Generalizing Detached Self-Reference and the Semantics of Generic *One*

FRIEDERIKE MOLTMANN

Abstract: In this paper I will give an analysis of what I call 'generalizing detached self-reference' within a general account of reference to the first person. With generalizing detached self-reference an agent attributes properties to a range of individuals by putting himself into their shoes, or simulating them. I will show that generalizing detached self-reference plays an important role in the semantics of natural language, in particular in the English generic *one* and in what syntacticians call arbitrary PRO.

The pronoun *one* as in (1a) is a generic pronoun that bears a particular semantic connection to the first person and is of significant philosophical interest:

(1a) One can see the picture from the entrance.

Generic *one* is a pronoun that, I will argue, expresses generalizing detached self-reference. It is a *first-person oriented generic pronoun* in the sense that it does not stand for the speaker's actual person, but rather for a range of individuals that the speaker identifies with or simulates. In this paper, I will develop a semantic analysis of generic *one* within a general account of detached self-reference. Detached self-reference is more familiar from attitude reports such as *I imagine being Napoleon*, and my account will apply to those as well.

The analysis I will develop is what I will call an attitudinal account of detached self-reference, an account that assigns a central role to notions of self-attribution and pretend self-attribution of properties and to the notion of simulation. The account is based on a view on which it is not propositions that play the central semantic role for the semantics of sentences, but rather 'attitudinal objects', objects of the sort 'John's belief that S or 'John's assertion that S'. The analysis I will develop captures intuitions according to which generic *one* involves a 'detached' (or 'objective' or 'attenuated') self, while avoiding objectifying a 'detached' self. The attitudinal account of detached self-reference accounts for a range of semantic and pragmatic properties of generic-*one* sentences that are rather independent of their truth conditions.

I would like to thank the two anonymous referees for their comments as well as John Campbell, Mike Martin, Chris Peacock, Ken Safir, and Robert Stalnaker for discussions of material in this paper.

Address for correspondence: Institut d'Histoire de Philosophie de Sciences et de Technique, 13 rue du Four 75006, Paris, France.
Email: fmoltmann@univ-paris1.fr

Mind & Language, Vol. 25, No. 4 September 2010, pp. 440–473.

Generic *one*, as in (1a) has an empty, unpronounced counterpart, namely what generative syntacticians call 'arbitrary PRO' (PRO_{arb}), which corresponds to the empty subject position in (1b) and which has exactly the same semantic function as generic *one* (though a complementary syntactic distribution):

(1b) It is nice PRO_{arb} to see the picture from the entrance.

The first-person connection of generic *one* or its empty counterpart can manifest itself in various ways. One such way consists in drawing a generalization on the basis of the speaker's own, perhaps unique, experience or action. Thus, (1a) can be truthfully uttered on the basis of the speaker's own, perhaps unique experience of seeing the picture from the entrance, while at the same time making a generalization: for anyone *x*, *x* can see the picture from the entrance. Similarly, (1b) is naturally used as an expression of the speaker's own evaluation of his seeing the picture from the entrance, while at the same time making a generalization: for anyone x, x's seeing the picture from the entrance is nice for *x*. With this particular manifestation of the first-person connection, an agent generalizes a self-ascription of a property by abstracting from the particuliarities of his own situation and thus ascribing the property to anyone else—or rather anyone the agent can assume is as normal as he himself. Both of these components, the first-person connection, in whatever way it may manifest itself, and the generalization, are part of the meaning of generic *one* (or arbitrary PRO), or so I will argue.

The first-person connection consists in a certain detached form of self-reference, which makes generic *one* particularly interesting philosophically. The first-person connection of generic *one* does not necessarily consist in an ascription of the predicate to the speaker's own person. With generic *one*, rather, the speaker may just project himself onto others or in fact simulate entirely counterfactual conditions. Thus, even (1a) could be uttered truthfully by someone unable to see, as long as that speaker is able to project himself onto those who can. First-person-oriented generic sentences involve self-reference in a quite different way than the most familiar cases of *de se*-interpreted pronouns. In the familiar cases, self-reference with pronouns interpreted *de se* is reference to the actual person. With generic *one*, by contrast, the self-reference may be 'detached' from the actual person and thus may be self-reference while identifying with another person or the set of people meeting a certain condition. Generic *one* appears to involve reference to an impersonal self—the 'objective self' in the sense of Nagel (1983, 1986) or what Williams (1973) calls the 'attenuated self', a self that is dissociated from the actual person with her various experiences and activities, and that can instead take the point of view of anyone else (or no point of view at all, as Nagel would say).

In some cases, also non-generic first-person pronouns or their empty counterparts, in particular what linguists call 'controlled PRO', may involve such 'detached self-reference'. This is possible especially in contexts of propositional attitudes such as imagination and desire, as in *I imagine PRO being Napoleon* or *I want PRO to be Napoleon*. I will propose an account of both detached self-reference and generalizing detached self-reference which will not involve positing an 'objective' or 'attenuated'

self as a semantic value of the relevant pronouns, but rather makes crucial use of a primitive notion of pretence or simulation. In the former case, this notion will play a role in an attitude that is a self-attribution or rather pretend self-attribution of a property (following Lewis' (1979) account of attitudes *de se*). In the latter case, it will modify what generic *one* or arbitrary PRO range over, namely individuals *as* entities the speaker identifies with or simulates, (thus a notion of 'reference under a perspective').

Besides the particular kind of first-person-based genericity illustrated in (1a, b), the first-person connection of generic *one* or arbitrary PRO can manifest itself in a second way, namely by the use of a sentence as an expression of a generalization *aimed at* a first-person ascription on the part of the *addressee*. This is the case with deontic sentences as in (2), whose purpose is generally meant to influence the addressee's practical reasoning:

(2) One should not lie.

The first-person connection of (2) on that use will be an immediate possible self-application of a rule which the addressee is to employ in his practical reasoning.

Generic *one* (as well as its empty counterpart) is a pronoun that always leads to generic sentences and involves a particular kind of self-attribution (or pretend self-attribution), be that on the part of the speaker himself or on the part of the addressee whose acceptance of the sentence the speaker intends and with whom the speaker may thus identify. Recognizing the first-person-orientation of generic *one* or its empty counterpart can explain a number of peculiarities of sentences containing them that go beyond their usual truth conditions, including the use of such sentences in philosophical discourse about subjectivity, self-knowledge, and consciousness.

I will first discuss one of two important uses of first-person-oriented generic pronouns, namely a use on which such pronouns lead to generalizations based on the relevant agent's self-ascription or pretend self-ascription of a property, a use that also plays a central role in the philosophical literature on the 'self'. I will then give a general account of detached self-reference with attitudes *de se* and show how on the basis of that the semantics of generic *one* can be developed. In the final section, I will discuss the second important use of generic *one*, having to do with the role of generic *one* in practical reasoning.

1. Some General Linguistic Facts About First-Person-Oriented Pronouns

First some general clarifying linguistic remarks about generic *one* are in order. I will call generic *one* a *first-person-oriented pronoun*, because it generally does not just stand for the speaker, but for other individuals as well, though as individuals the speaker identifies with.

Generic *one* differs from a first-person pronoun such as *I* in English in that it need not relate to the speaker as the first person, but in embedded contexts

relates to whoever is the agent of the attitude or speech act, for example John in (1c):

(1c) John thinks that one can see the picture from the entrance.

While generic *one* (with its empty counterpart) is the most important first-person-oriented generic pronoun in English, it is at the same time a somewhat problematic expression. Generic *one* in American English is increasingly replaced by *you* and its uses show a significant range of instability. The present interest, though, is not so much the linguistic details of a particular pronoun in English, but rather an important semantic 'strategy' that is most obviously expressed by generic *one* but that is also involved in a great variety of other expressions or uses of expressions across languages in general. The very same analysis that I propose for generic *one*, in fact, is meant to apply also to arbitrary PRO, whose occurrences are much less problematic. In fact, generic *one* alternates with arbitrary PRO, which occurs in those contexts in which an empty pronominal element, rather than an overt noun phrase is required. This can be seen first from the fact that arbitrary PRO and generic *one* can co-vary in contexts like (3a, b) (that is, they take the same semantic values under the relevant assignments), and second from the fact that arbitrary PRO may act as the antecedent for *one,* as in (4a), or the reflexive *oneself*, as in (4b):

(3a) PRO_{arb} to live a great life is to realize *one*'s true potential.
(3b) PRO_{arb} to have been diagnosed with a serious illness means that *one* cannot easily get new insurance.
(4a) It is nice PRO_{arb} to see *one*'s parents.
(4b) It is great PRO_{arb} to be able to teach *oneself.*

For generic *one*, moreover, there is, even in American English, at least one context in which its occurrences remain stable and unproblematic. These are the uses of generic *one* in the philosophical literature on subjectivity, consciousness, self-knowledge, and simulation, uses which illustrate extremely well some of the most important semantic properties of that expression. Some examples from the relevant philosophical discourse are given below:

(5a) *One* cannot be mistaken about the content of *one*'s own mental states.
(5b) [PRO_{arb} Speaking of oneself] is incompatible with not knowing that the subject *one* is speaking of is *oneself* (Anscombe 1975).
(5c) *One* cannot be presented to *oneself* as an object in introspection.
(5d) *One* can predict other people's actions by putting *oneself* into their shoes.

In the examples in (5), *one* (or its empty counterpart) is in fact the best expression to use, not easily replaceable by another. This is because the generalizations expressed in (5) are based on an irreducible first-person ascription: they express 'intuitions'

or conceptual truths made on the basis of essential first-person attributions of properties.[1]

One might think that Nagel's (1974) sentence (6), which involves arbitrary PRO, is a counterexample to the generalization:

(6) There is something it is like [PRO to be a bat] (Nagel, 1974, p. 435).

