

ANDREW SNEDDON

MORAL RESPONSIBILITY: THE DIFFERENCE OF STRAWSON,
AND THE DIFFERENCE IT SHOULD MAKE*

Accepted: 18 February 2005

ABSTRACT. P.F. Strawson's work on moral responsibility is well-known. However, an important implication of the landmark "Freedom and Resentment" has gone unnoticed. Specifically, a natural development of Strawson's position is that we should understand being morally responsible as having externalistically construed pragmatic criteria, not individualistically construed psychological ones. This runs counter to the contemporary ways of studying moral responsibility. I show the deficiencies of such contemporary work in relation to Strawson by critically examining the positions of John Martin Fischer and Mark Ravizza, R. Jay Wallace, and Philip Pettit for problems due to individualistic assumptions.

KEY WORDS: externalism, Fischer, individualism, moral psychology, moral responsibility, Pettit, Ravizza, Strawson, Wallace, Wolf

What I call the individualist assumption in the philosophy of psychology is the assumption that whatever we find it significant to name – whatever patterns of feeling and behavior we find salient enough to use in explanations of ourselves and others – must also pick out complex events, states, or processes with respect to some theory of the functioning of individual organisms (for example, neurophysiology) (Scheman, 1996).

1. INTRODUCTION

My topic is moral responsibility. I will examine the familiar question, "What is it for an agent to be morally responsible?" However, I am going to address this question in conjunction with another, less familiar one: "What are the methodological options for studying moral responsibility?" My reasons for introducing this methodological concern are simple: I think both that it is important and that it has been overlooked. My overall aim is to show that P.F. Strawson's seemingly well-understood account of moral responsibility has methodological implications that have been widely overlooked.

*Thanks to an audience at Carleton University for helpful discussion.

2. WHERE TO START?

So, what is it for an agent to be morally responsible? What can we say as a starting point to this sort of inquiry? Like much contemporary work on moral responsibility, I shall follow Strawson to begin. The common view of Strawson's work on moral responsibility has two aspects. The first is that he connected moral responsibility to the deployment of "reactive attitudes". The second is that he articulated a particular way in which our moral notions and determinism could be compatible.

Strawson asks whether the thesis of determinism is compatible or incompatible with moral discourse. Determinism is part of a thoroughly objective account of the world. To approach phenomena in this spirit is to adopt a thoroughly objective attitude towards them. Strawson encourages us to think of as many kinds of interpersonal relationships as possible, then to think of the kinds of importance we attach to the attitudes and intentions directed towards us by the others in these relationships, and then to think of our own reactive attitudes (Strawson, 1974, p. 6). These reactive attitudes include personal ones, such as resentment, and more general ones that Strawson thinks include characteristically moral ones. Once we have listed the appropriate relationships and attitudes, the question to ask is whether the acceptance of the thesis of determinism could lead us always to look on everyone exclusively with the objective attitude (Strawson, 1974, p. 11). Doing so would mean giving up the subjective engagement of which we have reflectively framed an account. Since morality is part of this subjective engagement, an affirmative answer to the question entails that determinism is incompatible with moral categories, including moral responsibility. Strawson's answer, for both the personal and more general reactive attitudes, is that this is not logically inconceivable, but that it is practically so.

There is a second aspect to Strawson's conclusion. On occasion, we do adopt the objective attitude towards others. For example, when we find out that someone is mentally incapacitated in specific ways, we suspend our attitudes of resentment, and we cease to deploy the apparatus of moral responsibility in connection to the conduct of this person. The way we suspend our subjective engagement is very important. We do not take up the objective attitude as a result of conviction of the truth of determinism. Instead, the adoption of such an attitude is a consequence of the giving up of our subjectively engaged perspective. Determinism, or particular applications of this thesis, is never the cause of our suspension of our normal attitude. Further, we abandon our normal perspective for specific reasons in specific cases. We do not give it up wholesale as a result of a general theoretical conviction.

