

Understanding the Intentions Behind the Referential/Attributive Distinction

Megan Henricks Stotts

This is a post-peer-review, pre-copyedit version of an article published in *Erkenntnis*. The final authenticated version is available online at: <https://doi.org/10.1007/s10670-016-9821-y>.

Abstract:

In his recently published John Locke Lectures, Saul Kripke attempts to capture Keith Donnellan's referential/attribution distinction for definite descriptions using a distinction between general and specific intentions. I argue that although Kripke's own way of capturing the referential/attribution distinction is inadequate, we can use general and specific intentions to successfully capture the distinction if we also distinguish between primary and secondary intentions. An attributive use is characterized by the fact that the general intention is either the primary or only designative intention, whereas a referential use occurs when a specific intention is either the primary or only designative intention. Along the way, accounts of the referential/attribution distinction offered by John Searle and by Kepa Korta and John Perry come in for criticism as well, and we'll also discuss Michael O'Rourke's dual-aspect uses of definite descriptions.

1. Introduction

Keith Donnellan's (1966) distinction between referential and attributive uses of definite descriptions (hereafter, the R/A distinction) has attracted extensive philosophical attention for fifty years. Donnellan began with a seemingly simple observation: a speaker can use the description 'Smith's murderer' in two different ways when saying, "Smith's murderer is insane." A speaker who has just discovered Smith's heinously murdered body and does not know who did the deed would be likely to use the description 'Smith's murderer' with the intention to designate whoever happens to fit that description. Such a use is *attributive*. On the other hand, an attendee of Jones's trial for the murder of Smith might observe Jones's erratic behavior and utter the same sentence while intending to use 'Smith's murderer' to designate *Jones*, whom she sees in front of her and has in mind. A use of that sort is *referential* (Donnellan, 1966, pp. 285–6).

There is no consensus about how to understand the R/A distinction. Saul Kripke (1977, 2013), who extends the distinction to include analogous cases involving proper names, offers an account of it in terms of general intentions (*i.e.*, intentions to refer to the conventional, semantic

referent of a designator) and specific intentions (*i.e.*, intentions to refer to some particular individual). I will argue that Kripke's two kinds of intentions are crucial for capturing the R/A distinction, but we also need to distinguish between primary and secondary intentions. Roughly, one intention is *primary* with respect to another (which is *secondary*) if the first intention partially caused the second to exist. Using this distinction, I will argue for two main claims:

- In an attributive use of a designator, the speaker's intention to designate whichever individual is the designator's conventional, semantic referent is either her only designative intention, or it is primary with respect to her other designative intentions.
- In a referential use of a designator, the speaker's intention to designate a particular individual she has in mind is either her only designative intention, or it is primary with respect to her other designative intentions.¹

I will show that the key cases that support these claims are also the cases on which Kripke's account of the R/A distinction falters. I will also argue that two other accounts of the R/A distinction are unsuccessful: John Searle's (1979) account, which bears similarities to the account I'm proposing, and Kripke and John Perry's (2011) account, which is able to deal with some but not all of the cases that cause trouble for Kripke. Along the way, we will also discuss Michael O'Rourke's (1998) notion of a dual-aspect use of a definite description.

2. *The Teacher Case*

With the R/A distinction in mind, consider a new case. A grade school teacher needs to arrange her students for a class picture. She begins by asking them to line up from tallest to shortest, saying, "The tallest student in our class should line up first." The teacher intends to designate whichever student is denoted by 'the tallest student in our class.' But suppose she also believes

that a particular student, Henry, is the tallest student in the class. This belief, along with her intention to designate whoever is the tallest, leads her to also intend to designate Henry.

Designating Henry is a foreseen and *intended* consequence of designating the tallest student. We might reconstruct the teacher's thought process in the following way: "I will use this description to designate whoever is the tallest. But that's Henry, of course, so I will use the description to designate Henry [by using it to designate whoever is the tallest]." This is the *Teacher Case*.