(6) does not imply that the speaker himself is able to project himself onto a bat, at least on the relevant understanding of the sentence. Thus, (6) seems to violate the generalization about generic *one*. However, an alternative analysis of (6) is available, which is plausible on independent grounds and which I will make use of on other occasions later. On that analysis, the evaluative expression *something it is like* involves an implicit argument acting as the evaluator or judge, the kind of argument that can be made explicit by a *for*-phrase, as in (7):

(7) There is something it is like for a bat_i [PROi to be a bat].

A term denoting such an evaluator will then be able to bind the empty element PRO, which is hence 'controlled' rather than 'arbitrary' PRO, that is, it acts just like a bound variable, bound by the implicit term for the evaluator.[2]

Before turning to further data concerning the first-person-oriented genericity expressed by generic *one*, a few words are necessary about the semantic status of generic *one* as a pronoun.[3] Generic *one* always occurs in generic sentences. As such, it can occur in two apparently distinct ways: as genericity-inducing, as in the first occurrence in (8a), and as a bound variable, as in the second occurrence in (8a):

(8a) *One* sometimes thinks *one's* life is too short.

Nonetheless, as I have argued in Moltmann (2006), in both occurrences generic *one* is best taken to be an expression that introduces just a variable, a variable that

---

[1] Interestingly, instead of generic *one*, the non-generic pronoun *I* can be used for philosophical purposes (generalizing first-personal knowledge) in just the same way (as was pointed out to me by François Recanati):

(1) I cannot be mistaken about the contents of my own mental states.

I can generally not replace *one* in other contexts (physical possibility, deontic predicates, generalizing from one's own experiences, conditionals). The reason, it appears, is that generalizing *I*-sentences, as one may call them, are possible only when they express metaphysical truths, based on first-person knowledge. They cannot be used to express empirical self-knowledge or to make deontic statements. Clearly, the generalization is not part of the semantics of generalizing *I*-sentences themselves, but follows from the fact that such sentences, with their literal meaning, are meant to express a metaphysical truth, which by nature is generalizable. The generalizing effect thus is a conversational implicature. The use of *I* in such contexts necessarily implies what *one* would express in those same contexts.

[2] Nagel (1974) is in fact interested not just in individual subjective experiences—for some x 'what it is like for x to be a bat'—but in *types* of subjective experiences, which means first-person-based genericity.

[3] For a linguistically much more detailed discussion of what follows see Moltmann, 2006.

is to be bound by a generic quantifier Gn, a quantifier that does not represent an explicit noun phrase, but rather a silent feature in sentence-initial position. Thus the logical form of (8a) will be as below:

(8b) Gn x x sometimes thinks that x's life is too short.

There are two main reasons for that. First, two different occurrences of generic *one* (PRO_{arb}) may covary, without either being in a position in which it could act as a quantifier binding the other, as in (3a, b), and also in (9a, b):

(9a) If *one* is French, *one* is European.
(9b) PRO_{arb} to be happy is PRO_{arb} to live a good life.

Second, the generic quantifier associated with genericity-inducing *one* always takes wide scope over any other quantifier or operator in the sentence regardless of where in the sentence generic *one* occurs, unlike quantificational noun phrases:

(10a) *One* cannot always recognize everyone.
(10b) Not everyone recognizes *one* at night.

(10a) cannot mean 'not always is there mutual universal recognition', but only 'for anyone x: x does not always recognize everyone'. Similarly, (10b) cannot mean 'some people fail to recognize anyone at night', but only 'for anyone x, not everyone recognizes x at night'.

The generic quantifier Gn in (8b) is meant to be just the kind of quantifier used in the linguistic literature on genericity (see Krifka *et al.*, 1995; Cohen, 2002; Greenberg, 2005 among others). While it is agreed that this quantifier should allow for exceptions and has modal force (ranging not just over actual individuals), there is a lack of agreement as to how the quantifier actually should be defined. A plausible way of understanding it given in the literature, a way also fairly suited for present purposes, is to take it to be a combination of a universal quantifier ranging over possible worlds, restricted by some contextually given accessibility relation R (relating certain 'normal' worlds to the actual world), as well as a universal quantifier ranging over individuals, possibly restricted by a contextually given condition C (Greenberg, 2005):

(11) $\forall w \forall x (wRw_o \,\&\, x \in D(w) \,\&\, C(x) \rightarrow P(x))$

Here D(w) is the domain of w. Note that both R and C may impose rather severe restrictions, so that the generic quantifier may range over just certain actual individuals in a limited context. Including other possible worlds accounts for the fact that generic sentences in general are not just about actual individuals.

Using a universal quantifier rather than a quantifier like 'most' explains entities not meeting the generalization count indeed as exceptions (which would not be the case with a quantifier like 'most'). It is the contextual restrictions both on the worlds and the individuals that capture the fact that there *can* be exceptions (Greenberg, 2005).

The contextual restriction *C* of the generic quantifier is needed to account for why *one* in (12), for example, may range only over the students in a particular contextually given class:

(12) One has to hand in the essay tomorrow.

Of course, the logical form in (11) does not yet capture the particular first-person-oriented meaning of generic *one*. To account for that, the simple variable 'x' in (11) will later be replaced by a more complex expression, with an additional component representing the 'first-person-orientation'.

The first-person-oriented meaning is needed not only to account for the first-person-based genericity observed in (1)–(4), but also for the observation that genericity-inducing *one* is subject to severe restrictions on which predicates it can go with. For example, the predicates in the following sentences are hardly acceptable with generic *one*, though they are fine in other generic sentences:

(13a) ?? One has a nose.
(13b) The typical person has a nose.
(14a) ?? One lives in a big city.
(14b) People live in big cities.

(13a) is impossible under normal circumstances; whereas (13b) is acceptable as well as true. Also (14a) is hardly acceptable, whereas (14b) is fine, though false. Sentences like (13a) and (14a) are possible under special circumstances, though, for example when expressing a law or condition with an intended immediate self-application or as a description of a realization of a requirement, uses which in fact belong to a different 'strategy' for using generic *one*, which I will come to shortly (see also Moltmann, 2006). The point is that such sentences cannot be used in the same way as (13b) and (14b), which do not require special circumstances.

Other examples in which generic *one* is unacceptable except under special circumstances are those in (15):

(15a) ?? One has at least one passport.
(15b) ?? One has parents.
(15c) ?? One has a body.
(15d) ?? One breathes.
(15e) ?? One was nervous.
(15f) ?? One went home.

Again such predicates are fine in other generic sentences with suitable choices of generic NPs such as *the typical person* or *people*:

(15a′) The typical person has at least one passport.
(15b′) The typical person has parents.
(15c′) The typical person has a body.
(15d′) The typical person breathes.
(15e′) People were nervous.
(15f′) People went home.

The restriction on predicates is not strict, though: all the examples in (15a)–(15f) can be made acceptable when understood (as far as possible) as a fulfilment of a script or plan or as a state of affairs that is also a fulfilment of a general requirement or norm. But these would be different uses of generic *one*, conforming to a different, second strategy (see below).

There are also contexts in which generic *one* imposes no restrictions on the predicate whatsoever. First, when occurring as a bound variable, generic *one* imposes no restrictions:

(16a) Sometimes one forgets that one has a nose.
(16b) One can doubt that one has a soul.

Second, generic *one* imposes no restrictions on the predicate in indicative conditionals:

(17a) If one lives in a big city, one lives in a city.
(17b) If one has a nose, one can breathe.
(17c) If one is nervous, one should take a tranquillizer.
(17d) If one is human, one has a soul.
(17e) If one has a significantly increased temperature, one has a fever.

There are thus no strict conditions on which kinds of predicates generic *one* can co-occur with. Given this, rather than trying to rule out unacceptable examples by imposing lexical restrictions on predicates acceptable with generic *one*, it is best to explore under what circumstances generic *one* is possible. There are two types of 'semantic strategies' that can be distinguished with which generic *one* is made possible in a sentence:

Strategy 1: Inference from the First Person (first-person based genericity)
Generic *one* (and arbitrary PRO) is licensed in a (simple) sentence establishing a generalization based on a first-person application of the predicate or first-person 'pretend application' of the predicate

Strategy 2: Inference to the First Person (first-person targeted genericity)
Generic *one* (and arbitrary PRO) is licensed in a (simple) sentence stating an (already established) generalization that is to allow for an immediate application to the first person in the reasoning relevant in the context.

That is, a sentence with generic *one* is acceptable if it either expresses a generalization on the basis of the first person (or the pretending first person) or else it is meant to be immediately applicable to the first person. These two strategies are not entirely independent, as we will see: at least with predicates of moral evaluation (*It is wrong* PRO_{arb} *to do* X), generic *one* and arbitrary PRO are licensed both on the basis of Inference from the First Person and Inference to the First Person.

The first-person-orientation of generic *one* and arbitrary PRO concerns not only the speaker, but also, in embedded contexts, whoever may be the described agent of the reported attitude or speech act:

(18a) John said that one can see the picture from the entrance.
(18b) John said that it is nice PRO_{arb} to see the picture from the entrance.

Crucially, Strategy 1 does not require the predicate to actually hold of the first person, but rather the relevant agent may just pretend to meet the conditions the predicate expresses.

The latter will in particular explain the possibility of generic *one* in indicative conditionals (Section 4).

2. Strategy 1: Inference from the First Person

2.1 Inference from the Actual First Person

With regard to Strategy 1, I will first discuss cases where first-person-based genericity is in fact based on an actual first-person application. I will then turn to the particular issues arising with a pretend first-person application.

