

Self-Reflexive Thoughts

Gilbert Harman

Princeton University

### **1 Are Self-Reflexive Thoughts Possible?**

What is it, or what would it be, for a thought (or experience) to be self-reflexive —a thought (or experience) that was intrinsically about itself? Derek Parfit once suggested to me an example of the following sort.

Alice has insomnia. She has trouble falling asleep and part of the problem is that she worries about it and realizes that her worrying about it tends to keep from falling asleep. It occurs to her that thinking that she will not be able to fall asleep may be a kind of self-fulfilling prophecy. Perhaps she even has a thought that might be expressed like this: “I am not going to fall asleep because of my having this very thought.” This thought (perhaps correctly) attributes to itself the property of keeping her awake.

Could there be such a self-reflexive thought? Someone might suggest that there could not be such a thought, because the supposition that there could be such a thought leads to contradiction via a version of the liar paradox.

Assume there can be self-reflexive thoughts.

Then one might have the thought of the form, “This very thought is not true.”

That thought would be true if and only if it were not true.

But that consequence is self-contradictory.

So, there could be no such thought.

So, there cannot be self-reflexive thoughts.

One obvious objection to this argument is that, even if there could not be a thought attributing falsity to itself, there might be other self-reflexive thoughts.

A more basic objection to the argument is that it assumes all instances of the following truth schema are correct.

(T) The thought that- $\phi$  is true if and only if  $\phi$ .

Kripke (1975) argues that the moral of the liar paradox is in part that (T) does not hold for all cases. Applying one of his points to the case of thought, it seems that if Alice is to have any concept of truth at all, she should be able to have the thought that her friend Harry is currently thinking something that is not true, not realizing that Harry is thinking that she, Alice, is currently thinking something true. It seems absurd to suppose that Alice and Harry could not have those thoughts, and their having those thoughts is

incompatible with (T).

Assume (T) holds in full generality.

Then, Alice's thought is true if and only if Harry's thought is not true.

And Harry's thought is true if and only if Alice's thought is true.

So, Harry's thought is not true if and only if Alice's thought is not true.

So, Alice's thought is true if and only if Alice's thought is not true, which is self-contradictory and so impossible.

The assumption that (T) holds in full generality leads to a contradiction.

Therefore, (T) does not hold in full generality.

This means that the liar paradox does not provide a good reason to suppose that thoughts or experiences cannot be about themselves. And Parfit's example seems a possible case of such a self-reflexive thought.

## **2 What Conceptual Resources Are Required?**

What conceptual resources are required to have self-reflexive thoughts? Obviously, a

thinker must have whatever resources are required to have thoughts at all. If thought requires the use of concepts, a thinker must have some concepts to think with. But does a thinker have to have the concept of a thought in order to have a self-reflexive thought?

It might seem that if Alice is to have a thought about a thought at all she needs to have a concept of a thought. And it might be argued that to have such a concept one must have some sort of conception of thinkers and their thoughts, a theory of mind, perhaps. And it might also be argued that very young children and at least some animals lack the relevant theory of mind and for that or some other reason fail to have the sort of concept of a thought that is required in order to have thoughts about thoughts, including self-reflexive thoughts that are about themselves. Similarly, it might be argued that very young children and at least some animals cannot have self-reflexive experiences.

A number of issues arise here. One is this. Consider a young child who has not yet acquired a theory of mind and who does not have a concept of a thought. The child does have thoughts but is too young to be able to think that others have thoughts and so perhaps unable to think that it has thoughts, because (it might be argued) the child does not have the conception of a thought as a thought of one or another thinker. The child has no concept of a thinker, that is of a creature that can think. Still, it seems that the child might have trouble falling asleep, have some sense that its worrying is part of the problem, and even have the thought “this is going to make it hard for me to fall asleep,” where the “this” refers to that very thought, although (by hypothesis) not categorized as “this thought,” because the child does not have the concept of a thought yet.

A different, if related, issue is what it counts as having a self-reflexive thought.

One way to respond is to say that a thought is self-reflexive if it contains a reference to or representation of itself. But consider Alice seeing herself in a mirror, not recognizing herself, thinking that the smiling person she sees is thinking a happy thought. Perhaps in this way Alice has a thought that is contingently or accidentally about that very thought, the thought that that very thought is a happy one. But, depending on our purposes, we might not want to count such a thought as self-reflexive, because it is not a thought *de se* to use David Lewis's (1979) terminology.

