

Chapter 1

Adequacy Conditions and a Prototheory

What is it to know who someone is? There should be some enlightening general answer to this question, however crude. Speakers of everyday English often claim to know who such-and-such a person is or ascribe such knowledge to others. Philosophers concerned with the logic of knowing, with the natural functions of singular terms, or with the nature of doxastic states appeal to the notion of "knowing who" as if it were sufficiently clear to go on with.¹ But until recently, the literature contained no more than a half-hearted discussion of the notion. In the 1975 pilot article on which this book is based, we tried to fill this lacuna; here we make a far more thorough attempt to explicate "knowing who" and related locutions. In later chapters we bring our analysis to bear on related philosophical issues.

Some obvious questions: Can we provide an interesting and useful taxonomy for the different sorts of answers that might, under various circumstances, be given to queries of the form "Who is N?"? Given the (much-conceded) multiplicity of tests for "knowing who" appropriate to different sorts of situations, is there any single canonical paraphrase appropriate to all instances of "S knows who N is"? In the end can we give a general theory of "knowing who" that illuminates the traditional issues that have been supposed to hang on the notion?

We concern ourselves mainly with expressions of the form "S knows who N is," though several cognate locutions deserve individual treatment ("S doesn't know who N is," "S remembers who N is," "S wonders who N is," "S told T who N is," etc.). It is important to bear in mind throughout that in such formulas "N" can be replaced by terms of several different sorts: proper names, demonstratives and other pronouns, and singular and plural definite descriptions.

1 Multiple Tests and Privileged Facts

It is a truism that a well-reasoned answer to the question "Does S know who N is?" may rest, in different situations, on widely disparate grounds; different circumstances demand different tests for knowing who someone is. Although such a test normally consists of eliciting an appropriate singular term from S

in response to "Who is N?," the standard of "appropriateness" differs dramatically.

It is easy to imagine situations in which an informant S's response to "Who is N?" prompts a repetition of the question; that is, S might respond to our original query with "N is M," "M" being a singular term distinct from "N,"² and we might well pursue the matter by asking "But who is M?" So far as grammar is concerned, this process could continue indefinitely, were we sufficiently loath to credit S with knowledge of N's identity. However, it seems that if S really does know who N (= M, = ...) is in the ordinary sense of that expression, the serial questioning would (from the epistemological point of view) have to stop; at some point in the chain of questions, S would have provided us with contextually conclusive evidence (whether or not we see fit to accept that evidence) that S does indeed know who N is. After all—surely—if S knows *enough* facts about someone, then S cannot fail to know who that person is.

This datum gives rise to the suspicion that there is some privileged group of identifying facts about a person such that to cite a member of that group is to provide one's audience with conclusive grounds for saying that one knows who that person is; more strongly, it is natural to suspect that the ability to cite a member of the privileged group *constitutes* knowing who that person is. Several candidates for privileged feature suggest themselves: It might be held that one knows who N is if one can produce N's name, address, and occupation; or that one knows who N is if one knows individuating details of N's physiognomy and/or has a record of N's fingerprints; or that one knows who N is if one can physically locate or lay hands on N on demand.

Even if such hopes as these are unrealistic, can some such group of identifying facts *carry more weight* in certifying that S knows who N is? It certainly seems so. We would be much more likely to proclaim that S knows who N is if S were able to give us N's physical description, address, and occupation than if S were merely to offer the (unique) number of hairs on N's head. And N's social security number is surely a better guide to N's identity than is the (also unique) number of hackberry trees visible from his maternal grandmother's billiard-room window. This much seems beyond doubt; but are these disparities formally significant? Can we distinguish logically or syntactically between facts that count *vis-à-vis* N's identity and facts that do not? (This seems far too much to hope for.) Or are these data merely epiphenomena of the passing interests and predilections of twentieth-century middle-class American academics? Or—to aim somewhere in between—do they reflect pervasive though contingent features of normal human endeavor in general?

Let us quickly assess our strongest hypothesis: that there is a type of individuating fact such that S's citing a fact of that type about N is conclusive grounds for saying that S knows who N is. (Alternatively, it might be said

that a chain of questions of the form "Who is N?" must or can always terminate in the citing of such a fact.) It is fairly easy to see that no hypothesis of this strong sort will do. If "D" is an identifying description of the allegedly preferred type, it should not make sense (or at least it should be palpably pointless) to ask, "Who is D?" But given any candidate whatever, it makes perfect sense to ask the relevant question unless some specific feature of context makes it otiose. To see this, consider the three candidates mentioned above: name, address, and occupation; physiognomy-*cum*-fingerprints; and physical location. Anyone can imagine circumstances in which it would be not only sensible but vital to ask, "But who is this Irving Smedley, the office boy who lives at 851 Amalgam Lane?" or, "But who is the man whose photograph and fingerprints we have here?" or, "But who is this scruffy character that we have just heaved into the drunk tank?" If these questions make sense, as they surely do, then possession of identifying facts of any of the three sorts we have mentioned is insufficient to guarantee knowing who someone is. And if having received a convincing answer to these questions, one could further ask, "But who is ———?," where the blank is filled by the immediately preceding answer, then our chain of questions need not come to an end on meeting an answer of one of our three types. It seems, then, that no type of identifying description is autocratically privileged in the way we have tried to suppose and that chains of "Who is ———?" questions may terminate in different ways on different occasions—we expect or demand different sorts of answers to "Who is N?"

