

Immunity to Error through Misidentification and Non-attributive Self-reference

Ted Lougheed (ted_lougheed@carleton.ca)

Institute of Cognitive Science, Carleton University

1125 Colonel By Drive, Ottawa, ON K1S 5B6 Canada

Abstract

Recent empirical literature (Jeannerod & Pacherie, 2004; Mizumoto & Ishikawa, 2005) purports to challenge the thesis that certain forms of self-awareness are immune to errors of misidentification with respect to the first-person (IEM). I argue, first, that these studies do not present a challenge to the IEM thesis, and furthermore that IEM is indicative of a fundamental distinction between two ways of being self-aware—a distinction that has real consequences for empirical studies of self-awareness. In the final section of the paper I suggest that the non-attributive self-reference (NSR) thesis better explains what is special about the distinction than IEM does by itself.

Keywords: self-awareness; self-as-subject; non-attributive self-reference; immunity to error through misidentification; consciousness; proprioception; attribution errors

Introduction

Recent empirical research into the self-attribution of actions (Jeannerod & Pacherie, 2004) and illusions of body ownership (Mizumoto & Ishikawa, 2005) purports to challenge the thesis that certain forms of self-awareness are immune to error through misidentification with respect to the first-person (hereafter, IEM). Briefly, IEM is the view that for certain judgments, a person cannot make the judgment, claim that the judgment is about her, and be wrong. Exactly what falls under the label of “certain judgments” is a point on which there is disagreement, and I will argue that the only judgments that count as IEM are judgments about the owner of experience—the one that is having the experience. IEM does not apply to judgments about others, a point that will become important when I examine the empirical challenges to IEM.

IEM has been the basis for making a fundamental distinction between two senses of self. First is the sense of self as an object of experience (hereafter, SAO-awareness). SAO is comprised of the particular components of experience that I attribute to myself; in other words, the object(s) of current experience that I take to be part of me. Second is the sense of self as the subject of experience (hereafter, SAS-awareness), that is, the sense that there is something—namely, *me*—doing the experiencing, owning or authoring (Shoemaker, 1968; Evans 1982). I examine the empirical challenges to IEM, and argue that while the research may challenge some claims about IEM with respect to certain kinds of self-knowledge, we still have IEM with respect to ownership of the experience.

Nevertheless, there is good reason to suggest that IEM does not fully capture what is distinctive about SAS-awareness.

awareness. In particular, it does not capture the fact that SAS-awareness does not divide experience—that when I am aware of SAS, there is no alternative subject of experience with whom I could confuse. In the final section of the paper, I argue that the thesis of non-attributive self-reference, or NSR (what Brook, 2001 calls *non-ascriptive* self-reference¹) explains whatever immunity may exist and, because it also explains the non-experience-dividing quality of SAS just mentioned, tells a better story about what is special about SAS-awareness than bare IEM does by itself.

Two Senses of Self: SAS and SAO

Self-as-object or SAO can be defined as the set of features or properties—for example our bodily appearance or mental states and the like—that I identify with myself. When I think of SAO, I am thinking of a particular object in the world, namely a human being named Ted, with a certain height, weight, hair color, and so forth. Aspects of my physical appearance, political opinions, taste preferences, propensity to particular emotions, and other particular characteristics are all features that might compose one's sense of SAO at a given time.

Our sense of SAO is constantly changing; at any given moment I may only be conscious of a small subset of self-related properties. I specify “at any given moment” because I cannot possibly be conscious of the features that I associate with myself all at once, given the limitations of working memory. SAO-awareness changes over time as I come to attribute different features to myself. The properties that I attributed to myself as a child are not the same ones that I attribute to myself now, so my SAO may be different in almost every respect. Indeed, the features I attribute to myself can change daily, or from one moment to the next. I do not have a static conception of myself as an object that I call upon whenever I think about myself. At one moment I may have a rich concept of myself, as when caught in a moment of deep introspection while standing in front of a mirror; at other times, such as in the middle of a dream, I may have only a hazy concept of myself consisting of one or two significant features. These features need not accurately reflect the real properties of my bodily or mental states—they are simply the parts of my experiential field that I take to be part of me, rather than someone or something else.

