THE JOURNAL OF PHILOSOPHY

VOLUME CIII, NO. 10, OCTOBER 2006

INTRASUBJECTIVE INTENTIONAL IDENTITY*

Some forty years ago, Peter Geach posed what he called the problem of *intentional identity*:[1]

(1) Hob thinks a witch blighted Bob's mare, and Nob thinks she killed Cob's sow.

One reading of this sentence, the *Geach* reading as I shall call it, can be true even if there are no witches, and even if neither Hob nor Nob has mistaken any particular person for a witch. So understood, the sentence seems to say that Hob's and Nob's attitudes are in some sense *about the same witch*, even though she does not (necessarily) exist. This reading, Geach argued, lies beyond the scope of conventional intensional formal languages for propositional attitudes.[2] If Geach is right about that, and if such languages mark the limit of what we understand about the logic of attitude ascriptions, then there is something about their logic that escapes us. Further, if Geach is right, and if such languages can be used to represent any proposition that your preferred semantical framework deems intelligible, then there is something about the content of propositional attitudes that lies beyond the scope of your preferred semantical views.

* For helpful comments and discussion, I wish to thank Lisa Downing, Sam Fleischacker, Dorothy Grover, Bill Hart, David Hilbert, Nick Huggett, Peter Hylton, Mitzi Lee, Abe Roth, Marya Schechtman, and an anonymous referee for this JOURNAL.

[1] Geach, "Intentional Identity," this JOURNAL, LXIV, 20 (October 26, 1967): 627–32. For simplicity I have replaced Geach's 'Nob wonders whether…' with 'Nob thinks that …'.

[2] By a conventional intensional formal language for the attitudes, I mean (roughly) one (a) that contains—in addition to the usual individual constants, predicates, truth-functional connectives, and apparatus for quantification and identity—operators or predicates for belief, desire, and other attitudes; and (b) for which all quantifier expressions carry existential import. I say "roughly" to allow for certain variations that do not affect the argument here. Often I will refer to these as conventional (or standard) intensional languages. Classic versions in a possible worlds semantical framework are analogous to those provided for alethic modality in Saul Kripke, "Semantical

0022-362X/06/0310/481–502

I shall review Geach's argument in the next section. My primary aim in this paper is to argue that Geach's problem is considerably more interesting—and more disturbing—than we realized. I argue that three strategies for deflecting the philosophical impact of Geach's puzzle—here called *Skepticism*, *Reductivism*, and *Reasoned Indifference*—are woefully inadequate. A fourth strategy posits nonexistent entities (Hob's and Nob's witch, for instance) as new semantical determinants. Yet this strategy, which I here call (*Intentional*) *Realism*, is generally regarded as something to be avoided at all costs. So we are in a nasty predicament.

Let me begin with an overview of the four strategies. The *Realist* strategy is the idea that the truth conditions of the Geach reading of (1) are best expressed by means of sentences of the form of (2) or (3).[3]

(2) $\exists\alpha$ [Hob believes (α is a witch & α blighted Bob's mare) & Nob believes (α killed Cob's sow)]

(3) $\exists\alpha\exists\beta$ [Hob believes (α is a witch & α blighted Bob's mare) & Nob believes (β killed Cob's sow) & $\alpha \approx \beta$]

Here, the quantifiers '$\exists\alpha$' and '$\exists\beta$' are interpreted referentially (rather than substitutionally), but they lack existential import. That is, they range over entities that in an appropriate sense "do not really exist." In some versions of this approach, the quantifiers in (2) and (3) are said to range over *intentional objects*, and for convenience I will adopt that term here, although I wish to be as noncommittal as possible about the nature of the relevant quantifier domains. The analysis given by (2) posits intentional objects that are global in the sense that they can "appear" in the attitudes of more than one person. The analysis given by (3) posits intentional objects that are local in the sense that they can appear in the attitudes of only

---

Considerations on Modal Logic," *Acta Philosophica Fennica*, XVI (1963): 83–94. For early work on this sort of approach, see Jaakko Hintikka, *Knowledge and Belief* (Ithaca: Cornell, 1962). For possible worlds versions with an applicative (function and argument) semantics, see David Lewis, "General Semantics," *Synthèse*, XXII (1970): 18–67; and Richard Montague, "The Proper Treatment of Quantification in Ordinary English," in Hintikka, Michael Moravcsik, and Patrick Suppes, eds., *Approaches to Natural Language* (Boston: Reidel, 1973), pp. 221–42. For versions that utilize logically structured propositions, see Lewis, "General Semantics," section V; and Scott Soames, "Direct Reference, Propositional Attitudes, and Semantic Content," *Philosophical Topics*, XV (1987): 47–87. There are important syntactic and semantic differences among such languages, yet none that will affect the arguments to be advanced here.

[3] Geach's sentence poses at least two distinct, but related, questions: (a) What is the formal syntax of the sentence? (b) What sort of formal language is required simply in order to state the truth conditions of the sentence? Sentences (2) and (3) are offered as an answer to the second question, not as an answer to the first. The suggestion that (2) or (3) correctly state the truth conditions for (1) would seem compatible with a number of theories as to its formal syntax.

one person, and posits a counterpart relation, $\approx$, among the local intentional objects.

For a variety of reasons, most contemporary philosophers find the Realist approach exceedingly unpalatable. It can be argued (a) that it violates a useful criterion of ontological commitment; (b) that it is ontologically extravagant; (c) that by countenancing nonexistent objects, it faces well-known logical difficulties; (d) that we lack a decent metaphysical theory of the postulated intentional objects; (e) that we lack principles of existence (or subsistence), identity, and predication for the postulated intentional objects; (f) that by postulating a domain of intentional objects mediating the relation between mind and world, the approach takes us too close to the problematic epistemology of Descartes and Locke; and (g) that we lack a theory of the intentional relation expressed implicitly by the double in-quantification in (2), and explicitly by the predicate '$\approx$' in (3). (Notice that (g) means that this approach cries out for a theory of mind-to-mind intentional relations, in addition the various mind-to-world intentional relations we usually countenance.) I hope I have not left anything out.