One might at this point suggest that "S knows who N is" is multiply ambiguous. If we are asked, "Who is that man cutting his toenails in the conservatory?," we might naturally respond, "What do you mean? Do you want to know what his name is? Or what he does? Or what he's doing at this party?" And if we are asked, "Do you know who that man is?," it is even more natural to reply, "In one sense I do, but in another sense I don't" (e.g., I can tell you his name and occupation but not what he did to earn an invitation to this exclusive affair).³ But several considerations militate against conceding that "S knows who N is" has different *senses* on different occasions. (i) It would be hard to *locate* the difference. There is neither any apparent syntactic ambiguity nor any independent evidence that "know" itself is polysemous. The most plausible option would be to say that "who" is ambiguous, but ambiguity in an interrogative pronoun would be hard to spell out. (ii) The alleged ambiguity would be not merely multiple but monstrous. We illustrate its profusion and offer a catalog of cases in chapter 2. (iii) The putative ambiguities would be widely proliferated. They would infect all the cognate locutions, as well as others that are somewhat more loosely related, such as "S remembers who N is," "S revealed who N is," "S made a guess as to who N is," and so on. (iv) Ambiguities should not be multiplied beyond necessity in any case. Our only motive to date for

positing ambiguity here is the diversity of tests, and no one but a crass verificationist would take that to signal semantic ambiguity without further examination.

There do seem to be cases in which it is simultaneously both true and false to say of a person that he knows who someone is, i.e., in which the most natural and nonmisleading response to "Does S know who N is?" is, "He does and he doesn't." Consider the plight of the copy boy at the *Daily Planet* who catches occasional glimpses of Clark Kent. It is true that the copy boy knows who Kent is, in that the boy can produce Kent's name, address, staff position, journalistic accomplishments, salient personal characteristics, preference in sandwich condiments, and the like, but in an obvious way the boy also does *not* know who Kent is; he does not know that Kent is really the Man of Steel, scion of Krypton, and savior of Truth, Justice, and the American Way. (Less imaginative examples suffice to make the point: In one way, plainly, we know who the Dalai Lama is; he is the exiled leader of Tibetan Buddhism. But we do not know where he came from and what he did before ascending to his present religious position.) Yet we can maintain the presumed univocity of "S knows who N is" by regarding it as elliptical and positing a *hidden parameter*. If we can find a plausible domain for such a parameter, we can account neatly for the "yes and no" cases without yielding to the ambiguity hypothesis and hold that "knowing who" is merely *relative*.

Fortunately there is a candidate, however vague and unilluminating, captured by a fairly colloquial idiom: S may know who N is for some purposes but not for others. (Or S may know who N is for all practical purposes.) The answer S would give to "Who is N?" (which would determine our verdict as to whether S knows who N is) depends not just on the state of S's knowledge about N but (over and above that) on his interest in N, his purposes at the time regarding N, or, we might say, his *project* vis-à-vis N. Thus our copy boy knows who Clark Kent is for the everyday purposes of getting about the newspaper business but not for the purpose of getting immediate help for Lois Lane (who is being eaten by a giant squid) or for that of reciting an honor roll of patriots, crime fighters, and heroes. And we know who the Dalai Lama is for the purpose of telling the recent history of Buddhism but not for the purpose of writing a book of success stories. These locutions seem natural enough; so let us say that our parameter is represented in each case as a *purpose* or *project*, typically (but not necessarily) that of S, the knower. Note that the purpose mentioned or presupposed in a knowing-who attribution need not be one that either S or the speaker is actually acting on, even tacitly so.

We say much more about "projects" in succeeding chapters. For now let us formulate our preliminary analysis of "S knows who N is," first having glanced at a few earlier ideas suggested by some philosophers' writings on related topics.

2 Adequacy Conditions

What should we expect from a formal analysis or regimentation of “S knows who N is”? On our view, the following.

Requirement 1 The theory should meet the demand with which this chapter began, viz., it should yield a general answer to the question of what it is to know who someone is. (We do not expect that the general answer will be terribly illuminating, since our notion of purpose or project is necessarily so inclusive, but our hypothesized regimentation should display whatever is nontrivially common to all cases, or all standard cases, of “knowing who.”)

Requirement 2 Our proposal should do this by capturing all the felt implications of ordinary knowing-who locutions (and, needless to say, by officially neglecting all those sentences intuitively not implied by the ordinary locutions). Of course, we have to start by offering rough approximations, in the sense that we may have to accept some unobvious consequences and/or neglect some apparent consequences, so long as the respective consequences or nonconsequences are close to being “don’t cares.” Any such mildly counterintuitive results are refined away later.

Requirement 3 Our theory should show exactly where our posited parameter fits into the analysis of “knowing who,” i.e., just how the parameter functions semantically in relation to the other basic ingredients of “S knows who N is.”

Requirement 4 The theory should display any actual ambiguity that *may* turn up. (Our defense of univocity in section 1 above is presumptive only. Although we reject the claim that each test for “knowing who” determines a different sense of “S knows who N is,” we later find grounds for admitting one source of syntactic and semantic ambiguity.)

Requirement 5 It is clear that, superficially, the position marked by “N” in “S knows who N is” is not purely referential. Substitution of coreferring singular terms into that position can change truth value. Consider a detective’s best friend, Boris, who has (unbeknownst to the detective) committed the nasty murder that the detective happens to be investigating. In this case the expressions “Boris” and “the murderer” corefer. But they do not substitute *salva veritate* into the context “The detective knows who ——— is,” since the detective (of course) knows who Boris is but does not know who the murderer is. This referential opacity may be purely superficial; but it is the task of a canonical regimentation to pinpoint its source and perhaps to relate it to the operation of the purpose parameter. We make an attempt at this in section 4 of this chapter.

