presuppositions and scope. *The Journal of Philosophy*, 104(2):71–106.

Russell, B. (1905/2005). On denoting. *Mind*, 114(456):873–887.

Salmon, N. (2002). Demonstrating and necessity. *Philosophical Review*, 111(4):497–537.

Salmon, N. (2006). Terms in bondage. *Philosophical Issues*, 16:263–274.

Salmon, N. (2008). That *F*. *Philosophical Issues*, 141:263–280.

Sauerland, U. (2003). Unpronounced heads in relative clauses. In Schwabe, K. and Winkler, S., editors, *The Interfaces: Deriving and Interpreting Omitted Structures*, pages 205–226. John Benjamins.

Schlenker, P. (2005). *Minimize Restrictors!* (notes on definite descriptions, condition C, and epithets. In Maier, E., Bary, C., and Huitnik, J., editors, *Proceedings of Sinn und Bedeutung*, volume 9, pages 385–416.

Stockwell, R. P., Schachter, P., and Partee, B. (1973). *The Major Syntactic Structures of English*. Holt, Rinehart and Winston.

Swanson, E. (2005). Pronouns and complex demonstratives. *Manuscript*.

Szabó, Z. G. (2000). Descriptions and uniqueness. *Philosophical Studies*, 101:29–57.

Szabó, Z. G. (2005). The loss of uniqueness. *Mind*, 114(456):1185–1222.

Vendler, Z. (1967). *Linguistics in Philosophy*. Cornell University Press.

Vikner, C. and Anker, P. (2002). A semantic analysis of the English genitive. *Studia Linguistica*, 56:191–226.

Wettstein, H. (1984). How to bridge the gap between meaning and reference. *Synthese*, 84:63–84.

Wolter, L. (2003). Demonstratives, definite descriptions, and definiteness.

Wolter, L. (2006). *That's that; the Semantics and Pragmatics of Demonstrative Noun Phrases*. PhD thesis, University of California at Santa Cruz.

Wolter, L. (2007). Situation variables and licensing by modification in opaque demonstratives. *Proceedings of Sinn und Bedeutung*, 11:612–625.

Wolter, L. (2009). Demonstratives in philosophy and linguistics. *Philosophy Compass*, 4:451–468.

Yang, H. S.-F. (2005). *Plurality and Modification in Mandarin Nominal Phrases*. PhD thesis, University of Texas at Austin.