In attitude contexts, the first-person-orientation of generic *one* is particularly transparent, for example when a generic-*one* sentence is embedded under an epistemic predicate:

(19a) John found out that one can see the picture from the entrance.
(19b) John confirmed that one can see the picture from the entrance.

For (19a) to be true it is sufficient that John has had the experience of seeing the picture from the entrance (while assuming he is relevantly normal). Similarly, for (19b) to be true, all John needs to have done is having gone to the entrance and having seen the picture from there (while taking himself to be relevantly normal). Generic-*one* sentences differ thus from universally quantified and other generic sentences, such as those below:

(20a) John found out that people can see the picture from the entrance.
(20b) John confirmed that everyone can see the picture from the entrance.

For (20a) and (20b) to be true, John has to have made sure in other ways that people other than himself can see the picture from the entrance.

The same point is made by the contrast between (21a) and (21b):

(21a) John found out that one gets sick when one eats these mushrooms.
(21b) John found out that people get sick when they eat these mushrooms.

For (21a) to be true, it suffices that John found out that he himself got sick once, so that assuming he is sufficiently normal, he can generalize on that basis the possibility of getting sick to other people that have also eaten the mushrooms. By contrast, for

(21b) to be true John has to find that out empirically that other people than himself get sick after eating the mushrooms.

Even the non-factive attitude verb *imagine* displays the difference:

(22a) John imagines a way one could cross the river.
(22b) John imagines a way people / a person could cross the river.

For (22a) to be true it is sufficient that John imagines himself crossing the river. By contrast, John's imagination needs to involve other people than himself in (22b).

There are other semantic contexts than those expressing physical possibility that allow for generic *one* by making use of first-person-based genericity, namely sentences describing experiences in certain types of situations or the frequency of acts or experiences:

(23a) One feels tired after such a long day.
(23b) One sometimes thinks one's life is too short.

The relevant (and most natural) reading of (23a) is one on which the speaker, by uttering the sentence, expresses his own state of tiredness. The speaker at the same time, though, generalizes by abstracting a type of situation from the situation he finds himself in. Similarly, (23b) naturally expresses a generalization of the speaker's own occasional thoughts of his life's shortness.

The first-person-orientation of generic o*ne* and arbitrary PRO manifests itself also in the ability of those pronouns to serve in an immediate description of a first-person experience, as in the following examples:

(24a) I find that one can easily forget one's own past experiences.
(24b) I now know how it is [PRO_{arb} to be treated like a king].

The embedded sentences in (24a, b) naturally serve as direct descriptions of a first-personal psychological state, though of course the generalizing force is there as well. (24a) and (24b) thus differ markedly from (25a) and (25b), where the attitude described takes as its immediate source third-person observations, or else has a derived content, obtained only inferentially from a first-person experience:

(25a) I find that people can easily forget their past experiences.
(25b) I now know how it is for a person to be treated like a king.

First-person-based genericity, using the appropriate pronouns, is clearly the only way of generalizing irreducibly subjective experiences as types of experiences, a point made quite apparent by the linguistic structure of the philosophical examples in (5).

While first-person-oriented pronouns provide the most suited expressions for generalizing essential first-personal situations, such as experiences, actions, or intentions, they can also be used for generalizing states of affairs that are not essentially first-personal and thus serve a wider purpose, for example in (26):

(26) One can be listed in the phone book without an address.

First-person-based genericity involves the ability of abstracting from the particularities of one's own person and situation, judging oneself to be normal in relevant respects, and generalizing to anyone meeting the same conditions. This kind of generalizing self-attribution of properties as a form of abstraction involves distinguishing between relevant and irrelevant features of a given person and her situation. Self-reference of this sort can be called *generalizing detached self-reference*.

Generalizing detached self-reference may also be viewed as *generic simulation*, where simulation is to be understood in the sense of Simulation Theory. Simulation Theory (Goldman, 1989, 1992; Gordon, 1986, 1995; Heal, 1989) has been developed as a theory about how people ascribe propositional attitudes to others and predict or explain their behaviour. Simulation Theory is fundamentally a first-person approach to the attribution of attitudes to others and to the prediction and explanation of their behaviour. According to Simulation Theory, third-person ascriptions of attitudes and explanations and predictions of actions are based on first-person ascriptions: by pretending to be the other person or taking the other person's point of view—in other words by *simulating* the other person. Simulation Theory thus contrasts with the more traditional view about the ascription of attitudes and the explanation / prediction of actions of others, the Theory Theory. The Theory Theory is a third-person approach to the attribution of attitudes to others and to the prediction and explanation of their behaviour: according to the Theory Theory, propositional attitudes and actions are attributed directly, on the basis of a tacit theory about other people's behaviour, a theory based in particular on observations about how other people behave.

First-person-based genericity then would be genericity based on simulation, rather than theory about others: a property is attributed to anyone (meeting relevant restrictions) on the basis of the speaker's attributing that property as if to himself, while abstracting from the particularities of his own person and situation.

The kind of simulation discussed in the literature on Simulation Theory is of course not generic simulation, but specific simulation: the attribution of properties to a specific person in whose situation the agent puts himself. In the case of generic simulation, the relevant intentional agent generalizes simply his own situation, abstracting from the features of his situation that are particular to himself. The agent then does not need to make any adjustments to adopt another person's point of view, but simply needs to abstract from what is specific to his own situation in order to attribute the property in question to others.

2.2 Inference From the Simulating First Person

In general in fact, the application of the predicate to the first person is not obligatory. Instead the generalization may be made on the basis of the speaker simply identifying with everyone meeting certain conditions (which the speaker himself need not meet). As mentioned, (1a) might be understood that way. Particularly convincing examples making the point in two different ways are (27a) and (27b):

(27a) One can see me from the entrance
(27b) One can solve the equation.

(27a) is perfectly acceptable without of course the speaker being able to see himself from the entrance. By uttering (27a) the speaker will just identify with (or put himself into the shoes of) people standing at the entrance, while referring to his actual person with *me*.

(27b) can be uttered appropriately and truthfully by someone who lacks the relevant mathematical ability. That speaker just emphasizes with anyone (actual or possible) with sufficient ability.

Detached self-reference with generic *one* occurs also in the following variation of an example by Mach (reported in Perry 1996). In the situation in question, Mach is looking at a mirror in a bus without recognizing the man he sees there as himself. In this case, (28) is perfectly acceptable:

(28) Mach noticed how one looks with unkempt hair and shabby clothes.

In (28) Mach in fact identifies with the person he sees in the mirror without making an actual identification with his actual person. Here again generic *one* expresses generalizing detached self-reference without implying self-reference to the agent's actual person (unlike with ordinary *de se*-interpreted pronouns).

Detached self-reference manifests itself also in the fact that generic *one* when used so as to include the speaker does not imply reference *de se* in the traditional sense.[4] Consider (29a, b):

(29a) Sometimes one thinks that through great deeds one can become a hero without realizing that one oneself can become a hero that way.

(29b) Sometimes one thinks one should help others without realizing that one oneself should help others.

Even though at first sight (29a) and (29b) seem impossible, indicating a standard *de se* reading, the sentences can become acceptable in a situation in which the speaker simply fails to apply the generalization he believes in to his own person (and thus fails to draw the practical consequences from it).[5]

Later, in Section 4, I will come to yet another, particularly important context of generic *one* with detached self-reference, namely indicative conditionals.

3. Detached Self-Reference

3.1 Detached Self-Reference with First-Person Pronominal Elements

Generic *one* expresses self-reference that may detach itself entirely from the relevant agent's actual person. Generic *one* allows for a self-attribution of a property while

[4] See Castañeda, 1966, 1967; Evans, 1982; and Perry, 1979 for relevant discussions.

[5] Note that the examples show that the construction *he himself* or *one oneself* does not as such indicate self-reference, but self-reference together with reference to the actual person (*de se* in the familiar sense).

the agent identifies himself with someone else or a collection of individuals that does not include himself. The possibility of detached self-reference as such is not a peculiarity limited to generic *one*. In certain contexts it is also available with non-generic pronominal elements, and the conditions on the possibility of detached self-reference with those elements can shed light on how detached self-reference in general should be analysed. Let us therefore take a closer look at the relevant cases.

A well-known example of detached self-reference, discussed by Williams (1973), is (30), with non-arbitrary or controlled PRO, the empty subject of infinitive or gerundive clausal complements:

(30) I imagine PRO being Napoleon.

In (30), the empty pronominal element PRO is controlled by *I* (that is, it is coreferential with *I*). PRO in this case does not refer to the speaker's actual person; but rather contributes simply to an identification of the speaker with Napoleon, possibly without thereby transferring any of the speaker's properties to Napoleon. With the utterance of a sentence like (30) the speaker simply expresses a projection of himself onto Napoleon.

Detached self-reference with controlled PRO is not restricted to contexts of imagination, but is also available with desire, as in (31):

(31) I want PRO to be Napoleon.

The question then is, under what conditions is detached self-reference possible?

One important generalization about detached self-reference is that it is also available with predicates other than those expressing identity with a specific individual. Detached self-reference is available, for example, in (32), which expresses identification with a kind of individual or an arbitrary instance of such a kind:

(32) I imagine PRO being a soldier in a war.

The imagination reported in (32) may consist in a projection of the agent onto the role of a solder in a war, without thereby carrying over any of the actual properties of the agent. This is also possible with non-sortal predicates as in (33):

(33) I imagine PRO swimming across the channel.

The imagination reported in (33) may consist in a projection of the agent onto a situation of swimming across the channel, in which the agent does not have any of the particular features he actually has. The effect of such detachment from the speaker's actual person is that the content of the imagination reported in (32) and in (33) comes out generic: the speaker imagines how it is to experience swimming across the channel in general.