The three deflective strategies aim to skirt all of those difficulties. *Reductivists* argue that with some ingenuity we *can* after all express the truth conditions for the Geach reading of (1) in conventional intensional formal languages.[4] The idea is that a proper theory of the pronoun in the Geach reading of (1) will reveal it to be synonymous (at least relative to a context of use) with another English sentence in which the pronoun has been replaced by an appropriate definite, indefinite, or universal noun phrase. The latter sentence is to be *well-behaved* in the sense that its truth conditions can be expressed within the framework of standard intensional languages. *Skeptics* have argued in various ways that sentence (1) simply lacks the reading Geach claims for it.[5] The *Indifferent* do not care. This reaction can be perfectly reasonable, so it is possible to speak of *Reasoned Indifference*. When you have a well-supported theory, recalcitrant data are sometimes best simply ignored, especially if they are of little or no independent interest. No one has

[4] See, for instance, Daniel C. Dennett, "Geach on Intentional Identity," this JOURNAL, LXV, 11 (May 30, 1968): 335–41; and Montague, "The Proper Treatment of Quantification in Ordinary English." For an interesting and extremely well-developed Reductivist approach, see Jeffrey C. King, "Intentional Identity Generalized," *Journal of Philosophical Logic*, XXII (1993): 61–93, and "Anaphora and Operators," *Philosophical Perspectives*, VIII, *Logic and Language* (1994): 221–50.

[5] Dennett, "Geach on Intentional Identity"; and Montague, "The Proper Treatment of Quantification in Ordinary English" utilize a combined strategy. Reductive analyses are provided for certain versions of the Geach reading; skeptical arguments are then advanced against putative readings that do not succumb to the reduction.

ever argued that Geach's sentence is central to our philosophical, scientific, or other important concerns. It is hard to see what significant role the sentence could play in anything we really care about. So if the Skeptical and Reductivist strategies prove to be indefensible, and with a host of philosophical problems and puzzles facing the Realist strategy, the most reasonable course is surely to express passing interest in Geach's sentence, and then simply to *ignore* it as a bit of anomalous data.

My main aim in this paper is to try to sharpen and intensify the conflict between the Realist strategy and its deflective alternatives. Earlier I cataloged the most pressing problems for the Realist approach. I want to balance the scales by explaining what is deeply unsatisfying with the reactions of Skepticism, Reductivism, and Indifference. At the most salient level, I will be occupied with Indifference. I will argue that certain intentional identity sentences play crucially important roles in common-sense psychological explanations of human behavior. I shall also argue that these intentional identity statements are *essential* in performing these roles. That is, other, well-behaved attitude ascriptions cannot take over these functions. If successful, these arguments will ultimately show that all three deflective strategies are highly problematic. My hope is that the arguments will not only provoke renewed defenses of Skepticism, Reductivism, and Indifference, but stimulate a more serious consideration of Realism and how it could be provided with adequate philosophical foundations.[6]

I. OVERVIEW OF GEACH'S PUZZLE

Readings of Geach's sentence (1) can be distinguished by entailment patterns. On one reading, Geach's sentence entails both that a real witch exists, and that Hob and Nob have beliefs about the same real object. This is the so-called *de re* reading of the sentence, and it can be represented in conventional intensional formal languages by a sentence like (4):

(4) $\exists x[x$ is a witch & Hob believes (x blighted Bob's mare) & Nob believes (x killed Cob's sow)]

[6] It is not entirely clear that problems (a) through (g) are insuperable. On problem (a) see, for instance, William P. Alston, "Ontological Commitments," *Philosophical Studies*, IX (1958): 8–17; and Susan Haack, "Quantifiers," in *Philosophy of Logics* (New York: Cambridge, 1978), pp. 39–55. On problems (c) and (d), see Terence Parsons, *Nonexistent Objects* (New Haven: Yale, 1980). On problems (c), (d), (e), and (g), see my "Intentional Identity and the Attitudes," *Linguistics and Philosophy*, XV (1992): 561–96, and "A Perspectivalist Semantics for the Attitudes," *Noûs*, XXIX, 3 (1995): 316–42.

I mention the *de re* reading of (1) only to set it aside. I shall also set aside a neighboring *de re* construction carrying only the second of the two entailments I just mentioned:

(5) $\exists x$[Hob believes (x is a witch & x blighted Bob's mare) & Nob believes (x killed Cob's sow)]

An altogether different reading of (1) carries neither of the two entailments. This reading can be true in situations like the following:

Story 1. Bob's mare is sick, and Hob and Nob were called in for a joint diagnosis. "A witch blighted her," Hob says. "Yes," Nob replies, "and I'll bet she killed Cob's sow." In reality, Bob's mare has a congenital defect that is beginning to take its toll.

Story 2. Yesterday's newspaper carried a story about a witch named Helga who is on something of a rampage. Hob and Nob don't know one another at all, but each reads the newspaper story. Hob thought to himself, "I'll bet Helga blighted Bob's mare." Nob thought, "I'm sure Helga killed Cob's sow." But Helga doesn't really exist.

I shall call this the *Geach reading* of the sentence, or the *Geach sentence* for short. (One might also call it the *de dicto* reading, but that term would prejudge certain theoretical issues. The most we can say is that Geach's reading is *not* the *de re* one.)

Notice that the first conjunct of the Geach sentence falls easily within the scope of conventional semantical theories. For sentence (6) below can be analyzed as (7), often called the *de dicto* analysis of (6).

(6) Hob believes a witch blighted Bob's mare.
(7) Hob believes $\exists x$(x is a witch & x blighted Bob's mare)

This is important. That Hob has a belief about something that does not exist causes no problem at all for conventional theories. It is rather the intersubjective import of Geach's sentence that causes the problem for the standard accounts. To see this, notice first that we would normally formalize (8) in first order logic as in (9).

(8) A witch blighted Bob's mare and she killed Cob's sow.
(9) $\exists x$(x is a witch & x blighted Bob's mare & x killed Cob's sow)

That is, we would construe the pronoun 'she' as functioning like a variable, bound by the quantifier phrase 'a witch'.[7] The problem is that there is no way to combine the techniques of analysis we have

[7] The bound-variable analysis of (8) is even more apparent in Montague-style grammars for English. See, for instance, Montague, "The Proper Treatment of Quantification in Ordinary English."

applied to (6) and to (8) to arrive at a plausible analysis of the Geach sentence. For instance, (10) will not do:

(10) Hob believes $\exists x$[x is a witch & x blighted Bob's mare & Nob believes (x killed Cob's sow)]

The last occurrence of 'x' in (10) is bound by the quantifier, but to achieve this we have placed 'Nob believes' in the scope of 'Hob believes'. This analysis of the Geach sentence does not entail that Nob believes anything at all, so it is clearly unacceptable. An additional problem with (10) is that it entails that Hob has beliefs about Nob and about what Nob believes. As story 2 shows, the Geach sentence can be true even if Hob has no beliefs at all about Nob. (Sentences (4) and (5) above likewise construe the pronoun in (1) quantificationally, but these are *de re* constructions and do not capture the Geach reading.)

