Social Neuroscience
Publication details, including instructions for authors and subscription information:
http://www.informaworld.com/smpp/title~content=t741771143

Simulation trouble

Shaun Gallagher [a]
[a] University of Central Florida, Orlando, FL, USA

Online Publication Date: 01 January 2007
To cite this Article: Gallagher, Shaun (2007) 'Simulation trouble', Social Neuroscience, 2:3, 353 - 365
To link to this article: DOI: 10.1080/17470910601183549
URL: http://dx.doi.org/10.1080/17470910601183549

SOCIAL NEUROSCIENCE, 2007, 2 (3–4), 353–365

Psychology Press
Taylor & Francis Group

Simulation trouble

Shaun Gallagher

University of Central Florida, Orlando, FL, USA

I present arguments against both explicit and implicit versions of the simulation theory for intersubjective understanding. Logical, developmental, and phenomenological evidence counts against the concept of explicit simulation if this is to be understood as the pervasive or default way that we understand others. The concept of implicit (subpersonal) simulation, identified with neural resonance systems (mirror systems or shared representations), fails to be the kind of simulation required by simulation theory, because it fails to explain how neuronal processes meet constraints that involve instrumentality and pretense. Implicit simulation theory also fails to explain how I can attribute a mental or emotion state that is different from my own to another person. I also provide a brief indication of an alternative interpretation of neural resonance systems.

When it comes to explaining how we understand other people some of the very best contemporary philosophers, psychologists, and neuroscientists are simulationists. Rather than appealing to a theoretical use of folk psychology, they appeal to their own experience as a measure of others' experience. A short list of simulationists—for example, Vittorio Gallese, Alvin Goldman, Robert Gordon, Jane Heal, Susan Hurley, and Marc Jeannerod—however, already suggests some problems. Not everyone on this list understands simulation in the same way. In effect, there are different simulation theories, and although it is important to distinguish them, in this paper I will not be able to address each and every variation. Rather, I will present arguments against what I take to be the main elements of simulation theory. My primary target in this paper, however, is an implicit version of simulation theory based on the recent neuroscience of resonance systems (mirror neurons, shared representations).

Presenting arguments against simulation theory (ST) does not mean that I favor its main competitor, theory theory (TT). I do not think that the theory of mind (ToM) explains our primary and pervasive way of understanding others any more than the concept of simulation does. The alternative to both of these approaches is not something that I will directly argue for here (see Gallagher, 2001, 2004, 2006, for the alternative). But since part of the argument that I develop against simulation feeds into and depends on that alternative, a quick summary may be helpful.

The alternative approach is based on evidence from developmental psychology, neuroscience, and phenomenology, and is composed of three aspects. (1) Starting in infancy, we understand conspecifics based on a direct perception of their emotion-expressive movements, intentional actions, gestures, facial expressions, etc. (in developmental psychology this is referred to as ''primary intersubjectivity''; Trevarthen, 1979). ''Direct''

Correspondence should be addressed to: Shaun Gallagher, Philosophy and Cognitive Sciences, University of Central Florida, Orlando, Florida 32816, USA. E-mail: gallaghr@mail.ucf.edu

Part of the research for this paper was done at the Center for Subjectivity Research at the University of Copenhagen where I was Visiting Professor and supported by the university's Research Priority Initiative on Body and Mind.

An earlier version of this paper, entitled ''La perception d'autrui en action,'' was presented as a public lecture in the series *Fondements cognitifs de l'interaction avec autrui*. Collège de France (February 2006). My thanks to Alain Berthoz, Jean-Luc Petit, and Alvin Goldman, and to three anonymous reviewers for help in clarifying some of these issues.

www.psypress.com/socialneuroscience DOI:10.1080/17470910601183549

perception here means perception without some further cognitive or inferential step that goes beyond what is perceived, for example, in an attempt to grasp hidden mental states.[1] Thus, in a non-mentalizing way, I am able to see meaning, intention, and emotion in the actions of others, and in their gestures and facial expressions. This provides some limited understanding of others without recourse to the concept of mental states. (2) Around the age of 1 year we begin to develop capacities for shared attention and, on this basis, begin to understand others in terms of their involvements in pragmatic contexts (secondary intersubjectivity; Trevarthen, 1980; Trevarthan & Hubley, 1978). As developing agents we do not stand to the side as third-person observers, we engage in second-person interactions, and this is something we continue to do throughout our normal social life. Extending this idea, the pragmatic situations, social practices, and socially defined roles through which our interactions with others happen turn out to do much of the work of shaping our understanding of others. We understand what they are doing and what their intentions are because their actions are set within environments that help to define meaningful actions. Specific contexts make sense out of their actions, gestures, and expressions. In such cases, which are most cases, we do not try to get into the other person's mind, we try to get into their world, or more precisely, into a world that we already share with them. (3) These abilities for intersubjective perception and pragmatic interaction are carried forward and enhanced, starting around the age of 3 or 4 years, by a developing narrative competency that allows us to use proto-narratives (Lewin, 2005) or narrative practices to frame our everyday interpretations of others (Gallagher & Hutto, in press; Hutto, 2004). We begin to use narratives to put persons and contexts together in ways that generate the more subtle and sophisticated understandings of others that we have as adults. This kind of narrative practice is not based on ToM or folk psychology as understood in TT or ST, but may in fact be the basis for our use of folk psychological concepts (like beliefs, desires, reasons) in those rare cases when we encounter puzzling behavior and attempt to explain it (Hutto, in press). Most often, however, our narrative competency combined with our perceptual abilities and action contexts provides sufficient scaffolding for our understanding of others.

This alternative view, of which I have just given the most minimal account, suggests that these perceptual, emotional, and pragmatic capacities for interaction, together with narrative competency, are, at the personal or behavioral level, sufficient to explain pervasive everyday social understanding. As I have just suggested, this does not rule out the possibility that in rare cases we do take a theoretical stance, or that we sometimes use simulation routines to solve puzzling cases. But these are the rare cases. Ordinarily, in our everyday encounters in the pragmatic and social contexts that characterize our lives, we gain a perceptual grasp of another's contextualized actions, gestures, and expressions, and we understand their speech acts as meaningful and intentional, without looking beyond such meanings to their mental states. And if, instead of directly interacting with others, we are called upon (or we call upon ourselves) to think more deeply about them, our tendency is not to call up a theory or a simulation, but to call up a narrative framework to facilitate our understanding of them.[2]

ARGUMENTS AGAINST EXPLICIT ST

Simulation theory (ST), as an approach to ToM, has been developed in several versions. For the purposes of this paper I will distinguish between explicit and implicit versions of ST. I acknowledge that there are hybrid versions (e.g., Gallese & Goldman, 1998; Goldman, 2006; Jeannerod & Pacherie, 2004) and ''radical'' versions (e.g., Gordon, 2004), and although an argument against one version will not necessarily work against all of them, I think that much of what I argue against explicit and implicit versions also applies to the other models.

