

# PROCEDURAL MEANING AND DEFINITE DESCRIPTIONS\*

RAMIRO CASO

Universidad de Buenos Aires - CONICET - GAF

ramirocaso05@yahoo.com.ar

## Abstract

The present work explores the possibility of conciliating the truth-conditional relevance of referential uses of definite descriptions with the assignment of a univocal linguistic meaning to these constructions. It is argued that conciliation is possible if we reject the thesis, central to the debate between Russellians and ambiguity theorists, according to which referential uses are truth-conditionally relevant if and only if they constitute referential meanings. We sketch a framework within which the denial of that thesis has theoretical content, by drawing on the conceptual resources of Relevance Theory and on a pragmatic conception of reference, following Strawson (1950). The linguistic meaning of definite descriptions is analyzed as a procedural meaning (Blakemore 1987) that is semantically underdetermined with respect to both referential and attributive readings, and a pragmatic strategy for understanding this ambiguity is sketched.

KEY WORDS: Semantic underdeterminacy; Truth-conditional pragmatics; Speaker reference.

## Resumen

El presente trabajo explora una posible conciliación de la relevancia para las condiciones de verdad de los usos referenciales de descripciones definidas con la postulación de un significado lingüístico unívoco para tales construcciones. Se argumenta que dicha conciliación es posible si se abandona la tesis, central en el debate entre russellianos y teóricos de la ambigüedad, según la cual los usos referenciales son relevantes para las condiciones de verdad si y sólo si constituyen significados referenciales. Se esboza un marco dentro del cual la negación de esta tesis tiene contenido teórico, apelando a los recursos conceptuales de la Teoría de la Relevancia y a una concepción pragmática de la referencia en la línea de Strawson (1950). Se analiza el significado de las descripciones definidas como un significado procedimental (Blakemore 1987), semánticamente subdeterminado respecto de una lectura referencial o atributiva, y se esboza una manera pragmática de entender esta ambigüedad.

PALABRAS CLAVE: Subdeterminación semántica; Pragmática de condiciones de verdad; Referencia del hablante.

1. Two questions a theory of definite descriptions should answer are the following:

\* Special thanks are due to Michael Devitt, Eleonora Orlando, Juan Manuel Berros, Justina Díaz Legaspe, and Laura Skerk for their insightful comments.

- (i) Are referential uses truth-conditionally relevant? That is, given a sentence  $S$  containing a definite description  $d$ , ( $Sd$ ), do the truth conditions of an utterance of ( $Sd$ ) vary according to whether  $d$  is used referentially or attributively?
- (ii) Do referential uses constitute referential meanings? That is, is the determiner (or the description as a whole) lexically ambiguous?

Ambiguity theorists, such as Michael Devitt (2004, 2007a,b), answer both questions affirmatively: referential uses are truth-conditionally relevant, and they constitute referential meanings. Russellians, on the other hand, such as early Neale (1990), answer both questions negatively: definite descriptions have only quantificational meanings, and referential uses are truth-conditionally irrelevant. Following some of the insights of truth-conditional pragmatics, I would like to advance the idea that it is possible to combine the truth-conditional relevance of referential uses with a univocal account of descriptions.

What both Russellians and ambiguity theorists have in common is that they accept the following biconditional:

- (1) Referential uses of definite descriptions are truth-conditionally relevant if and only if they constitute referential meanings

Actually, it is the acceptance of (1) what forces uniform answers to questions (i) and (ii) above. However, this acceptance can be contested, and thus non-uniform answers to those questions are possible. In particular, we'll try to give some plausibility to the joint claim of (2) and (3):

- (2) Referential uses of definite descriptions are truth-conditionally relevant
- (3) Definite descriptions are lexically univocal

The rejection of (1) makes sense only under a specific account of the relation between truth-conditions and pragmatics: one according to which truth-conditions are partly pragmatically constructed; that is, they are determined both by a linguistically encoded meaning and by the occurrence of pragmatic processes operating on that meaning. In adopting this stance, I'm in line, *modulo* particular differences, with authors such as Bach (1999, 2001), Carston (1988), Neale (2004), Powell (2001b), Recanati (1993), and Sperber and Wilson (1995), among others. The position that I want to present here presupposes then a non-traditional notion of semantic content:

one more closely related to the notion of linguistically encoded, context-independent meaning than to the notion of truth-conditions. Once this conception is accepted, and the distinction between semantics and pragmatics is re-drawn accordingly, it is possible to propose a univocality account of definite descriptions that nonetheless makes referential uses, though pragmatically explained, truth-conditionally relevant.