Alternative proposals abandon the idea that the pronoun in the Geach sentence is quantificational. One of these supposes that it is as a *pronoun of laziness*: a pronoun that goes proxy for a definite description constructible from its antecedent and the surrounding discourse. On this proposal one could take the Geach sentence to have the same truth conditions as (11), where the definite description takes narrowest possible scope:

(11) Hob believes a witch blighted Bob's mare, and Nob believes that *the witch that Hob believes blighted Bob's mare* killed Cob's sow.

The intended reading of (11) can be expressed in standard intensional languages as (12):

(12) Hob believes $\exists x$(x is a witch & x blighted Bob's mare) & Nob believes $\exists x$[Hob believes (x is a witch & x blighted Bob's mare) & $\forall y$(Hob believes (y is a witch & y blighted Bob's mare) $\supset$ $y = x$) & x killed Cob's sow]

This, of course, is an attempt to implement the Reductivist strategy. It should be clear that this particular proposal fails. Sentence (12) is true only if Nob has beliefs about Hob and what Hob believes; yet in story 2 the Geach reading of (1) is true even though neither Hob nor Nob knows of the other's existence.[8] Other Reductivist hypotheses can be advanced, and as the discussion unfolds I will be discussing

[8] On pronouns of laziness, see Peter Geach, *Reference and Generality* (Ithaca: Cornell, 1962), and "Referring Expressions Again," *Analysis*, XXIV (1964): 73–92. In "Intentional Identity," Geach considers and rejects the laziness analysis of (1), largely for the reasons stated above.

some of these—implicitly for most of the paper, but explicitly at the end. The prospects for Reductivism, we shall see, are rather dim.

If the truth conditions of the Geach sentence cannot be represented in conventional intensional formal languages for the attitudes, no proposition expressible in these languages would be what the Geach sentence asserts. Given a few plausible assumptions, it would follow that when the Geach sentence is true, the propositions believed by Hob or Nob, or both, are not expressible in these languages.[9] In that case, conceptions of propositions normally thought to provide adequate foundations for theories of content for the attitudes would have been shown to be insufficient for that purpose.

Let us take stock. Unless the Reductivist strategy can be successfully implemented, the truth conditions of the Geach sentence—and arguably the propositions Hob or Nob, or both, believe when it is true—cannot be expressed in languages that otherwise seem as powerful as one could wish for. Problems (a) through (g) eagerly await us if we try to extend the language and the underlying theories of content in the way the Realist suggests. No wonder, then, that people have reacted with skepticism about the linguistic intuitions that generate the problem. Or that some have argued that the Geach sentence ultimately does reduce, by principles governing the pronoun, to a sentence that falls within the scope of standard intensional formal languages and associated theories of propositional content. Or that the sentence generates only passing interest, thereafter to be largely ignored.

II. INTRASUBJECTIVE CASES

When he introduced the puzzle, Geach provided no hint as to what roles intentional identity statements might play in natural language: what purposes they might serve, what utility they might have. He simply argues that we cannot account for our linguistic intuitions about the sentence. Pursuing the problem in this fashion, it quickly becomes hard to see it as anything more than a logical curiosity—and

[9] Call a proposition expressible in conventional intensional formal languages for the attitudes an *ordinary* proposition, and a proposition that is not ordinary in this sense an *extraordinary* one. Then, as a first approximation, the following assumptions would suffice to establish, from the fact that the Geach sentence expresses an extraordinary proposition, that Hob or Nob, or both, believes an extraordinary proposition when the Geach sentence is true: (i) that the Geach sentence contains two belief ascriptions with Hob and Nob, respectively, as subjects; (ii) that belief ascriptions express relations between persons and propositions; and (iii) that the syntactic rules used to generate the Geach sentence cannot generate sentences expressing extraordinary propositions from sentences expressing ordinary ones.

one safely ignored. Some later illustrations of intentional identity, such as this example from Mark Richard, seem to fit that view of the significance of the problem:[10]

(13) Hob thinks the unicorn ate the petunias, but Nob thinks it probably didn't.

Considered in isolation from other kinds of cases, examples like Geach's (1) and Richard's (13) lend considerably plausibility to the strategy of Reasoned Indifference. Yet these illustrations create a false impression of the importance and difficulty of the problem. Intentional identity comes in both intersubjective and intrasubjective forms, and the latter play important roles in explaining human behavior, while still posing all the problems of philosophical analysis we find in Geach's intersubjective example. Take the case of Grandma, for instance.

Story 3. Grandma is stomping off toward the barn with her rifle. She looks very upset. "What's wrong, Grandma?" you ask. "It's a snake. A dad-blasted snake and it's eating the chickens. I'm going to hunt it down and shoot it." But there is no snake. The missing chickens have run off to play with the chickens on a neighboring farm.

We ask: Why is Grandma headed toward the barn, gun in hand? Within the framework of commonsense psychology, (14) provides a perfectly good explanation:

(14) Grandma thinks there's a snake in the barn, and she wants to shoot it.

The reading we want here is not the *de re* reading given by (15), nor the neighboring *de re* construction (16), since Grandma's belief is clearly not about any particular snake, nor even about any particular real object to the effect that it is a snake in the barn.

(15) $\exists x$[x is a snake & Grandma believes (x is in the barn) & Grandma wants (Grandma shoots x)]

(16) $\exists x$[Grandma believes (x is a snake & x is in the barn) & Grandma wants (Grandma shoots x)]

Nor will sentence (17) do, since it does not entail that Grandma wants anything, but only that she *believes* she does.

(17) Grandma believes $\exists x$[x is a snake & x is in the barn & Grandma wants (Grandma shoots x)]

[10] Richard, *Propositional Attitudes: An Essay on Thoughts and How We Ascribe Them* (New York: Cambridge, 1990), pp. 4–5.