[1] This concept is developed in phenomenologists like Scheler (1923/1954), who calls it ''primary perception'' (e.g., p. 10) and Merleau-Ponty (1964).

[2] Narratives that frame our social interaction are different from the kind of theories to which TT appeals. The latter involve a reference to unobservable entities, that is, mental states which are theoretically postulated entities comparable to the black holes of astrophysics. If Bruner (1986) is right, narrative is a particular mode of thinking that relates to the concrete and particular; it takes context to be primary in the determination of meaning. In contrast, the theoretical is concerned with the abstract and general, and, in this sense, abstracts away from the particular context in favor of propositional attitudes. See Gallagher and Zahavi, Chapter 9 (in press) for further discussion.

According to explicit versions of ST, simulation is a conscious or introspective process in which I imagine myself in the other's situation and use the model (the simulation) that is generated to predict the other's mental states. Goldman has been a good representative of this position.[3] He argues that simulation is explicit insofar as it involves a conscious introspective use of the imagination to conceptually manipulate propositional attitudes (beliefs, desires). ''When a mindreader tries to predict or retrodict someone else's mental state by simulation, she uses pretense or imagination to put herself in the target's 'shoes' and generate the target state'' (Goldman, 2005a). This process involves three steps:

1. First, the attributor creates in herself pretend states intended to match those of the target. In other words, the attributor attempts to put herself in the target's ''mental shoes.''
2. The second step is to feed these initial pretend states, e.g., beliefs, into some mechanism of the attributor's own psychology . . . and allow that mechanism to operate on the pretend states so as to generate one or more new states, e.g., decisions.
3. Third, the attributor assigns the output state to the target . . . e.g., we infer or project the decision to the other's mind (Goldman, 2005b, pp. 80–81).

The first step—''the attributor creates in herself pretend states intended to match those of the target''—seems troublesome in that the simulator apparently already has some idea of what's going on with the other person. Where does that knowledge come from and why isn't that already the very thing we are trying to explain? Some theorists who combine TT and ST suggest that folk psychology provides some general rules about how people think and behave in certain situations, and that this is what the simulationist can use to generate the pretend mental states needed for the simulation process (e.g., Currie & Ravenscroft, 2003). Alternatively, however, one might appeal to subpersonal mirror resonance processes (discussed below), although one then faces the problem of how to translate these processes into a conceptual grasp of propositional attitudes. It is clearly not enough to suggest that a belief feels different from a desire because it is generated by different subpersonal processes, which are themselves generated by differential activations induced by our perception of the other (cf. Goldman, 2002, pp. 11–12). One needs to say something about the particular content of the belief used in the simulation, and not just that it is a belief rather than a desire.

There are both logical and phenomenological arguments to be made against this explicit version of ST. Some of these arguments have been around for a long time and have been directed against earlier versions of ST, developed under the heading of ''inference from analogy.'' Gilbert Ryle, for example, argued that the logic of simulation isn't correct because the idea of imputing to a variety of others what is true of my simulated action ignores the diversity of their actions. ''[T]he observed appearances and actions of people differ very markedly, so the imputation to them of inner processes closely matching [one's own or] one another would be actually contrary to the evidence'' (Ryle, 1949, p. 54). A similar objection was raised by Max Scheler (1923/1954). If I project the results of my own simulation onto the other, I understand only myself in that other's situation, but I don't understand the other. In other words, given the large diversity of motives, beliefs, desires, and behaviors in the world, it is not clear how a simulation process based on my own relatively narrow experience (or relatively unique circumstance) can give me a reliable sense of what is going on in the other person's mind, or in their behavior.

Another objection is based on developmental considerations. Infants are capable of understanding the intentions of others. At 18 months of age they are capable of recognizing and completing another person's failed intention (Meltzoff, 1995). Even earlier, infants seem capable of parsing intentions (see Baird & Baldwin, 2001; Baldwin

[3] I'm using Goldman as a representative of explicit simulation, and this was certainly his view in his earlier papers (e.g., 1989, 1992). In his most recent work, Goldman (2006) defends a hybrid version of ST. He distinguishes between high-level and low-level mind reading. I take this distinction to be compatible with the distinction between explicit and implicit versions, respectively. Low-level simulation is ''simple, primitive, automatic, and largely below the level of consciousness'' (p. 113). High-level simulation includes one or more of the following features: ''(a) it targets mental states of a relatively complex nature, such as propositional attitudes; (b) some components of the mindreading process are subject to voluntary control; and (c) the process has some degree of accessibility to consciousness . . . . The low-level prototype is the mirroring type of simulation process. The high-level prototype is the kind that uses pretense or E[nactment]-imagination'' (p. 147).

& Baird, 2001). Whether some animals are capable of understanding intentions is still a matter of debate. In any case, if infants are capable of understanding the intentions of others, then, as Scheler also suggests, it is nonetheless unlikely that they are capable of the complex cognitive processes needed for explicit simulation.

A third argument that can be made against explicit ST is what I call the *simple* phenomenological argument. According to Goldman and other defenders of the explicit version of ST, simulation is not only explicit but also pervasive. That is, we use it all the time, or at least it is the default way of understanding others. Goldman (2002, pp. 7–8) calls this a moderate claim:

> The strongest form of ST would say that all cases of (third-person) mentalization employ simulation. A moderate version would say, for example, that simulation is the *default* method of mentalization . . . I am attracted to the moderate version . . . . Simulation is the primitive, root form of interpersonal mentalization.