Ideally, a philosophical proposal for the semantical analysis of a natural-language construction would have syntactic plausibility in that it could plausibly be incorporated into an independently motivated grammatical theory of that construction's syntactic formation and behavior and thereby illuminate its striking grammatical features. In practice this rarely happens, for the classical aim of regimentation is basically that of finding a philosophically useful paraphrase of the target construction, couched in a perspicuous symbolic idiom whose formulas wear their truth conditions on their sleeves. Once found, such regimentations are often hyperbolically characterized as expressing the target constructions' *logical forms*; but the association of target sentences with their alleged logical forms is justified only as a handy translation that preserves the sentences' most salient inferential properties—most are not even as impressive as Russell's theory of descriptions. What *makes* these symbolic paraphrases the logical forms "*of*" their respective target sentences, let alone what they have to do with anyone's actual production and interpretation of those target sentences in speech, is left unexplained.⁴ In this and the next chapter, we avoid speaking of logical form, for our initial concern is simply that of finding perspicuous paraphrases that satisfy the five desiderata we have mentioned. If taken as proposals about *logicogrammatical* form, our suggested regimentations would impose a heavy demand for syntactic apparatus that is (to say the least) unlikely to be met. However, it is a major task of chapters 3–5 to reduce our explanatory debt by showing how one might divide the work of our regimentations between the syntactic and the semantic components in a truth-theoretic account of propositional-attitude ascriptions so that the successor regimentations are sufficiently simple and similar to their English analysanda as to be at least *prima facie* plausible candidates for genuine logical forms. In thus diminishing the explanatory deficit and trying for enough equity to make the balloon payment acceptable, we do not, of course, eliminate the deficit entirely, nor do we wish to minimize the importance of eventually paying it off. On our view, a proposal about the logical form of an English sentence that makes even comparatively minimal demands on syntactic implementation, in the absence of an account of what the analysans actually contributes to the grammatical production of the sentence, is blind.

3 Hintikka's Proposal

To the best of our knowledge, the only theorist actually to have offered an analysis of "S knows who N is" before our 1975 article was Jaakko Hintikka in *Knowledge and Belief*.⁵ (His main objective in making his proposal was to put his analysis to technical use in providing semantical foundations for quantification into epistemic contexts, but he also purported to be giving a canonical rendition of the ordinary notion of "knowing who" (see p. 132),

and he says that he “cannot find any fault with” his analysis (p. 143). Let us test Hintikka’s proposal against the requirements we set out above and briefly examine a similar but more elaborate hypothesis suggested by some remarks of David Kaplan; then we proffer our own prototheory and discuss its merits.

Hintikka begins by offering a general explication of “S knows who is such that p,” where “p” is an open sentence expressing an attribute presumed to individuate⁶

$$(1) (\exists x)K_S p.$$

Certainly this is a natural hypothesis. It is tempting to agree that S knows who is such that p just in case there is a person whom S knows to be such that p, i.e., just in case there is some particular person known by S to have the individuating attribute in question. Now, taking our attribute to be that of *being N* (letting our open sentence be “ $x = N$ ”), we can give a more specific analysis of “S knows who N is” (= “S knows who it is that is N”)⁷:

$$(2) (\exists x)K_S(x = N).^7$$

Appealing as (2) is, however, it has a fatal defect as it stands, apparently noticed by Hintikka on pp. 144–145. The truth of (2), at least on a naive understanding of its components, is readily conceded to be necessary for that of “S knows who N is,” but it is hardly sufficient. For (2) is evidently implied by

$$(3) K_S(N = N).$$

Presuming that S is a noncretin and knows that N exists, S knows that N is self-identical. But it does not follow that S knows who N is. Our detective knows that the murderer exists and (for what it is worth) that the murderer is self-identical, but does not know who the murderer is. If (2) is to be saved, we must find a way of interpreting (2) so that it is not implied by (3), and that means sacrificing some of our preanalytical trust in existential generalization.

Hintikka plausibly suggests restricting existential generalization of a name “N” occurring within the scope of an epistemic operator to cases in which the knower knows who “N”’s referent is. This restriction, Hintikka says, is no mere ad hoc evasion of what he takes to be nasty consequences of Quine’s theory of opacity; he believes it captures a sound intuition to the effect that “a conclusion in which the identity of at least one individual is assumed to be known can be drawn only from premises at least one of which embodies the same assumption” (p. 150).⁸ If the restriction is accepted, it neatly blocks the troublesome inference from (3) to (2); such an inference would be legitimate only if (3) were to be conjoined with a canonical representation of “S knows who N is”—but of course that representation is

just (2) itself, and so the resulting argument would be (explicitly) question begging.

This response to our particular difficulty itself corresponds to an intuitive piece of ordinary reasoning, as Hintikka points out on pp. 148–149: S's ability to produce a true identity statement concerning N is plainly insufficient for saying that S knows who N is. S's knowledge that N is M (where "M" is some further singular term denoting N) does not constitute knowledge of who N is unless S (antecedently) *knows who M is*. (Our detective does not know who the murderer is, even if he knows that the murderer is the man whose footprints were found in the fishtank, unless he knows who the latter man is.)

There are, then (Hintikka urges against Quine), quantifiers that legitimately bind variables across epistemic operators such as "S knows that," but they range only over persons *known to S* (persons x such that S knows who x is) (pp. 155–156). An existential quantifier of this type is read as "Of some man known to S, S knows that" (p. 155). What, then, does Hintikka's analysis of "S knows who N is" come to? Expression (2) is evidently to be read as "Of some man known to S, S knows that that man is (identical with) N." But for a person to be "known to" S is just for S to know *who* that person is. So what (2) really says is, "Of some person such that S knows who that person is, S knows that that person is N." And since (2), thus understood, essentially contains an occurrence of a knowing-who locution, it is circular to offer (2) as a general analysis of "knowing who." This is not to deny that (2) is correctly said to be *equivalent* to our analysandum, but *qua* analysis, (2) does not meet our first requirement for a theory of knowing-who locutions—it fails to display whatever may be nontrivially common to all cases of knowing who someone is. (We do not mean to imply that Hintikka believes otherwise.)

It does seem that if we know who some person M is (for purpose P) and if we know that $N = M$, we know who N is (for purpose P); the converse of this is trivial. The circularity that vitiates (2) thus guarantees uninteresting satisfaction of requirement 2 (the demand that felt implications be captured). Requirement 3, however, is ignored. Because of the circularity of the proposed analysis, we are told nothing about our teleological parameter's role in the structure of "knowing who"—indeed, we are not even told that there is such a parameter—nor does Hintikka's discussion require (as it should) that there exist either a hidden parameter or an ambiguity.