Furthermore, I want to suggest that we might want to count Alice as having a self-reflexive thought even if it does not contain any sort of specific reference to or representation of itself, at least in the way that Alice's worry about insomnia might contain a specific reference to itself. Here are four other related ways in which a thought might count as self-reflexive.

1. A thought may make use of a "token reflexive" notion.
2. A thought may have informational or propositional content that necessarily involves the thought itself.
3. A thought may present its content in a perspectival way.
4. A thought may function in relation to other sorts of thoughts in a self-reflexive way.

### **Thoughts about Oneself**

Suppose Betty has the following thought.

(1) I am less than six feet tall.

Such a thought is true if and only if the person thinking the thought is less than six feet tall. The truth conditions of the thought instance refer to that very thought. So, it might be said that there is a sense in which the content of the thought is a claim about the thought and so a sense in which the thought is about itself.

Following Reichenbach (1947), we can say the thought has a “token reflexive” aspect. The truth conditions of a token or instance of the thought are conditions that are in part about that token. We might suppose that the thought element “I” has the (Reichenbachian) content “the thinker of this thought,” i.e., “the thinker of the thought containing this very element”.

Of course, in *reporting* Betty’s thought (1) we would say something like

(2) Betty thinks that she is less than six feet tall.

We would not normally say,

(3) Betty thinks that her thought’s thinker is less than six feet tall.

Still, Betty’s thought might be self-reflexive in the way that Reichenbach’s analysis suggests even if we would not report it as self-reflexive. We certainly recognize a

difference between a first person thought and a third person report of that thought.

Consider the ambiguity of

- (4) Betty thinks that she is less than six feet tall and so does Carl.
- (4-1) Betty thinks that she is less than six feet tall and Carl thinks she is less than six feet tall.
- (4-2) Betty thinks that she is less than six feet tall and Carl thinks that he is less than six feet tall.

In the second interpretation (4-2), what Betty and Carl have in common is that they both think (1). Or they both think something that might be translated as (1).

To use Kaplan's (1979) terminology, in the first interpretation (4-1) Betty and Carl's thoughts have the same (propositional) "content" but different "characters" and in the second interpretation (4-2) their thoughts have different propositional contents but the same character.

It might be suggested that the "propositional content" of a thought can be represented as the set of possible worlds in which that content is true. Quine (1968) and Lewis (1979) respond that thoughts like (1) have what Lewis calls a "de se" character that is better represented as a set of centered worlds, centered on the individual with the thought. Or better: centered on the individual with the thought at the time of the thought. Or, perhaps even centered on that thought token! But, if the content of the thought is centered on the thought token, isn't that a respect in which the thought is about the

thought token?

### **Visual Experiences**

This discussion about the way the content of a thought is centered on the thinker, and maybe even the thought token, suggests a possibly related comparison with the perspectival point of view characteristic of visual experience. There is a sense in which perceptual experience represents its point of view and another sense in which it does not.

Consider ways in which visual experiences might be self-reflexive. Many and perhaps all visual experiences have a certain perspectival content. That is they present or represent what is perceived from a certain point of view— “from here.” The point of view is in a certain way manifested in experience, which is not to say that the point from which the scene is presented as viewed is normally itself seen. People do not normally see their own eyes, unless they are looking in mirrors, and they presumably never see their visual experiences, at least in the sense in which they see external objects. Things visual experience presents as seen are presented as having certain properties “from here”. One thing is represented as *to the left of* or *in front of* another *as viewed from here*, for example. Things are presented as at various relative distances, some as *nearer* and some as *farther* from here (Harman, 1990).

So, in one respect a visual experience does not normally represent its point of view—it does not represent its point of view as something seen. But in another respect such an experience does normally represent its point of view—it represents objects as related in certain ways to that point of view. Using terminology suggested by Perry

(1998) for a related case, we might say that the experience “concerns” a certain “unarticulated” point of view even if it is not “explicitly about” that point of view.

Furthermore, a visual experience provides *information* about its own cause. A veridical experience of seeing a tree carries the information that the experience was caused in part by perception of a tree of a certain sort, at a certain distance, etc. A perceiver can use this information to negotiate the environment and to react to changes in the environment, so the information is we might say *functionally* part of the content of experience.