One might think that (33) is perhaps generalizing detached self-reference with arbitrary, rather than controlled PRO. But if this were so, then the generic reading should also be available with other attitude verbs that do not as such allow for detached self-reference as in (31) and (32). This is not the case, however. *Remember* is a verb that does not allow for detached self-reference of the sort in (31) and (32), and it does not allow for a generic reading of a gerundive complement either:

(34) I remember PRO swimming across the channel.

Unlike with *imagine*, a simple gerundive complement of *remember* as in (34) must relate to a specific event in the past. In his memory about himself, the speaker may also set aside any other properties he may have that won't play a role in his identification with a role in the event in the past. But this does not lead to a generic effect. The reason is that *remember* simply does not allow for detached self-reference, and thus in particular does not allow for generalizing detached self-reference.

To get a generic reading with *remember*, a more complex construction must be chosen:

(35) I remember how it is [PRO to swim across the channel].

Here, crucially, the clause is embedded under a predicate of evaluation *how it is*. This predicate, as was pointed out earlier in connection with Nagel's sentence (6), involves an implicit evaluator *for one*, which is what carries the first-person-oriented genericity. This means that (35) involves (bound-variable) arbitrary PRO, rather than controlled PRO.

We can thus conclude that detached self-reference (and the particular case of generalizing detached self-reference) is possible with controlled PRO only with suitable attitude verbs such as verbs of imagination and desire, namely just those verbs expressing future and present 'projection'. It is not possible with verbs like *remember* because such verbs do not express projection, but rather express an epistemic attitude directed towards the agent's actual past. Before I turn to what this means for the analysis of generic *one*, one further remark about the availability of detached self-reference is in order.

Detached self-reference in contexts of imagination is available also with the overt first-person pronoun *I* in the subject position of a tensed clause:

(36a) I imagine that I am George Bush.
(36b) I wish that I was Napoleon.

With *I*, detached self-reference is available also within the antecedent of certain counterfactual conditionals:

(37) If I were you, I would participate in the race.

(37) is about the speaker's decisions when identifying with the addressee. It is not about the speaker's decision when counterfactually having certain properties of the addressee. The latter is what (38) is about:

(38) If I were like you, I would participate in the race.

Nor is (37) about an impossible identity. Setting aside the issue of counterfactuals with impossible antecedents, (37) certainly is not equivalent to (39):

(39) If I were identical to you, I would participate in the race.

Note that the subjunctive in the antecedent of (37) is obligatory.

Does *I* then display the same possibilities for detached self-reference as controlled PRO? No: *I* in contexts of imagination does not generate the generic reading that PRO can generate:

(40) I imagine that I am swimming across the channel.

(40) carries over the concept the speaker has of himself; it lacks the generic effect that (33) has. Note also that if detached self-reference was a general possibility for *I*, we would not expect the subjunctive to be obligatory in counterfactuals.

We can thus conclude that the cases of detached self-reference with *I* in contexts of certain attitudes and in counterfactuals should receive special treatment (involving perhaps 'self counterparts' that consist in people other than the speaker himself).

The difference between *I* and controlled PRO is important: analysing controlled PRO as a variable is plausible both philosophically and linguistically, but not so for *I*. The analysis developed in the next section is meant to give a unified account of controlled PRO, arbitrary PRO, and generic *one* as variable-introducing expressions, but it does not concern itself with detached self-reference with *I*.

3.2 An Attitudinal Account of Detached Self-Reference

The point of departure for my semantic analysis of first-person oriented generic pronouns will be the possibility of detached self-reference with controlled PRO. In the case of controlled PRO, the detachment can be traced to the attitude in question (an attitude of imagination or desire), and it is this feature that, I propose, should then be carried over to the analysis of generic *one* and arbitrary PRO.

Crucially, in the content of a propositional attitude *de se*, the speaker should not figure himself as a component. That the content should not depend on a particular agent is also reflected in the way identity of content is expressed linguistically: pro-sentential quantifiers like *something* and nominalizations such as *the same belief* allow abstracting from the particular agent in question. Thus, on one reading, the following inferences are intuitively valid:

(41a) John believes that he is the winner.
Bill believes that he is the winner.
John and Bill believe the same thing.

(41b) John believes that he is the winner.
Bill believes that he is the winner.
John and Bill have the same belief.

Lewis (1979) takes attitudes *de se* to consist in a self-attribution of a property, thus avoiding a propositional content in which either an agent or a particular mode of presentation of a self plays a role. This view corresponds to a semantic analysis on which an attitude verb like *expect* takes a property, not a proposition, as an

argument (cf. Chierchia, 1990). Thus, (42a), which has an obligatory *de se* reading, would be represented as in (42b), rather than as in (42c), which corresponds to an interpretation *de re*:

(42a) John expects [PRO to win].
(42b) expect(John, λx[win(x)])
(42c) expect(John, ∧ he will win)

This account does not require making attitude verbs ambiguous as to whether they take properties or propositions as arguments. Lewis himself in fact suggested that propositional attitudes always involve a self-ascription of a property. In the 'propositional' cases, it would consist in a self-location in a possible world: a proposition construed as a property of a possible world is attributed to the agent's own world, thus locating the agent in the worlds of his belief or other attitude. Later I will propose a generalized view of *de se* as well, though of a somewhat different kind.

My motivation in adopting Lewis' account is somewhat different from Lewis' own motivation. The present interest lies in giving an account of detached self-reference. What is crucial in the present context is that Lewis' account allows distinguishing different kinds of self-ascriptions; that is, different kinds of attitudes *de se* will involve different kinds of self-ascriptions of properties. Detached self-reference is made possible by certain propositional attitudes (desire, as in (31), and imagination, as in (32)), but not others (such as thinking or remembering (as we saw with (34))). Taking *de se* contents to be properties can account for the role of the kind of propositional attitude for enabling detached self-reference. An attitude such as imagination or desire when taking a property as its content consists not in an actual self-ascription of a property, but rather a pretended or simulated self-ascription. By contrast, an attitude of thinking or remembering when taking a property as its content consists in an actual self-ascription of a property. That is, it depends on the kind of attitude how self-ascription of a property is to be understood. This *attitudinal account* of detached self-reference reduces detached self-reference to attitudes that consist in a pretend or simulated self-ascription of a property. In this way the attitudinal account avoids objectifying an attenuated or objective self.

How does this account for the possibility of generalizing detached self-reference with controlled PRO? Generalizing detached self-reference can simply be traced to the flexibility of the relevant notion of pretence. Pretending to have a property can mean either of two things:

[1] projecting one's actual person onto having the property, that is, taking one's own actual person counterfactually to have the property (thus making minimal adjustments given the properties the actual person already has).

[2] projecting oneself onto just anyone having the property, abstracting away from all the particularities of one's own person, that is, pretending to be just anyone having the relevant properties, entirely disregarding one's own actual properties.

[2] is a form of generic pretence or *generic simulation*, and it is precisely that which is involved in generic *one* too. I will discuss the notion of generic simulation in greater detail later when analysing the semantics of generic *one* in indicative conditionals (Section 4).

An account of the content of attitudes *de se* as properties is not complete unless it also gives a specification of the truth conditions of such contents. While propositions are entities that have their truth conditions essentially, properties do not have truth conditions, but need to be supplemented by what they are to be predicated of to yield a truth value.

It is not only *de se* contents in particular attitude reports that are truth-conditionally complete. Also, anaphoric reference generally consists in reference to a truth-conditionally complete object. Thus, when (42a) is continued by *that will be the case*, the speaker can only mean that John will win, not that he himself will win. That is, the truth conditions of a particular belief content can never involve predicating the property of another agent, say, a speaker who is making reference to that content. The same holds when someone agrees with, say, a claim *de se*. When John claims that he is the winner, and Bill agrees, Bill can only have agreed with the claim that John has won, not that he, Bill, has won.

To obtain the truth conditions of a *de se* content, one might be tempted to invoke some kind of strategy associating the relevant individual with the property that is that content. However, there is a simpler account available: there is already a single object that has just the right truth conditions that a *de se* content intuitively has. In the case of (41a) this is the object that is 'John's belief that he is the winner', and in the case of (42a) it is 'John's expectation that he will win'. John's belief that he is the winner (or John's expectation that he will win) is not a proposition, but rather what I call an *attitudinal object* (Moltmann, 2003). John's belief that he is the winner is not just a proposition because it has properties that relate specifically to the belief mode and the agent (it is something that can be 'strong', 'justified', or 'unexpected', unlike a proposition), and it is not a belief state, because, like propositions, but unlike belief states, it has truth conditions (Moltmann, 2003). Unlike the property λx[winner(x)], John's belief that he is the winner is truth-conditionally complete: it is either true or false.

There are philosophical as well as linguistic arguments for preferring attitudinal objects to propositions (or at least to take them as more primitive than propositions), which I will only briefly mention. General philosophical reasons in favour of attitudinal objects are that attitudinal objects, being intentional objects, go along with an intentionalist view of truth and reference, the view that truth is not possible without intentionality (Moltmann, 2003), and moreover with the truth-directedness of proposition-like objects, a notorious problem for structured propositions (Jubien, 2001). Linguistic reasons favoring attitudinal objects include observations to the effect that prosentential expressions such as *that* or *something* take attitudinal objects rather than pure propositions as values (Moltmann, 2003). Attitudinal objects (or kinds of attitudinal objects) arguably should replace propositions for the various purposes for which they have been invoked (Moltmann, 2003).

There are also specific motivations from generic-*one* sentences for using attitudinal objects: attitudinal objects allow representing the particular first-person access that generic-*one* sentences encode (Section 3.3).

What is most important in the present context is that attitudinal objects, based on an attitude *de se*, are clearly truth-conditionally complete objects, whereas the content of an attitude *de se* (a property) is truth-conditionally incomplete. Lewis' account of attitudes *de se* provides thus a further motivation for attitudinal objects: attitudinal objects provide the missing truth-conditional part of that account.