Hintikka does remark on the multiplicity of tests for "knowing who" (p. 149n),⁹ but he dismisses this phenomenon as being irrelevant to his theoretical purposes:

No matter how the criteria for the truth of statements of the form ' $(\exists x)K_S(x = N)$ ' are chosen (within the limits of our normal logic), the truth of the two statements ' $(\exists x)K_S(x = M)$ ' and ' $K_S(N = M)$ ' accord-

ing to these criteria entails the truth of $(\exists x)K_S(x = N)$ according to the *same* criteria....

As we have conceded, this claim seems unexceptionable, but it is totally unilluminating as regards the role of the varying "criteria" in the semantics of "knowing who."

Hintikka's system as a whole does not permit the derivation of "S knows who M is" from the conjunction of "S knows who N is" and " $N = M$," so it acknowledges the referential opacity of "knowing who." Hintikka also gives a compelling general account of opacity in terms of a person's idiolectic referring expressions' having different referents in different possible worlds. We are not yet persuaded by this account,¹⁰ but we would not go so far as to insist that Hintikka's analysis cannot meet requirement 5.

4 Kaplan and Privileged Names

The failure of (2) to meet more than one of the preceding four requirements is certainly adequate grounds for looking further. Let us therefore turn to an analysis reconstructed from a hint given by Kaplan in "Quantifying in."¹¹ Kaplan writes:

One might understand the assertion, "Ralph has an opinion as to who Ortcutt is" as a claim that Ralph can place Ortcutt among the leading characters of his inner story, thus that Ralph believes some sentence of the form $\ulcorner \alpha = \text{Ortcutt} \urcorner$ with α vivid. (p. 136)

Kaplan goes on to assimilate the view alluded to in the quotation to Hintikka's, though Kaplan's and Hintikka's respective ways of explaining their basic notions differ so widely that it is difficult to assess this remark. It is easy, however, to expand Kaplan's hint into an account of "knowing who" (albeit one that Kaplan himself might disavow) by replacing "has an opinion as to" and "believes" with "knows" and then spelling out the notion of *vividness* to some extent. Minimally, then, our new analysis of "S knows who N is" is: For some referring expression α which is a vivid name, S knows-true the result of concatenating α , "=", and "N" in that order ("N" being a schematic letter here). (As we shall see, vivid names are not always public linguistic items; "=" and "N" may well be morphemes of the "language of thought" in which S's "inner story" is told.) Formally (using Kaplan's Quine quotes):

$$(4) (\exists \alpha) K_S \ulcorner \alpha = N \urcorner.$$

It is easy to see the superficial similarity of this formula to Hintikka's. And a vivid name turns out, as Kaplan seems to suggest on p. 137, to be *something like* a singular term which refers to an individual known to S, but we believe the differences here are more striking than the similarities.

The most important difference for our purposes is that (4), *qua* analysis of “S knows who N is,” *prima facie* avoids the vitiating circularity of (modified) (2). Kaplan characterizes the notion of a vivid name without reference to “knowing who” (though we later argue that his characterization does not help solve our immediate problems). Actually, we think that in giving an analysis of “*knowing who*” (rather than of “*having an opinion as to who*”) Kaplan (and we) would want to add at least two more conditions: (i) that our vivid name α actually denote N (that α have N as its semantic referent), and (ii) that S’s use of α not be merely accidental—that there be some appropriate causal connection between N’s in fact having the name α and S’s mentally using α to refer to N. We can add these conditions to (4) simply by invoking Kaplan’s defined predicate “**R**”:

$$(5) (\exists \alpha)(\mathbf{R}(\alpha, N, S) \ \& \ K_s \ulcorner \alpha = N \urcorner).$$

Kaplan explains “**R**” as follows:

R (α , N, S) is true (read “ α represents N to S”) just in case (i) α denotes N, (ii) α is a name of N for S, and (iii) α is (sufficiently) vivid. (p. 138)

Clauses (i) and (ii) merely state the two conditions we have just added. But clause (iii) is obscure, and Kaplan’s roundabout introduction of his term “vivid name” is by his own admission somewhat fanciful, despite the fact that the notion bears the whole weight of his final theory of quantifying in.¹² Let us try to get at what it comes to.

Kaplan introduces the notion of vividness by analogy with the vividness of a picture. The main ingredients of this more literal kind of vividness seem to be *clarity* and *detail* (p. 134), as opposed to blur, fuzziness, shadow, and obscurity. (The vividness of a picture may vary slightly, though not radically, according to the special interests of the person contemplating it.) Somewhat similarly, a vivid name is a “conglomeration of images, names, and partial descriptions which S employs to bring N before his mind” (p. 136). This conglomeration, we take it, plays the role of a singular term in S’s “language of thought.”¹³ Since a vivid name is “made” partly of mental images and definite and indefinite descriptions, it is, like a picture, open to varying degrees of clarity and detail. A vivid name is, then, such a mental conglomeration of “sufficient” clarity and detail. The trick is to see how far we can go in treating a vivid name as we would a purely linguistic item. What sorts of properties can a vivid name have in common with an ordinary (linguistic) singular term?

A vivid name can *denote* a particular individual (p. 137). The following definition seems to capture what Kaplan has in mind here (though he does not offer it explicitly): A vivid name α denotes N iff (i) the (linguistic) names in α actually denote N in the ordinary sense, (ii) the descriptions in α are in fact true of N, and (iii) the eidetic part of α depicts N more or less accurately.

Obviously, a name's being vivid does not guarantee its denoting, as Kaplan points out (p. 137). In our formula (5), S is intended to have a name for N whose descriptive content is cautious enough to be accurate but still detailed enough to be vivid.