Some versions of conceptual role semantics and some versions of functionalist accounts of psychological states and events claim that the content of such states and events is determined by the functional or conceptual role of aspects of those states and events (for discussion see Greenberg and Harman, forthcoming).

In addition, Sean Kelly (2005), relying on Merleau-Ponty (1962), argues that experience often has a normative aspect, providing a felt tendency toward acting in a way that would provide a better view of one or another aspect of the object. He suggests that something similar explains differences in the way an object looks depending on how far away it is. If this is correct, it indicates a sense in which the content of visual experience is self-reflexive in indicating how that very experience might be improved.

The informational content of the visual experience that  $\phi$  might be or include something like “This very experience is the result of perceiving that  $\phi$ .” Searle (1983) observes that someone perceives that  $\phi$  only if his or her experience is caused in a relevant way by what makes it the case that  $\phi$ . The experience is a successful perception only if

that condition is met. Searle takes the “intentional content” of an experience to be its success condition, so he takes the intentional content of an experience to be self-reflexive. (This use of “intentional content” is meant to cover content of all sorts of mental states and events, not just the content of intentions in the ordinary sense. I will discuss that narrower topic, the content of intentions, in the following section.)

It might be objected that an imaginative experience of a leafy tree can be a success in the relevant sense—imagining that there is a tree here— without the experience’s arising from an encounter with such a tree. But the imaginative experience might have the same content as a visual experience. So, the objection continues, the content of the visual experience cannot include the condition that the experience arises from an encounter with a tree.

One possible reply is that two experiences might share certain content without having exactly the same content. There is a “difference in feel” between imagining something and perceiving it, and it is not unreasonable to suppose that the difference has something to do with what a person takes to be the cause of the experience.

Another reply (possibly at odds with the first) is that the relevant sort of imagining is imagining seeing something. In this view, the imagining contains the content contained that might be part of a perceptual experience but contains it bracketed, as it were, and marked as merely imagined.

Furthermore, if Searle’s self-reflexive intentional content plays a functional role in how a person reacts to perceptual experience, that provides a further respect in which the content of the experience is self-reflexive.

Again it might be objected that, since a person can have perceptual experiences without having any concept of perceptual experiences, it does not make sense to suppose that the experience is self-reflexive. But first, as suggested by Perry (1998), the content of one's experiences and thoughts may have "nonarticulated" aspects. And, second, the content of perceptual experiences (and thoughts) may have nonconceptual aspects, where the content can be attributed to the experience (or thought) in virtue of the way the experience or thought functions in a person's psychology (Kelly, 2000).

### **Intentions**

I now want to discuss the narrower sort of "intentional content," the content of intentions to do something. There are reasons to treat intentions to do something as self-reflexive. The point is clearest for positive intentions, which are to be distinguished from negative intentions and conditional intentions. Roughly speaking, positive intentions are ones envisioned leading to the intended result, as when Doris intends to go to today's Faculty Meeting.

Negative intentions, like Edward's intending not to go to today's Faculty Meeting are different. Edward does not think of his intention as responsible for his not going to the meeting, because it would normally be enough if he merely did not have the positive intention to go to the meeting.

Conditional intentions are an intermediate case. They are envisioned as leading to the desired result only on the condition that a certain condition is met. Frank intends to go to the meeting if his class ends before 3 pm.

We might say that an agent's positive intention to  $\psi$  is successful only if the agent intentionally  $\psi$ s, where that involves the agent's  $\psi$ -ing in something like the way in which the agent intended to  $\psi$ . Reflection on cases suggests that, if an agent's intention to  $\psi$  is a positive intention, the agent intentionally  $\psi$ s only if the agent  $\psi$ s because of that intention. (Harman 1976, Searle 1983). So the success condition associated with a positive intention to  $\psi$  includes that the agent  $\psi$ s because of that intention—the intention leads the agent to  $\psi$ . And, as with perceptual content, if the intentional content of an intention is given by its success conditions, then the intentional content of an intention is self-reflexive—the intention is the intention to  $\psi$  in consequence of having that very intention.

Another way to argue for a similar conclusion is this. Given that an agent intentionally  $\psi$ s only if the agent  $\psi$ s in the way the agent intends, the principle that an agent intentionally  $\psi$ s only if the agent  $\psi$ s as the result of intending to  $\psi$  would be explained by the assumption that the way the agent intends to  $\psi$  involves the agent's intention's leading to the agent's  $\psi$ -ing.