Not every kind of attitudinal object has truth conditions, but, generally, an attitudinal object is associated with conditions of satisfaction. Whereas a thought or a belief has truth conditions, a desire or expectation does not have truth conditions, but rather *fulfilment conditions* and an imagination has, one may say, *conditions of representational accurateness* (at least if it is about existing objects).

Attitudinal objects, unlike propositions, are mind-dependent objects. However, it is possible to have an attitude of sharing or of evaluation towards the attitudinal object of another (*Mary shares John's belief that S, Mary considers John's belief that S true*). Moreover, besides particular attitudinal objects, such as John's belief that S, there are kinds of attitudinal objects, the kind of thing *the belief that* S stands for (Moltmann, 2003). Kinds of attitudinal objects do not involve a particular agent, but have instances that do. Kinds of attitudinal objects arguably are also what *the same thing* and *the same belief* in the conclusion in (41) refer to. That is, *the same thing* and *the same belief* in (41) make reference to an object of the sort 'the belief that he is the winner', an object constituted by the property of being the winner and the belief mode.

Attitudinal objects have two sorts of properties: properties that relate to the attitudinal mode (*John's belief that S is understandable / is unshakable / is well-grounded*) and properties that relate to the content (*John's belief that S is true / is about Mary*). These two kinds of properties impose adequacy conditions on how to construe attitudinal objects: attitudinal objects depend both on an attitudinal mode and on propositional constituents. The attitudinal mode can be seen as providing the 'glue' among the propositional constituents so as to guarantee that the attitudinal object have satisfaction conditions. This naturally goes along with a (neo)Russellian view on which attitude verbs express multigrade relations taking the various propositional constituents in the relevant order as arguments (Jubien, 2001; Moltmann, 2003). A particular attitudinal object then will depend on such a multigrade relation and n propositional constituents; a kind of attitudinal object will simply depend on a multigrade attitudinal relation and n propositional constituents. Thus, if the logical form of an attitude report such as John believes that Mary likes Bill is (43a), the attitudinal object that is John's belief that Mary likes Bill is as given in (43b) and the kind of attitudinal object that is the belief that Mary likes Bill is as given in (43c):

(43a) believe(John, LIKE, Mary, Bill)
(43b) f(John, BELIEVE, LIKE, Mary, Bill)
(43c) f_k(BELIEVE, LIKE, Mary, Bill)

Here I made use of a function f mapping an agent a, a multigrade attitudinal relation R, and n propositional constituents $X_1, \ldots, X_n$ (for some number n) to a particular attitudinal object ($f(a, R, X_1, \ldots, X_n)$) and a function f_k mapping a multigrade attitudinal relation and n propositional constituents (for some number n) to a kind of attitudinal object ($f_k(R, X_1, \ldots, X_n)$). The nature of this composition as well as questions about the propositional constituents involved should not further concern us in this paper (see Moltmann, 2003 for a proposal). In the following discussion of *de se* contents, I will make use of only very simple kinds of propositional contents, which consist of just a property, thus disregarding any internal composition the property may have (and that may thus be represented as different propositional constituents).

The truth conditions, or rather satisfaction conditions, of attitudinal objects involving a *de se* content are very simple:

(44) For an agent a, an attitudinal relation R and a property P, $f(a, R, P)$ is satisfied iff P holds of a.

This also accounts for detached self-reference. John's imagination of being Napoleon is, intuitively, not true or false. But it has conditions of representational adequacy, and these conditions say that John's imagination is representationally adequate just in case John is in fact Napoleon, which is impossible. Of course, the content of John's imagination as such is not impossible: the content is just a property, a property that is being self-ascribed. Desires are attitudinal objects that have fulfillment conditions, and the fulfillment condition of John's desire to be Napoleon is that John himself be Napoleon, a condition which simply can never be satisfied. The content of the desire as such is not contradictory: the content is just a property, a property that is being self-ascribed.

There is a different view in the recent literature on truth conditions of sentences with detached self-reference. Recanati (2007) argues that in the case of detached self-reference with imagination, the speaker's actual self does not enter into the truth conditions of the content. For Recanati, the 'cognitive content' of the imagination is given by a property (as on Lewis' account), and thus is truth-conditionally incomplete. The truth conditions are given by an Austinian proposition, a proposition in which for (30) only Napoleon, not the imaginer's self plays a role. This view cannot be correct, though, as a general account of detached self-reference. The intuitions about the fulfillment conditions of desire with detached self-reference, as in (31), are very clear. Here the agent's actual self must enter into the truth conditions: the desire described by (31) is fulfilled just in case the agent is in fact Napoleon, not just if Napoleon exists (or whatever the truth conditions of the corresponding Austinian proposition are). It is thus the attitudinal object rather than an Austinian proposition that gives the right truth-conditional completion for attitude reports involving detached self-reference.

3.3 The Attitudinal Account Applied to Generic *One*

What enabled detached self-reference with attitudes like imagination and desire was a notion of pretence or simulation which modified an attitude of self-ascription.

The same kind of notion can account for detached self-reference with generic *one*. Detached self-reference with generic *one* is independent, though, of the attitude whose content the generic-*one* sentence may form since generic-*one* sentences can provide the content of any kind of propositional attitude. Instead the notion of pretence or simulation is involved in the semantic values of generic *one*. The main intuition is this: an agent applies a predicate to a value *d* of generic *one* 'as if *d* was the agent himself' or better: the agent applies a predicate to *d* on the basis of him 'identifying with' *d* (or 'simulating' or 'projecting himself onto' *d*). This identification of course does not require the agent himself to be among those values. It just requires that the agent project himself onto anyone meeting the relevant contextually given conditions.

While with imagination and desire, the concept of pretence or simulation was a modifier of an epistemic or doxastic attitudinal relation of self-ascription (pretend to believe / to know), with generic *one* it is a modifier of the identity relation (and thus yields the identification relation). Nonetheless it is the same notion, and it is this notion that makes it possible to dispense with objectifying a detached or attenuated self.[6]

What role does the identification relation play in the semantics of generic *one*? What is crucial about the role of the notion of identification is that the predicate applies to any value *d* of generic *one* not as such but on the basis of the agent identifying with *d*. If a predicate applies to an entity *d* on the basis of that entity *d* having some property *P*, then, one can also say, the predicate applies to *d* qua having *P*. The latter notion of an entity *d* qua *P* is expressed explicitly by constructions such as *John as a teacher*, as in *John as a teacher is not very rich*. The effect of *as a teacher* as a modifier of *John* is to condition the application of predicates to John in a certain way. One of the conditions, and that is the one relevant for first-person-based generic *one*, is that *d*'s being *P* constitutes an epistemic ground for *d* falling under the predicate in question. Thus, John's being a teacher is grounds for John not being very rich. An example such as *John as a teacher is thirty years old* does not make sense since John's age cannot in any way be based on his being a teacher.

Reference with an expression like *John as a teacher*—just like quantification with generic *one*—involves an object under a perspective. In linguistic semantics, reference to objects under an aspect has become an important concept, especially in lexical semantics (Moltmann, 1997; Asher, 2006, ms), though there is no agreement

[6] In Moltmann, 2006, I proposed a simpler analysis that does not itself involve a notion of simulation or pretence. It involves simply a complex variable, an ordered pair which consists of an 'ordinary' variable and the property of being identical with the relevant intentional agent, a complex variable of the sort <x, λy[y = z]>. The variable 'z' will then later be bound by the lambda operator defining the meaning of the entire sentence as a property. The problem with this account is that the property of being identical with the relevant agent does not actually hold of a relevant value of 'x', unlike in the constructions *John qua being a teacher* and *John as teacher* which require John to be a teacher if they are to refer successfully in extensional contexts.

as to the right formal treatment of that notion. The semantics of generic *one*, if the present view is right, constitutes another phenomenon of reference to objects under a perspective (or rather quantification over such objects). What is special about generic *one* is only the particular perspective it involves: the aspect in the case of generic *one* consists in the property of being identified with by the relevant agent.

There are two particularly plausible kinds of accounts of *as*-phrases as in *John as a teacher*, both of which could be used for analysing generic *one*. First, the *as*-phrase might be taken to apply to the semantic value of *John* as a generalized quantifier, mapping it onto the set of those properties that John has based on his being a teacher.[7] Alternatively, *John as a teacher* might be taken to stand for a different entity from John, let's say, for a 'qua object' in the sense of Fine (1982) or an 'aspect' in the sense of Asher (2006). If an object under a perspective is conceived as an ontologically distinct entity, then such an entity, say 'John qua being a teacher', should fulfil the following conditions: (1) it exists as long as John is a teacher, (2) it is identical to another object x just in case x is 'composed', as a qua object, of John and the property of being a teacher, and (3) it has a property P just in case John has P and John's having P is in some way based on John's being a teacher; in other words, John's being a teacher must in some way be relevant for the application of P.[8]

For the purpose of the semantics of generic *one*, I will choose the second approach to reference under a perspective, in part just for the sake of formal simplicity. Generic *one* then ranges not over entities d meeting the contextual condition C, but over entities d qua λy[Izy], where λy[Izy] is the property of being a y such that z identifies with y. The variable 'z' is the variable that is to stand for the relevant agent, which means it will be bound by the lambda-operator defining the property involved in the self-ascription.

Bound-variable *one* differs from genericity-inducing *one* in that it does not introduce functional expressions of the sort 'x qua λy[Izy]', but rather simple variables. Thus, a sentence like (45a) is to be analysed as in (45b):

(45a) One can see one's face in the mirror.
(45b) λz[Gn x can see in the mirror(qua(x, λy[Izy]), ιy[face(y, x)])]

Generic-*one* sentences thus express properties which as contents of a propositional attitude or illocutionary act are self-ascribed by the relevant agent, just like sentences with *de se*-interpreted pronouns. In the next section, we will see that there are also some fundamental differences between pronouns interpreted *de se* and generic *one*, namely with respect to the role self-ascription plays.