A vivid name can also bear the appropriate genetic relation to its putative referent; let us say that a vivid name α is a name of N for S iff there is an appropriate sort of causal chain connecting S's (mental) tokening of α (i) to the actual states of affairs depicted by the descriptive content of α and (ii) to the events of "dubbing" in which N acquires the (linguistic) names in α . If a mental name both denotes N and is a name of N for S and is still "sufficiently" vivid, it *represents* N for S and (Kaplan says) puts S "en rapport" with N as an individual (p. 138). According to Kaplan, S must thus be en rapport with N if we are truly to say of S that he knows or believes such-and-such of N, i.e., if we are to quantify in and proclaim a *relation* between S and N.¹⁴

Let us now try to plug all this back into our most recent proposed analysis of "knowing who." According to that proposal, S knows who N is just in case S has a name that represents N to him and he knows-true some (mental) sentence consisting of that name concatenated with the mental version of " $= N$."

This analysis seems to meet requirement 1; it provides a general answer to our original question. It also satisfies requirement 5: From " $(\exists \alpha) (R(\alpha, N, S) \ \& \ K_S \ulcorner \alpha = N \urcorner)$ " and " $\ulcorner N = M \urcorner$ " we cannot derive " $(\exists \alpha) (R(\alpha, M, S) \ \& \ K_S \ulcorner \alpha = M \urcorner)$," because the occurrences of "N" and "M" in our " K_S " clauses lie in effect inside quotation. It is alleged that S knows-true some *sentence* or at least some "sentence" of his language of thought; no (pure) reference is made in those clauses to N or to M.

An analysis based on (5), however, has serious failings as well. To begin with, it does not tell us much about our parameter or even about *any* approach to the "yes and no" cases described in section 1. The vividness of a name, we are told, is to some extent relative to special interests, but, understandably, not very drastically. Some pictures just *are* clearer and more detailed than others; Kaplan's own examples of relativity (p. 134) are, though certainly sound, a little contrived. The proponent of (5) fails to exhibit the *full-scale* relativity of "knowing who" to purposes, since the relativity evident in (5) is minimal.¹⁵

More important, the entailment relations of "S knows who N is" are not well represented by (5). Here are some of the disparities:

(i) Expression (5) asserts the existence of vivid names and entails the existence of *mental sentences*. It is not clear that these consequences are shared by "S knows who N is"; could S not know who N is even if there were no language, mental or otherwise? This objection is not serious, however. First, we should not easily grant that S *could* know who anyone was in such impoverished circumstances, lacking an internal system of representation.

Second, any faintly plausible analysis that we can think of makes reference to some such items, including that which we shall devise. Whether or not it is metaphysically wrong to saddle “knowing who” with such an ontology, it is indispensable for the present, at least as a heuristic device.

(ii) Suppose that S has a name α for N that both denotes N and is a name of N for S; suppose further that S knows N personally. Given what Kaplan has said on p. 136, it follows that α represents N for S and presumably that S knows-true “ $\alpha = N$ ”. In short, it follows (on the analysis under discussion) that S knows who N is. But that should not follow, for we sometimes have a representing name for someone whom we know personally, the name being as vivid as one might like, and still utterly fail to know who that person is for some key purposes (cf. the case of the copy boy and Clark Kent). It might be replied that, from the initial conditions hypothesized here, it does follow that S knows who N is *for some purpose or other*. This much seems to be true, but trivially so—for, given that we know that a particular person exists at all, we know who that person is for some boring purpose or other; stated categorically as it is, without any reference to the teleological parameter, the Kaplan-style analysis is at best misleading on this point.

(iii) It is not clear how we are to take the requirement that S’s representing name for N “robustly and clearly delineate” N’s nature or character. We know who Thomas Edison was, for standard purposes: He was the man who invented incandescent light bulbs. Does it follow that we have a representing name for Edison that robustly and clearly delineates him? In whatever intuitive sense we attach to this phrase (having to do, again, with detail and completeness in particular), it seems not—for all we know about him is that he invented incandescent light bulbs.

(iv) A related matter: Kaplan requires (see again p. 136) that N play a “major role” in S’s inner story. We assume that this means S must have fairly intimate knowledge of N (cf. Kaplan’s remark about Julius Caesar; he speaks of being “well acquainted” with Caesar) or at least that N has some prominence in S’s thoughts from time to time. But neither of these conditions is really necessary for S’s knowing who N is.

Unless we have misconstrued Kaplan, it looks as if vivid names are not exactly what we are after. His emphasis on detail and completeness of information seems misplaced for our purposes. Also, the notion of vividness itself is too vague. However, we provisionally exploit Hintikka’s and Kaplan’s practice of invoking privileged or preferred singular terms, as this still seems to be the most promising line to take in catching the intuitive appeal of (2) while blocking the trivial inference from (3). It might be objected at this point (e.g., by Quine) that such an inegalitarian attitude toward referring expressions drags in a metaphysically repugnant form of essentialism; but this is not so—we confess to an inegalitarian weighting of

individuals' properties but only relative to someone's purpose or project, and not even Quine could object to so interest-bound a notion of "essence."

One further objectionable feature of (5) as an analysis of "S knows who N is" is that the privileged status of S's representation is made relative to the subject in particular. We want "knowing who" to be relative to projects, but in our Kaplan-style analysis it is also relative to the idiosyncrasies of S's personal inner story. On our view, these features of S should be irrelevant. Relative to a particular project, it is easy to frame a question-and-answer test for whether an *arbitrarily* chosen subject knows who some N is. For example, relative to the desirable task of throwing the murderer in jail, we can say in advance what test *anyone* would have to pass in order to count as knowing who the murderer is: A subject would have to be able to tell whether he had the right human body in handcuffs (name, address, and occupation have only contributory relevance in this case). It would be strange to say that, once we are holding our parameter fixed, different tests might still be required for different subjects—we can think of no reason why this could ever happen. We conclude that a name is privileged (in a way useful to us) relative to the teleological parameter but irrespective of the identity of S, the knower.