Furthermore, this self-reflexive content of a positive intention is functional in the sense that, in order to reason to a positive intention to do something, an agent must be able to suppose that the resulting intention will indeed lead to the result intended. If the agent is convinced that the decision cannot affect what will happen, then the agent cannot reason practically about it and so cannot form the positive intention in question. So, conceptual role or functional theories of content might attribute self-reflexive content to

positive intentions.

Positive intentions are different in this respect from desires and hopes. You can want something or hope for something without envisioning your desire or hope being responsible for its happening.

(Of course, desires and hopes are—usually? always?—*de se*. One desires something for oneself or for something related to oneself. The Reichenbachian content of such desires treats them as token reflexive.)

Even though there is a sense in which positive intentions refer to or concern themselves as leading to the intended outcome, this does not mean that such a reference is articulated in the intentions. It is therefore quite possible for young children to form intentions even if they have no explicit concept of intentions or of self-reference.

Unlike positive intentions, negative intentions, like the intention not to go to the Faculty meeting, do not work by getting the intender to bring about the intended state of affairs. This raises the question why someone would even bother to form such a negative intention. One reason is that a person might do so in order to settle the matter as to whether he or she will go to the meeting. And that answer suggests a more general principle that would apply to all intentions—positive, negative, and conditional: *To intend that  $\phi$  is to intend that that very intention settles it that  $\phi$ .* In the case of positive intentions, an intender can settle it that the intended result will happen only by intending to bring that result about. But that is not required for negative and conditional intentions. It is often said that the difference between intention and belief has to do with “direction of fit.” We might speculate that one such difference between these types of mental is that,

whereas the content of an intention that  $\phi$  is something like, “My being in this mental state settles it that  $\phi$ ,” the content of a belief that  $\phi$  is something like “I am in this mental state because of something that settles it that  $\phi$ .”

### **KK and TT Theses**

Turning to a somewhat different topic, I observe that, if beliefs can be self-reflexive, a certain sort of objection to the so-called KK thesis and other related principles seems to fail.

(KK) If one knows that  $\phi$ , one knows that one knows that  $\phi$ .

(TT) If one thinks that  $\phi$ , one thinks that one thinks that  $\phi$ .

The objection is that such theses lead to infinite regresses and imply that if one knows or thinks anything, one knows or thinks infinitely many distinct things. Betty knows that she is less than six feet tall, so she knows that she knows that she is less than six feet tall, so she knows that she knows that she knows that she is less than six feet tall, and so on. But she does not have the time or the mental space for all this knowledge and all those thoughts, so the (KK) and (TT) theses cannot be true.

That particular objection can be avoided if knowledge and thoughts can be self-referential. Suppose that in knowing that  $\phi$ , one knows something of the form, “ $\phi$  and I know this,” where *this* refers to the whole conjunctive thought. And suppose that in

thinking that  $\phi$ , one thinks something of the form, “ $\phi$  and I think this,” where *this* refers to the whole conjunctive thought.

Consider the following possible self-reflexive knowledge:

(7) I am less than six feet tall and I know this.

Betty’s knowledge that she knows (7) is not distinct from what she knows in knowing (7) so there is no unwelcome regress.

Of course, if Betty’s thought is in words, the words are “I am less than six feet tall” not “I am less than six feet tall and I know this.” So, what does it mean to say that she has that more complex thought?

One possible answer is that her thinking the words, “I am less than six feet tall,” constitutes her having the thought expressed in (7). Her tokening those words functions as having the thought expressed in (7), for example, allowing her to answer such questions as “Do you know whether you think that you are less than six feet tall?”

### **Speaker Meaning and Common Knowledge**

A related point arises from Grice’s (1957) initial account of speaker meaning appealed to a self-reflexive intention: “*A* meant<sub>NN</sub> something by *x*” is roughly equivalent to “*A* uttered *x* with the intention of inducing a belief by means of the recognition of this intention” (384).

In his original discussion, Grice added, “This seems to involve a reflexive paradox,

but it does not really do so.” Later, as various complications were noted, Grice (1969) and Schiffer (1972) replaced the self-reflexive analysis with ones involving a series of intentions, with later intentions referring to the earlier intentions. This led to issues about the existence of the potentially infinite regress of intentions required, issues that could have been avoided by staying with self-referential formulations.

It seemed that additional clauses were needed to avoid counter-examples, but as clauses were added it became unclear how speakers could have such complicated intentions. Again it would have been better to have kept the original analysis with its self-reflexive intention.