Similar cases can occur when someone wants to designate whatever satisfies a certain description but also has a firm belief about which individual does satisfy it. Imagine a presenter of an award, who peeked inside the envelope and already knows who won, describing the prize associated with the award to the audience. He says, "The winner will receive \$5000." He intends to designate whoever will receive the award, but because of what he saw inside the envelope, he also intends to designate the specific person whose name is inside. Similarly, think of statements about the future president of the United States before the election. In a discussion about the U.S. procedure for changing presidents, I might say, "The new U.S. president will take the oath of office in January," intending to talk about whoever wins the 2016 election. But imagine that I believe that Hillary Clinton will win, and therefore form an intention to make a claim about her in particular. Importantly, I am not suggesting that people in such situations *always* or *must* form the second intention. In my examples, I am stipulating that the speakers do form the second intention, and all I need is the claim that we *sometimes* use descriptions with these sorts of intentions.

How should we categorize the Teacher Case (and others like it)? Is it attributive, referential, or something else? The teacher intends to pick out whoever fits the description she used. Yet, she also intends to pick out someone in particular whom she has in mind: Henry. We

can make sense of the case by thinking about what sort of outcome would fulfill the teacher's overall plan in uttering the description. Obviously, if Henry lines up first and he really is the tallest, both of the teacher's intentions will be satisfied. But now imagine that Henry is *not* actually the tallest student—Rosa has recently surpassed him. Because the teacher's larger goal is to arrange the students for a class picture, she needs them to be in the actual order of their heights. In that case, if Henry ended up first in line, the teacher would be disappointed when she noticed that Rosa was actually taller—her overall plan would then be best served by *Rosa* lining up first, even though she had also intended to designate Henry. Like the speaker who says, "Smith's murderer is insane" after discovering Smith's body, the teacher's main goal is to designate whomever the description denotes. For this reason, the Teacher Case is attributive.

But because the Teacher Case is quite different from paradigmatic attributive uses, one might wonder whether it is better thought of as in some sense both attributive and referential. Here it will be helpful to consider O'Rourke's notion of a dual-aspect use. In a dual-aspect use, "both a particular item and a descriptive condition are essential to the utterance" (O'Rourke, 1998, p. 266). He describes a case in which the detective investigating Jones's involvement in Smith's death has concluded that a bust of Heidegger in Jones's study is the murder weapon. She picks it up and says to Jones, "The murder weapon is heavy." In this case, the detective intends to designate that particular bust of Heidegger, and she intends to designate whatever item was the murder weapon. Both of these are essential to the detective's plan: she only wants to designate the bust because she thinks it is in fact the murder weapon, and she is only bothering to designate whatever is the murder weapon because she thinks that description denotes the particular bust in Jones's study (O'Rourke, 1998, p. 267).

However, the Teacher Case is not dual-aspect. The descriptive condition is essential to what the teacher is trying to achieve, but the particular individual (Henry) is not. If anyone other than the tallest person in the class ends up at the front of the line, the teacher will not have what she needs: she will not be in a position to easily arrange the students for the class picture. On the other hand, if someone other than Henry ends up first in line, that will be just fine—as long as that person is the tallest in the class. The Teacher Case does not have the same structure as O'Rourke's example. So, we can retain our conclusion that the Teacher Case is attributive.

3. Capturing the R/A Distinction

To capture what the Teacher Case reveals about the R/A distinction, we need Kripke's distinction between general and specific intentions. According to Kripke, a general intention is a speaker's standing intention to designate the conventional, semantic referent of a designator whenever she uses it (so, for a definite description, it will be an intention to designate whichever object fits the description). The teacher has a standing intention to designate whoever is the tallest student in the class when she uses 'the tallest student in our class.' A specific intention is a speaker's intention on a particular occasion to designate an individual that she has in mind. The teacher has a specific intention to designate Henry (Kripke, 1977, p. 264; 2013, p. 119).²

We need to modify Kripke's notion of a general intention because it runs into problems with ambiguous designators. Kripke (2013, p. 119) would say that my general intention with the name 'Michael Jordan' is the following: when I utter 'Michael Jordan,' I will designate the man Michael Jordan who was an incredible basketball player. However, imagine that I know another person named 'Michael Jordan.' This would seem to mean that I have the general intention to

designate *that* man whenever I utter ‘Michael Jordan’ as well. As a result, I would have two inconsistent intentions. Kripke’s way of understanding general intentions, then, isn’t quite right.