If simulation is both explicit and pervasive, then one should have some awareness of the different steps that one goes through as one consciously simulates the other's mental states. But there is no phenomenological evidence for this; there is no experiential evidence that I use such conscious (imaginative, introspective) simulation routines when I interact with or come to understand another person. That is, when we consult our own common experience of how we understand others, we just don't find such processes, and that is puzzling if they are supposed to be explicit and pervasive. This is not to say that we never use simulations, but that itself is telling. If we are confronted with some strange or unaccountable behavior, I may try to understand the other person by running a simulation routine. But I can easily become aware that I am in fact taking this approach, and it is all the more apparent when I do this, simply because it tends to be the exception, and it tends to stand out in its rarity. This clearly tells against the idea that I employ simulation in the usual everyday circumstance, and it goes against any claim that in fact we use simulation routines all the time, but phenomenology constantly and consistently misses this fact. What phenomenology does show is that most of our encounters are not third-person puzzles that require explanation of the other person, which in turn is generated by first-person simulations. Rather, most of our encounters are second-person interactions in which I easily have a sense of what is going on with the other person based on our common pragmatic or socially contextualized interactions, with no cognitive simulation required.

The explicit simulation theorist might argue that we use simulations so much that they become habitual, and for that reason we should not expect them to stand out in our experience. Running simulation routines is like driving a car. Just as we are not explicitly aware of all of our driving habits, in the same way we are not aware that we are simulating. According to the simple phenomenological objection, however, if such processes stay at the personal level, they would remain accessible to conscious reflection, or at least they would become apparent, as unworkable habits, in problematic situations when our habitual strategies break down. Not only does this not happen, but it turns out that in problematic situations when our habitual strategies break down we often find that we employ the very sort of simulation process described by explicit ST, which again suggests that simulation is the exception rather than the rule.

IMPLICIT SIMULATION

ST can easily counter the simple phenomenological argument by moving to the implicit version of simulation. Implicit ST conceives of simulation as a subpersonal process, and this approach has gained more ground in recent years by appealing to the neuroscientific evidence involving subpersonal activation of mirror neurons, shared representations, or more generally, resonance systems (e.g., Fadiga, Fogassi, Pavesi, & Rizzolatti, 1995; Rizzolatti, Fogassi, & Gallese, 2000; Rizzolatti, Fadiga, Gallese, & Fogassi, 1996). On the one hand, implicit versions of ST do not succumb to the simple phenomenological objection since if simulation is subpersonal, it is not something of which we would be aware. On the other hand, the strictly implicit version of ST is actually an argument against the explicit version of ST. That is, if our understanding of others is in fact mediated by an implicit and automatic simulation process, then we have little need for the more explicit version. Indeed, to the extent that implicit ST explains the phenomenological scarcity of explicit simulation, it would support the simple phenomenological argument against explicit

simulation as the default model of social understanding. Along this line Gallese (2005, p. 102), who I will take as the clearest representative of implicit ST, states: ''Whenever we face situations in which exposure to others' behavior require a response by us, be it active or simply attentive, we seldom engage ourselves in an explicit, deliberate interpretive act. Our understanding of a situation most of the time is immediate, automatic, and almost reflex like.'' This, of course, doesn't rule out hybrid models that combine implicit and explicit elements (see Goldman, 2006).

Implicit ST appeals to what is now familiar social neuroscience. In broad terms, one's motor system reverberates or resonates in one's encounters with others. My motor system is activated when I perceive another person performing an intentional action, for example. Mirror neurons in the pre-motor cortex and in Broca's area of the human brain are activated both when the subject engages in specific instrumental actions and when the subject observes someone else engage in those actions (Rizzolatti et al., 1996, 2000). Also, specific overlapping neural areas (shared representations), in parts of the frontal and parietal cortexes, are activated when I engage in intentional actions, when I observe some other person engage in that action, and when I imagine myself or another person engaged in that action (e.g., Grèzes & Decety, 2001). These subpersonal mechanisms are said to constitute a simulation of the other's intentions (Gallese, 2001; Gallese & Goldman, 1998). Thus Gallese contends that ''when we observe actions performed by other individuals our motor system 'resonates' along with that of the observed agent ... action understanding heavily relies on a neural mechanism that matches [simulates], in the same neuronal substrate, the observed behaviour with the one [the observer could execute] ...'' (2001, pp. 38–39).

The neuroscientific evidence for these processes is very strong, and it is not my intention here to deny the neuroscience. But is it appropriate to characterize these processes as simulations, as Gallese goes on to do? The implicit simulation hypothesis is just this: understanding others is achieved by simulating the other's action ''with the help of a motor equivalence between what the others do and what the observer does'' (2001, p. 39). This is a subpersonal process generated by ''automatic, implicit, and nonreflexive simulation mechanisms ...'' (Gallese, 2005, p. 117). He refers to his model as the ''shared manifold hypothesis'' and distinguishes between three levels (Gallese, 2001, p. 45):

1. The *phenomenological level* is the one responsible for the sense of similarity ... that we experience anytime we confront ourselves with other human beings. It could be defined also as the *empathic* level.
2. The *functional level* can be characterized in terms of simulation routines, *as if* processes enabling models of others to be created.
3. The *subpersonal level* is instantiated as the result of the activity of a series of mirror matching neural circuits.

Recall the Ryle–Scheler objection against explicit ST, that since it is based on a projection from experience that is narrowly my own, no inference about the experience of the other person (which may be greatly diversified from mine) is justified. Does a subpersonal simulation lock us up within our own first-person system? Defenders of the implicit version of ST have a good, albeit partial, response to this. There is growing consensus that mirror neurons (and shared representations) are *neutral*—neither first nor third person. They are activated both for my own action and for observation of the other's action: thus, activation of the system simulates the intentional action but not the agent (deVignemont, 2004; Gallese, 2005; Hurley, 2005; Jeannerod & Pacherie, 2004).

In this case, however, the complete subpersonal simulation process, like its explicit cousin, involves a multi-step process. First, we perceive the other's behavior; this is followed immediately by activation of shared representations—in neutral mode registering what Jeannerod and Pacherie (2004) call ''naked intentions'', and this is followed by a determination of agency (i.e., a specification of who did the action—me or the other person. This final step is accomplished by what Georgieff and Jeannerod (1998) have called the ''Who'' system, which distinguishes self-action from other-action and which may involve activation in the inferior parietal lobule (Blakemore & Frith, 2003; Decety, 2005; Jackson & Decety, 2004).