Overall, the lesson is that once we pay proper attention to the interpersonal domain in which the practices of moral responsibility have their home, it will become clear that determinism is no threat to moral responsibility. Morality and determinism are compatible because the objective domain which is the appropriate home of the discourse of determinism is, in a certain sense, irrelevant to morality. As Strawson puts it,

... questions of justification are internal to the structure [of human attitudes and feelings] or relate to modifications internal to it. The existence of the general framework of attitudes itself is something we are given with the fact of human society. As a whole, it neither calls for, nor permits, an external 'rational' justification (Strawson, 1974, p. 23).

3. HOW TO PROCEED?

Strawson does not explicitly address the question of what it is to *be* morally responsible – his position is first and foremost about what it is to attribute responsibility – but a reasonable interpretation of his position on the present topic is not hard to find. According to Strawson, to hold someone responsible is to deploy the reactive attitudes towards that person. When excusing conditions are present, we suspend the reactive attitudes and refrain from holding the person responsible for his/her actions. Presumably then, in the first place, to be responsible is to be an apt candidate for the reactive attitudes. Put another way, to be morally responsible is to fit into the social practices governing the deployment of the reactive attitudes. In short, it is to acquire a *social competence*.

This is the point at which my methodological question becomes important. How should we understand the acquisition and possession of a social competence? There are at least two broad ways of approaching this general question:

- A. A social competence can be explained solely in terms of the intrinsic properties of the individual agent. I shall call this, following the epigram from Scheman and established usage in philosophical psychology, the *individualistic approach*. I shall use “competence_I” to denote individually construed competences.
- B. Alternatively, a social competence can be explained in terms of the way the agent fits into his/her context. Instead of solely using intrinsic properties of the agent, relational ones would also be used in the explanation of the competence in question. I shall call this the *externalistic*

approach, and I shall use “competence_E” to denote externalistically construed competences.¹

Here is one important difference between individualistic and externalistic approaches to moral responsibility: if moral responsibility is a competence_I, then if one imaginatively holds the intrinsic properties of a morally responsible individual constant but varies the properties of his/her environment, especially the social properties, then the individual will always be morally responsible. No change in solely in environmental properties can affect his/her status as morally responsible. By contrast, if moral responsibility is a competence_E, then performing the same sort of thought experiment can result in changes to the agent's responsibility. Changes solely to the agent's environment, and not to the intrinsic properties of the agent, can change the agent from morally responsible to not morally responsible. Moreover, if one holds environmental properties constant, then two agents with very different intrinsic properties could both be morally responsible. In the case of at least some externalistic competences, *no* particular intrinsic properties of an agent may be necessary or sufficient for their realization. Moral responsibility might be such a competence_E.

As it happens, it is very revealing to use this very general methodological taxonomy to categorize extant accounts of moral responsibility. So far as I can tell, *almost all* major contemporary theorists pursue the individualistic approach (Watson, 2001; Fischer, 1999).² That is, they all theorize about moral responsibility in such a way that changes solely to an agent's environment and not to the intrinsic properties of the agent cannot affect the agent's status as morally responsible. By contrast, Strawson's own position, developed and extended to the issue of being morally responsible, is externalistic. Given the influence of Strawson's position, it is very interesting and peculiar that contemporary work that is explicitly influenced by Strawson diverges from his position in this way. We shall later look at the positions of John Martin Fischer and Mark Ravizza, R. Jay Wallace, and Philip Pettit to examine their individualism.

At this point, no arguments have been provided to support either an individualistic or an externalistic approach to moral responsibility. I take this lack of pre-judgment to be appropriate given the recent rise of externalism to counter-balance the individualistic tendencies in philosophy

¹ In a similar vein, Bernard Berofsky (1995, p. 34), speaking of abilities, argues that our everyday concept is not to be understood solely in terms of “internal states” of agents.

² Gary Watson (2001) takes seriously the idea that moral responsibility might have an irreducibly externalistic aspect; refer to Fischer (1999) for further discussion.

of mind.³ As a preliminary step in determining whether moral responsibility is a competence_I or a competence_E, let's look at some non-moral examples.