To accommodate ambiguity, we can think of general intentions as *procedures* that generate intentions on particular occasions, rather than as ordinary intentions themselves. For example, I have a procedure of using the name ‘Michael Jordan’ to designate the man who was an incredible basketball player. Each time I use that procedure, I generate an intention to designate that man with that particular utterance of the name. However, if I knew someone else named ‘Michael Jordan,’ I would have a second procedure with that name. When I followed that second procedure, I would generate an intention to designate that other individual.

Unlike with intentions, there is no problematic inconsistency in having two incompatible procedures, so long as one never attempts to follow both of them simultaneously. For instance, I have multiple procedures for cooking eggs: one procedure is to fry them, and another is to poach them. Clearly there is no problematic inconsistency involved in having both procedures. So, moving forward, we will use Kripke’s notion of a *specific intention*, but we will speak of a person’s *general procedure* with some term, and of intentions that result from an application of that general procedure on a particular occasion, rather than in terms of general intentions.³

Using this terminology about intentions, what can we say about the R/A distinction? Let’s tackle the attributive use first, by returning to the attributive version of the ‘Smith’s murderer’ case. When the speaker does not know who killed Smith but makes an inference from the state of Smith’s body before saying, “Smith’s murderer is insane,” she intends to designate whoever murdered Smith. In other words, she has formed an intention as specified by her general procedure with the description ‘Smith’s murderer.’⁴ And because she does not know who murdered Smith, she does not have anyone in mind, so the intention formed in accordance

with her general procedure with ‘Smith’s murderer’ is her only designative intention.⁵ Intentions resulting from a speaker’s general procedure clearly play an important role in attributive uses.

Now, consider the Teacher Case again. The teacher intends to designate whoever is the tallest in the class when she utters ‘the tallest student in our class.’ But interestingly, the teacher also has a specific intention: she intends to designate Henry. Importantly, the teacher’s specific intention is not on the same footing as the intention resulting from her general procedure with ‘the tallest student in our class.’ The teacher first forms an intention to designate whoever satisfies ‘the tallest student in our class.’ Then that intention, along with her belief that Henry is the tallest, leads her to form a second intention. She believes that by designating the tallest student she will designate Henry, and this leads her to intend to designate Henry. We might say that the intention she formed in accordance with her general procedure with ‘the tallest student in our class’ is *primary* with respect to her specific intention to refer to Henry.⁶

This notion of one intention being primary with respect to another (which is accordingly a *secondary* intention) is not limited to communicative contexts. Think of a person who intends to cheer up her friend and believes that she could do this by sending flowers or by cooking her friend dinner. She also believes that sending flowers would be quicker than cooking dinner. So, she forms an intention to send flowers. Her intention to cheer up her friend is primary, and her intention to send flowers is secondary.

For another example, think of a teenager who wants to clean his room because the mess bothers him. Then, he realizes the clean room will please his parents and develops an intention to please them. He now intends to please his parents, but he wouldn’t have formed the intention to please them on this occasion if he hadn’t already had the intention to clean his room. The

teenager's intention to clean his room is primary with respect to his intention to please his parents.

Notice that in the flowers example, the satisfaction of the secondary intention is a means to satisfying the primary intention. But in the tidy teenager example, the satisfaction of the secondary intention is instead a likely result or consequence of the satisfaction of the primary intention. This shows that the relationship between the intentions themselves is independent of the relationship between the *events that would satisfy* those intentions.