It will pay to stop and consider Jeannerod and Pacherie's claims about naked intentions. They assume that an articulation at the level of the neural activations between those activations responsible for (1) registering in the perceiving system the ''naked'' intention in an action, and (2) registering the agent for the action, means

that there is an articulation at the level of experience between the perception of intention and the experience of agency. ''We can be aware of an intention, without by the same token being aware of whose intention it is . . . something more than the sole awareness of a naked intention is needed to determine its author'' (2004, p. 140). The question is this: if in fact the brain can process information about intentions without assigning the intentions to a particular agent, is it legitimate to say that our experience is similarly articulated? Jeannerod and Pacherie suggest that it is:

> When the naked intention one is aware of yields an overt action, the extra information needed to establish authorship may be found in the outside world. The question ''Is this intention mine?'' would then be answered by answering the question: ''Is this my body performing the corresponding action?'' (p. 140).

Phenomenologically (experientially), however, intentions in almost all cases come already fully clothed in agent specification. The ''who'' question does not come up at the level of experience, because the neural systems have already decided the issue. The wonderful thing about the ''Who system'' is that it is completely neurological and subpersonal—and the results of its activation are hardly ever experientially manifested as ''making a decision about who did the action.'' Rather, the results of its activation are experientially manifested as ''X's action'' where X is either you or me. Indeed, our direct perception is highly reliable with regard to discriminating between self and non-self. Pathologies and oddly arranged experiments may reveal ''who'' problems, but in normal ecological behavior it is generally clear whose intention/action it is. Philosophers like Wittgenstein (1958), Shoemaker (1968), and Evans (1982), provide good reasons why the identification question—''Someone is intending to pick up the apple, is it me?''—just doesn't come up.

On this score, I want to suggest that there is no necessary isomorphism between the phenomenological, the functional, and the neuronal levels. So, if the neuronal processes can be defined as involving a step-wise process, this does not mean that a step-wise process necessarily shows up at the level of experience, and *vice versa*. If we can distinguish a step-wise process in experience, that does not mean that we will find exactly the same step-wise structure in the underpinning functional or neuronal processes. This is tied into the concept of multiple realizability. For the purposes of this paper, I will call any claim to the contrary the fallacy of supposed isomorphism.[4]

PROBLEMS WITH IMPLICIT SIMULATION

This brings us to the first set of questions about implicit ST. Specifically I want to argue that an alternative interpretation of neural resonance is possible. Implicit ST interprets neural resonance as simulation, but it could easily be interpreted as part of the neuronal processes that underlie intersubjective *perception* rather than simulation. That is, the articulated neuronal processes that include activation of mirror neurons or shared representations may underpin a non-articulated direct perception of the other person's intentional actions, rather than a distinct mental process of simulating their intentions. This claim requires that we conceive of perception as a temporal phenomenon, and as an enactive sensory-motor phenomenon. Let us examine this possibility.

First, mirror neurons fire 30–100 ms after appropriate visual stimulation (Gallese, private correspondence). This short amount of time between activation of the visual cortex and activation of the pre-motor cortex, raises the question of where precisely to draw the line between perceptual processes and something that would count as a subpersonal simulation. Even if it is possible to draw a line between activation of the visual cortex and activation of the pre-motor cortex, this does not mean that this line distinguishes between perception and simulation as a step-wise process. The simulation interpretation of these processes seems to read into the neuronal level the step-wise processes of simulation, first identified as explicitly conscious simulation processes.

Let us be careful and clear about this. Gallese and the implicit simulationists are not claiming that the step-wise *neuronal* processes (sensory activation of visual cortex followed by mirror system activation) generate a step-wise *conscious* process of perception plus simulation. In this regard they are not committing the fallacy of

[4] I do not mean to rule out all isomorphisms. This only goes to the claim that there is a necessary isomorphism between neuronal, functional, and phenomenological levels.

supposed isomorphism. Gallese contends that the simulation stays implicit, that the mirror system activation itself can be read, functionally, as part of a step-wise simulation process. Nonetheless, I think two issues can be raised in this regard.

The first issue is this: deciding that mirror neurons function as simulations depends on taking a step-wise model that was developed at the explicit, conscious, or personal level, and looking for that step-wise model at the neuronal level. On various versions of ST, simulation involves a step-wise process that begins with perception and ends with some form of inferential understanding. We first see an action that we need to understand; we then simulate it in our own mind or motor system; we then attribute agency for the action, and infer something about the other's experience. But even if neuronal processes that involve information flow from sensory cortex to pre-motor cortex take some time (as much as 100 ms), it is not clear that we should identify these step-wise processes as perception plus simulation, rather than a temporally extended and enactive perceptual process. That is, at least in terms of temporal parameters, the fact that at the neurological level, S (sensory processing) is followed by M (activation of mirror neurons) does not mean that one should think of this as perception followed by simulation.

If we think of perception as an enactive process (e.g., Hurley, 1998; Noë, 2004)—as involving sensory-motor skills, rather than as just sensory input/processing; as an active, skillful, embodied engagement with the world rather than the passive reception of information from the environment—then it may be more appropriate to think of the resonance processes as part of the structure of the perceptual process when perception is of the action of conspecifics. Fogassi and Gallese, despite their simulationist interpretation, put this point clearly: ''perception, far from being just the final outcome of sensory integration, is the result of sensorimotor coupling'' (2002, p. 27). Mirror activation, on this interpretation, is not the initiation of simulation, it is part of a direct intersubjective perception of what the other is doing.

At the phenomenological level, when I see the other's action or gesture, I see (I *directly perceive*) the meaning in the action or gesture. I see the joy or I see the anger, or I see the intention in the face or in the posture or in the gesture or action of the other. I see it. I do not have to simulate it. And I immediately see that it is *their* action, gesture, emotion, or intention, and it is extremely rare that I would be in a position to confuse it with my own. Although Jeannerod and Pacherie (2004) defend a version of ST, they nicely express the phenomenological alternative: ''Perception and action are closely integrated and when we visually perceive actions, we seem to be immediately sensitive to the distinctive properties of intentional behavior'' (p. 139). This may be due to the underlying articulated neuronal processes, but it is not clear at all why we should think of these processes on the simulationist model.

Of course the simulationist can accept the phenomenology (''Yes indeed, that is what *seems* to happen'') and still hold to the interpretation that the specific subpersonal processes involve simulation. But what precisely justifies this interpretation? After all, what happens on the neurological level is simply a complex sequence of neuronal activations. If we look at those processes from a functionalist perspective already framed by ST, then we tend to read those processes as involving simulation. If, in contrast, we look at those processes from a phenomenological level that suggests a direct perception of the other's intentions, then we tend to read those processes as perceptual without simulation. Can the simulationist offer any convincing evidence that the activation of resonance processes is in fact a simulation?