First, a competence_I. I take tying one's shoelaces to be an individualistic competence. When we learn to tie our laces, what we acquire are skilled finger movements. If we hold the intrinsic properties of a competent lace-tier constant and vary their environment, their lace-tying competence will not be affected. They might find themselves in environments where there is no need for this competence (e.g., because there are no laces to be tied), but this does not entail that the individual is not a competent lace-tier.

Competences_E are more complex. Lace-tying would be an externalistic competence if there were socially specific standards that, at least in part, determined which manipulations of laces counted as lace-tying and which did not. In such a case, holding the intrinsic properties of the individual constant and varying the environment could affect their status as a competent lace-tier: what counts as lace-tying in one social context would be idiosyncratic lace-knotting in another context, not really lace-tying as it is recognized there.⁴

As an example of a real-world competence, consider the ability to play tennis. Andre Agassi is a competent tennis player (to say the least). So is Andy Roddick. By this I mean that they both excel at meeting whatever criteria of excellent participation in tennis there are. But their competences are manifest in different ways. There are all sorts of strokes that both players must be able to perform to be competent. Some of these might be necessary for competence, some might be merely sufficient – i.e., helpful,

³ Debate between individualists and externalists started over content. Individualists contended that mental content was determined by intrinsic properties, typically physical ones, of individuals. Thought experiments by Hilary Putnam⁶ and Tyler Burge⁶ questioned this. These arguments suggested that content was at least partly determined by aspects of agents' environments. Debate has since turned to whether cognitive processes might also be widely distributed. This idea is currently more controversial than the idea that content is externally determined, but support is mounting from a variety of directions. Philosophers, often working in conjunction with the sciences of the mind, have offered theoretical support. Wilson (1994, 1995, 2001, 2004), Clark and Chalmers (1998), Rowlands (1999), Hurley (1998) are representative of this sort of case for (using the terminology of Clark and Chalmers) "active externalism". Support from a broadly engineering perspective has come from Brooks (1991). Clancey (1997a,b). Hutchins (1995) offers evidence from an observational approach. Sneddon (2002) has argued that there is some reason to think that psychopathology will support externalism about some cognitive processes. On a slightly different note, Paprzycka (2002) argues for changes to psychological theories to include ordinary and widespread externalist rationalizing explanations of everyday action.

⁴ Maybe tying one's laces *is* an externalistic competence even for us. Little children who secure their shoes in unwieldy knots are said not to know how to tie their laces, even if these knots successfully hold their shoes on their feet.

but not vital. The fact that Agassi and Roddick excel in different aspects of tennis supports this impression. Agassi's success hinges, in large part, on his return of serve. By contrast, Roddick is an average returner at best. His excellence rests on his serve. Insofar as we mean to emphasize what is excellent about their abilities when we describe Agassi and Roddick as (colloquially) able tennis players, their competences are realized by different causal processes.

The Agassi–Roddick example presents an interpersonal comparison of the realization of competence, but the same case can be made with an *intrapersonal* example. Martina Navratilova is another competent tennis player. Her success at Wimbledon was accomplished with a serve-and-volley strategy: follow one's serve by running to the net and trying to control play to one's advantage with volleys rather than groundstrokes. Such competence in tennis hinges on two types of strokes, clearly executed in different ways. The serve is the stroke that initiates play. The player has a very high degree of control of the location of the ball at the time of contact because s/he throws it into the air him/herself. The server has as much opportunity as possible to plan the pace and trajectory of the ensuing shot. A volley could hardly be more different. It is hit typically when the player is close to the net. The ball has not bounced on his/her side before contact. Because the player is so close to his/her opponent, and because the opponent has, in all likelihood, hit the ball rather hard, a volley is often a reflex shot. It can consist of little more than a stab in the direction of the ball. This is in marked contrast to the consciously planned execution of the serve. Assuming that there is significant difference between the planned execution of a movement and the reflex execution of a movement, it is safe to conclude that Navratilova's tennis competence is realized by a variety of distinct causal processes. Whether any particular one is necessary is very difficult to say. It certainly cannot be assumed at the beginning of inquiry, but instead must be demonstrated.