SAS-awareness, by contrast, is the sense I have of myself as the *subject* of particular experiences, that is, the

¹ In the interests of clarity and cross-disciplinary understanding, I have chosen the term *non-attributive* instead of *non-ascriptive*, as the term *ascription* is less common across disciplines.

entity who has the experiences. I make a distinction between SAS—the notion of a single entity common to all experience—and SAS-awareness, which is the state of being conscious of said notion. SAO-awareness is sufficient for SAS-awareness, but SAS-awareness is insufficient for SAO-awareness; it is possible for me to be SAS-aware yet fail to attribute any objects of experience to me. The point is not that SAS-awareness is a particular experience *in addition to* SAO-awareness; it is rather that SAS-awareness makes possible the identification of particular features as *mine*. To have SAS-awareness is to appear to myself *as* myself, as Brook (2001) explains. That is, to experience “*x* as subject” is for *x* to appear to me as *x* rather than *y*. Consider a particular conception that I have of myself at the time of writing. It is not enough to be aware of a constellation of specific features, because I must also recognize that those features are *my* features, and not those of somebody else or the external world.

To appreciate the significance of the SAS/SAO distinction, it is helpful to consider a case in which no such distinction exists. Consider a being—a mollusk, perhaps—that is conscious of its own body, and capable of distinguishing its body from external objects, yet fails to conceive of that body as “*me*.” Such a being would be capable of autopoiesis, taking a special interest in protecting its body yet not conceiving of that body as belonging to a common subject of experience (the “I” or first-person).

We need not speculate beyond our own species to appreciate the significance of SAS-awareness, however. It is possible for me to be conscious of a part of my body yet fail to recognize that it is *my* body part. I may see a partial reflection of myself in a mirror (where all I can see is an arm), for example, and not realizing that I’m looking at a mirror, think that I am seeing someone else’s arm.

IEM: Definition and Brief History

Immunity to error through misidentification with respect to the first-person (IEM) is the claim that I cannot make a judgment about myself, take the judgment to be about me, and be wrong. IEM only applies in one direction, meaning that it does not apply to judgments about others. I can make a judgment about someone else, take the judgment to be about me, and be wrong. This point is often lost in discussions of IEM, and it must be kept in mind when I consider the empirical challenges, below.

When I make a judgment involving SAS, it is IEM because there is no identification component involved (Shoemaker, 1968; Evans, 1982). That is, it is neither necessary nor sufficient to identify certain properties in order to identify SAS, so there is no logical possibility of an error of misidentification. For example, when I say, “I feel hungry,” I cannot be wrong about *who* is feeling hungry, even if I am wrong about being hungry. If, on the other hand, I look to my right, see a wrist with a watch on it, and declare “I am wearing a watch,” I could be in error about who is wearing the watch (e.g., if the view of the rest of my

arm is obscured and the wrist belongs to someone else). The latter kind of error is relatively easy to produce under experimental conditions, especially when participants need to rely on indirect visual feedback about their own bodies (Jeannerod & Pacherie, 2004). Even when a participant knows that a body part is not her own, a feeling of ownership over that part can be induced, as in the famous “rubber hand illusion” pioneered by Botvinick and Cohen (1998). In this illusion, the participant’s arm is hidden while a rubber hand is placed on the table, oriented in the same direction as her real hand but laterally displaced. When strokes are applied synchronously to both real and rubber hands, the participant reports feeling the touch in the location of the rubber hand. The illusion breaks down when strokes are applied asynchronously, demonstrating that the effect is not simply a product of empathy.