This brings us to the second issue, which I think counts as an argument against the implicit version of ST. What theorists of implicit simulation (Gallese, Jeannerod, Pacherie), and even critics of implicit ST, like Saxe (2005), call ''simulation'', is not simulation in any genuine sense of the word. Consider, first, two definitions of ''simulation'' offered by the Oxford English Dictionary. (1) Simulation is an imitation, in the sense of something *not real*—counterfeit; to simulate means to feign, to pretend. I will call this the pretense definition. (2) Simulation in the sense of a simulator: a model (a thing) that we can *use* or do things with so we can understand the real thing. I will call this the instrumental definition. Both senses of the term appear to be ubiquitous in the literature of both explicit and implicit ST.[5] Consider the following characterizations (italics are mine).

[5] I think it is clear that, given the history of ST, the definition of simulation relevant to ST is first worked out in accounts of explicit ST, and then is uncritically used in accounts of implicit ST.

Instrumental def. Simulation means "*using* one's own evaluation and reasoning mechanisms as a model for theirs . . . " (Dokic & Proust, 2002, p. viii).

Pretense def+instrumental def. Simulation involves "pretend states" where, "by pretend state I mean some sort of surrogate state, which is *deliberately adopted* for the sake of the attributor's task . . . In simulating practical reasoning, the attributor *feeds* pretend desires and beliefs into her own practical reasoning system" (Goldman, 2002, p. 7).

The surrogation or pretense, however, is of a precise kind. Bernier (2002) makes this explicit as an essential element found in ST.

Instrumental def+pretense def. "According to ST, a simulator who runs a simulation of a target would *use* the resources of her own decision-making mechanism, in an 'off-line' mode, and then the mechanism would be fed with the mental states she would have *if she was in the target's situation"* (Bernier, 2002, p. 34)

For ST, a simulation is not simply a model that we use to understand the other person—theoretical models would suffice if this were all that was required. Even the fact that the model is constituted in our own mechanisms is not sufficient. Rather, I must use the model "as if" I were in the other person's situation. As Gallese puts it:

Pretense def. "[O]ur motor system becomes active *as if* we were executing that very same action that we are observing" (2001, p. 37).

Gordon (2005, p. 96) locates this pretense right at the neuronal level:

Pretense def. The neurons that respond when I see your intentional action, respond "*as if* I were carrying out the behavior . . .".

The concept of simulation, as defined by ST, then, clearly needs to meet these two conditions: it is a process that I control in an instrumental way (in the explicit version it is "deliberately adopted"), and it involves pretense (I put myself "as if" in the other person's shoes).

These characterizations are, as I've suggested, ubiquitous in ST.

Instrumental and pretense defs combined. Mental simulation is a cognitive model, "*ability* or *heuristic* or *methodology"* (Jacob, 2002, who cites Gordon for the latter term)—by which I "*engage* in pretense," *put myself* in someone else's shoes, *compare* my experience to their experience, and *predict* their mental state, emotion, or behavior. We *use* ourselves as a model . . . I *create* in myself some pretend beliefs . . . and so forth.

This is the way simulation is characterized not only by theorists of explicit simulation, but also by theorists of implicit simulation. The pretense condition mentioned by Gallese is accomplished in a simulation considered to be "an interactive *model* of what cannot be known in itself" (2003). At the subpersonal level, the brain in a stepwise fashion is *modeling* the intentional action of others. Gordon (2004), p. 1) suggests that on the "cognitive-scientific" model, "one's own behavior control system is employed as a *manipulable model* of other such systems. (This is not to say that the 'person' who is simulating is the model, rather, only that *one's brain can be manipulated to model other persons*)."

Thus, according to ST, simulation involves the instrumental use of a first-person model to form third-person "as if" or "pretend" mental states. For subpersonal processes, however, both of these characterizations fail. (1) If simulation is characterized as a process that I (or my brain) instrumentally uses or controls, if this is what simulation is, then it seems clear that what is happening in the implicit processes of motor resonance is not simulation. We, at the personal level, do not *do* anything with the activated brain areas—in fact, we have no instrumental access to neuronal activation, and we can not use it as a model. Nor does it make sense to say that at the subpersonal level the brain itself is *using* a model or methodology, or *initiating* pretend states, or that one set of neurons makes use of another set of neurons as a model in order to generate an understanding of something else. In precisely the intersubjective circumstances that we are considering, these neuronal systems do not take the initiative; they do not activate themselves, but are activated by the other person's action. The perception of the other person's action automatically activates in our brain the same areas that are activated when we engage in similar action. The other person *has an effect on us.* The other

elicits this activation. This is not a simulation, but a perceptual elicitation. It is not us (or our brain) *doing* it, but the other who does this to us.[6]

Furthermore, (2) in subpersonal processes there is no pretense, and this is the case whether we consider neuronal processes as vehicles or in terms of the content that they might represent. As vehicles, neurons either fire or do not fire. They do not pretend to fire. More to the point, however, what these neurons represent or register cannot be pretense in the way required for ST. As we saw, the mirror system is neutral with respect to the agent; there is no first- or third-person specification involved. In that case, they do not register *my* intentions as pretending to be *your* intentions; there is no "as if"—there is no neuronal subjunctive—because there is no "I" or "you" represented.

In response to this criticism it might be argued that no one is claiming that there is pretense in the firing of neurons, but rather that the neuronal processes underpin a personal-level experience of the other that qualifies as a pretense state.[7] But if, as we saw, implicit simulation theorists like Gallese contend that the simulation stays implicit, and that the mirror system activation itself can be read, functionally, as a simulation process, then on the ST definition of simulation, those subpersonal processes themselves must involve pretense. Alternatively, if the claim is that certain neuronal processes generate pretense states at the personal level, then the simulation becomes explicit and other objections hold.

Goldman and Sripada (2005) acknowledge a discrepancy between the ST definition of simulation and the working of subpersonal mirror processes. "Does [Gallese's] model really fit the pattern of ST? Since the model posits unmediated resonance, it does not fit the usual examples of simulation in which pretend states are created and then operated upon by the attributor's own cognitive equipment (e.g., a decision-making mechanism), yielding an output that gets attributed to the target . . .". To address this discrepancy they propose a minimal definition of simulation:

> However, we do not regard the creation of pretend states, or the deployment of cognitive equipment to process such states, as essential to the generic idea of simulation. The general idea of simulation is that the simulating process should be similar, in relevant respects, to the simulated process. Applied to mind reading, a minimally necessary condition is that the state ascribed to the target is ascribed as a result of the attributor's instantiating, undergoing, or experiencing, that very state. In the case of successful simulation, the experienced state matches that of the target. This minimal condition for simulation is satisfied [in Gallese's model] (Goldman & Sripada, 2005, p. 208).