Our focus can be made even more specific with an *intra-shot* example. Volleys are hit in a variety of ways. Most are largely reflex, but if your opponent hits a looping, weak shot to you, you actually have a great deal of time in which to plan your volley. Hence even volleying competence is realized by distinct causal processes.

As it happens, the lesson that has unfolded in the various tennis examples holds generally for biology. Hugh LaFollette and Niall Shanks point out that an important commonplace of biological inquiry is that commonality of biological function between members of different species does not entail commonality of cause of these functions (LaFollette and Shanks, 1996, p. 100). They can be realized in different ways. I take the present considerations of the realization of competences between individuals to be a related issue.

Overall, it is appropriate to see “tennis competence” not as a term that picks out a single ability, but rather as one that is realized by a variety of individually construed abilities. Moreover, it is reasonable to think that this competence_E does not require any particular abilities in order to be realized. Turning our attention from experts to beginners helps to make this clear. Think of the burgeoning abilities of novice tennis players. For beginning competence, what is required is to be able to get the ball over the net. This is accomplished in a variety of ways, none of which is individually necessary for tennis competence_E. This notion can be generalized to more expert performance: although the state of the professional game at time *T* might be such that a powerful serve is necessary for success, this has certainly not been the case at previous times *T*−*n*, and it might well not be the case at future times *T* + *n*. The players at these different times all count as competent, yet their intrinsically construed abilities differ.

So, the question is which is more similar to moral responsibility: lace-tying or tennis? I’m inclined to think that being morally responsible is more like being able to play tennis than it’s like being able to tie one’s shoelaces. Like tennis, to be morally responsible is to have an array of abilities. In the case of moral responsibility, one has various sensitivities that one applies to the conduct of oneself and others, and I am not sure that any of these in particular are necessary for all social practices that could define moral responsibility. My argument for seeing moral responsibility as a competence_E will be to develop Strawson’s position. To assess the reasons for taking an individualistic approach to moral responsibility, I will examine the positions of Fischer and Ravizza, Wallace, and Pettit.

3.1. Strawson: Being Responsible

The question is, “What is it for an agent to be morally responsible?”. Putting this more precisely, we can ask, “What makes it true of someone that the criteria of being morally responsible are met?” The typical individualistic strategy deployed to answer this question is to turn to psychology: an individual meets the criteria of being morally responsible when they have certain psychological states/processes/abilities. A thoroughly Strawsonian account of being responsible, however, must be explicitly opposed to such an account. The reason is that explanations in terms of psychological states/events/processes are a specific variety of causal explanation. Causal notions, however, have their natural home in the domain of determinist discourse, which is external to the participant perspective within which the reactive attitudes are deployed. A Strawsonian account of being an apt candidate for the reactive attitudes must use other resources.

There is an obvious response one might make at this point. Perhaps it is reasonable to see theoretical psychology as belonging to the objective perspective, but there is also folk psychology. Since this is pre-theoretical discourse, perhaps it is reasonable to see it as most naturally belonging to the participant perspective, not the objective one. After all, Strawson himself locates the reactive attitudes within this perspective, and “attitude” is a psychological notion. So, instead of trying to provide criteria of moral responsibility in terms from some particular corner of the psychological sciences, instead try to formulate them using folk psychological notions.

This way of proceeding faces a dilemma. Since psychological sciences belong to the objective perspective, and since the present idea is to try to use folk psychological notions to formulate criteria of moral responsibility, it is reasonable to ask about the relationship between folk psychology and scientific psychology. There are two possibilities here:

- (1) Suppose that it is very close. In this case, it looks as if folk psychology gets a principled interpretation *within* the objective perspective. This robs the participant perspective of independent resources to provide a psychological account of the criteria of moral responsibility. Strawson’s original arguments work against folk psychology that has been incorporated by scientific psychology.
- (2) Suppose that it is not very close. This would preserve folk psychological notions independent of scientific psychological ones. But now it is reasonable to ask why we should understand these folk psychological notions as designating facts about the mind. If principled scientific psychology has no place for these notions, then it is not at all clear that they carve the mind at its joints, so to speak. Securing the independence of folk psychology from scientific psychology robs folk psychology of its authority. This possibility does not undermine the authority of the participant perspective altogether. It does it only for participant discourses that have a principled cousin in the objective perspective: e.g., folk physics and scientific physics, folk biology and scientific biology. In particular, moral discourse seems to have no such objective counterpart, so the authority of this participant perspective discourse is not threatened.