I assume Grice’s worry about self-reflexive paradox is the worry that self-reference leads to paradox. But, as already noted, considerations about such paradoxes do not imply that there cannot be thoughts about themselves. So, Grice’s original formulation of his analysis of speaker meaning can be kept, avoiding the enormous difficulties that arose from trying to formulate a version without thoughts about themselves.

A similar point applies to the kind of analysis offered by Lewis (1969) and Schiffer (1972) of what it is for it to be common knowledge that  $\phi$  among a group of people. Their basic idea was to suppose it is common knowledge that  $\phi$  in a group of people  $G$  if and only if: (1) everyone in  $G$  knows that  $\phi$ , (2) everyone in  $G$  knows (1), (3) everyone in  $G$  knows (2), etc.

As in discussions of Grice’s analysis of speaker meaning, this sort of analysis of common knowledge led to worries about how many such iterations an ordinary person

can handle. A better analysis of this sort would appeal to self-referential knowledge with a relatively simple clause that implies all the iterations, taking it to be common knowledge that  $\phi$  in a group of people  $G$  if and only if everyone in  $G$  knows “ $\phi$  and everyone in  $G$  knows this,” where the final “this” refers to the conjunction in question.

### **Appeals to Higher Order States**

There are reasons to suppose that all or most psychological states and events are self-reflexive in one or another way. That result may be bothersome for various philosophical theories that attempt to explain one or another notion in terms of higher-order thoughts. For example, theories that would distinguish values from mere desires by appeal to second order desires might have difficulty (Harman 1993).

Similarly for theories that identify conscious experiences as those one has a higher-order thought about or perception of (e.g., Rosenthal 2002, Caruthers 2000, Lycan forthcoming). Such theories face the following dilemma. If they allow that reflexive thoughts and experiences count as the relevant sort of higher-order thoughts or experiences, there will be difficulty in supposing that any thoughts or experiences are not conscious—perhaps there will be no unconscious beliefs, no unconscious intentions, and no unconscious perceptual experiences. Even blindsight will count as conscious experience, because of the perspectival character of blindsight.

Your stomach apparently regularly tests its Ph-level. The resulting representation is perspectival and functions as a self-reflexive representation in the way it regulates the amount of stomach acid produced in the stomach to which it attached rather than

something else. Presumably that does not count as a conscious representation because it is a subpersonal rather than personal representation—it is an experience of your stomach, not an experience of yours.

What about unconscious but self-reflexive intentions and blindsight? Someone might argue that they are subpersonal representations like representations of acidic level in the stomach and so they are not counter-examples to the claim that conscious states and experiences of the person are states for which there are corresponding higher-order thoughts or experiences. But this would be to identify almost all personal states and experiences with conscious states and experiences and the appeal to second-order representations would be doing no work.

I said there was a dilemma for higher-order theories of consciousness. I just described the horn on which having a self-reflexive thought or experience counts as having the relevant higher-order thought or experience. If that horn is accepted, all or almost all thoughts and experiences count conscious, which is presumably an undesirable result.

The other horn is to deny that a self-reflexive thought or experience counts as the relevant higher-order thought or experience. What's needed, it might be said, is a higher-order thought that articulates a reference to the lower-order thought.

The problem with this horn is to find evidence for the existence of relevant higher-order thoughts or experiences which are not the thoughts or experiences themselves. In as much as the evidence is that conscious thoughts and experiences are thoughts and experiences one knows one has and can report to others, it would seem that, if higher-order states and experiences are needed to explain such self-knowledge and ability to self-

report, it is unclear why the relevant higher-order states and experiences cannot be the original states and experiences themselves. What is added to the explanation if the relevant higher-order states and experiences are distinct from the original states and experiences?

In response to this dilemma, it might be suggested that the relevant self-reflexive states and experiences (or the distinct higher-order states and experiences) must be *poised for free use in reasoning and rational control of thought*. Block (1995) calls states and experiences meeting that condition “A-conscious” (“access conscious”) and observes that many theorists identify consciousness with A-consciousness. But, like appealing to a distinction between personal and subpersonal states and events, this added criterion might be what’s doing the work in distinguishing conscious from unconscious states and events.