If this is a necessary condition, it is clearly not a sufficient one, as Goldman (2006, p. 131 ff.) realizes, because on this minimal definition and without something further, it is not clear what would motivate me to ascribe the state that I was undergoing to someone else. Activation of the mirror system does not, on its own, necessarily involve mind reading or attribution to another agent. What one would have to add in order to get to the point of making this a simulation of the other's action, intention, or mental state are precisely the pretense and instrumental aspects that Goldman and Sripada are discounting.[8]

Furthermore, and importantly for implicit ST, if simulation were as automatic as mirror neurons firing, then it would seem that we would not be able to attribute a state different from our own to someone else. But we do this all the time. I can understand that the woman in front of me is enthusiastically and gleefully reaching to pick up a snake at the same time that I am experiencing revulsion and disgust about that very possibility. It seems straightforward to say that I see what she is doing and that I see she is doing it with enthusiasm and glee, but that my own feelings are quite

[6] It may seem contradictory to claim in the previous paragraphs that perception is enactive, or as Noë says, "perception is action," and in this argument to claim that the activation of the resonance system is the result of a passive elicitation, so that the motor aspect of perception does not involve our action, but is a case of us being affected by the other. I think that a fuller account of enactive perception has to be able to accommodate this passive, affective aspect of perception (see Gallagher, 2005).

[7] My thanks to an anonymous referee for this objection.

[8] This motivates Goldman's move to a hybrid model. Implicit mirroring may be considered a minimal condition for simulation, but for the full simulation that constitutes mind reading (the attribution of mental states to another) one requires higher-level processes of "classification" (which he characterizes as introspective, although not necessarily conscious introspection; 2006, p. 245 ff.) and "projection." Although pretense is not required for low-level mirroring, it is one possible way to realize high-level simulation and mind reading (p. 49).

different. Neither my neural states nor my feelings/cognitions match hers, and it thus seems that I can understand her actions and emotions (which are completely different from mine) without even meeting the minimal necessary condition for simulation, that is, matching my state to hers.

One might want to add to this minimally necessary condition for simulation the further stipulation that the simulating process must be ''concretely similar'' to the simulated process, where ''concretely similar'' means that there is a fine-grained (though not necessarily one-to-one) isomorphism for each step involved in the respective processes (Fisher, 2006). But this makes the problem worse if we are conceiving of the simulation as happening on a neuronal level. If my neural/mental/emotional state has to match the other person's neural/mental/emotional state, it is difficult to understand how we could attribute to someone else a mental or emotional experience that is different from our own.

In regard to these minimal definitions of simulation, also consider the difficulties involved if we were interacting with more than one other person, or trying to understand others who are interacting with each other. Is it possible to simulate the neural/mental/emotional states of two other people at the same time, or quickly enough, if in fact our simulations must be such that we instantiate, undergo, or experience, that very state, and/or that such simulations must be concretely similar? There would be an impossible amount of cognitive work or subpersonal matching required to predict or to understand the interactions of several people if the task involves simulating their mental states, especially if in such interpersonal interactions the actions and intentions of each person are affected by the actions and intentions of the others (see Morton, 1996, for a similar point).[9]

The woman reaching for the snake is a good example with which to make one more clarification. The term ''simulation'' is sometimes used by neuroscientists to refer to certain motor control processes for action planning. Efference copy sent through forward control mechanisms, for example, is said to constitute a simulation of an intended movement in order to compare it with an ongoing movement to predict its success. The brain runs this simulation to make fast nonconscious corrections to keep the action on track. This use of the term is entirely independent of ST, and the objections I am raising here do not apply to it.[10] In the context of motor control the term ''emulation'' is sometimes substituted for ''simulation'', and I'll follow that practice here.

In the case of seeing someone reach to pick up a snake, it seems possible that my understanding of the other's actions and their enthusiasm for those actions is initiated purely in enactive perceptual processes that would include (nonsimulative) activation of shared representations. Those same shared representations may further generate, or may help to constitute a motoric emulation of my own action of picking up a snake, which leads to my sense of revulsion and disgust at the very possibility. This emulation, however, is not an implicit simulation that contributes to my understanding of the other person; it is rather something that helps to constitute my own feel-

[9] The requirement that the simulation has to be concretely similar also raises problems for the instrumental and pretense conditions even for the explicit version of ST. If our simulation has to be concretely similar to the simulated state for it to be considered a simulation, assuming explicit instrumental control of our simulation process, how will we know how to run or control our simulation unless we already know in some detail what the other's state is like. And how do we come by that knowledge? If the answer is through simulation, then we have an infinite regress. In regard to explicit pretense, Fisher (2006), who models simulation as a reasoning process, rejects this aspect as inconsistent with simulation being concretely similar, for if we simulate a reasoning process, we are really reasoning, and not just pretending to.

[10] Some theorists, however, have appealed to these motor emulation processes as possible mechanisms involved in the simulation of another's action (e.g., Gallese, 2001; Hurley, 2005, pp. 181–188; see Iacoboni, cited in Millikan, 2005, p. 188, note 2). Our own motor system comparators are activated to simulate and thereby anticipate the other's action. The brain could be said to predict the other person's actions in this way. On this account the perception of the other's action is automatically informed by a subpersonal simulation; perception of action involves a loop through the motor control comparator. Can ST adopt this model of simulation? The problem, again, is that the pretense condition is not met; there is no ''*as if* it were I'' involved, and in that regard it fails to be the kind of simulation required by ST. If indeed the subpersonal emulation is neutral in regard to whose action is at stake, then it can be only a representation of an intentional action in my motor system, but not a representation of my own motor action *as if it were* the other's. It is not at all clear that, as Gordon (2005, p. 96) suggests, the neurons respond ''*as if* I were carrying out the [other's] behavior'' in any sense in which the ''as if'' registers subpersonally. Even assuming a ''Who system,'' a specification in my motor system that the action belongs to another is not equivalent to the specification ''*as if* I were carrying out the action.'' If this is a simulation of intentional action, it is, nonetheless, not the kind of simulation that ST needs; it may be nothing more than motor priming or emulation, or what Hurley calls mirroring (2005, p. 184).

ings about this action, which turn out to be very different from the ones I understand to belong to the other person. Just this contrasting situation, between her enthusiastic action and my repulsion, may motivate a more explicit hermeneutical process. That is, depending on circumstances, I may explicitly frame her action in a narrative that makes sense of it (e.g., this woman is, after all, the curator of the zoo's snake house, she is trained to handle snakes, and does it all the time), or if there is some puzzling element involved (e.g., this woman is my wife who I know to be afraid of snakes), I may even engage in one of those rare instances of explicit simulation (e.g., assuming I am not in a position to ask her what she is doing, I might try to imagine what she is thinking—perhaps that this is a toy snake, etc.).