We now propose our own analysis of "S knows who N is," proceeding more or less along Kaplan's lines. We invent our own category of privileged names, making use of Kripke's notion of a rigid designator. In chapter 2 we introduce a key modification, one that we take to codify our most important insight about "knowing who"; then in chapter 4 we formulate the final version of our theory.

5 "Knowing Who" and Important Names

Let us bluntly begin by stating our preliminary analysis, explaining it subsequently:

$$(6) (\exists \alpha) (\text{Impname}(\alpha, N, P) \ \& \ (\exists \beta) (\Delta(\beta, N) \ \& \ K_s \ulcorner \beta \text{ is } \alpha \urcorner)).$$

The second clause of this formulation is intended to say roughly that S knows of N that he is identical with the referent of α ; the first clause specifies α 's status as a privileged name.

The clause " $\Delta(\beta, N)$ " is to be read as " β referentially designates N." As we use that term, a referential designator is first a *rigid* designator. (According to Kripke, this is for it to "designate . . . its referent wherever [at every possible world in which] the object exists."¹⁶) Even though the same singular term may be used rigidly on one occasion but nonrigidly or flaccidly on another, we need not explicitly relativize rigidity to a speaker at a time; because " β " in (6) ranges over expressions of S's language of thought, the requirement that β rigidly designate N *in that language* is already sufficient to ensure that β rigidly designate N "for S." We now add, in keeping with Kripke's views

and with what we take to be the requirements of “knowing who,” that a rigid term β designates N in S’s language of thought only if there is an appropriate genetic connection of Kripke’s and Kaplan’s type between S’s (implicit) use of β and the dubbing of N with β . The idea of the second clause of (6), then, is that S knows-true a superficial identity sentence, one term of which is used (in S’s language of thought) to pick out N, regardless of N’s contingent properties, whatever they may be.^{17,18}

“Impname (α , N, P)” is read, “ α is an important name of N for purpose P.” As we have insisted, a name is privileged in our sense only relative to certain purposes; but when those purposes have been fixed, the name is privileged irrespective of the identity of S, the person using the name. Superficially, the intuitive content of “Impname” is something like this: “Impname (α , N, P)” is true just in case α denotes N and also is the sort of item that would be accepted as a reply by a person asking “Who is N?” with purpose P explicitly in mind. Saying this, of course, does not help explain how “Impname” is ultimately to be spelled out or how the use of an important name is related to the purpose or project. We address these matters in chapter 2.

Let us now check our preliminary analysis against the requirements we appealed to in section 2. Plainly, it satisfies requirement 1, as did the Kaplan-style proposal. It also satisfies requirement 3 by placing the teleological parameter explicitly. According to our analysis, “knowing who” is relative to project precisely because a singular term α , as a name of an individual N, is *important* to people variously, depending on their interests and aims. (That explanatory fact should not surprise anyone.)

What about felt implications (requirement 2)? Expression (6) is certainly sufficient for “knowing who.” If S knows that N satisfies ‘is α ’ (and appropriately answers ‘N is α ’ to the question of who N is), when α is an important name in the sense we have less than adequately described, then surely S knows who N is. And S knows that N satisfies ‘is α ’ if S has a *referential* designator β of N such that S knows-true ‘ β is α ’. So our only remaining question is that of whether satisfying (6) is *necessary* for S’s knowing who N is.

Like the Kaplan-style analysis, our own is committed to “names in one’s language of thought”; we continue to live with that for now for the reasons originally mentioned. Happily, our analysis avoids saddling “knowing who” with some of the counterintuitive consequences to which the Kaplan-style theory succumbed. Expression (6) entails none of the following: (i) that S’s important name for N provides any particular degree of detail or completeness considered in the abstract, (ii) that S has any particularly intimate knowledge of N and N’s doings, (iii) that S is well acquainted with N (even in the sense in which a historian can be “well acquainted” with Caesar), (iv) that N has any special prominence in S’s thoughts. As we shall see, having an important name of N may amount to knowing just one key fact about N—

possibly an obscure fact at that. We have also avoided another pitfall of the Kaplan-style analysis: Knowing someone personally does not guarantee knowing who that person is (for all practical purposes) in the sense of (6), although, depending on the actual workings of an important name, it may guarantee knowing who the person is for some easily specifiable purposes.

Finally, we can list a few intuitive theses about “knowing who” that have considerable preanalytical plausibility and show that (6) squares with them. (i) We have said that S’s knowing or being acquainted with N should not in every situation suffice for S’s knowing who N is, nor should it be required, and nothing in (6) does, in general, require it. Of course, some particular purposes P are such that one does need to know N personally in order to know who N is for P; this fact is accommodated by the concept of an important name. (ii) It is hard to maintain that elaborate factual knowledge is always required for “knowing who.” In some cases, naturally, certain sorts of facts are required, such as those concerning N’s occupation or N’s social status or whatever. But in other cases, it is not clear that any fact other than the superficial $\ulcorner \beta \text{ is } \alpha \urcorner$ itself is needed (we have more to say about this in chapter 2). Our analysis saves all these intuitions too. (iii) If S is asked, “Who is N?,” S’s answer “N is M” gives us reason to say that S knows who N is if and only if it is clear that S knows who M is.¹⁹ Our analysis preserves this fact. From “ $K_S \ulcorner N = M \urcorner$ ” we cannot derive (6), since “ $K_S \ulcorner N = M \urcorner$ ” does not guarantee that S has an important name of N for the relevant purpose(s). Expression (6) does follow, however, from the conjunction of “ $K_S \ulcorner N = M \urcorner$ ” and

$$(7) (\exists \alpha)(\text{Impname}(\alpha, M, P) \ \& \ (\exists \beta)(\Delta(\beta, M) \ \& \ K_S \ulcorner \beta \text{ is } \alpha \urcorner))$$

(“S knows who M is”)—the proof is trivial. (iv) “If you know *who* does something, you ipso facto know *that* someone did it,” Hintikka observes (p. 160). In our case, if you know who has the property of *being N*, you accordingly know that someone has that property, i.e., that N exists. This fact falls easily out of our analysis; if, for the relevant α and β , S knows-true $\ulcorner \alpha \text{ is } \beta \urcorner$, then presumably S also knows-true $\ulcorner (\exists x)(x = \beta) \urcorner$. (v) Finally, notice that (6) does not require S to have purpose P explicitly in mind, to act on or with P, or even to have thought about P, however vaguely. Purposes are mentioned independently of *whose* purposes they might be at any time.