### **3 Summary**

I began by considering whether self-reflexive thoughts or experiences are possible and observed that worries about semantic paradox provide no reason to doubt their existence. Next I considered whether self-reflexive thoughts or experiences required special conceptual resources and noted various ways in which such thoughts or experiences might be self-reflexive without containing articulated references to themselves. They might contain token reflexive notions. They might be perspectival through having *de se* content or through representing things from a certain point of view related to the relevant thought or experience. Their self-reflective content might function in a way that gives them a distinctive conceptual role. I suggested that positive intentions are self-reflexive in that such an intention involves a conception of itself as leading to the intended outcome. It may be that any intention is self-reflexive in involving a conception of itself as settling it

whether something is the case, and it may be that any belief is self-reflexive in involving a conception of itself as resulting from something that settles it whether something is the case. I discussed the way in which various attractive philosophical ideas are best expressed though appeal to self-reflective thoughts: KK and BB theses, Grice's account of speaker meaning, and Lewis and Schiffer's analysis account of common knowledge. Finally, I noted that, if most thoughts and experiences are self-reflexive, that poses a problem for a couple of philosophical appeals to higher order states—certain attempts to analyze values and certain attempts to explain what makes a state conscious.

## **Bibliography**

- Block, Ned, (1995). "on a Confusion about a Function of Consciousness." *Behavioral and Brain Sciences* 18, 227-287.
- Carruthers, Peter, (2000). *Phenomenal Consciousness*. Cambridge: Cambridge University Press.
- Greenberg, M. and Harman, G., (forthcoming). To appear in *The Oxford Handbook of the Philosophy of Logic*, Ernie Lepore and Barry Smith, eds. (Oxford: Oxford University Press).
- Grice, H. P., (1957). "Meaning," *Philosophical Review* 66, 377-388.
- Grice, H. P., (1969). "Utterer's Meaning and Intention," *Philosophical Review* 78, 147-177.
- Harman, Gilbert, (1976). "Practical reasoning," *Review of Metaphysics* 29, 431-463
- Harman, Gilbert, (1993). "Desired Desires," *Value, Welfare, and Morality*, Ray Frey and

- Chris Morris, eds. (Cambridge: Cambridge University Press: 1993), pp. 138-157.
- Kaplan, David, (1979). "On the Logic of Demonstratives," *The Journal of Philosophical Logic* 8, 81-98.
- Kelly, Sean, (2000). "Grasping at straws: motor intentionality and the cognitive science of skillful action." In *Heidegger, Coping, and Cognitive Science: Essays in Honor of Hubert L. Dreyfus*, vol. II, Mark Wrathall and Jeff Malpas (eds.), (Cambridge, MA: MIT Press), 161-178.
- Kelly, Sean, (2005). "Seeing Things in Merleau-Ponty," *Cambridge Companion to Merleau-Ponty*, ed. Taylor Carman (Cambridge: Cambridge University Press), 74-110.
- Kripke, Saul, (1975). "Outline of a Theory of Truth," *Journal of Philosophy* 72, 690-716.
- Lewis, David, (1969). *Convention*. (Cambridge, MA: Cambridge University Press).
- Lewis, David, (1979). "Attitudes *De Dicto* and *De Se*." *Philosophical Review* 88, 513-543.
- Lycan, W. G., (forthcoming). "The Superiority of HOP to HOT." *Higher-Order Theories of Consciousness*, R. Gennaro, ed. (John Benjamins: Amsterdam).
- Merleau-Ponty, Maurice, (1962). *Phenomenology of Perception*, tr. Colin Smith, (New York: Routledge & Kegan Paul).
- Perry, John, (1998). "Myself and I," in *Philosophie in Synthetischer Absicht*, ed. Marcelo Stamm (Stuttgart: Klett-Cotta), 83-103. Reprinted in Perry (2000).

Perry, John, (2000). *The Problem of the Essential Indexical and Other Essays*, Expanded Edition (Stanford, CA: CSLI Publication).

Quine, W.V., (1968). "Propositional Objects," *Critica* 2, 3-22.

Reichenbach, Hans, (1947). *Elements of Symbolic Logic* (New York: Free Press).

Rosenthal, David, (2002). "Explaining consciousness." In David Chalmers (ed.),

*Philosophy of Mind* (Oxford: Oxford University Press) 406-421.

Schiffer, Stephen, (1972). *Meaning* (Oxford: Oxford University Press).

Searle, John, (1983). *Intentionality* (Cambridge, England: Cambridge University Press).