CONCLUSION

It is not clear why we should think of the activation of resonance systems as a simulation process of the sort required by ST. This is not to deny that resonance processes are at work in our perception of the other person. Moreover, the nature of the resonance processes involved in such encounters makes our perception of other conspecifics different from our perception of objects and instruments. But it does not make our perception and understanding of others the result of an implicit simulation. In effect, simulation is a personal-level concept that cannot be legitimately applied to subpersonal processes. As a personal-level concept it signifies a relatively rare and not at all pervasive or default way of understanding others.

I have suggested that implicit ST is only one, and not necessarily the best, interpretation of the significance of motor resonance systems. I have argued that implicit resonance processes are not simulations in any sense that is useful for ST. Furthermore, if implicit ST were a good account of our primary and pervasive ability to understand others, it would count as an argument against explicit ST, since explicit simulations would be redundant in this case. Likewise, however, if our default mode of understanding others were based on explicit simulation, then the claims of implicit ST about the adequacy of motor resonance processes would be wrong. Goldman's view of implicit motor resonance processes is that they do not constitute simulations of a sort that would be sufficient to do the full job, but do generate some background information that is useful to initiate the explicit simulation process. I have argued, in contrast to both explicit and implicit ST, that implicit motor resonance processes are important enactive processes that contribute to the constitution of the perceptual access that we have to the intentions of others.

I do not claim, however, that we get a full account of human intersubjectivity in the idea that we have *perceptual* access to the intentions of others. Perceptual access to the other person's contextualized bodily movements, gestures, facial expressions, and so forth does give us a partial sense of what is going on with them, what they mean and what they feel. This, together with our interactions with others in pragmatic and social contexts, where those contexts and situations enrich our understanding even further, gives us a relatively reliable, but still relatively elemental understanding of them. There is much more to say about the role of language and narrative competency in a fuller account of intersubjectivity (see Gallagher, 2006; Gallagher & Hutto, in press; Hutto, 2003, 2004, in press; Lewin, 2005). Even in that larger story, however, the ToM approaches that emphasize either simulation, or the role of folk psychology as background theory, have a minimal role to play in our normal and everyday interactions.

REFERENCES

Baird, J. A., & Baldwin, D. A. (2001). Making sense of human behavior: Action parsing and intentional inference. In B. F. Malle, L. J. Moses, & D. A. Baldwin (Eds.), *Intentions and intentionality: Foundations of social cognition* (pp. 193–206). Cambridge, MA: MIT Press.

Baldwin, D. A., & Baird, J. A. (2001). Discerning intentions in dynamic human action. *Trends in Cognitive Science*, *5*(4), 171–178.

Bernier, P. (2002). From simulation to theory. In J. Dokic & J. Proust (Eds.), *Simulation and knowledge of action* (pp. 33–48). Amsterdam: John Benjamins.

Blakemore, S.-J., & Frith, C. D. (2003). Self-awareness and action. *Current Opinion in Neurobiology*, *13*, 219–224.

Bruner, J. (1986). *Actual minds, possible worlds.* Cambridge, MA: Harvard University Press.

Currie, G., & Ravenscroft, I. (2003). *Recreative minds.* Oxford, UK: Oxford University Press.

Decety, J. (2005). Perspective taking as the royal avenue to empathy. In B. F. Malle & S. D. Hodges (Eds.), *Other minds: How humans bridge the divide between self and other* (pp. 135–149). New York: Guilford Press.

deVignemont, F. (2004). The co-consciousness hypothesis. *Phenomenology and the Cognitive Sciences*, *3*(1), 97–114.

Dokic, J., & Proust, J. (2002). Introduction. In J. Dokic & J. Proust (Eds.), *Simulation and knowledge of action*. (pp. vii–xxi). Amsterdam: John Benjamins.

Evans, G. (1982). *The varieties of reference*. Oxford, UK: Oxford University Press.

Fadiga, L., Fogassi, L., Pavesi, G., & Rizzolatti, G. (1995). Motor facilitation during action observation: A magnetic stimulation study. *Journal of Neurophysiology*, *73*, 2608–2611.

Fisher, J. C. (2006). Does simulation theory really involve simulation? *Philosophical Psychology*, *19*(4), 417–432.

Fogassi, L., & Gallese, V. (2002). The neural correlates of action understanding in non-human primates. In M. I. Stamenov & V. Gallese (Eds.), *Mirror neurons and the evolution of brain and language* (pp. 13–35). Amsterdam: John Benjamins.

Gallagher, S. (2001). The practice of mind: Theory, simulation, or interaction? *Journal of Consciousness Studies*, *8*(5–7), 83–107.

Gallagher, S. (2004). Understanding interpersonal problems in autism: Interaction theory as an alternative to theory of mind. *Philosophy, Psychiatry, and Psychology*, *11*(3), 199–217.

Gallagher, S. (2005). A new movement in perception: Review of Alva Noë's *Action in Perception* (book review). *Times Literary Supplement,* 9 September 2005.

Gallagher, S. (2006). The narrative alternative to theory of mind. In R. Menary (Ed.), *Radical enactivism: Intentionality, phenomenology, and narrative* (pp. 223–229) Amsterdam: John Benjamins.

Gallagher, S., & Hutto, D. (in press). Primary interaction and narrative practice. In Zlatev, J. Racine, T. Sinha, C. & Itkonen, E. (Eds.), *The shared mind: Perspectives on intersubjectivity.* Amsterdam: John Benjamins.

Gallagher, S., & Zahavi, D. (in press). *The phenomenological mind*. London: Routledge.

Gallese, V. (2001). The ''shared manifold'' hypothesis: From mirror neurons to empathy. *Journal of Consciousness Studies*, *8*, 33–50.