This notion of one intention being primary relative to another has some intuitive pull, but we can say more. One intention is primary with respect to another intention at a particular time if the following criteria are satisfied:

- (1) The first intention at least partially caused the second intention to exist at that time.
- (2) The second intention did not partially cause the first intention to exist at that time.
- (3) The two intentions belong to the same person.⁷

The person's intention to cheer up her friend is primary because it partially *caused* her to form an intention to buy flowers, and not the other way around. Her intention to cheer up her friend, along with some beliefs, is what led her to intend to buy flowers. Similarly, the intention that the teacher formed in accordance with her general procedure with 'the tallest student in our class' is primary because that intention was part of what *caused* her to form her specific intention to designate Henry, and not the other way around. We cannot tell the story of why she intends to designate Henry without mentioning the fact that she intends to designate the tallest student in the class, but explaining why she intends to designate the tallest student does not require any mention of Henry—she wants to designate the tallest student because she needs to arrange the class for a picture.

It is worthwhile to note that in this case, in contrast to the non-linguistic cases, the satisfaction of the secondary intention is neither a means to nor a consequence of the satisfaction of the primary intention. Rather, the satisfaction of the primary intention *constitutes* the satisfaction of the secondary intention, assuming that Henry really is the tallest. To designate the tallest student in the class just *is* to designate Henry. Again, the relationship between the intentions is independent of the relationship between the states of affairs that would satisfy them.

So, to account for both the more typical attributive uses and the Teacher Case, I propose the following: in an attributive use of a designator, the speaker's intention formed in accordance with her general procedure is either her only designative intention (as in the attributive use of 'Smith's murderer'), or it is primary with respect to her other designative intentions (as in the Teacher Case).

Now for the referential use. The speaker who observes Jones's behavior during his trial has a specific intention to designate Jones, whom she has in mind. But she believes that 'Smith's murderer' denotes Jones, so she also forms an intention in accordance with her general procedure with 'Smith's murderer.' She has both a specific intention to designate Jones and a procedure-generated intention to designate whoever murdered Smith. Her specific intention is primary with respect to the intention generated by her general procedure because the only reason she followed her general procedure with 'Smith's murderer' on this occasion is that she wants to talk about Jones, whom she has in mind, and she believes that Jones fits the description 'Smith's murderer.'

Now consider one of Donnellan's own cases, which we'll call the *King Case*:

Suppose the throne is occupied by a man I firmly believe to be not the king, but a usurper. Imagine also that his followers as firmly believe that he is the king. Suppose I wish to see this man. I might say to his minions, "Is the king in his countinghouse?" I succeed in referring to the man I wish to refer to without myself believing that he fits the

description. It is not even necessary, moreover, to suppose that his followers believe him to be the king. If they are cynical about the whole thing, know he is not the king, I may still succeed in referring to the man I wish to refer to (Donnellan, 1966, pp. 290–1).

This use of ‘the king’ is clearly referential—the speaker has a particular person in mind (the usurper) and intends to designate him. But there is something odd going on. The speaker does *not* intend to use ‘the king’ to designate the person who satisfies the description. If the minions replied with, “No, the usurper has had him locked in the dungeon for the past five years, remember?”, the speaker would not have achieved anything he intended to do with his utterance. It seems that the speaker is not following his general procedure with ‘the king’ on this occasion—that is, he has not formed an intention to designate whoever is the king (*cf.* Bratman, 1987, pp. 88–9). The speaker refrains from following his general procedure with ‘the king’ because he can get something he wants more by doing so: information about the location of the usurper without raising suspicion.

We can account for both the King Case and the more typical referential cases by saying that in a referential use of a designator, the speaker’s specific intention to designate an individual she has in mind is either primary with respect to her other designative intentions, or it is her only designative intention on that occasion because she has chosen not to follow her general procedure.

Let’s also consider the dual-aspect use. In O’Rourke’s case, the detective has a specific intention to designate the bust of Heidegger that she is holding, in order to talk about *its* weight. That intention, along with her belief that the bust is the murder weapon, partially causes her to follow her general procedure with ‘the murder weapon’ and intend to designate its conventional denotation. But from another angle, we can say that the detective wants to designate whichever item actually is the murder weapon in front of Jones, the suspect, in order to observe his reaction.

That intention, along with her belief that the bust is the murder weapon, partially causes her to intend to designate the bust in particular. The two intentions are causally interdependent.