**2.** But first we have to take a look at the scope of Devitt's main argument against pragmatic explanations of referential uses of definite descriptions in (2007a). It purports to be an argument against all pragmatic explanations, insofar as "a pragmatic explanation of referential uses *must be based* on the view that a person using a description referentially in uttering a sentence conveys a singular proposition *while saying a general quantificational proposition*" (see 2007b, p. 49, and also 2007a, p. 18; the first emphasis is ours).

As I pointed out in (unpublished), I think this argument is sound against Gricean-Russellian positions. Now, this is not equivalent to saying that every pragmatic explanation lies within its range, since not every such explanation is coupled with a Gricean-Russellian core, as shown by Recanati's (1989), Powell's (2001b), and our own. So, at least in principle, univocality theories of definite descriptions not subject to Devitt's main argument seem possible.

In what follows, we'll attempt to sketch an explanation of referential uses which (a) is pragmatic in nature, and (b) doesn't seem to be subject to Devitt's main criticism of Gricean-Russellian positions. As we've mentioned, this criticism seems to rest crucially on the idea that any pragmatic explanation must do two things: first, find a suitable general proposition corresponding to the Russellian analysis of definite descriptions; second, derive the intended, singular proposition corresponding to the referential use, by means of Gricean machinery (or some other sort of pragmatic processes). The account we'll propose is not committed to either of these: insofar as it is an account that posits a linguistically encoded meaning for descriptions that is neither referential nor attributive, it does not assign to sentences containing such expressions any proposition as a matter of semantic value and, in particular, no general proposition from which to derive a singular one. As a consequence, it does not rely on a Gricean mechanism to obtain the intuitive truth-conditions of referential uses, which are determined by pragmatic processes of quite a different sort. Section 3 briefly addresses the issue of spelling out a univocal linguistic meaning for descriptions. Section 4 is an attempt to sketch a pragmatic explanation of reference determination.

3. To fix ideas, I'll be adopting a broadly relevance-theoretic framework.<sup>1</sup> On another occasion (unpublished) I suggested that we could understand referential uses as providing enough clues so as to enable the hearer to successfully identify the intended referent. We can generalize this idea, and think of definite descriptions in general as providing clues that help in the identification of referents. More technically, and drawing heavily upon the work of Powell (2001b) and others, we can say that a definite description linguistically encodes a *procedure for reference identification*. As Blakemore (1987) shows, it makes sense to distinguish between two kinds of meaning within the framework of Relevance Theory (RT): conceptual and procedural meaning.<sup>2</sup> Allowing for this distinction, we can say that definite descriptions encode both types of meaning. The conceptual meaning encoded by a description *d* is composed of both the idea of uniqueness of reference (given by the presence of a determiner plus a singular noun phrase), and the properties provided by the descriptive part of *d*. The procedural meaning encoded by *d* gives "instructions" on how to use this conceptual meaning in reference identification. Roughly put, this procedure amounts to something like: "look for an object *x* such that ( $\varphi x$ ) and satisfies the expected relevance".<sup>3</sup>

Obviously, we have to say something about this, since as it stands it is both obscure and inaccurate. When I advanced in (unpublished) the idea that reference assignment to definite descriptions ought to take into account the representation conversational participants have of the context of utterance (what I'd like to call, following RT, their *mutual cognitive*

<sup>1</sup> For a full exposition of Relevance Theory, see Sperber and Wilson (1995). For a concise one, the reader may check Wilson and Sperber (2002) and Carston (2004).

<sup>2</sup> Conceptual meaning is, roughly, meaning as usually understood in terms of reference, conceptual content, or other); procedural meaning, on the other hand, consists of constraints over how to retrieve propositional constituents from lexical items. That is, it consists of rules that guide the speaker/hearer in the interpretive process, in assigning a value to an expression. What in relevance-theoretic terms is the conceptual/procedural distinction is not that far from Kaplan's (unpublished) distinction between *rules of meaning* and *rules of use*, the idea of an expression linguistically encoding a procedure being akin to Kaplan's idea of an expression semantically encoding a rule of use: in roughly the same way as a demonstrative may be thought of as encoding a rule like "use 'that *F*' to refer to some maximally salient *F*", definite descriptions may be thought of as encoding a different, more elaborated, rule or procedure. The important point we'll try to make is that there is only one rule or convention for descriptions, which covers both referential and attributive readings, and not two distinct, separate meanings.