On the Realist approach, we would state the truth conditions of the Grandma sentence (14) along the lines of (18) or (19):

(18) $\exists\alpha$[Grandma believes (α is a snake & α is in the barn) & Grandma wants (Grandma shoots α)]

(19) $\exists\alpha\exists\beta$[Grandma believes ($\alpha$ is a snake & α is in the barn) & Grandma wants (Grandma shoots β) & $\alpha \approx \beta$]

As in the intersubjective application of the Realist approach, these quantifiers range only over intentional objects that, in an appropriate sense, do not exist. (When the Realist approach is applied as in (19), we assume that the intentional objects are local not only in the sense that they cannot appear in the attitudes of more than one person, but also in the sense that they cannot appear in more than one attitude-type of a single person.) All the earlier objections (a) through (g) to the Realist approach apply here—except that (g) now becomes: (g′) we lack a theory of the intrasubjective intentional relation expressed implicitly by the double in-quantification in (18), and explicitly by the predicate '$\approx$' in (19). So we reach for the three deflective strategies here as much as in the case of Hob and Nob.

Unlike Geach's (1) and Richard's (13), the Grandma sentence (14) is actually useful in explaining human behavior. Such intrasubjective intentional identity statements are perfectly natural explanations of human behavior in terms of belief and desire. (Indeed, similar intrasubjective intentional identity statements might be useful in explaining (other) *animal* behavior. You hold your closed hand out to your dog, the way you always do when you have a treat for him. Why is he pawing at your hand and trying to open it? "He thinks you have a treat in your hand, and he wants to eat it." Only you are teasing him; your hand is empty.) The Skeptical response—that sentences (1) and (13) lack the reading Geach claimed for them—seems utterly preposterous in the case of (14). The reaction of Reasoned Indifference—that we need not provide a philosophical theory of the problematic readings of (1) or (13), since they serve no clear explanatory or other purpose—seems completely implausible in the case of (14).

Yet perhaps we can explain Grandma's behavior by means of some *well-behaved* sentence. The most promising idea would be to try to explain her actions by means of a sentence in which the pronoun in (14) has been replaced by a descriptive noun phrase (NP) of some sort. I call this the *descriptivist* explanatory strategy, and I shall consider three kinds of descriptivist candidate explanations: those in which the pronoun has been replaced by (i) an indefinite NP (*a K*), (ii) a definite NP (*the K*), or (iii) a universal NP (*every K*), where *K* is any common

noun phrase (such as 'snake' or 'snake such that it is in the barn'). In each case, the common noun phrase *K* may be either (a) *plain* (that is, free of doxastic expressions such as 'believes' or 'thinks'), or (b) *doxasticized* (containing at least one doxastic expression). Thus 'snake in the barn' is plain, 'snake believed by Grandma to be in the barn' is doxasticized. I extend this terminology to cover NPs formed by combining common noun phrases *K* with various determiners such as 'a', 'the', and 'every'. Combining (i), (ii), and (iii) with (a) and (b), we have six kinds of descriptivist candidates to consider all together.

We begin by first considering descriptivist candidates that rely on *plain* NPs:

(20) Grandma believes there's a snake in the barn. Grandma wants to shoot *a snake in the barn.*

(21) Grandma believes there's a snake in the barn. Grandma wants to shoot *the snake in the barn.*

(22) Grandma believes there's a snake in the barn. Grandma wants to shoot *every snake in the barn.*

For short I shall refer to these as the *plain indefinite*, *definite*, and *universal* candidates, respectively. I shall postpone discussion of the universal candidate (22) for a later section, and focus for now on (20) and (21).

One problem with explanations like (20) and (21) is that they are subject to what I shall call the *Distribution Problem.* On the desired reading, the second sentences in (20) and (21) are expressed in standard intensional languages as (20a) and (21a), respectively:

(20a) Grandma wants $\exists x(x$ is a snake & x is in the barn & Grandma shoots $x)$

(21a) Grandma wants $\exists x[x$ is a snake & x is in the barn & $\forall y(y$ is a snake & y is in the barn $\supset y = x)$ & Grandma shoots $x]$

(We need the *de dicto* readings here, since the *de re* readings of (20) and (21) are false in story 3.) The trouble is that (20a) and (21a) are also false in the story. For assuming that desire distributes over conjunction—a principle I defend in the next section—(20a) and (21a) both entail (23):

(23) Grandma wants $\exists x(x$ is a snake & x is in the barn)

Yet (23) is false in the story. For (23) says that Grandma wants it to be the case that there is at least one snake in the barn. Grandma despises most snakes. She certainly does not want any in the barn. Since the intended readings of (20) and (21) entail (23), and (23) is false in story 3, it follows that (20) and (21) are false in story 3. So (20) and (21) cannot explain Grandma's behavior in that story. (More precisely, the argument is this: given the Distribution Principle defended in the next

section, and given that the intended readings of (20) and (21) fall in the scope of conventional intensional languages, these two sentences each entail (23), which is clearly false in story 3. So either the Distribution Principle is false, or the intended readings of (20) and (21) lie outside the scope of conventional intensional languages, or (20) and (21) are false in story 3 and cannot explain Grandma's behavior.[11])

We now turn to candidate explanations that appeal to *doxasticized* NPs. For short I shall refer to these as *doxasticized* indefinite, definite, and universal candidates:

(24) Grandma thinks a snake is in the barn. Grandma wants to shoot *a snake such that she thinks it is in the barn.*

(25) Grandma thinks a snake is in the barn. Grandma wants to shoot *the snake such that she thinks it is in the barn.*

(26) Grandma thinks a snake is in the barn. Grandma wants to shoot *every snake such that she thinks it is in the barn.*

Again, I want to postpone discussion of the universal candidate (26) and to focus for now on (24) and (25). These are subject to a doxasticized version of the Distribution Problem. For the intended (narrow scope) readings of the second sentences in (24) and (25) both entail (27):

(27) Grandma wants $\exists x[x$ is a snake & Grandma believes (x is in the barn) & Grandma shoots $x]$

Assuming again that desire distributes over conjunction, (27) entails (28):

(28) Grandma wants $\exists x[x$ is a snake & Grandma believes (x is in the barn)]