The observation that there is something special about “I” thoughts has been observed for several centuries, most notably in the form of Descartes’ “cogito” in his *Meditations* (1641/1996), where he famously declares that he cannot be in error about being a thinking thing, despite being able to doubt all other contents of experience. It was not until the early part of the last century, however, that the phenomenon of IEM was identified as such. Wittgenstein (1934/1960) was the first to explicitly note that certain predicates were immune to misidentification errors:

There are two different cases in the use of the word “I” (or “my”) which I might call “the use as object” and “the use as subject”... One can point to the difference between these two categories by saying: the cases of the first category involve the recognition of a particular person, and there is in these cases the possibility of an error, or as I should rather put it, the possibility of an error has been provided for. ...On the other hand, there is no question of recognizing a person when I have a toothache. To ask, ‘are you sure it is you who have pains?’ would be nonsensical. (pp. 66-7)

As there is no identification of a particular object (person) involved and thus no possibility of misidentifying the object, Wittgenstein took the lack of a referent as evidence against a physical self. I will not speculate on the metaphysical implications of IEM in this paper; suffice it to say that the absence of a referent (or representation, percept, etc.) need not imply the absence of underlying psychological processes (that I assume are physical).

IEM in its current form was first explicitly identified by Shoemaker (1968). Following directly from Wittgenstein’s distinction between subject and object, Shoemaker observes that identifying oneself as an object in the world cannot be the foundation for the use of the first-person pronoun as subject. He states that to identify himself as an object would require either finding something true of the object that he knows to be true of himself, or knowing that the object is related to himself in a way that only he could be related to the object. If neither condition holds, he would be led to an infinite regress of identifying an object with some further object, and that object with yet another object, and at no point would he relate an object with himself. Furthermore, the identification of some object with him would always

carry with it the potential for an error of misidentification; when he uses "I" as subject, there is no such possibility (Shoemaker, 1968).

Now because this fundamental knowledge of self as the subject of experience cannot be based upon identification of objects, Shoemaker concludes that there is no principled reason against other sorts of knowledge being identification-free. He supposes that there might be an entire class of psychological predicates that are identification-free. For example, "sees a blue cylinder near the center of the field of vision" would qualify as such a predicate, according to Shoemaker, because we do not need to make a further identification of *what* sees a blue cylinder in the center of the field of vision.

Gareth Evans (1982) extends identification-free knowledge to certain bodily predicates, specifically proprioceptive knowledge—the sense of where our body parts are located relative to one another. In other words, proprioceptive knowledge of my body parts is enough to guarantee that those parts are mine. On this view, when the knowledge that my arm is extended over my head comes to me via proprioception, it makes no sense to ask, "Someone's arm is over his head, but is it my arm that's over my head?" The claim that proprioceptive knowledge counts as identification-free is directly challenged by Mizumoto & Ishikawa (2005), and I consider this challenge below.

Brook (2001) claims that if I ever have IEM in identifying myself, it can only be when I am aware of myself *in virtue of having an experience*, and not on any other basis. On this interpretation, it is not enough for a judgment to be identification-free to be IEM—the judgment must occur with respect to what one is currently experiencing. I defend this interpretation of IEM, claiming that any predicate about an *experience* is IEM, whereas any predicate about a non-phenomenal state of affairs will fail to be IEM *in principle*—that is, as a matter of logical necessity. I can be wrong about who is wearing a watch, just in case it turns out that the watch is not fastened around *this body's* wrist. If I said instead, "I think that I am wearing a watch," I am describing my experience of thinking "I am wearing a watch", and cannot be in error about *who* is having that experience, although I could still be in error about wearing the watch, that is, the content of my experience. In sum, *I cannot be wrong about who is having this experience*.

While there may be cases of *de facto* IEM with respect to SAO, what is missing is the *logical* guarantee. Logically speaking, however unlikely, we have no difficulty supposing that someone could mistakenly attribute an identification-free predicate to herself, when the predicate applies to someone else. The fact that we never have access to the mental states of others is a contingency of our particular biology.

I defend Brook's claim that IEM applies only at the level of the structure of experience—specifically to SAS—

not the contents of experience. I maintain that neither mental nor physical predicates are logically IEM, and that logical IEM extends only to the experience itself. The subject's idea of "I" is not tied to any particular mode of experience, perceptual or otherwise, but a kind of realization generated by a separate cognitive faculty (SAS-awareness).