<sup>3</sup> The notions of *relevance*, *optimal relevance*, *expected relevance*, etc., are relevance-theoretic terms. For an understanding of those terms, we refer the reader to the expositions of Relevance Theory indicated in fn.23.

*environment*, or MCE), Devitt suggested that such a representation might include referential representations of objects.<sup>4</sup> I want to cash out this suggestion. I don't know exactly how to characterize a referential representation of an object; however, I take such a characterization to involve something like the following: if someone has a referential representation of an object, then she is in a position to have *de re* thoughts about it; or, somewhat equivalently, in a position to entertain object-dependent propositions about it. Now, a subject is in a position to have *de re* thoughts about some object only if she has a privileged epistemic relation to it. Other ways of making essentially the same point are demanding that she have a *de re* individual concept of that object, or that she have a *de re* mode of presentation of it. In any case, the epistemic relation that is at the basis of this possibility is a reminiscence of Russell's acquaintance relation; it doesn't matter right now how exactly we interpret it: having a causal-perceptual link to an object counts as having such a relation, and is enough for our purposes.

To interpret a description referentially is, then, to look for an object  $o$  such that we are in a position, within the particular MCE we are in, to entertain  $o$ -dependent propositions. To interpret a description attributively amounts to settle for thinking non-referentially of an object, regardless of whether we are also able to have *de re* thoughts about it or not. Let  $d$  and  $(\varphi x)$  be as before,  $(Sd)$  a sentence containing  $d$  in subject position,  $c$  an MCE, and  $D_c$  its domain that is, the set of objects accessible at (MCE). We can put the procedure encoded by  $d$  in a slightly algorithmic form as follows:

**PROCEDURE FOR DESCRIPTIONS.** *Interpret  $d$  referentially; for every  $i$  in  $F = \{x : x \in D_c \wedge (\varphi x)\}$ , taken according to a preference order  $P$  on  $F$ , construct the  $i$ -dependent proposition  $p_i$  that corresponds to  $(Sd)(i/d)$ ; if optimal relevance is achieved, then retrieve  $p_i$  and stop; else, interpret  $d$  attributively; construct the corresponding object-independent proposition  $q$ ; if optimal relevance is achieved, then retrieve  $q$  and stop; else stop<sup>5</sup>*

<sup>4</sup> In (unpublished) I was working provisionally within a different framework, and adopted the notion of Epistemic Context of Utterance (ECU), a relative of Skerk's (unpublished) homonym notion, as the representation that a conversational participant has of her Pragmatic Context of Utterance. I think the adoption of the clearer notion of MCE has solid theoretical and explanatory dividends over the adoption of either notion of ECU, but I'll not be arguing here for this position.

<sup>5</sup> Here, whether an object  $o$  satisfies  $(\varphi x)$  within an MCE does not depend on whether the object *actually* satisfies  $(\varphi x)$ , but on whether it does for at least some of the conversants.

One might be worried about the precedence that a referential interpretation has over an attributive one. Obviously, one has to go first; otherwise, it wouldn't be that

To see how this procedure works, we must take a look at some examples. Suppose my friend and I are attending a chess match. We both know that one of the players is Russian, the other, Ukrainian. Alas, we ignore their respective names –maybe they are too hard to remember or even pronounce. However, we can see them –so we are in a position to have *de re* thoughts about them–, and know which is which. In this setting, I say to my friend:

(4) The Russian player will move first

To retrieve the proposition I expressed in uttering (4), she will go through the procedure encoded by ‘the Russian player’. So first she will interpret this description referentially, and look within her MCE for a referentially represented object *o* such that it is both a chess player and Russian. Since there is one such object, and in this case only one such object, she will construct only one object-dependent proposition. Since this proposition is relevant enough (it provides a maximum of cognitive effects compatible with a minimum of processing effort), she’ll retrieve it as the proposition I expressed by means of (4).