Gallese, V. (2003). The roots of empathy: The shared manifold hypothesis and the neural basis of intersubjectivity. *Psychopathology*, *36*, 171–180.

Gallese, V. (2005). ''Being like me'': Self–other identity, mirror neurons and empathy. In S. Hurley & N. Chater (Eds.), *Perspectives on imitation I* (pp. 101–118). Cambridge, MA: MIT Press.

Gallese, V., & Goldman, A. (1998). Mirror neurons and the simulation theory of mind-reading. *Trends in Cognitive Sciences*, *2*, 493–501.

Georgieff, N., & Jeannerod, M. (1998). Beyond consciousness of external events: A ''Who'' system for consciousness of action and self-consciousness. *Consciousness and Cognition*, *7*, 465–477.

Goldman, A. (1989). Interpretation psychologized. *Mind and Language*, *4*, 161–185.

Goldman, A. (1992). In defense of simulation theory. *Mind and Language*, *7*, 104–119.

Goldman, A. (2002). Simulation theory and mental concepts. In J. Dokic & J. Proust (Eds.), *Simulation and knowledge of action* (pp. 1–19). Amsterdam: John Benjamins.

Goldman, A. (2005a). Mirror systems, social understanding and social cognition. *Interdisciplines.* (Retrieved 12 January, 2007 from: http://www.interdisciplines.org/mirror/ papers/3)

Goldman, A. (2005b). Imitation, mind reading, and simulation. In S. Hurley & N. Chater (Eds.), *Perspectives on imitation* II (pp. 80–81). Cambridge, MA: MIT Press.

Goldman, A. (2006). *Simulating minds: The philosophy, psychology and neuroscience of mindreading*. Oxford, UK: Oxford University Press.

Goldman, A. I., & Sripada, C. S. (2005). Simulationist models of face-based emotion recognition. *Cognition*, *94*, 193–213.

Gordon, R. M. (2004). Folk psychology as mental simulation. In N. Zalta (Ed.), *The Stanford encyclopedia of philosophy.* (Retrieved 12 January, 2007 from: http://plato. stanford.edu/archives/fall2004/entries/folkpsych-simulation/)

Gordon, R. M. (2005). Intentional agents like myself. In S. Hurley & N. Chater (Eds.), *Perspectives on imitation* I (pp. 95–106). Cambridge, MA: MIT Press.

Grèzes, J., & Decety, J. (2001). Functional anatomy of execution, mental simulation, and verb generation of actions: A meta-analysis. *Human Brain Mapping*, *12*, 1–19.

Hurley, S. L. (1998). *Consciousness in action*. Cambridge, MA: Harvard University Press.

Hurley, S. L. (2005). Active perception and perceiving action: The shared circuits model. In T. Gendler & J. Hawthorne (Eds.), *Perceptual experience*. New York: Oxford University Press.

Hutto, D. (2003). Folk-psychological narratives and the case of autism. *Philosophical Papers*, *33*(3), 345–361.

Hutto, D. (2004). The limits of spectatorial folk-psychology. *Mind and Language*, *19*, 548–573.

Hutto, D. (in press). *Folk psychological narratives.* Cambridge, MA: MIT Press.

Jackson, P. L., & Decety, J. (2004). Motor cognition: A new paradigm to study self other interactions. *Current Opinion in Neurobiology*, *14*, 259–263.

Jacob, P. (2002). The scope and limits of mental simulation. In J. Dokic & J. Proust (Eds.), *Simulation and knowledge of action* (pp. 87–109). Amsterdam: John Benjamins.

Jeannerod, M., & Pacherie, E. (2004). Agency, simulation, and self-identification. *Mind and Language*, *19*(2), 113–146.

Lewin, P. (2005). Understanding narratively, understanding alterity. *Human Studies*, *28*, 375–383.

Meltzoff, A. N. (1995). Understanding the intentions of others: Re-enactment of intended acts by 18-month-old children. *Developmental Psychology*, *31*, 838–850.

Merleau-Ponty, M. (1964). *The primacy of perception and other essays on phenomenological psychology.* Evanston, IL: Northwestern University Press.

Millikan, R. G. (2005). Some reflections on the theory theory-simulation theory debate. In S. Hurley & N. Chater (Eds.), *Perspectives on imitation* II (pp. 182–188). Cambridge, MA: MIT Press.

Morton, A. (1996). Folk psychology is not a predictive device. *Mind*, *105*, 119–137.

Noë, A. (2004). *Action in perception*. Cambridge, MA: MIT Press.

Rizzolatti, G., Fadiga, L., Gallese, V., & Fogassi, L. (1996). Pre-motor cortex and the recognition of motor actions. *Cognitive Brain Research*, *3*, 131–141.

Rizzolatti, G., Fogassi, L., & Gallese, V. (2000). Cortical mechanisms subserving object grasping and action recognition: A new view on the cortical motor functions. In M. S. Gazzaniga (Ed.), *The new cognitive neurosciences* (pp. 539–552). Cambridge, MA: MIT Press.

Ryle, G. (1949). *The concept of mind*. New York: Barnes & Noble.

Saxe, R. (2005). Against simulation: The argument from error. *Trends in Cognitive Sciences*, *9*(4), 174–178.

Scheler, M. (1954). *The nature of sympathy* (Trans. P. Heath). London: Routledge & Kegan Paul. (Originally published in 1923 as *Wesen und formen der sympathie*. Bonn, Germany: Verlag Friedrich Cohen)

Shoemaker, S. (1968). Self-reference and self-awareness. *Journal of Philosophy*, *65*, 555–567.

Trevarthen, C. (1979). Communication and cooperation in early infancy: A description of primary intersubjectivity. In M. Bullowa (Ed.), *Before speech*. Cambridge, UK: Cambridge University Press.

Trevarthen, C. (1980). The foundations of intersubjectivity: Development of interpersonal and cooperative understanding of infants. In D. R. Olson (Ed.), *The social foundation of language and thought: Essays in honor of Jerome S. Bruner* (pp. 316–341). New York: Norton.

Trevarthen, C., & Hubley, P. (1978). Secondary intersubjectivity: Confidence, confiding and acts of meaning in the first year. In A. Lock (Ed.), *Action, gesture and symbol: The emergence of language*. San Diego, CA: Academic Press.

Wittgenstein, L. (1958). *The blue and brown books*. Oxford, UK: Basil Blackwell.