Yet (28) is false. Grandma wants no such thing. Yes, she believes a snake is in the barn. But she does not really *want* there to be a snake such that she believes it to be in the barn. "There are some things you'd just as soon not know about, and this snake he's one of them. My heart, you know. But since he's there and I've got my chickens to think about, well I just have to go on and shoot it." In some versions of the story, there will be a different reason why (28) is false. Suppose Grandma knows that when it comes to snakes in the barn, she is very keen and rarely issues false positives. She knows that if

[11] I would suggest that (20) is simply false in the story. On the other hand, (21) seems ambiguous between a false reading given by (21a), and a true reading in which the definite description 'the snake in the barn' is wholly anaphoric on the indefinite description 'a snake'. On the latter reading, (21) is little more than a stylistic variant of (14), and like (14) it lies beyond the scope of conventional semantical theories.

she believes there is a snake in the barn, this is very likely true. She believes there is a snake in the barn, but strongly prefers the situation in which the barn is snake-free and in which she accordingly does not believe there is a snake in the barn (and she strongly prefers that situation to the one in which the barn is snake-free but in which she falsely believes there to be a snake). So she does not want it to be the case that she believes there is a snake in the barn. It seems that no definite or indefinite candidate explanation, whether plain or doxasticized, can successfully explain Grandma's behavior in story 3.

III. THE DISTRIBUTION PRINCIPLE

The argument I used to generate the Distribution Problem relied on the assumption that desire distributes over conjunction:

(29) *S* wants $(p \,\&\, q)$ / $\therefore$ *S* wants p

At first there seem to be counterexamples. One might suggest, for instance, that (30) does not entail (31), and similar examples are easy to come by.

(30) I want to put on a parachute and jump out of a plane.
(31) I want to jump out of a plane.

The desire expressed in (30) is conditional—I want to jump out of a plane, *but only if I (first) put on a parachute*—and (31) neglects to mention this, so (31) is false.

Or so one might argue. This argument against the Distribution Principle relies on the assumption that where the content of desire is conditional, it is false to ascribe a desire with that content unless you explicitly mention the condition. Since it is not clear how we should understand the logical form of conditional desire ascriptions,[12] we would do better to formulate this assumption by saying that argument form (32) is valid:

(32) I want $(p \,\&\, q)$
$\sim$ [I want $(\sim p \,\&\, q)$]
$\therefore \sim$ (I want q)

[12] '*I want A, but only if B*'. Does this have the form $WA \supset B$? Or $W(A \supset B)$? Should that material conditional be a counterfactual conditional? Or is this kind of conditional desire to be formalized with a binary conditional desire operator, as $W(A//B)$? How does tense figure in? '*If B, then I want A*'. Does this have the form $B \supset WA$? Or $W(B \supset A)$? Should that material conditional be a counterfactual conditional? Or is this kind of conditional desire to be formalized with a binary conditional desire operator, as $W(A/B)$? At least for the latter sort of conditional desire report, the problems and theoretical options seem analogous to those for conditional obligation, on which there is an extensive literature: for an overview and references, see Risto Hilpinen, "Deontic Logic," in Lou Goble, ed., *The Blackwell Guide to Philosophical Logic* (Malden, MA: Blackwell, 2001), pp. 159–82; see especially §8.5.

If argument form (32) is valid, one could then argue that since (30) and (33) below are both true, (34) is also true; hence (31) above is false.

(30) I want to put on a parachute and jump out of a plane.
(33) I do not want to not put on a parachute and jump out of a plane.
(34) I do not want to jump out of a plane.

The trouble with such arguments against the Distribution Principle is that they are easy to stand on their head. I want to put on a parachute and jump out of a plane—yes, but only if the ripcord is attached to the plane, and only if the plane does not lurch suddenly just as I am about to jump, and only if I do not land in a tree, and only if I am not hit on the way down by another airplane, and.... The problem is that if argument form (32) is valid, then (30) itself is going to be false—and it is going to be very difficult indeed to find a finite replacement.

The more general point, of course, is that most of our desires are conditional. I want to go rock climbing, but not if I break my neck in the process. You want to buy a new camera, but not if you will lose it the very next day. He really wants to eat this hamburger, but not if he will choke to death on it. In each case, the conditional desire can be expressed in terms of a conjunction: He wants to eat this hamburger, but *he does not want to eat this hamburger and choke on it.* It would be ludicrous to infer from the italicized conjunct that he does not want to eat the hamburger. The unconditionalized statement is true—he wants to eat this hamburger—and similarly for the other cases. So, I want to suggest, if you *really do* want to put on a parachute and jump out of a plane, then you do want to jump out of a plane. The Distribution Principle is valid.[13]

This is one of those many places where it is important to distinguish between information pragmatically imparted by a use of a sentence in a context, and the proposition semantically expressed by the sentence in the context.[14] Background information and context can af-

[13] The Conjunction Principle is different: *S* wants *p*, *S* wants *q* / ∴ *S* wants (*p* & *q*). This inference is clearly invalid. You can want your brother to come to Thanksgiving dinner and want your sister to come to Thanksgiving dinner without wanting them both to attend.

[14] See, for instance, Paul Grice, "Logic and Conversation," in Donald Davidson and Gilbert Harman, eds., *The Logic of Grammar* (Encino, CA: Dickenson: 1975), pp. 64–75; Saul Kripke, "Speaker's Reference and Semantic Reference," in Theodore Uehling, Peter French, and Howard Wettstein, eds., *Midwest Studies in Philosophy*, Volume 2, *Contemporary Perspectives in the Philosophy of Language* (Minneapolis: Minnesota UP, 1977), pp. 64–75, especially §3; Jon Barwise and John Perry, *Situations and Attitudes* (Cambridge: MIT 1983), especially pp. 258–64; Nathan Salmon, *Frege's Puzzle* (Cambridge: MIT, 1986); Soames, "Direct Reference, Propositional Attitudes, and Semantic Content," especially §v, and "Substitutivity" in Judith Jarvis Thomson, ed., *On Being and Saying: Essays in Honor of Richard Cartwright* (Cambridge: MIT, 1987), pp. 99–132.

fect how we revise our understanding of a person's psychological states when we accept a new report of her desires. The therapist hears his depressed patient say, "I want to jump out of a plane. I want to overdose on pills." The skydiving instructor, after a presentation to recruit new students, hears someone say, "I want to jump out of a plane; I want to feel the exhilaration." The therapist comes to believe the patient desires to jump out of a plane without a parachute; the skydiving instructor comes to believe the prospective student wants to jump out of a plane while wearing a parachute. Yet in both cases, the unadorned report of desire is perfectly true. The content of the unadorned report, together with background information and context, can provide the audience with information that goes beyond the content of the report itself.