Now, suppose that within my friend’s MCE there is more than one Russian player. After all, she knows plenty of them, and more than one might even be perceptually available at the time of utterance. In this case, the procedure commands my friend to take into account each Russian player of which she has a referential representation. This is not done blindly, however, but following some preference order relation on the class of Russian players accessible at *c*. In this case, the most relevant one, that is, the first one to be accessed, will be the player we are both seeing, or the Russian player sitting at the table and not, for example, the other Russian player sitting right next to us.

---

much of a procedure. Our choice of evaluating first referential interpretations is motivated by relevance considerations: first, from a processing effort point of view, a great deal of our descriptions seem to be referential, so that, based on frequency, referential interpretations should go first (it just makes sense to evaluate first the kind of interpretation that is more frequent); second, from the point of view of cognitive effects, referential interpretations are more relevant than attributive ones, insofar as a piece of information about some particular individual or other has more cognitive effects than a general information about the world. For example, the information that *that* man is the burglar that broke into my house last night has more cognitive effects than the information that there was exactly one burglar that broke into my house last night: I may call the police, be cautious around him, form the expectation of recovering my stolen goods...

But suppose I utter (4) when we first hear of a Russia-Ukraine chess match, and before any player is selected for the event: suppose I'm trying to make a prediction. My friend is still acquainted with several Russian players, so that she can have *de re* thoughts about plenty of them. However, none of them is relevant enough; that is, no object-dependent proposition she can form will meet the expected relevance. When she runs out of object-dependent propositions to consider, the procedure dictates that she must interpret the description attributively, and retrieve the corresponding general proposition, maybe more explicitly expressed by:

- (5) The Russian player, whoever he might turn out to be, will move first

So she retrieves the corresponding quantificational proposition as the one I expressed by means of (4).<sup>6</sup>

Now, there is a third possibility: the procedure halts without delivering any proposition. This happens when neither the constructible object-dependent propositions, nor the quantificational one, achieve optimal relevance. One usual case is given by (6):

- (6) The vase had a crack on it

Suppose my friend and I are looking for vases. We enter into a shop, and find ourselves in front of a shelf full of china vases. None of them draw my attention at all, but, completely unaware of it, I slightly delay my glance upon a particular one. She wrongly interprets this as my noticing that particular vase. In leaving the store, she utters (6). I go through the procedure. First, I try to interpret the description referentially. Since no vase is relevant enough, I retrieve no object-dependent proposition.<sup>7</sup> The

<sup>6</sup> Again, one might think that, in understanding utterances of sentences containing definite descriptions in situations in which it is highly unlikely that the speaker have any particular object in mind, the better procedure is to start with the quantificational reading. We could have stated the procedure the other way around, starting with the quantificational reading, and moving to the referential one if the quantificational reading is not relevant enough. The order here is really immaterial, and makes no real difference. However, it makes some sense to put the referential reading first: as a matter of fact, the majority of our uses of definite descriptions are referential ones, so it makes sense to evaluate those readings first.

<sup>7</sup> Actually, what happens is that the preference relation on the class *F* of china vases generates one equivalence class containing all vases present. So maybe I'm able to dismiss the vases in the shop *in toto*, and I don't have to consider *every* vase-dependent

next step is to interpret the description attributively. However, since I know she's trying to convey something about some *particular* vase or other we saw, the quantificational proposition doesn't achieve optimal relevance either. So I don't retrieve any proposition at all, and the procedure halts, resulting in my failure to interpret my friend's utterance.<sup>8</sup>

4. We have adopted a very audience-oriented stance in dealing with our problem. Now, certainly it is one thing to say how a hearer tells that a description refers to some object or other, and quite another to say what makes it the case that a description refers to its referent. While the first question is epistemic in character, the second is a constitutive one. However, there is a justification for the approach we've taken: establishing what the hearer needs to do in order to ascertain reference yields a definite constraint on what is needed to explain reference determination by the speaker. We've seen that reference is resolved by means of the linguistic meaning of the description –that is, the procedure given above–, and general pragmatic abilities that enable the hearer to identify the intended referent. All the speaker has to do to make successful reference by means of a description is to take those elements into account. Our answer to the constitutive question is, briefly, the following: what makes it the case that a description refers to an object is just its being so used by the speaker.