3.4 Two Kinds of First-Person Orientation of Propositional Contents

Both the content of first-person-oriented generic sentences and the content of sentences with a pronoun interpreted *de se* have been construed as properties. In

[7] This is in fact what is proposed in Moltmann, 1997.
[8] The first two conditions are found in Fine, 1982; the third condition certainly needs to be added (see also Moltmann, 1997).

both cases, the properties are self-ascribed when acting as contents of propositional attitudes and as contents of assertions. However, the two kinds of contents differ fundamentally in their truth conditions and with respect to acceptance and agreement on the part of other interlocutors.

A content of a sentence with a *de se*-interpreted pronoun can be considered true or false only relative to the agent of the relevant propositional attitude or assertion. Joe's assertion that he himself is a hero aims at making the addressee accept that Joe is a hero, not that he, the addressee, is a hero. If Joe asserts that he himself is a hero and the addressee agrees, the addressee agrees that Joe is a hero, not that he, the addressee, is a hero. That is, accepting the assertion of a sentence with a pronoun interpreted *de se* means accepting the truth-conditionally complete object that involves the agent of the assertion.[9] We have seen that this is also the case with prosentential anaphora: when Bill says *that's true* in regard to Joe's assertion that he himself is a hero, Bill means that it is true that Joe is a hero, not that he, Bill, is a hero.

This is different for the content expressed by a sentence with a first-person-oriented generic pronoun. Here it is whoever accepts or just grasps the content of the sentence that matters: accepting or even just taking into consideration Joe's assertion that one can see the picture from the entrance requires self-applying the property expressed by *one can see the picture from the entrance*. Thus, when Mary agrees with Bill's assertion that one can see the picture from the entrance, she reapplies the property that is the content of Bill's assertion to herself, rather than subscribing to the result of applying that property to Bill. She also applies the property to herself when she disagrees with Bill's assertion. Thus first-person-oriented generic sentences require the property they express to be self-applied by whoever grasps the content of the sentence (agreeing with it, disagreeing with it, or just taking it into consideration). This also holds for the corresponding attitudinal object: if Bill shares John's knowledge that one can see the picture from the entrance, this suggests that Bill puts himself into the shoes of someone who sees the picture from the entrance and thus that Bill applies the property expressed by *one can see the picture from the entrance* to himself. By contrast, if Bill shares John's knowledge that he himself is a hero, then Bill believes that John, not that he, Bill, is a hero. The same holds if Bill accepts John's assertion that he himself is a hero. The attitudinal object that corresponds to a generic-*one* sentence needs to be accessed itself in a first-personal way in order to be evaluated as true or false. By contrast, the attitudinal object that corresponds to a *de se* content requires no specific first-person access.

The first-person access that generic-*one* sentences require does not influence the truth conditions of such sentences, though; it only ensures that the self-identification indicated by generic *one* as the basis for applying the predicate concerns the assessor himself. The truth conditions of a generic-*one* sentence should be the same for the

[9] Stalnaker (1981) in fact takes the behaviour of sentences with *de se*-interpreted pronouns to be grounds for rejecting Lewis' account of such sentences as expressing properties. Stalnaker instead takes them to express propositions, just like any other sentences.

interlocutors, at least as long as they agree as to what the domain of quantification of the generic quantifier is.[10]

The following is a simple way of accounting for the difference in first-person-orientedness formally. 'Propositional' contents in general are two-place relations with one argument position for a sequence of parameters of evaluation (such as a world, a time, and an agent) and a second argument position for an agent. If an agent (in a world at a time) stands in a propositional attitude to such a content, then he self-applies this relation with respect to both argument positions (as well as with respect to the world and time he finds himself in). Even if the sentence contains no first-person related pronominal element, its content can be construed formally that way: as long as the sentence is true in a world w at a time t, then for any agent a the triple $<w, t, a>$ and a will stand in the relation expressed by that sentence.

This can be stated in terms of conditions on attitudinal objects. An attitudinal object now depends on an attitudinal relation R_1, an agent a, and a two-place relation R_2 between an agent and an agent. That there is such an attitudinal object requires that the relation be self-ascribed:

(46a) For an attitudinal relation R_1, and agent a and a two-place relation R_2, the attitudinal object $f(R_1, a, R_2)$ exists iff $R_1(a, R_2)$ (which means a self-ascribes R_2 with respect to both argument positions in the way R_1).

Then we have:

(46b) An attitudinal object $f(R_1, a, R_2)$ can be grasped (understood / assessed / (dis)agreed with) by an agent d only if d ascribes R_2 to a and *himself* (in some way or another).

That is, the agent of the relevant propositional attitude or illocutionary act self-ascribes the 'propositional' content with respect to both of its argument positions, given (46a); but only the second argument position is filled in by that same agent. An example in which the two kinds of self-ascriptions are involved is the attitudinal object that is 'John's fear that one can see him from the entrance', with a *de se* reading of *him*. For John to have that fear, John self-ascribes a two-place property, in such a way that he identifies with anyone y in the domain (which does not include John himself) and self-ascribes the property of being an x such that y can see x from the entrance. An agent a grasps the attitudinal object that is John's fear that one can see him from the entrance in case a identifies with anyone y in the domain and ascribes to John the property of being an x such that y can see x from the entrance.

What are the truth conditions of an attitudinal object with a two-place propositional content? The self-application of the first argument position of the two-place

[10] Generic-*one* sentences give rise to intuitions of relative truth in the sense of MacFarlane (2005) and others. See Moltmann (to appear) for an extensive discussion.

propositional content is not part of the truth conditions of an attitudinal object; it only constitutes a condition on the existence of an attitudinal object and on grasping it. An attitudinal object of the form $f(R_1, a, R_2)$ is true just in case $R_2(a, a)$, that is, in case the two place-content holds of the agent and himself. Obviously, the first-person-orientation of generic *one* cannot be accounted with a standard notion of a (mind-independent) propositional content that consists in truth conditions; it requires the notion of an attitudinal object. It manifests itself for that reason in various linguistic intuitions that one might have considered pragmatic rather than semantic.

I now turn to the most important argument for detached self-reference being involved in the semantics of generic *one*, namely the behaviour of generic *one* in indicative conditionals.

4. Strategy 1a: Generic *One* in Indicative Conditionals

Generic *one* can occur in indicative conditionals, where, as was noted earlier, both in the antecedent and in the consequent generic *one* accepts predicates that are generally problematic in non-conditional sentences:

(47a) If one is young, one has lots of energy.
(47b) If one has a nose, one can smell.
(47c) If one has to take an exam, one gets nervous.
(47d) If one is 2 meters tall, one is taller than John.

I will argue that what licences generic *one* in indicative conditionals is also Strategy 1, but in this case it specifically involves detached self-reference.

Why is generic *one* possible in conditionals without imposing any semantic restrictions on the predicate? This follows simply from the analysis of conditional generic-*one* sentences below and the fact that the identification relation 'I' allows for complete detachment from one's own actual person:

(48) Gn x (P(qua(x, λz[Iyz]) → Q(qua(x, λz[Iyz]))

(48) says that for any x, if P were to hold of x on the basis of the agent y identifying with x, then Q would hold of x on the basis of the agent y identifying with x.

The reason why generic *one* is acceptable in such conditionals is because of the degree of detachment that the identification relation 'I' allows: According to (48), if anyone x has the property P in virtue of the relevant agent identifying with x, then x also has Q in virtue of the agent identifying with x. This crucially involves generic simulation allowing complete detachment from the properties of the agent's own person: the agent pretends to be just anyone x having the antecedent property and it is said that then the consequent property holds of x as well.

The semantics of conditional generic-*one* sentences in (48) implies a very simple way of verifying such sentences in a first-personal way: conditionals with generic

one can be verified just on the basis of the agent himself pretending to have the antecedent property (abstracting from all of his actual properties) and then seeing whether he, under pretence, would also have the consequent property. No other individuals in fact need to be involved for the verification of an indicative conditional with generic *one*.

This account of generic *one* in conditionals receives support from an independently motivated view of conditionals in general. A common view about (indicative) conditionals is that conditionals involve hypothetically adding the antecedent to one's present state of accepted information and then verifying that the consequent follows from the resulting information state (cf. Ramsey, 1931; Gaerdenfors, 1986). Given this view, the hypothetical addition of the antecedent to one's stock of beliefs, the attitude of acceptance, is, as Recanati (2000) has emphasized, a form of simulation: it is simulated belief. Thus, according to the view in question, indicative conditionals as such involve a form of simulation. On the present account of conditional generic-*one* sentences, it is just that the agent engages in generic simulation as well, identifying himself with anyone meeting the condition expressed by the antecedent.

A particularly interesting feature of conditional generic-*one* sentences is the choice of the mood of the antecedent. The conditional can be an indicative conditional even if the speaker does not fulfill the condition expressed by the antecedent, as in (49a, b):

(49a) If one is an angel, one is neither human nor divine.
(49b) If one is a Martian, one is not susceptible to human disease (Safir, 2000).

Thus, by using generic *one* in an indicative conditional, the agent does not just abstract from the particularities of his own situation, but rather can take the point of view of an entirely different situation, a situation in which he has quite different properties than he actually has. Given any standard semantics of indicative conditionals, this means that the pretence of having those counterfactual properties takes place directly, without the agent making any comparison between his actual situation and a counterfactual situation.