Here is a (neo-)Strawsonian⁵ account of being responsible. The reactive attitudes are deployed within the participant perspective, as opposed to the objective perspective characterized by causal discourse. Since we have

⁵ I will henceforth refer to this as Strawsonian, but it should be taken as, strictly, a development of his account of attributing moral responsibility, and hence as *neo*-Strawsonian.

reason to think that these attitudes stand in no need of justification in terms not at home in the participant perspective, it seems that a Strawsonian account of being responsible should locate the criterion of this notion within this perspective. The criteria of being an apt candidate for the reactive attitudes should not be causal, but *participatory*, or, to use a more familiar word, pragmatic.⁶ That is, it is public, behavioral facts that make it true of someone that they are morally responsible. In particular, it is important to see how someone interacts with others, or, to put it even more broadly, how someone fits into the social context in which s/he finds him/herself. As such, this clearly cannot be an individualistic account of what it is to be morally responsible. Such an account of moral responsibility appeals not to the individualistically construed concepts of some psychological or neurological theory, but instead to non-individualistic evidence about the relations between people. In short, a thoroughly Strawsonian account of being morally responsible is externalistic.⁷

To show that one is morally responsible, a person must demonstrate that s/he is an appropriate target for resentment, gratitude, etc. This is done through participation in the very practices through which these attitudes are exercised. More specifically, one demonstrates social competence by acting in accordance with, e.g., rules, principles, expectations, etc. With regard to expectations, one demonstrates social competence interpersonally, and hence publicly. In Strawsonian terms, one does this by acting sensitively to the deployment of the reactive attitudes, and as if using these attitudes oneself. Whether one actually experiences the feelings seems to be beside the point, since whether one does or not is inaccessible to others.

Overall, this will sound odd to many, but it shouldn't. It sounds dubious if we imagine a heretofore isolated adult trying to gain entry into a social group constituted by certain moral practices. The real world, however, is rarely like this. First, no person exists completely outside of already existing social groups. Second, the groups that do exist are not isolated from each other. Moral practices overlap; there is much in common around the world with regard to the sorts of attitudes that are deployed within various participant perspectives. Consequently, any individual who did find him/herself in the very odd situation of trying to demonstrate one's competence in moral practices would have much to use from his/her background. Third, the *exceedingly* vast majority of people do not find themselves in a position where their moral participatory competence must be demonstrated all at once. Instead, we start as children and hone our participatory skills as we mature. We have, literally, years to demonstrate that we are apt candidates

⁶ See Wallace (1994), 9 for like characterization.

⁷ The remarks just made about folk psychology show why one cannot respond to this view by using explicitly externalistic folk psychological notions.

for the reactive attitudes. Our existing practices reflect this fact. In general, we do not hold children responsible for their actions to the same degree as we do adults. But we do hold them "sort of" responsible. We routinely deploy modified versions of the reactive attitudes towards children. The extent of such deployment depends, fairly reliably, on the level of maturity of the recipient. We even explicitly deploy some reactive attitudes in a teaching mode. For instance, it is common for adults to thank children in more explicitly overt ways than they thank other adults. The rationale is clearly instructive: it is a way of helping to bring up these children to be full participants in the moral domain.

Fourth, in practical contexts involving adults, actual demonstration of competence is of secondary importance. When adults from different social groups meet, their long period of education is implicitly recognized. The most common way of behaving towards strangers is to assume that they are apt candidates for the reactive attitudes. In practice it is *failure* to fit in that strangers must demonstrate, not competence with regard to the moral practices that constitute the participatory perspective. We assume that people are morally responsible when the conditions that would defeat seeing someone as responsible are not filled. Both the meeting and the avoiding of defeating conditions, such as mental incapacity, are done with publicly observable, interpersonal performances. To construe the absence of the satisfaction of defeating conditions in terms of intrinsically construed capacities of the agent is an individualistic interpretation of publicly observable and in large part contextually individuated competencies. A common practice is to *assume* intrinsically construed capacities, but we *encounter* competence_E.