Empirical Challenges to IEM

It may seem difficult, if not impossible, to imagine what conditions would need to hold for IEM to be false. However, for IEM to be empirically testable, we need to be able to specify what would need to be true for IEM to fail with respect to SAS. IEM would fail just in case I could have an experience, claim that it is my experience, but be wrong about who is experiencing it. If I *think* that I am having an experience of a blue mug on the table in front of me, but in fact it is John having that experience, and *I am not actually having the experience*, then SAS would not be IEM. It does not seem to be possible to claim someone else's experience as one's own, because the very nature of having an experience makes it one's own, even if the source is someone else.

We can imagine a scenario where one is having a shared experience—that is, my nervous system is hooked up in such a way that a copy of John's sensory inputs is sent to my brain. John goes about his daily business, while I sit in a lab somewhere. There is an intuitive appeal to supposing that the experiences are *actually* John's, because we identify John with the physical body that is going about its daily activities—standing in line at the coffee shop, walking to work, writing papers, smelling roses, and so forth. There is a sense in which these experiences are *also* mine, however, because I am in fact experiencing them. There seems to be something inherent to experience that guarantees that the possessor recognizes it as her own experience, to the extent that she possesses a sense of self. In the complete absence of the distinction between self and other, the question of who is having what experience can't even be formulated. I am not suggesting, contrary to many conceptions of self-awareness, that to be aware is necessarily to be *self-aware*. I am suggesting that the possession of a sense of self guarantees mental self-attribution of all experience.

I examine two empirically-supported claims against IEM in this section, specifically that some judgments involving reference to SAS violate IEM. Jeannerod and Pacherie (2004) claim that the fact we can make errors regarding the agency behind an observed action demonstrates that we can make errors of misidentification with respect to SAS; meanwhile, Mizumoto and Ishikawa (2005) claim that under certain conditions, proprioception is not IEM. While I do not accept Evans' claim that proprioception is logically IEM, I will show that Mizumoto and Ishikawa's experiment fails to challenge the claim that proprioception is *de facto* IEM.

Claim: Agency is not IEM

Jeannerod and Pacherie (2004) claim that we do not have IEM with respect to agency (i.e., authorship of actions). They identify agency as a subset of subjectivity (i.e., SAS). If they are right, then we do not have IEM with respect to SAS in certain cases. I argue that while their results may show that our actions are not IEM, we still have IEM with respect to SAS. My argument is based on a rejection of the premise that agency is part of SAS.

Jeannerod and Pacherie's claim is based on the following considerations:

- 1) We are aware of the intentions of ourselves and others non-inferentially, that is, we perceive them directly via the process of action simulation. The same parts of the brain that are active when we plan our own actions are active when we observe someone else's actions.
- 2) We are aware of intentions as "naked" intentions, that is, there is nothing inherent to awareness of an intention to identify whose it is. An additional step is needed to identify the owner of the intention.
- 3) The process by which we attribute an intention to an owner is not entirely reliable, so it is possible to misattribute in certain circumstances.

According to the authors, action simulation is the primary mechanism by which we self-attribute. On the action simulation theory, simulated and actual actions are indistinguishable at the neural level save for the execution. A person is better at simulating her own action than simulating that of another, because she is far more familiar with her own actions. However, by manipulating the accuracy of simulation, one can cause a person to identify herself as the source of someone else's action.

The authors claim that the Alien Hand paradigm (Nielsen, 1963; Fournier & Jeannerod, 1998) demonstrates how such errors of misidentification can be caused in mentally healthy people, and that we have a tendency to over-attribute actions to ourselves. In this experiment, the participant places her arm through a hole in a box, and is instructed to draw a line towards the top of a page. A small window in the top of the box allows the participant to observe her line-drawing. Unbeknownst to the participant, the box is fitted with a mirror so that the participant actually sees the hand of the experimenter (both the participant and the experimenter wear a glove to mask obvious differences between the hands).