To appreciate the implications of our stance on the constitutive question, it might be useful to draw a parallel with a relatively well-understood pragmatic *phenomenon*: that of generalized conversational

---

proposition I could construct. The same could apply to the case of my prediction: no Russian player is particularly salient, so maybe my friend can dismiss them *in toto*, bypass the construction of all object-dependent propositions, and simply go for the quantificational one.

<sup>8</sup> We note in passing that the referential interpretation is not mandatory whenever some relevant object is perceptually available. Suppose I say, without trying to be discrete, and with respect to someone clearly perceptually available to all conversational participants:

The person sitting at my right is going to be sorry by the end of the night

Given the setting, I could have said just as well:

She will be sorry by the end of the night

pointing at that person. The processing effort demanded by 1. is greater than the one demanded by 2., and in no case the cognitive effects of 1., interpreted referentially, will be greater than those of 2.. So in 1., the description won't be so interpreted. Surely, cases like this one are cases of extremely uncooperative speakers: it is hard to imagine any natural example in which a speaker would use 1. instead of 2..



implicatures. It is generally agreed that an utterance of (7) effectively conveys something like (8):

- (7) John has three children  
 (8) John has exactly three children

One might ask what makes it the case that ‘three’ means here *exactly three* –as opposed to just meaning *at least three*. There are two ways of understanding this question. One way is the following: what makes it the case is the fact that ‘three’ is used by the speaker to mean *exactly three*. Alternatively, it could be understood that what makes it the case is the fact that, as a constitutive matter, ‘three’ means *exactly three*. But ‘three’ just doesn’t mean, as a constitutive matter, *exactly three*. All there is to its meaning so is its being used by the speaker to mean so. “‘Three’ means *exactly three*” is just shorthand for “The speaker uses ‘three’ to mean *exactly three*”. This explanation is pragmatic: the speaker just relies on the linguistic meaning of ‘three’, together with the conversational context, and the hearer’s ability to identify the content of her communicative intention.

Roughly the same goes for referential uses. The question about what makes it the case that a description refers to an object may be understood in two ways. If it is understood as equivalent to ‘what makes it the case that a description is effectively used by a speaker to refer to some determinate object’, then the explanation is purely pragmatic, and is given in terms of the speaker’s intention to refer to that object –which is part of her communicative intention, and thus a pragmatic intention, not a semantic one–, and her presumption of the hearer’s ability to identify the content of her referential intention. On the other hand, if the question is understood in its constitutive sense, then the answer is: nothing. Descriptions do not refer by themselves, but are used to make reference by speakers. Reference is hereby taken to be a pragmatic issue, not a semantic one.

**5.** We set out to sketch a pragmatic explanation of referential uses not subject to Devitt’s main criticism of Russellian positions. We saw that such a criticism seemed to rest crucially on the idea that any pragmatic explanation must find a suitable general proposition corresponding to the Russellian analysis of definite descriptions, and proceed to derive the intended, singular proposition corresponding to the referential use, by means of Gricean pragmatic processes. The present account is not committed to either of these: insofar as it is an account that posits a

linguistic meaning that is neither referential nor attributive, it does not assign to sentences with definite descriptions any proposition as a matter of semantic value and, in particular, no general proposition from which to derive a singular one; as a consequence, it does not rely on a Gricean mechanism to obtain the intuitive truth-conditions of referential uses, which are determined by pragmatic processes of quite a different sort.<sup>9</sup>

Since we are going revisionist on traditional wisdom about descriptions and meaning, there is one methodological consideration that might justify the endeavour: if we have a theory of descriptions that (a) posits a univocal meaning, (b) accounts for referential and attributive uses in terms of pragmatic principles independently motivated, and (c) is empirically adequate (that is, it accommodates the relevant empirical data), then this theory is to be preferred to theories that neglect linguistic evidence like Gricean-Russellian theories, which deny the truth-conditional relevance of referential uses, and to theories that posit lexical ambiguities to account for a *phenomenon* otherwise explainable by means of pragmatic processes. After all, under such a revision of a theory of descriptions, we seem to be able to preserve the best of both types of account: the truth-conditional significance of referential uses, together with a univocality account of meaning.<sup>10</sup>