According to our criteria for primary intentions, a primary intention must partially cause its secondary intention to exist, but the secondary intention must *not* partially cause the primary intention to exist. Because both of the detective's intentions fail to meet these criteria, neither of her intentions is primary.

So, in dual-aspect uses, the speaker has both a specific intention and an intention formed in accordance with her general procedure with the relevant designator, but neither intention is primary. Although (or rather, because) both intentions meet the first criterion for being primary, neither of them meets the second criterion: each partially causes the other to exist, so neither of them is such that its existence is not partially caused by the other.

4. Kripke Criticized

Now we'll turn toward criticism of other attempts to capture the R/A distinction. We'll leave the dual-aspect use aside and focus on others' difficulties in capturing certain referential and attributive uses. Kripke claims that attributive uses are characterized by the fact that only the general intention is present, whereas referential uses involve both a general intention and a separate specific intention (Kripke, 1977, p. 264). This characterization is inadequate: because it focuses only on which intentions are present and not on which intentions are *primary*, it fails to capture the Teacher Case or the King Case.⁸

We have seen that the Teacher Case is attributive, but it involves more than just the general intention. The teacher intends to designate the individual that satisfies 'the tallest student in our class,' but she also intends to designate Henry, whom she has in mind. Because the

Teacher Case is attributive but involves both the general intention and a separate specific intention, Kripke cannot accommodate it.⁹

The King Case, on the other hand, is referential, and yet the general intention is not involved. The speaker does *not* intend to designate whoever actually is the king; his only designative intention concerns the usurper. In his own discussion of the King Case, Kripke says that the notion of what is “officially regarded” as the denotation of ‘the king’ replaces the usual denotation of ‘the king’ in the speaker’s general intention (Kripke, 2013, p. 122).

Leaving aside the rather *ad hoc* nature of this move, there are also related cases that it does not cover. Imagine that a boy is enthusiastically practicing the piano around dinnertime. His father says fondly to his sister, “Go pry Liberace away from his piano and tell him it’s time for dinner.” The father is going against his general procedure about how to use ‘Liberace,’ and the boy is not officially regarded as the referent of ‘Liberace.’ Nonetheless, the father has successfully used the name ‘Liberace’ to refer to his son. This kind of use is not terribly common, but it is certainly a part of our linguistic practice.¹⁰ It is a clear case of a referential use that lacks the general intention, so it is just as much a problem for Kripke as the Teacher Case.

5. Searle

Searle’s (1979) discussion of the R/A distinction is also worth considering in light of its similarities to the approach on offer here. Searle argues that we can understand the phenomena behind the R/A distinction in terms of primary and secondary *aspects* under which reference is made. However, this approach fails to accommodate many everyday referential uses.

Searle (1979, p. 194) tells us that “whenever a speaker refers he must have some linguistic representation of the object ... and this representation will represent the object referred

to under some *aspect* or other.” An aspect seems to be more or less a way of presenting an individual—for example, as Smith’s murderer, as the man behaving erratically while being questioned, or as the son of Mr. and Mrs. Jones.

In attributive cases, the speaker uses a designator that expresses the aspect under which she actually intends to refer to the individual. This may be because it is the only aspect under which she can refer to that individual, or because it is the only aspect of the ones available to her that is relevant to her current purposes (Searle, 1979, p. 198). For instance, the person who encounters Smith’s heinously murdered body utters ‘Smith’s murderer,’ intending to refer under the aspect of being Smith’s murderer, and nothing else.

In referential cases, the speaker uses a term that expresses one aspect under which she can refer to the individual (the secondary aspect), but she is primarily thinking of the individual under some other aspect that is available to her (the primary aspect) (Searle, 1979, pp. 195–6). The trial attendee uses ‘Smith’s murderer’ to refer to Jones, even though she is primarily thinking of him as the person behaving erratically while being questioned.