When the participant is asked to draw the line, the experimenter mimics the action. Midway through the task, the experimenter makes the line diverge from straight according to a set number of degrees up to 10°. Even though the line starts to diverge, the participant still believes that she is the one drawing the line, so long as the divergence is not too great. Furthermore, she attempts to correct for the path of the line, but does not realize that she is doing so. When confronted with the discrepancy, participants would commonly confabulate, claiming for instance that they must

have been tired or inattentive (Jeannerod & Pacherie, 2004).

As Jeannerod and Pacherie explain, when we make a determination about who is performing an action, we rely on visual information so long as the discrepancy is not too large. So long as our intention is roughly congruent with the perceived action, we attribute the action to ourselves. The authors claim that the feeling of ownership is tightly bound to the feeling of agency, such that when we make a mistake with respect to the latter, we also make a mistake with respect to ownership. In other words, we make an error of misidentification with respect to SAS.

There are a number of problems with the authors' conclusion. First, their argument depends on the highly dubious suggestion that we are aware of intentions non-inferentially. They conflate perception of action with awareness of intention but are not justified in doing so based on the evidence they provide. Second, being wrong about the author of an action is not the same as being wrong about the owner of an intention. Clearly, the participant is confused about the former, and does not have IEM with respect to actions (this should not be surprising), but actions, as movements of objects of experience, are part of SAO, not SAS. The participant is not confused about her intention just because she mis-attributes the action; it just happens that someone else's action tracks that intention accurately. The problem of mis- or over-attribution reduces to a mere problem of causation, and does not introduce any difficulties for IEM.

Finally, even if the participant was mis-attributing the action of another to herself, this would not count as a counter-example to IEM. Shoemaker never claimed that IEM applied in the reverse direction—that a person could not make a judgment about someone else, think that it was a judgment about someone else, and yet think the judgment is about her. IEM only applies to judgments with respect to the self.

Claim: Proprioception is not IEM

Mizumoto and Ishikawa (2005) attack IEM from a different angle, claiming that proprioception is not IEM because it is tightly (perhaps inextricably) bound to visual perception. They claim that misidentification errors with respect to vision are thereby also errors with respect to proprioception. I argue that, far from challenging the view that proprioception is IEM, Mizumoto and Ishikawa's findings actually *support* that view, and bolster the central distinction we make between SAS and SAO.

The authors base their claim on the results of a study on the "body-swap" illusion, whereby the participant is made to feel as if she has swapped bodies with another participant. This experiment is supposed to show that a touch is taken to be located where it is seen, not where it is in fact felt. The authors argue that this experiment demonstrates that judgments based on proprioception are not IEM. The participant wears a head-mounted display that relays visual information from stereo cameras, mounted on the

head of another participant. The two participants sit in chairs some distance apart. The view from the HMD shows only the one participant sitting in a chair. One of two manipulations is done to try to “trick” the participant into reacting as if she were the other. In the first manipulation the experimenter taps the shoulders of both participants at the same time, and then asks the first participant to report on how she felt. The second manipulation uses a non-verbal measure. In this one, the experimenter pretends to hit the participant and records whether or not she reacts (Mizumoto & Ishikawa, 2005).

On the first manipulation, all participants reported feeling as if they were the other participant. On the non-verbal measure, one of three participants reacted to the threat of being hit, and Mizumoto and Ishikawa report that the other two “almost” reacted.

Mizumoto and Ishikawa take these results to show that misleading visual information can cause errors in judgments based on proprioception. Because an identification component (in this case, misleading visual stimuli) is involved, and the judgment is based on proprioception, they claim that proprioceptive judgments are therefore subject to errors of misidentification. This claim hinges on the assertion of the *de facto* inseparability of visual and proprioceptive information. They suggest that single-mode proprioception is impossible, because there is always an “imagined” visual space associated with the proprioception. So, although it is *logically* possible to be IEM with respect to proprioceptive data alone, it is never *actually* possible because the visual information is intimately linked with proprioception (Mizumoto & Ishikawa, 2005).

There is a problem with the authors’ claims. Mizumoto and Ishikawa interpret the IEM thesis correctly, but over-extend what they are actually justified in claiming about proprioception. Mizumoto and Ishikawa acknowledge that it doesn’t make sense to ask, “Someone is being tapped, but is it really me who is being tapped?” (2005), but they persist in claiming that proprioception is not immune because the judgment based upon proprioception is in error.