Another reason for accepting the Distribution Principle looks to the way desires are processed in instrumental reasoning, and is much more speculative. Suppose that this processing proceeds by subgoaling.[15] In particular, suppose that if a person desires that $(p \,\&\, q)$, the processing of that desire automatically generates subgoals p and q. If subgoals—at least those introduced in this specific way from conjunctive desires—are themselves desires, then the Distribution Principle must hold.

IV. UNIVERSAL NOUN PHRASE CANDIDATES

I claimed earlier that intentional identity statement (14) provides a good, common-sense explanation of Grandma's behavior in story 3, and that no well-behaved attitude ascription provides as good an explanation.

(14) Grandma thinks there's a snake in the barn, and she wants to shoot it.

In section II, I argued that proposed explanations that appeal to narrow scope definite or indefinite NPs, whether plain or doxasticized, fail because they fall prey to the Distribution Problem. I postponed discussion of proposed explanations that replace the recalcitrant pronoun with a narrow scope *universal* NP, such as plain (22) or doxasticized (26).

(22) Grandma thinks a snake is in the barn. Grandma wants to shoot *every snake in the barn.*

(26) Grandma thinks a snake is in the barn. Grandma wants to shoot *every snake such that she believes it to be in the barn.*

[15] See Elaine Rich, *Artificial Intelligence* (New York: McGraw-Hill, 1983), pp. 25–53; and Keith J. Holyoak, "Problem Solving," in Daniel N. Osherson and Edward E. Smith, eds., *Thinking* (Cambridge: MIT, 1990), pp. 117–46.

On the readings we want, the truth conditions of (22) and (26) are represented by (22a), and (26a), respectively:

(22a) Grandma believes $\exists x$(x is a snake & x is in the barn). Grandma wants $\forall x$[(x is a snake & x is in the barn) $\supset$ Grandma shoots x].

(26a) Grandma believes $\exists x$(x is a snake & x is in the barn). Grandma wants $\forall$x[Grandma believes that (x is a snake & x is in the barn) $\supset$ Grandma shoots x)]

The Distribution Problem is obviously no difficulty whatsoever for (22a) and (26a). So universal NP candidates like (22) and (26) would seem to allow us to explain Grandma's behavior without intentional identity statements like (14). For those who wish to pursue the strategy of Reasoned Indifference, this greatly increases the interest and importance of the universal NP candidates. Yet (22) and (26) are themselves subject to other serious objections. (Some of these objections will apply with equal force against the definite and indefinite NP candidates considered earlier.)[16]

The Specificity Problem. This is best explained by an illustrative example:

> *Story 4.* Grandma has a pet garter snake whom she calls "Jake." All too often, Jake slips out of his cage and slithers off to the barn to eat bugs. She thinks Jake is probably in his cage, but she acknowledges that there's also a fair chance he's out in the barn. If Jake is in the barn, she certainly doesn't want to shoot him. Grandma mistakenly believes there is a wild snake in the barn, and that it's after the chickens. She wants to shoot it.

Here, intentional identity statement (14) still provides a good explanation of Grandma's behavior. Yet (22) is false: it is simply not the case that Grandma wants to shoot every snake in the barn, on the reading given by (22a). A slight variant of story 4 will show that the doxasticized version (26) will not do, either. Suppose Grandma *does* believe that Jake is in the barn, but she certainly does not want to shoot *him*. Intentional identity statement (14) will still explain Grandma's behavior in the variant version, but (26) will be false on the reading given by (26a).[17]

[16] Notice that (22) is equivalent to "Grandma believes there's a snake in the barn. Grandma wants it to be the case that if there is a snake in the barn, then she shoots it." So the arguments against the universal NP candidate (22) apply equally well against the conditionalized version. Similar remarks apply to (26) and its conditionalized version.

[17] Notice that "I dropped a coin and it fell into the gutter" is true even if you dropped two coins, and only one of them fell into the gutter. Similarly, sentence (14) is true even if Grandma thinks there are two snakes in the barn, and she wants to shoot only one of them.

Of course we might be able to supplement the descriptive content in the relevant NPs, to home in on the right snake in story 4 and its variant:

(35) Grandma thinks a *wild* snake is in the barn. Grandma wants to shoot every object such that it is (or she believes it to be) a *wild* snake in the barn.

Yet in an extended version of story 4, in which Grandma likes wild garter snakes and thinks that there could (also) be one of these in the barn, we would have to complicate the explanation further:

(36) Grandma thinks a snake *such that it is wild but not a garter* is in the barn. Grandma wants to shoot *every snake in the barn such that it is wild but not a garter.*

If we try to explain Grandma's behavior along these lines, we must include in our explanation each such qualification, as Grandma's beliefs and desires about snakes become more nuanced. Some might be sanguine about the prospects for solving the Specificity Problem in this way. But there is good reason to be skeptical. More often than we would like to think, our attitudes toward objects are determined by features of our mental life that we are insufficiently aware of for them to enter into the contents of our beliefs and desires: our moods, emotions, the degree of stress we are under, whether we are well-rested or tired, our general level of happiness, background noise or other distractions, associations inaccessible to consciousness, and so on. Suppose you are meeting someone for the first time. Your attitudes toward that person, including some persistent desires, can be determined to a large extent by these sorts of factors. In that case, you might have persisting desires about that person, even though these factors—for instance, that at the time of meeting you were in a very good (or dreadful) mood that you were calm and at peace (or totally stressed out), or that the person unconsciously reminded you of a childhood classmate you loved (or hated)—do not enter into the content of any of your beliefs and desires. Love and hate probably work more in this way than we would ever like to admit.

Suppose that this is how it is with Grandma and her snake: the one she wants to shoot. What distinguishes the snake she wants to shoot from other snakes is simply that at the time she came to believe it existed, she was in a certain mental state. Suppose she is so unaware of being in that state that none of her attitudes has as any part of its content that she was in that state at that time. Here, it will not do to supplement the universal noun phrase in (22) or (26) with predicates like 'wild', or 'chicken-eating' or anything else of this sort.