6 Opacity and Referential/Attributive Ambiguity

All this is very encouraging. However, we have not measured our analysis against requirements 4 and 5, concerning possible ambiguities and opacity, respectively. And here a difficulty emerges. From (6) and “ $N = M$ ” we can easily derive (7), since all positions containing “N” in (6) are plainly trans-

parent. If (6) is an accurate analysis of "S knows who N is," then how can we save our intuition that "knowing who" is referentially opaque?

The answer to this is roundabout and requires some antecedent discussion of possible ambiguities. We believe, in fact, that "S knows who N is" is subject to one ambiguity that no one else has pointed out. It is most clearly perceptible when "N" is a definite description, say, "the murderer." Temporarily adopting Donnellan's referential/attributive distinction (see note 1) and assuming its syntactic/semantic reality, let us examine the consequences of reading this description first referentially and then attributively²⁰ in "S knows who the murderer is." On the referential reading, "the murderer" serves merely to pick out the individual we are talking about and not essentially to describe that individual in any way (thus, on such a reading, we can *nontrivially* say, "The murderer murdered someone," or wish that the murderer had not murdered anyone without desiring that the predicate calculus contain falsehoods). "The murderer" would be canonically represented either by a Russellized description taking wide scope or (on a more drastic understanding of Donnellan's distinction) by an unstructured referring expression, an individual constant. On the attributive reading, however, "the murderer" takes narrow scope; a sentence containing it thus construed is about *whoever* did the murder, whoever that might turn out to be. "The murderer" would be Russellized narrowly in the canonical representation, leaving behind a *general* statement with no variables bound from outside the relevant sentence operator. In an important way, therefore, attributive definite descriptions are not really singular terms (referring expressions) at all. We assume for purposes of this chapter and the next that the referential/attributive contrast is sufficiently familiar, hoping that our heuristic use of it will not be troubled by the various difficulties that have attended various attempts at refining it.²¹

On the referential reading, then, "S knows who the murderer is" is to be understood as saying that S knows who *that person*, the one whom we are picking out, is, for whatever purposes are in question, regardless of that person's having committed any murders. On the attributive reading, by contrast, "S knows who the murderer is" is equivalent to "S knows who did the murder," which says that there is some person whom S knows to have committed the murder. This latter (general) statement may well be false even when it is true that S knows who the murderer (referential use) is. It entails, for example, that S knows that someone did the murder; our referential reading does not.

In order to codify such facts, it appears that we must posit a second, alternative formal representation of "S knows who N is," where "N" is a definite description "the F." "S knows who N is" is then said to be syntactically and semantically ambiguous. We propose to Russellize the attributive description in the way suggested above:

$$(8) (\exists!x)(Fx \ \& \ (\exists\alpha)(\Delta(\alpha, x) \ \& \ \text{Impname}(\alpha, x, P) \ \& \ K_s \ulcorner F\alpha \urcorner)).$$

(For a particular replacement of the schematic letter "F" such as "invented bifocals," $\ulcorner \alpha$ invented bifocals \urcorner would be the result of concatenating α with whatever predicate translates "invented bifocals" in S's language of thought—i.e., a Mentalese ·invented bifocals· in Wilfrid Sellars's sense.)

Expression (8) seems to us to capture the way in which the attributive sense of "S knows who the F is" differs from the referential sense explicated by (6). The added implications we have mentioned are present in (8); and the analysis captures the fact that knowing who the F is in the attributive sense is not, as it is in the referential sense, knowing an apparent *identity*—it is, rather, knowing a (unique) predication. To see this, imagine our familiar question-and-answer situation. To ask "Who is the murderer?" in the attributive sense is precisely to ask, "Who did the murder?" An identity statement would be an inappropriate response to such a question; what is called for is some claim of the form, "M did the murder." Accordingly, (8) ultimately ascribes to S knowledge of a predication rather than knowledge of an identity.

Notice that the teleological parameter functions even in the attributive case, which is, in a way, preanalytically much less mysterious than the referential sense. (On the face of things, the only mysterious feature of "knowing who" in the attributive sense is the quantification into an epistemic context—"S knows of *some particular person* that that person is (uniquely) F"). But S may know, and at the same time not know, who did the murder, depending on purposes. For the purpose of writing history books, S may know, in that he knows that the murderer is the man named Boris Flammenwerfer, a chicken-sexer who hails from Berlin; yet S may still not know who the latter individual is for the purpose of laying hands on his person, throwing him into a cell, bringing him to trial, and executing sentence. (Flammenwerfer may long since have decamped incognito.) Thus, even in the attributive sense, S knows who the F is (S knows which individual is F) only relative to a project.