It is clear that the participants are not mistaken about who is being tapped, even if they are mistaken about who is about to be hit. In the latter case, however, they are not confusing someone else’s proprioception for their own, or wondering if the proprioception they are feeling is their own, which is what it would mean to misidentify the proprioception.

Judgments based on proprioception can be wrong, but we can’t be mistaken about the fact that those judgments are *ours*. If proprioception were not immune, then we would be able to have awareness of the proprioception of another and attribute it to ourselves in the same way I attribute the visual percept of a hand to myself. Even if vision and proprioception are tightly linked, there is a component of the judgment about which the participant cannot be in error, namely, *who* is feeling the tap. Mizumoto and Ishikawa have not presented any data that can challenge the claim of IEM

with respect to proprioception.

Also, while the authors claim that the results demonstrate a tight link between vision and proprioception, the experiment actually serves to demonstrate that participants can distinguish the source of visual perception (SAO) from proprioceptive feelings, which we know are ours in virtue of having, and therefore part of reference to SAS. Thus the findings of the authors actually support the claim that proprioception is IEM, and the distinction between SAS and SAO.

Non-attributive Self-reference (NSR)

IEM is often heralded as *the* key feature of reference to SAS. While I agree that IEM is significant, I argue that the thesis of non-attributive self-reference spells out *why* IEM is significant, and provides some clues as to the psychological mechanism of self-awareness.

The NSR thesis is as follows: when we refer to SAS alone, we do not attribute any particular properties to ourselves. All the experience that is required to refer to SAS is given in the structure of experience itself, rather than the content. It is neither necessary nor sufficient to be aware of any particular object of experience in order to be aware of SAS (Brook, 2001). We are of course frequently aware of SAO when we are aware of SAS. The point is that reference to SAS cannot be *reduced* to awareness of SAO.

We are not aware of ourselves in the same way that we are aware of objects of experience. For instance, for me to identify a car *as* a car, I need only be aware of some object of experience, namely a certain collection of colors, edges, sounds, and possibly smells and other sensory data that I associate with “car-ness”. But it is not sufficient, or even necessary, for me to identify certain properties to identify myself – I need only be having an experience, *any* experience whatsoever, to be aware of SAS. To use another example, if I think or say “I am seeing the color red,” I need not have in mind *any* particular properties of myself. All I can say about myself is that I am seeing the color red, but this does not provide me with any additional information about myself.

To identify a given subject as *my-self* due to the possession of some property, I would have to know that *I* observe it (Shoemaker, 1968). Because this self-knowledge is the grounding for my identification of self *as* myself, that knowledge can’t also be grounded on the identification of a particular object, or else we would have an infinite regress of self-predication.

If we grant that self-reference must at bottom be non-attributive (and I do not see how it could be otherwise) then IEM logically follows. Being non-attributive, SAS cannot be “experience-dividing.” (Brook, 2001; Bennett, 1974) That is, the subject (SAS) has no features with which to distinguish it from others. As Brook explains, “no representation of mine is made different from any other representation of mine by the fact that it makes me aware of myself as its subject... What we cannot do is compare

[SAS] to, contrast it with, one object *rather than another*." (2001) It is in virtue of the fact that I cannot contrast SAS with any object that I have IEM with respect to SAS. In other words, what it is for x to be "me" is just for x to be placed in a certain relationship with the structure of experience. This relationship holds both synchronically and diachronically, since when we are aware of SAS, it is always presented in the same manner.

Contrast the experience of a tree. I can identify one tree from another by various features, such as the shape of the leaves or the hardness of the bark, but there are no analogous features by which I can distinguish one subject of experience from another. I am only ever aware of one subject (namely, me)—I am only aware of others as objects, and if I were to become aware of another subject in the same way that I am aware of myself, there would be nothing to distinguish that subject from myself. IEM holds with respect to SAS precisely because there are no features to distinguish one subject from another, and thus no question of identity.