One might try to solve the Specificity Problem by introducing into the universal NP a complex predicate that singles out the right snake by referring to the episode in which the belief was first acquired, perhaps by referring to the time of that episode:

(37) Grandma believes there is a snake in the barn. There is a time *t* such that Grandma first came to have this belief at time *t*, and Grandma wants it to be the case that she shoots *every object x such that she first came to believe at t that x is a snake in the barn.*

One problem with (37) is that this explanation will work only if Grandma has the concepts of time and belief, since these concepts figure in the content of her desire according to (37). Such an assumption might work in the case of Grandma, but it will not work if we are explaining the behavior of a very small child. A second problem emerges if Grandma is mistaken about the time at which she first acquired the belief in question. Suppose she first acquired a belief that a snake (Jake, for instance) was in the barn at 10:00 a.m., but did not want to shoot that snake. At 11:00 a.m. the same day, she first acquired a belief that there was a "second" snake in the barn, and because of her mood and other such factors, she acquired a persistent desire to shoot that snake. Suppose she gets it backwards: she believes, of 10:00 a.m., that this is when she first acquired the belief about the snake she wants to shoot; and she believes, of 11:00 p.m., that this is when she first acquired the belief about the snake she does not want to shoot. In that case, (14) will still explain why Grandma is headed toward the barn with her rifle, but (37) will be false. A third problem is that (37) seems at best a highly misleading explanation of Grandma's behavior. For it suggests that what is important *to Grandma* about the snake, her reason for wanting to shoot it, is that she first became aware of its existence at a certain time.

Clearly, the Specificity Problem poses similar difficulties for the definite and indefinite candidates we considered earlier.

The Problem of Targeted Action. Another problem with the universal NP replacements emerges when we consider cases like the following one.

Story 5. When Grandma wants to clear the barn of wild snakes, she takes her new rifle because it reloads more quickly and doesn't have such a mean recoil. She does this even if she believes there is just one wild snake in the barn. "You might think there's just one," she says, but you can never tell when you might find more." When on the other hand she wants to kill just one particular wild snake she believes to be there, she takes her grandfather's rifle which for sentimental reasons she enjoys using more than the new one. On the present occasion, Grandma takes her grandfather's rifle. "I'll get him with this for sure," she says. She believes there is just one snake, but it is compatible with her beliefs that there are others.

The problem is that for any given kind *K*, an agent may act quite differently when the agent wishes to do something to *every instance* of kind *K*, and when she wishes do it to *a particular* instance of kind *K* that she believes to exist. This difference can persist even when the agent currently believes there is only one instance of kind *K*, provided that the agent realizes that she may later revise her beliefs to allow that there is more than one *K*. Intentional identity statement (14) will explain Grandma's behavior in story 5. Universal NP candidates like (22) will not:

(22) Grandma thinks a snake is in the barn. Grandma wants to shoot *every snake in the barn.*

This simple illustration takes the value of *K* to be the property of *being a snake in the barn*, but clearly the problem will persist with other (nondoxasticized) values of *K*. (It may be difficult to argue that the Problem of Targeted Behavior applies to candidates in which *K* is doxasticized in a way that refers to Grandma's present beliefs. Yet below we shall see there is an independent argument against the doxasticized versions.)

The Problem of Case-by-Case Strategies. In some situations, the agent might not have adopted any rule to the effect that every instance of kind *K* is to be shot, because she wants to decide on a case-by-case basis. Sometimes it can be hard to foresee every possibility. "If I find a snake that's not a garter, I'll kill it. But not if it has a collar around its neck indicating that it's someone's pet. But then again, if the collar says the snake belongs to Alvin, I'll shoot it. Unless Alvin apologizes. Unless he apologizes with his fingers crossed...." Suppose that Grandma has adopted a policy that she is going to treat the *K* instances on a case-by-case basis. She believes that a *K* snake is in the barn, and she has decided to shoot "this one." In that case, (14) still provides a good explanation. But because she has decided to treat *K* instances on a case by case basis, sentence (38) will be false:

(38) Grandma thinks that a (*K*) snake is in the barn. Grandma wants it to be the case that she shoots every *K* snake in the barn.

(It is tricky to argue that this problem applies to candidates in which *K* is doxasticized in a way that refers to Grandma's present beliefs. Yet below we shall see there is an independent argument against the doxasticized versions.)

The Practical Inference Problem. This objection applies only to the doxasticized versions. In practical reasoning, we sometimes form intentions on the basis of our beliefs and desires. Grandma does this when she forms an intention to shoot the snake on the basis of her

belief that there is a snake in the barn, and her desire to kill it. Let us try to home in on the inference pattern that Grandma relies on in making her inference. The logic of practical inference is of course far less clear than we would like it to be, yet schema (39) represents a form of inference that seems correct in respects that interest us here:

(39) *S* believes that there exists at least one *F*.
S desires that it will be the case that every *F* is *G*.
S believes that for any *x* that is *F*, the most effective means of seeing to it that it will be the case that *x* is *G* is to see to it that *S* does *A* to *x*.
∴ *S* intends to see to it that *S* does *A* to some *F*.

No doubt we would need to add certain provisos to allow for overriding desires and various other factors, but this will not affect the argument I wish to make. Moreover, schema (39) is stated in a form that includes the relevant propositional attitude in each step, as we would if we were explaining or predicting *S's* practical inference. In an actual logic of practical inference, it is less clear that these operators would be included in the premises of the inference.[18] This issue also does not affect the present argument.

Let us see what happens when we try to work with (26) as the basis for our explanation of how Grandma reasons in accord with schema (39).

(26) Grandma thinks a snake is in the barn. Grandma wants to shoot *every snake such that she believes it to be in the barn.*

Here we run into a problem. For (40) below is not an instance of schema (39), since we do not have a consistent value for '*F*' throughout.