Searle’s portrayal of referential uses is problematic because it requires that speakers always think of the individual they have in mind under some determinate aspect. Searle (1979, p. 196) is explicit about this, saying, “If nothing satisfies the primary aspect the speaker didn’t have anything in mind, he only thought he did.” But often when we have something in mind, the aspect under which we are thinking of it is indeterminate. For instance, the trial attendee may have Jones in mind when she says ‘Smith’s murderer,’ but it may not be clear to her whether she is thinking of him as the man behaving erratically in front of her, the man who is on trial, or the man in the blue suit. In this and other everyday referential cases, it is hard to see how anyone could identify one determinate aspect under which the speaker is thinking of the individual.¹¹

Now imagine that the trial attendee *does* have Jones in mind under a particular, determinate aspect: as the man in the blue suit. If it turns out that Jones's suit is gray, not blue, Searle would say that the attendee didn't really have Jones in mind. It is as if referential uses were really silent attributive uses: according to Searle, the speaker is still referring to whoever fits some description, but it is a description within her mind rather than one she utters. But having in mind isn't supposed to work this way—the speaker is thinking directly about a particular individual, and *not* just thinking about whoever happens to satisfy some description. So, because speakers in referential uses don't always think of individuals under determinate aspects, and even when they do, the accuracy of the aspect is not essential to referential success, it is better to cash out the R/A distinction in terms of primary and secondary intentions rather than primary and secondary aspects.¹²

6. Korta and Perry

We'll consider one more account of the R/A distinction, developed by Korta and Perry. Their account bears some similarities to Kripke's, is able to accommodate the King Case more naturally, but still runs aground on the Teacher Case.

As a part of their overall view that utterances have multiple levels of truth conditions, Korta and Perry (2011, pp. 5–7) argue that the R/A distinction tracks which levels of truth conditions a speaker intends to convey.¹³ The speaker could say, "Smith's murderer is insane," intending to convey a belief that is true if and only if there is a unique person who murdered Smith and that person is insane. In that case, the speaker intends to convey the truth conditions at the level that just takes into account the meaning of the sentence in English, and the use is attributive. This is quite similar to our notion of an intention formed in accordance with the

speaker's general procedures with the relevant terms. Then Korta and Perry say that the speaker could intend to convey that belief and thereby also to convey a belief that is true if and only if *Jones* is insane. This involves an intention to convey the level of truth conditions that also takes into account the fact that 'Smith's murderer' denotes Jones, and the use is referential (Korta & Perry, 2011, p. 100). This is quite similar to the notion of a specific intention.

Unlike Kripke, Korta and Perry have a natural way of accounting for the King Case, quite similar to the one suggested in Section 3. They allow that a speaker can convey the belief that Jones is insane even if she knows that Jones did not murder Smith—in other words, one can use a false sentence to convey something true. In that kind of case, the speaker only really intends to convey the true belief about Jones (Korta & Perry, 2011, pp. 99–101). This sounds a lot like my claim that only a specific intention is present.

Nonetheless, Korta and Perry are unable to account for the Teacher Case. The teacher intends to convey a command that will be satisfied if and only if the unique person who is the tallest in the class lines up first, and thereby to convey a command that will be satisfied if and only if *Henry* lines up first. This fits Korta and Perry's criteria for a referential use, but we have seen that the Teacher Case is attributive. Again, we need the distinction between primary and secondary intentions to capture the R/A distinction without miscategorizing the Teacher Case.

7. Conclusion

We have seen that accounts of the R/A distinction offered by Kripke, Searle, and Korta and Perry are inadequate. Kripke's notions of general and specific intentions are the key to capturing the R/A distinction, but matters are more complicated than Kripke acknowledges. An attributive use is characterized by the fact that the intention formed in accordance with the speaker's general

procedure with the designator is either her primary or only designative intention, whereas a referential use occurs when a specific intention to designate some particular individual is either the speaker's primary or only designative intention. And a dual-aspect use occurs when a speaker has both kinds of designative intentions, but neither is primary with respect to the other.