<sup>9</sup> Bach (2007) seems to think that such proposals are implausible, for reasons given in Neale (1990, pp. 110-112). I find nothing on the cited pages that amounts to an argument against such positions: only some worries regarding the plausibility of Recanati's particular account. I have to concede that this account is not entirely plausible; however, this doesn't cast doubt over other accounts of underdeterminacy. After all, a great deal of advancement has been made in the understanding of pragmatic processes since the late '80s and early '90s. On the other hand, I have to say that I find Bach's own proposal a little bit puzzling: if definite descriptions have only attributive meanings, then the construction, in referential uses, of the intended singular proposition must appeal to some pragmatic process that completely bypass the linguistically encoded meaning, and substitutes it with a vastly different one. This is what Bach calls *standardization*. I'm not quite convinced by Bach's claim that standardization is not in fact a semantic convention in the sense of Devitt (2004). In any case, it is more elegant a solution to posit a meaning that falls short of attributive and referential uses, and construct from that point. Incidentally, Bach himself proposes a kind of underdeterminacy account for complex demonstratives, attributing them a sort of procedural meaning that falls short of quantificational and referential uses. Strikingly, he considers it quite plausible, while he rejects it for definites.

<sup>10</sup> As Powell (2001a) shows, this kind of treatment is smoothly extensible to complex demonstratives, so that a unified theory of attributive and referential uses for both kinds of expressions seems plausible. This by itself constitutes a significant linguistic payoff.

## References

- Bach, K. (1999), "The Semantics-Pragmatics Distinction: What It Is and Why It Matters", in Turner, K. (ed.), *The Semantics-Pragmatics Interface from Different Points of View*, Oxford, Elsevier, pp. 65–84.
- (2001), "You Don't Say", *Synthese*, 128, pp. 15-44.
- (2007), "Referentially Used Descriptions: A Reply to Devitt", *European Journal of Analytic Philosophy*, 3, pp. 33-48.
- Blakemore, D. (1987), *Semantic Constraints on Relevance*, Oxford, Blackwell.
- Carston, R. (1988), "Implicature, Explicature, and Truth-Theoretic Semantics", in Kempson, R. M. (ed.), *Mental Representations: The Interface between Language and Reality*, Cambridge, Cambridge University Press, pp. 375-409.
- (2004), "Explicature and Semantics", in Davis, S. and Gilion, B. (eds.), *Semantics: A Reader*, Oxford, Oxford University Press, pp. 817-845.
- Caso, R. (unpublished), "Referential Uses and Non-Gricean Explanations", 2008. Presented at the Workshop on Referential Descriptions, Facultad de Filosofía y Letras, Universidad de Buenos Aires, March 28th, 2008.
- Devitt, M. (2004), "The Case for Referential Descriptions", in Reimer, M. and Bezuidenhout, A. (eds.), *Descriptions and Beyond*, Oxford, Clarendon Press, pp. 28-305.
- (2007a), "Referential Descriptions and Conversational Implicatures", *European Journal of Analytic Philosophy*, 3, pp. 7-32.
- (2007b), "Referential Descriptions: A Note on Bach", *European Journal of Analytic Philosophy*, 3, pp. 49-53.
- Kaplan, D. (unpublished), "What Is Meaning? Explorations in a Theory of Meaning as Use". Brief version – draft 1, ms., UCLA.
- Neale, S. (1990), *Descriptions*, Cambridge, Mass., The MIT Press.
- (2004), "This, That, and The Other", in Reimer, M. and Bezuidenhout, A. (eds.), *Descriptions and Beyond*, Oxford, Clarendon Press, pp. 68-182.
- Powell, G. (2001a), "Complex Demonstratives", *UCL Working Papers in Linguistics* 13, pp. 43–73.
- (2001b), "The Referential-Attributive Distinction: A Cognitive Account", *Pragmatics and Cognition*, 9, pp. 69-98.
- Recanati, F. (1989), "Referential/Attributive: A Contextualist Proposal", *Philosophical Studies*, 56, pp. 217–249.
- (1993), *Direct Reference: from Language to Thought*, Oxford, Blackwell.

- Skerk, L.(unpublished), “Descripciones definidas y referencia singular”, 2007.URL = <http://accionfilosofica.com/misc/1174588991crs.doc>.
- Sperber, D. and Wilson, D. (1995), *Relevance: Communication and Cognition*, Oxford, Blackwell.
- Wilson, D. and Sperber, D. (2002), “Relevance Theory”, *UCL Working Papers in Linguistics* 14, pp. 249–287.