By allowing an agent to simulate counterfactual conditions directly, without comparison of his own actual situation, the behaviour of generic *one* in conditionals supports one of two views of simulation, namely that of Robert Gordon, as opposed to that of Alvin Goldman. On Goldman's (1989, 1992) view, an agent when simulating another person puts himself into the other person's situation and, making minimal adjustments, sees what *he himself* would do or what mental states he would be in, under the circumstances of that other person. Only then will he attribute that same behaviour or those same attitudes that he would display to the other person. On that view, one would expect the subjunctive in sentences such as (48a, b): the agent puts himself counterfactually into the situation of anyone meeting certain conditions and then sees whether he would, in that situation, also satisfy the consequence. On Gordon's (1986, 1995) view, by contrast, an agent, when

simulating another person, 'imaginatively transforms himself into the target', by taking *directly* the perspective of the other person, in order to predict the behaviour of that person or to attribute mental states to him. On that view, the simulator does not try to see how he himself would be in the other person's situation, but rather focuses on how the world would be from the very perspective of that other person. With Gordon's view of simulation, the agent's own situation will play no role in conditionals when the agent takes the position of anyone meeting the conditions of the antecedent; no comparison of the agent's actual situation with a counterfactual situation characterized by the antecedent of the conditional takes place, and thus no corresponding adjustments. Therefore no subjunctive mood is to be expected.

This notion of simulation, it seems, is also needed to account for the way generalizing self-reference as such is expressed in natural language. Generic *one* as well as the generic reading of controlled PRO with *imagine* shows that the notion of simulation at stake allows total and immediate abstraction from the agent's own situation: the agent projects himself onto a general condition he may not fulfill himself and sees what (in general) is the case when in fact fulfilling this condition. That is, it is Gordon's notion of simulation that is operative in the generalizing self-reference or first-person-oriented genericity that is expressed by generic *one*.

Generic *one* is also possible in counterfactual conditionals:

(50) If one were an angel, one would be neither human nor divine.

The question such counterfactual conditionals raise is, what enables the use of the subjunctive when in such conditionals the subjunctive is in fact not required? The appearance of the subjunctive indicates that in counterfactual conditionals *one* is associated with a contextual restriction that is incompatible with the counterfactual condition expressed by the antecedent. By contrast, in indicative conditionals there is either no contextual restriction or else the contextual restriction is not in conflict with the content of the antecedent.

Conditionals with generic *one* express a generalization on the basis of the first person, with the first person being used for a generalizing inference rather than providing an experience or action from which the generalization is made. With generic-*one* sentences, the agent pretends to meet the condition expressed by the antecedent and then sees whether he also meets the condition expressed by the consequent. What is crucial for the use of *one* in conditionals is the recognition that an inference involving the agent himself is generalizable, rather than the recognition that a particular kind of experience or action is generalizable. Unlike in the latter case (that is, with non-conditional generic-*one* sentences), generic *one* in conditionals is not a required choice: an indefinite like *someone* or *a person* in the antecedent and a definite pronoun in the consequent would have the same overall effect: (51a) and (51b) appear to be equivalent:

(51a) If one is 2 meters tall, one is tall.
(51b) If someone / a person is two meters tall, he is tall.

Generic *one* in conditionals allows verifying a conditional in a particular way: on the basis of an inference involving the first person. The same generalizing inference, though, could have been made on the basis of an arbitrarily chosen third person, and hence the near-equivalence of conditional generic-*one* sentences and conditionals with indefinites such as *someone* or *a person*.

Conditional generic-*one* sentences are also (near-)equivalent to conditional sentences with universal quantifiers:

(51c) Everyone is such that if he is two meters tall, he is tall.

However, there are use-related differences regarding the corresponding conditional with generic *one*: conditional generic-*one* sentences display an 'easy' first-person-based verification strategy, whereas conditionals with a universal quantifier require making sure, in some way or other, that the generalization holds for each individual in the domain of the quantifier.

5. Generalizing Detached Self-Reference and the Objective Self

Natural language expressions such as controlled PRO and generic *one*, I have argued, may express detached self-reference: self-reference while the agent identifies with others. This identification may be specific simulation of another person or generic simulation of anyone meeting certain conditions (generalizing detached self-reference). Both kinds of simulation are possible with controlled PRO, but generic simulation is the only option with generic *one* (and arbitrary PRO). In the present semantic analysis of those expressions, attitudes consisting in a pretend self-ascription of properties (imagination, desire) as well as the notion of simulation (that is, the identification relation I) play a central role.

The notion of detached self-reference is of interest not only to cognitive scientists (simulation theory) and (now) semanticists; it has also played a central role in some purely philosophical work. Most notably, the possibility of self-reference with potentially complete detachment from the actual person has been pursued by Thomas Nagel (1983, 1986) for a great range of philosophical issues, in epistemology, philosophy of mind as well as ethics. Nagel's notion is that of the 'objective self'. The objective self is the self that can dissociate itself entirely from the actual physical and psychological person and thus may take an objective point of view of the world. Nagel gives various descriptions of the objective self, such as:

> The picture is this: essentially I have no particular point of view at all, but apprehend the world as centerless. As it happens, I ordinarily view the world from a certain vantage point, using the eyes, the person, the daily life of TN [Thomas Nagel] as a kind of window. But the experiences of TN are not the point of view of the true [objective] self, for the self has no point of view

> and includes in its conception of the centerless world TN and his perspective among the contents of that world (Nagel, 1986, p. 61);

and:

> How do I abstract the objective self from the person TN? By treating the individual experiences of that person as data for the construction of an objective picture. I throw TN into the world as a thing that interacts with the rest of it, and ask what the world must be like from no point of view in order to appear to him as it does from his point of view. For this purpose my special link with TN is irrelevant. Though I receive the information of his point of view directly, I try to deal with it for the purpose of constructing an objective picture just as I would if the information were coming to me indirectly. I do not give it any privileged status with respect to other points of view (Nagel, 1986, p. 62).

Nagel does not actually take the objective self to be a separate object distinct from the actual person, but rather appears to consider it just one aspect of one and the same entity of which the actual person is another (and thus in the sentence 'I am TN' the two aspects, the objective self referred to by 'I' and the actual person referred to as 'TN', are identified as belonging to one and the same entity). The notion of the objective self remains a difficult one, though, in need of further clarification (see Stalnaker, 2003 for a critical discussion).

But still, given the notion of the objective self as it stands, first-person-oriented pronouns, one might say, are a linguistic manifestation of that notion in that they involve self-reference with full detachment from one's own actual person: first-person-oriented pronouns lead to the expression of a generalization generated by the relevant agent's generalizing detached self.

In fact, the involvement of the objective self in the semantics of generic *one* is nicely illustrated by Nagel's (1986, p. 9) own use of the pronoun. When taking the point of view of the objective self, Nagel uses generic *one*:

> Withdrawing into this element [the objective self] one detaches from the rest and develops an impersonal conception of the world and, so far as possible, from the elements of self from which one has detached. That creates the new problem of reintegration [. . .]. One has to be the creature whom one has subjected to detached examination, and one has in one's entirety to live in the world that has been revealed to an extremely distilled fraction of oneself.

In the next paragraph, however, when Nagel takes the point of view of the actual person ((generic) non-detached self-reference), he switches to a generic use of *we*:

> It is necessary to combine the recognition of our contingency, our finitude, and our containment in the world with an ambition of transcendence, however limited may be our success in achieving it.

We does not have the first-person-orientation of generic *one*. Instead it is a pronoun used generically or non-generically with the condition that the speaker (the speaker's actual person) be included among its semantic values. *We,* as in the passage cited, is not subject to the same restrictions on predicates as generic *one* (*today we have passports* is fine, for example), and it is not licensed in the same environments as generic *one* (*we can solve the equation*, for example, is unacceptable if the speaker knows that he himself couldn't).[11]

Thus the detached self-reference involved in the semantics of controlled PRO and generic *one* (with Strategy 1) seems to match Nagel's notion of the objective self. However, whereas in Nagel's work the objective self is given some quasi-objectual status (as one aspect of an entity of which an actual person is another aspect), the present semantic analysis made use only of attitudes and relations involving the actual person, attitudes of pretend self-ascription and relations of identification or simulation.

6. Strategy 2: Inference to the First Person

Let me finally turn (rather briefly) to the second strategy that licences generic *one*: Inference to the First Person. Inference to the First Person does not start out with a particular experience or action of the speaker, but rather with a generalization that has been established independently. This generalization, however, is presented with the intention to be at least potentially applied in a first-person way by whoever accepts the sentence, in particular the addressee. Given the analysis of generic *one* as a quantifier ranging over entities qua being identified with by the relevant agent, the relevance of this identification will now be not of an epistemic nature (providing the first-person epistemic basis for the generalization), but rather of an inferential practical sort. Given that in usual circumstances—that is, on the non-pretend use—an agent identifies with himself, it is the self-application of the predicate (by whoever accepts the sentence) that will make up the practical relevance. If a predicate Q applies to an entity x qua being identified with by the relevant agent, then the agent's identifying with x will be relevant in one way or another (epistemically, practically, or otherwise) for Q holding of x. The application of a predicate Q to an entity x qua P thus is to be understood more

[11] As was pointed out by an anonymous referee, *we* also allows for a use involving generic simulation, namely when *we* refers to the team / group / country / faction that the speaker identifies with, for example in a situation of team competition. *We found an objection* can be true even if the speaker did not actually find the objection, but identifies with the group representative that did.

generally as: *x* being *P* is relevant (epistemically, practically, or otherwise) for Q holding of *x*.

As with Strategy 1, there is a range of contexts in which generic *one* can be licensed by the second strategy. Most importantly, Strategy 2 is used in deontic sentences:[12]

(52a) One is not allowed to enter the room.
(52b) One should not lie.
(52c) One should be respectful toward the elderly.