IEM is often used to support the claim that we are aware of ourselves in more than one way, as SAS and SAO. NSR gives a more complete story, and subsumes IEM as a particular consequence of NSR. My intention is not to claim that we should ignore IEM, but that we should recognize IEM as an indicator of NSR. We can understand the distinction as that between interesting phenomenon and underlying psychological mechanism.

A useful parallel is the relationship between an optical illusion and the psychological mechanism driving the illusion. As cognitive scientists, we are ultimately interested in understanding *why* we see the illusion and not the illusion itself. We may use the illusion as an interesting example to draw attention to the psychological story we want to tell, but we wish to keep separate, if only for terminological clarity, the illusion from the psychological mechanism driving the illusion. We do not ignore the illusion, but acknowledge that the illusion is indicative of a particular mechanism, perhaps going so far as to incorporate the illusion into an experimental study. Likewise, I do not propose to discard IEM, but to recognize it as indicator of NSR, and potential tool for empirical study of NSR. By isolating instances where IEM applies and where it fails to apply, we may be able to identify the psychological and even neural correlates of NSR.

Conclusion

IEM remains a viable idea despite attempts to challenge it via empirical means. The observation that certain judgments are IEM, however, has overshadowed the observation that reference to SAS is non-attributive—that it does not depend on awareness of particular objects of experience. NSR better explains the significance of the SAS/SAO distinction as it relates to the study of underlying psychological mechanisms than IEM can do by itself.

I do not mean to suggest that IEM should be ignored, to

be replaced wholesale by talk of NSR. The observation of IEM has value as a tool for demarcating the boundaries of our SAS-awareness. We should be able to exploit IEM under experimental settings to help us identify correlated mechanisms, which would allow insight into the underlying structure of experience.

If the distinction between IEM and NSR seems a trifle at this stage, it is due solely to the fact that we still know so little about the psychological mechanism(s) of self-awareness, in particular SAS-awareness. This paper represents only a small piece of a much larger project to unravel the psychology of SAS. The claim that reference to SAS is non-attributive raises a host of perplexing questions, such as *how* it is possible to refer to self without attributing anything *to* it, and other questions that I don't have the space to address here.

Acknowledgments

This research was supported by funds from Carleton University, including the Hamlin Graduate Fellowship, the David and Rachel Epstein Fund, and the Dr. Thomas Betz Memorial Fund. I would like to thank Andrew Brook, Raj Singh, and Frederique de Vignemont for providing me with helpful comments and criticism on previous drafts.

References

- Bennett, J. (1974). *Kant's Dialectic*. Cambridge: Cambridge University Press.
- Botvinick, M. & Cohen, J. (1998) Rubber hands 'feel' touch that eyes see. *Nature*, 391, 756.
- Brook, A. (2001). Kant, self-awareness and self-reference. In A. Brook and R. C. DeVidi (Eds.), *Self-reference and Self-awareness* (pp. 9-30). Philadelphia: John Benjamins Pub. Co.
- Descartes, R. (1641/1996). *Meditations on First Philosophy*. Trans. John Cottingham. Cambridge: Cambridge University Press.
- Evans, G. (1982). Self-identification. In *Self-reference and Self-awareness* (pp. 95-141).
- Fournier, P. & Jeannerod, M. (1998). Limited conscious monitoring of motor performance in normal subjects. *Neuropsychologia*, 36, 1133-1140.
- Jeannerod, M., & Pacherie, E. (2004). Agency, simulation and self-identification. *Mind & Language*, 19(2), 113-146.
- Mizumoto, M., & Ishikawa, M. (2005). Immunity to error through misidentification and the bodily illusion experiment. *Journal of Consciousness Studies*, 12(7), 3-19.
- Nielsen, T.I. (1963). Volition: a new experimental approach. *Scandinavian Journal of Psychology*, 4, 225-230.
- Shoemaker, S. (1968). Self-reference and self-awareness. In *Self-reference and Self-awareness* (pp. 81-93).
- Wittgenstein, L. (1934/1960). *The Blue and brown books*. New York: Harper & Row.