(40) Grandma believes that there exists at least one *snake in the barn.* Grandma wants it to be the case every *snake such that she believes it to be in the barn* is killed by her. Grandma believes that the most efficient way to see to it that any *snake in the barn* is killed by her is for her to shoot it. Therefore, Grandma intends to shoot some *snake in the barn.*

Rather, we would need (41):

(41) Grandma believes there is at least one *snake such that she believes it to be in the barn.* Grandma wants it to be the case that every *snake such that she believes it to be in the barn* is killed by her. Grandma believes that any *snake such that she believes it to be in the barn* could most effectively be killed by her if she shoots it. Therefore, Grandma intends to shoot some *snake such that she believes it to be in the barn.*

[18] See John Broome, "Practical Reasoning," in José Luis Bermúdez and Alan Millar, eds., *Reason and Nature: Essays in the Theory of Rationality,* (New York: Oxford, 2002), pp. 85–112.

In other words, the doxasticization of desire we see in (26) infects the whole of Grandma's practical inference. This is not just implausible, but wildly so. We would have one (plain) form of practical inference for reasoning about *real* snakes, and another (doxasticized) form for reasoning about snakes that are believed to exist, but do not.

Here is a much simpler account of Grandma's practical reasoning:

(42) Grandma believes that there is a snake in the barn, and she wants to kill it. Grandma believes that she could most effectively kill it by shooting it. Therefore, she intends to shoot it.

Obviously the practical inference schema on which (42) relies involves intentional identities. Yet I conjecture that without intentional identity it will be impossible to provide a plausible account of practical inference. All the more reason for thinking we cannot ignore it.

The Tracking Problem. (This objection applies to definite, indefinite, and universal candidates, whether plain or doxasticized.) We want to be able to explain Grandma's behavior as she tracks the snake (which does not really exist). *Now* Grandma is going up into the hayloft with the rifle. Sentence (43) explains this stage of her tracking behavior perfectly well:

(43) Now Grandma thinks a snake (or the snake) is up in the hayloft, and she (still) wants to shoot it.

If we are attempting to explain her behavior by appealing to universal candidates, we would need something like the following:

(44) Now Grandma thinks that a snake is up in the hayloft. Grandma wants to shoot every snake such that it is up in the hayloft.

As Grandma hunts, her beliefs change: she acquires new beliefs and loses some old ones. "It was in the cow's stall. But when I opened the barn door I heard him slither up into the hayloft." To explain Grandma's behavior by means of universal, definite, or indefinite candidates, we must assume that the content of Grandma's desire changes with each slight alteration in her beliefs about the snake. Explaining Grandma's behavior by means of an intentional identity statement like (14) carries no such commitment. This is an important difference. Grandma's beliefs purport to track the snake. To explain her behavior in terms of her beliefs and desires, these mental states must be coordinated around the "figure" of the snake. The descriptivist approaches we have been considering would appear quite plausible if we assumed that this coordination is achieved (only) through the descriptive content of the relevant beliefs and desires: as Grandma's beliefs about the snake shift, so does the content of her desires, in such a way

that the desires and beliefs share certain parts of their descriptive content. Yet this is a substantive assumption about mental processing—and one that seems highly implausible. A more reasonable picture assumes that such beliefs and desires are coordinated not by descriptive content, but by certain causal or functional relationships among aspects of the relevant mental representations. (A system of causally linked tags, for instance, would seem much more efficient than content-matching, when content on the side of belief is rapidly changing.)

V. CONCLUSION

In the opening section I distinguished four possible reactions to Geach's puzzle about intentional identity: Skepticism, Reductivism, Indifference, and Realism. Most of my argument in this paper aimed to defend two theses about the general framework of commonsense psychology:

(a) In this framework, we often explain human (and perhaps other animal) behavior by means of intentional identity statements.
(b) Omitting these statements from the framework would significantly diminish its explanatory scope and power.

The examples offered here in support of these two claims illustrate how implausible is the Skeptical claim that intentional identity statements lack the readings that Geach claimed for them. The arguments show also that the strategy of Reasoned Indifference is in serious trouble, since these sentences do play an important role in commonsense psychology, a role that cannot be played by the well-behaved sentences of the framework.

If correct, the arguments also critically undermine the Reductivist strategy. As I mentioned earlier, Reductivist accounts typically take the form of a hypothesis about the nature of the problematic pronoun in intentional identity statements. According to one, considered briefly earlier, the pronoun is a pronoun of laziness that goes proxy for a narrow scope definite description built up out the antecedent and surrounding discourse. A more liberal hypothesis claims that the pronoun functions as variable bound by a "substitutional quantifier" phrase that takes (narrow scope) definite descriptions as substituends. A third hypothesis proposed by Jeffrey King claims that the pronoun is a "context dependent quantifier."[19] According to this theory, the pronoun functions semantically like a quantifier phrase whose characteristics are determined in a rule-governed way from appropriate features of context. These theories all share the thesis that the truth conditions for (tokens of) intentional identity statements can be ex-

[19] See King, "Intentional Identity Generalized," and "Anaphora and Operators."

pressed in terms of well-behaved propositional attitude statements.[20] To the extent that the theories rest on the underlying idea that intentional identity is to be understood not in terms of causal or functional relations but in terms of identity of descriptive content, they can be seen as vestiges of descriptivist theories of reference and intentionality. Yet if no well-behaved propositional attitude statements can, without a significant loss of explanatory power, replace intentional identity statements in explanations of behavior, these Reductivist theories are in deep trouble. (Indeed, many of the candidate substitutes for intentional identity were not even true in the situations under consideration.)

These are of course plausibility arguments and are not intended to be absolutely conclusive. Yet too much discussion of intentional identity has proceeded without any attempt to consider how the sentences are actually employed. The unfortunate result has been that we sometimes underestimate the robustness of the intuitions supporting the recalcitrant readings, undervalue their theoretical interest, and misjudge the difficulty of providing a satisfying theoretical account. The Realist approach is, as I mentioned, not without serious problems of its own, yet the alternatives—at least those considered here—are not nearly as attractive as one might have supposed.[21]

WALTER EDELBERG

University of Illinois at Chicago

[20] I say "in terms of" because on the substitutional-quantificational interpretation of intentional identity, it is not claimed that some finite well-behaved sentence expresses the truth-condition. Rather, the intentional identity statement is claimed to be true if one or more of its infinitely many well-behaved substitution instances is true.

[21] For reasons of space, I have here limited my focus to how intrasubjective intentional identity figures in explanations of human behavior. In a subsequent paper, I plan to take up how intersubjective intentional identity figures in explaining the behavior of collectives.