¹ I use 'designative intention' to mean any kind of intention to pick out an individual with an utterance. Designative intention is the genus of which intentions to refer, denote, and demonstrate are species. I should also note that the issues under discussion here are closely related to the question of whether the R/A distinction is semantically significant. I will not explicitly enter into that debate, but I want to acknowledge that Kripke (1977, p. 255; 2013, pp. 122–3) takes his way of accounting for the R/A distinction to show that Donnellan does not give us a reason to think the distinction is semantically significant. My way of accounting for the distinction is compatible with that aim, but this paper should also be of interest to those who find arguments for the distinction's semantic significance convincing. For examples of such arguments, see Wettstein (1981), Reimer (1992), Amaral (2008), and Devitt (2004). Sennet (2002), on the other hand, presents considerations that point back toward seeing the R/A distinction as pragmatic, despite the sort of argument Reimer makes. Bach (2004), Salmon (2004), and Nunberg (2004) also provide developments of Kripke's idea that the distinction is a pragmatic phenomenon.

² These definitions differ slightly from Kripke's own. For Kripke (1977, pp. 263–4), the general intention involves "conventions of [the speaker's] idiolect." I think that insofar as an idiolect has conventions, they will be the conventions of the larger linguistic group, so I leave idiolects aside. Also, my claim that a specific intention involves having an object in mind is just one legitimate interpretation of Kripke's remarks about specific intentions.

³ My notion of a procedure bears significant similarities to Bratman's (1987, pp. 56–7) notion of a general policy. I should also note that Ray (1980, p. 445, n. 7) alludes to the ambiguity problem we have been discussing. Cashing out general intentions in terms of procedures avoids another issue that Ray mentions in the same footnote: "general intentions are of the wrong kind for initiating an action." If general intentions are really procedures that generate ordinary intentions when they are followed, then there is no problem about the connection to initiating action.

⁴ A speaker's procedure for using a definite description is derived from her procedures with each of the words included in the description. We will not say more about that complication here.

⁵ Or, as Kripke (1977, p. 264) puts it, the specific intention here just *is* the general intention. As Kanterian (2011, p. 369) notes, this claim is a little mysterious if specific and general intentions are supposed to be two fundamentally different kinds of intentions. However, our replacement of general intentions with general procedures that generate ordinary intentions on particular occasions removes the air of mystery.

⁶ I am indebted to Wettstein (1984, p. 69, 81 n. 23) for the idea of distinguishing between primary and secondary designative intentions. The role of that distinction in Wettstein's paper suggests that some of the discussion here might have interesting implications for indexicals.

⁷ This causal characterization of the distinction between primary and secondary intentions comes from McKinsey (1978, pp. 176–7), whom Wettstein (1984) cites. McKinsey himself cites Castañeda (1971). I say “caused the second intention to exist” rather than just saying “caused the formation of the second intention” to allow for cases in which the relative priority of two intentions changes over time. For instance, imagine someone intends to drive through her grandmother's childhood hometown because she believes it is the quickest way to her destination. That intention causes her to form an intention to please her grandmother, because she foresees the consequence and finds it desirable. Then she finds out that there is a faster route that avoids her grandmother's hometown, but her formerly secondary intention to please her grandmother is strong enough to keep the intention to drive through her hometown in existence. The intention to please her grandmother then becomes primary.

⁸ Interestingly, Kripke (2013, p. 119) does say that in a referential use, the speaker's “primary intention” may be to talk about the individual she has in mind. He doesn't develop the idea of primary intentions beyond this mention in connection with referential cases, and it doesn't appear in his main statement of how he handles the R/A distinction (Kripke, 2013, pp. 122–3). Furthermore, the way in which he mentions it doesn't allow him to capture the Teacher or King cases. However, it is interesting to see that idea make an appearance.

⁹ Ray (1980, p. 445, n. 7) also gives what he takes to be an attributive use that contains both a general and specific intention as a counterexample to Kripke's account of the R/A distinction, although without more development it is hard to see exactly how his case is supposed to work.

¹⁰ We might also use the name ‘Einstein’ for someone who is very smart (or sarcastically for someone who isn't), ‘Sherlock’ for someone who is being unusually inquisitive, or ‘Cinderella’ for someone who is complaining about chores.