Examples with arbitrary PRO are those with indirect questions below:

(53a) I will ask John since he knows what PRO_{arb} to do in such a situation.
(53b) John knows how PRO_{arb} to behave oneself.

Whereas with Strategy 1, the speaker's own experience leads to the generalization expressed by the generic-*one* sentence, in the present cases the speaker presents an internalized, but already established, generalization: a law, general requirement, or general recommendation. The generalizations expressed in (52a-c) crucially play a role in the speaker's reasoning for his actions, or better, are meant to play a role in the addressee's reasoning. For example, if the addressee accepts (52a), then this is likely to prevent the addressee from entering the room. That is, deontic generic *one* sentences are generally uttered with the intention that they play a future role as premises in the addressee's practical reasoning.

The reason why generic-*one* sentences making us of Inference to the First Person are so suited for governing an agent's practical reasoning is that they allow for an immediate first-person application by anyone who accepts them.

With generic-*one* sentences using Strategy 2, detached self-reference seems available too, as when a teacher says to his students:

(54) One has to hand in the essay today.

Thus, also with Strategy 2 self-application may take place in a detached way, at least on the part of the speaker.

Generic-*one* sentences using Inference to the First Person target potential actions on the part of the relevant intentional agent. They do so because they naturally act as premises for practical reasoning, that is, reasoning whose conclusion is an action (or at least the description of an action) on the part of the relevant agent. To make this more precise, consider a typical practical syllogism, involving the first person:

(55a) I intend to do E.
<u>I think that unless I do X, I cannot bring about E.</u>
I do X.

[12] One might have a different view about some normative sentences, namely those expressing ethical principles. Kant's categorical imperative would make them instances of Strategy 1, Inference from the First Person.

The second premise can be replaced by one that makes use of generic *one* instead of *I*:

(55b) I intend to do E.
I think that unless one does X, one cannot bring about E.
I do X.

Generic *one* is justified in sentences using Strategy 2 because such sentences are meant to lead to inferences, in particular practical inferences, in which self-reference is essential. Generic-*one* sentences are the only sentences truly suited for that purpose: they express generality and imply an immediate self-application by anyone who accepts them.

There are use-related differences between generic-*one* sentences and universally quantified and ordinary generic sentences that confirm the point:

(56a) One is not allowed to enter the room.
(56b) No one is allowed to enter the room.
(56c) A person is not allowed to enter the room.

In order to prevent the addressee from entering the room, a speaker would most naturally utter (56a), rather than (56b) or (56c). The addressee may even be excluded from the quantification domain: in (56b) because of an implicit contextual restriction and in (56c) because he is a legitimate exception. Laws and general advice are in fact typically formulated using generic *one*. The reason is that generic *one* carries both a general force and a first-person orientation. A deontic generic-*one* sentence commits an agent who accepts it to act in certain ways, having to make the sentence a premise in his practical reasoning.

The first-person connection of sentences with generic *one* or arbitrary PRO using Strategy 2 can also be seen from the suitability of such sentences for the expression of practical knowledge as in (57a, b), where generic *one* or arbitrary PRO could not be appropriately replaced by another generic NP, as in (58a, b):

(57a) I know what one can do.
(57b) I know what PRO_{arb} to do.
(58a) I know what people can do.
(58b) I know what people should do.

Strategy 2 is not always independent of Strategy 1. We have already seen examples, sentences expressing physical possibility, which could be licensed either by Strategy 1 or Strategy 2. With evaluative predicates, in fact, both Strategy 1 and Strategy 2 could be at play. With predicates of moral evaluation as in (59a, b), Strategy 2 will be clearly involved, whereas with predicates of some form of emotional evaluation, such as (60a, b), Strategy 1 will generally be the relevant one:

(59a) It is wrong PRO_{arb} to refuse to cooperate.
(59b) PRO_{arb} to help others is a virtue.

(60a) It is pleasant PRO_{arb} to walk in the park.
(60b) It is painful when one loses a parent.

At the same time, these generalizations are not strict: predicates of moral evaluation involve not only hypothetical practical reasoning, but also emotions of various sorts (I am outraged to have done X / that you have done X). Similarly, predicates of emotional evaluation will not just involve an inference from a first-person emotional state, but will also govern future practical reasoning. Evaluative predicates in general seem to involve both directions: inference from the first person (generalizing first-person mental states or acts) and inference to the first person (anticipating potential practical reasoning).

7. Conclusion

First-person-oriented genericity is a form of generalization that is associated with the meaning of a wide range of expressions, and most explicitly in English with generic *one*. It is clearly also a central cognitive notion. First-person-oriented genericity leads to self-locating beliefs and assertions, but it involves self-reference in a quite different way than familiar cases of pronouns interpreted *de se*, involving the notion of a self that can detach itself from the actual person. While detached self-reference is something familiar from contexts of imagination as well as the general philosophical work of Thomas Nagel, this paper has shown that detached self-reference, especially generalizing detached self-reference, also plays a central role in the semantics of certain kinds of generic sentences in natural language.

Institut d'Histoire de Philosophie de Sciences et de Technique IHPST
(Paris1/CNRS/ENS)

References

Anscombe, E. 1975: The first person. In S. Guttenplan (ed.), *Mind and Language. The Wolfscombe Lectures*. Oxford: Oxford University Press, 45–65.

Asher, N. 2006: Things and their aspects. *Philosophical Issues*, 16, pp. 1–23.

Asher, N. ms: *A Web of Words: Lexical Meaning in Context*. Ms. Toulouse: IRIT.

Castañeda, H-N. 1966: He*: A study in the logic of self-consciousness. *Ratio*, 8, 130–57.

Castañeda, H-N. 1967: The logic of self-attributions to others. *American Philosophical Quarterly*, 4, 85–100.

Chierchia, G. 1990: Anaphora and attitudes *de se*. In R. Bartsch *et al*. (eds), *Language in Action*. Dordrecht: Foris Publications, 1–31.

Cohen, A. 2002: Genericity. *Linguistische Berichte*, 10, 59–89.

Evans, G. 1982: Self-identification. In J. McDowell (ed.), *The Varieties of Reference*. Oxford: Oxford University Press, 205–33.

Fine, K. 1982: Acts, events, and things. In W. Leinfellner *et al.* (eds.): *Sprache und Ontologie. Proceedings of the Eighth Wittgenstein Symposion*. Vienna: Hoelder-Pichler-Tempsky, 97–105.

Gaerdenfors, P. 1986: Belief revision and the Ramsey Test for conditionals. *Philosophical Review*, 95, 81–93.

Goldman, A. I. 1989: Interpretation psychologized. *Mind & Language*, 4, 161–85.

Goldman, A. I. 1992: In defense of simulation theory, *Mind & Language*, 7, 104–19.

Gordon, R. M. 1986: Folk psychology as simulation. *Mind & Language*, 1, 158–71.

Gordon, R. M. 1995: Simulation without introspection or inference from me to you. In M. Davies and T. Stone (eds), *Mental Simulation. Evaluations and Applications*. Oxford: Blackwell, 53–67.

Greenberg, Y. 2005: *Manifestations of Genericity. Outstanding Dissertations in Linguistics*. New York: Routledge.

Heal, J. 1986: Replication and functionalism. In J. Butterfield (ed.), *Language, Mind, and Logic*. Cambridge: Cambridge University Press, 135–50.

Jubien, M. 2001: Propositions and the objects of thought. *Philosophical Studies*, 104, 47–62.

Krifka, M., Pelletier, F., Carlson, G., ter Meulen, A., Chierchia, G. and Link, G. 1995: Genericity. An introduction. In G. Carlson and F. Pelletier (eds), *The Generic Book*. Chicago: University of Chicago Press, 1–124.

Lewis, D. 1979: Attitudes *de dicto* and *de se*. *The Philosophical Review*, 88, 513–43.

McFarlane, J. 2005: Making sense of relative truth. *Proceedings of the Aristotelian Society*, 105, 321–39.

Moltmann, F. 1997: *Parts and Wholes in Semantics*. Oxford: Oxford University Press.

Moltmann, F. 2003: Propositional attitudes without propositions. *Synthese*, 135, 70–118.

Moltmann, F. 2006: Generic *one*, arbitrary pro, and the first person. *Natural Language Semantics*, 13, 257–81.

Moltmann, F. to appear: Relative truth and the first person. *Philosophical Studies*.

Nagel, T. 1974: What is it like to be a bat?. *Philosophical Review*, 83.4, 435–50.

Nagel, T. 1983: The objective self. In C. Ginet and S. Shoemaker (eds), *Knowledge and Mind*. Oxford: Oxford University Press, 211–32.

Nagel, T. 1986: *The View from Nowhere*. Oxford: Oxford University Press.

Ramsey, F. P. 1931: General propositions and causality. In *The Foundations of Mathematics*. London: Routledge and Kegan Paul, 237–55.

Recanati, F. 2000: Varieties of simulation. In J. Dokic and J. Proust (eds), *Simulation and Knowledge of Action*. Paris: Bibliotheque CREA, 151–71.

Recanati, F. 2007: *Perspectival Thought*. Oxford: Oxford University Press.

Safir, K. 2000: Keeping one's reference constant. Ms, Rutgers University.

Schiffer, S. 1987: *Remnants of Meaning*. Cambridge, MA: MIT Press.

Stalnaker, R. 1981: Indexical belief. *Synthese*, 49, 129–51.

Stalnaker, R. 2003: On Thomas Nagel's objective self. In R. Stalnaker (ed.), *Ways a World Might Be*. Oxford: Clarendon Press, 253–87.

Williams, B. 1973: Imagination and the self. In *Problems of the Self*. Cambridge: Cambridge University Press, 26–45.