One further detail should be mentioned:²² Knowing-who ascriptions in which "N" is replaced by definite descriptions are typically ambivalent as to whether it is the speaker or the subject who is to take responsibility for the *uniqueness* of an F, implied by the descriptor. It is perfectly acceptable to say that S knows who the F is when S satisfies (8) and also knows (as does the speaker) that there is at most one F. But it is also usually acceptable to say that S knows who the F is even when S may not know that there is only one F. (For example, if S, a Nicaraguan communist, has discovered that our agent in Nicaragua is a CIA spy, we may well (and properly) say, "S knows who our agent in Nicaragua is," even though S does not know that we have only one agent in Nicaragua.) To allow for the latter possibility, we have stated our

analysis (8) in a particularly weak form, by assigning the shriek wide scope. In some cases, however, a purist might insist that S does *not* know who *the* F is unless S actually knows that there is but one F; in such cases we might move (8)'s shriek inward, penetrating the knowledge clause, as in

$$(8^*) (\exists x)(Fx \ \& \ (\exists \alpha)(\Delta(\alpha, x) \ \& \ \text{Impname}(\alpha, x, P) \ \& \ K_S \ulcorner F! \alpha \urcorner)).$$

We shall not commit ourselves on the question of whether this further distinction betokens a real semantic ambiguity. The only evidence for saying that (8*) captures a genuinely distinct sense of (attributive) "knowing who" is that some speakers of English sometimes feel uncomfortable in ascribing such knowledge to a subject when the subject does not know that there is but one F. We believe that English is not determinate on this point. A possible explanation for this is that (i) uniqueness is usually *presupposed* all around (as a matter of fact, the vast majority of murders are the work of individuals rather than of committees), and (ii) in any case, uniqueness does not usually *matter* to speakers nearly so much as does the existence claim. So we do not choose between saying that (8) is correct to the exclusion of (8*), that (8*) is correct to the exclusion of (8), or that (8) and (8*) are both correct readings of an ambiguous attributive "knowing who" construction.

Expression (8), and a fortiori (8*), seem much stronger than (6), in the sense that they appear to require more epistemic activity on S's part in order for S to know who N is. But this is what we would have expected. "Knowing who" in the attributive sense does require more of the subject. In order to know who the murderer is, referentially speaking, S need only have an important name of him (of that person, regardless of his murdering or any of his other contingent attributes), vis-a-vis whatever project is in question. But attributively, S must know in addition some general facts about the world—e.g., that a murder took place and that whoever committed it fulfills such-and-such an identifying description—and only then go on to worry about knowing who the murderer is for practical purposes, such as locating and jailing him.

In expounding our referential/attribution ambiguity, we have thus far used only definite descriptions as examples. Demonstratives are by nature referential, or so it seems. Proper names are generally held to fare similarly, since the semantical function of a proper name is solely to pick out the appropriate referent (cf. note 16); but there are attributive uses of proper names,²³ and we reveal and examine further complexities in chapter 2.

It is time to return to requirement 5 and to explain the apparent opacity of "knowing who." In light of the ambiguity we brought out, we can do this fairly easily, for the most plausible cases in which truth value is not preserved through substitution of "coreferring singular terms" are those in which one of the alleged "singular terms" is actually a definite description (or possibly even an apparent proper name) *used attributively*. Since such expressions

count superficially as singular terms, they give rise to opacity in the accordingly superficial way. But opacity of *this* sort is simply a case of amphiboly; the description is not really (semantically) a singular term at all; thus we would not expect opacity of this relatively trivial sort to be reflected in logical form.

Does this treatment succeed in handling all cases of apparent opacity? It seems not. Even stipulating that all replacements of "N" are referential singular terms, we can doubt the truth of "S knows who M is" even when that sentence is, in fact, the product of substituting "M" for a coreferring term "N" in a true instance of "S knows who N is." For example, one might contend that our detective would say (and an impartial observer might agree) that obviously he knows who Boris is but does not know who the murderer is, even though "the murderer" is being used referentially here. Why is this?

If there truly are such cases, we believe that they are products of quite natural *parameter shift*. Certain singular terms carry with them suggestions of particular projects. The name "Boris," as the detective uses it, brings to mind Boris in his capacity as neighbor, best friend, confidant, and fly-tying companion to the detective, whereas the term "the murderer," even when used referentially to pick out Boris himself, plainly connotes someone whose smudged fingerprints we have, who was seen by a near-eyewitness fleeing the fatal scene in a bloodstained anorak, etc. We submit (for now) that in any such cases of apparent opacity over and above the referential/attributive ambiguity, what is going on is simply that two different values of the teleological parameter are being rung in. Nothing more recondite than this parameter shift is required to account for a change of truth value on substitution. If S knows who Cicero is for purpose P, then S knows who Tully is for P, even if S himself has never heard the name "Tully." (To see this, notice that *we* are the utterers of "S knows who Tully is." "Tully" is a name in *our* language, serving solely to indicate that person, Tully or Cicero, whatever he is called; it is not required or even expected that S would express himself in that way.) Our position, then, is that the apparent opacity of "knowing who" is real enough superficially (it occurs in either of the two ways we have mentioned), but in neither case does it penetrate to the level of logical structure—so we should not be surprised if "N" occurs everywhere transparently in (6).²⁴

This concludes the case in favor of our preliminary analysis. But the analysis still stands in need of modification, in light of a difficulty that first shows itself as merely technical but turns out to prompt a significant conceptual revision. We cannot discuss it until we say a bit more about important names in chapter 2, so we pass over it for now. Instead, let us note a few other *prima facie* drawbacks that we try to address later on.

First, we must provide our allusions to quasilinguistic items and "lan-

guages of thought" with some real substance; until we do so, these allusions are no more than a subterfuge designed to help us avoid commitment to the thesis that a knower must have and express himself in a natural language, such as English. Second, we have played fast and loose with the notion of a referential designator, which despite its currency and despite our quick attempt at definition, remains none too clear.

Third, our account so far is minimal, collapsing all the really difficult conceptual analysis into the allegedly primitive predicate "Impname," which is made to carry all the intuitive content of "knowing who." Toward giving the *meaning* of "S knows who N is," we have done relatively little. On the other hand, that is what one expects from a logical regimentation, and rightly so: It is better to have to deal with the meaning of an untidy primitive than with unknown logical grammar.²⁵