¹¹ On this point, *cf.* Wettstein (2004, p. 72). I should also note that Searle (1979, p. 196) allows that a “collection of aspects” (rather than just a single aspect) could be primary, but in order to remain consistent with his other claims he would have to require that any such collection of aspects be completely determinate as well. See, for instance, Searle’s (1979, pp. 197–8) claim that the primary aspect must be determinate in order for there to be a determinate statement that is made in a referential use.

¹² Searle (1979, pp. 202–3) acknowledges that his picture of the aspects under which we think of an individual when we refer is somewhat oversimplified. However, he still holds onto the idea that a speaker can specify the primary aspect under which she was thinking of her referent if the need arises.

¹³ Korta and Perry are building on Perry (1979) and Perry (2001/2012).

Acknowledgments:

For helpful feedback on drafts of this paper, I am grateful to Peter Graham, Meredith McFadden, Eliot Michaelson, Luis Montes, Courtney Morris, John Perry, Howard Wettstein, Monique Wonderly, and several anonymous referees.

References

- Amaral, F. S. (2008). Definite descriptions are ambiguous. *Analysis*, 68(300), 288–297.
- Bach, K. (2004). Descriptions: Points of reference. (In M. Reimer & A. Bezuidenhout (Eds.), *Descriptions and beyond* (pp. 189–229). Oxford: Clarendon Press.)
- Bratman, M. (1987). *Intention, plans, and practical reason*. (Cambridge: Harvard University Press)
- Castañeda, H.-N. (1971). Intentions and the structure of intending. *The Journal of Philosophy*, 68(15), 453–466.
- Devitt, M. (2004). The case for referential descriptions. (In M. Reimer & A. Bezuidenhout (Eds.), *Descriptions and beyond* (pp. 280–306). Oxford: Clarendon Press.)
- Donnellan, K. (1966). Reference and definite descriptions. *The Philosophical Review*, 75(3), 281–304.
- Kanterian, E. (2011). Kripke’s metalinguistic apparatus and the analysis of definite descriptions. *Philosophical Studies*, 156(3), 363–387.
- Korta, K., & Perry, J. (2011). *Critical pragmatics: An inquiry into reference and communication*. (New York: Cambridge University Press)
- Kripke, S. (1977). Speaker’s reference and semantic reference. *Midwest Studies in Philosophy*, 2(1), 255–276.
- Kripke, S. (2013). *Reference and existence: The John Locke lectures*. (New York: Oxford University Press)

- McKinsey, M. (1978). Names and intentionality. *The Philosophical Review*, 87(2), 171–200.
- Nunberg, G. (2004). Descriptive indexicals and indexical descriptions. (In M. Reimer & A. Bezuidenhout (Eds.), *Descriptions and beyond* (pp. 261–279). Oxford: Clarendon Press.)
- O'Rourke, M. (1998). Semantics and the dual-aspect use of definite descriptions. *Pacific Philosophical Quarterly*, 79(3), 264–288.
- Perry, J. (1979). The problem of the essential indexical. *Nous*, 13(1), 3–21.
- Perry, J. (2012). *Reference and reflexivity* (2nd ed.). (Stanford: CSLI Publications) (1st ed. published in 2001)
- Ray, R. (1980). Transparent and opaque reference. *Philosophical Studies*, 38(4), 435–445.
- Reimer, M. (1992). Demonstrating with descriptions. *Philosophy and Phenomenological Research*, 52(4), 877–893.
- Salmon, N. (2004). The good, the bad, and the ugly. (In M. Reimer & A. Bezuidenhout (Eds.), *Descriptions and beyond* (pp. 230–260). Oxford: Clarendon Press.)
- Searle, J. R. (1979). Referential and attributive. *The Monist*, 62(2), 190–208.
- Sennet, A. (2002). An ambiguity test for definite descriptions. *Philosophical Studies*, 111(1), 81–95.
- Wettstein, H.K. (1981). Demonstrative reference and definite descriptions. *Philosophical Studies*, 40(2), 241–257.
- Wettstein, H. K. (1984). How to bridge the gap between meaning and reference. *Synthese*, 58(1), 63–84.
- Wettstein, H.K. (2004). *The magic prism: An essay in the philosophy of language*. (New York: Oxford University Press)