There are two things worth noting in passing about this view of being responsible. First, agents who demonstrate their competence in the social practices that are characterized by deployment of the reactive attitudes count as responsible even when they are not actually being held responsible.⁸ Second, non-humans are not morally responsible because they routinely fail, completely, to demonstrate competence in the practices that provide the individuating conditions of the status of being responsible. However, once we find non-humans, whether artificial or natural, extra-terrestrial or terrestrial, that do demonstrate such competence, then they're in! We would have no principled grounds on which to deny that these were morally responsible beings.

⁸ However, given the omni-presence of social expectations, and given that we apply the social apparatus of holding responsible to ourselves, perhaps there is a sense in which we are always being held responsible. Either way, the property of being responsible persists, rather than coming and going with actual praise and blame.

3.2. *Getting Metaphysically Specific*

Although I have put the Strawsonian account of being responsible in terms of *pragmatic* criteria, it is possible, following recent work in externalist psychology, to say something about the metaphysics of this phenomenon. Robert Wilson addresses externalism about psychological states and processes by examining the nature of systems. It is a commonplace to see cognitive states, processes, and abilities as *realized* by cognitive systems. Wilson follows Sydney Shoemaker's distinction between core and total realizations to examine the phenomenon of realization. For a higher level property H and a system S in which it is realized, the core realization is, "... a state of the specific part of S that is most readily identifiable as playing a crucial causal role in producing or sustaining H." (Wilson, 2001, p. 8). The total realization of H is, "... a state of S, containing any given core realization as a proper part, that is metaphysically sufficient for H." (2001, p. 8). The system S in question will typically be specified through a *posteriori*, empirical inquiry. Wilson discusses pain, which is realized by our nociceptive system. Our nociceptive system is contained within our physical boundaries. However, Wilson contends that there are other systems that are not contained by such boundaries. Instead of containing these systems, organisms are themselves part of such wide systems. Using this possibility, Wilson identifies two sorts of externalist realization. "Wide" realization occurs when there is, "... a total realization of H whose non-core part is not located entirely within B, the individual who has H." (2001, p. 11). "Radically wide" realization involves, "... a wide realization whose core part is not located entirely within B, the individual who has H." (2001, p. 13).

More recently, Wilson has distinguished varieties of externalism in accordance with the two varieties of wide realization. When a property of an individual is widely realized, then it must be individuated in reference to the system that extends beyond the boundaries of the individual. This yields a position about taxonomy, which Wilson calls "taxonomic externalism" (2004, pp. 174–178). By contrast, properties with radically wide realizations are not really properties of the individual that one is examining, but are instead located beyond its physical boundaries. The associated view of externalism is, accordingly, "locational externalism" (2004, pp. 174–178).

Some examples are helpful here. The biological property of *being a predator* is one that is properly attributed to individual organisms, but one which they have by virtue of their role in a predator–prey system (Wilson, 2004, pp. 114–115). Accordingly, biologists should be taxonomically but not locationally externalist about predators. By contrast, certain research programmes in cognitive science describe cognitive tasks as being accomplished *between* individuals, or via individual – environment interaction.

Wilson discusses Edwin Hutchins's work on how navigational tasks are performed and Rodney Brooks' work designing robots as examples of such research programmes (2004, pp. 175–178). The contention is that the cognitive processes in question are actually located beyond the physical boundaries of the individuals participating in the systems, so we should be locationally externalist about them.

These metaphysical possibilities come with corresponding explanatory strategies. Entity-bounded realization calls for "constitutive decomposition" (2001, p. 11): the system in question is explained by making clear what parts of the organism stand in what relations to each other to function as a system. By contrast, wide varieties of realization call for "integrative synthesis," which locates the individual within a system made up by that individual and parts of its environment, and specification of the relations between these various parts of the system to explain how it functions and realizes such properties as H.

This conceptual framework can usefully be applied to the neo-Strawsonian position on responsibility developed here. Individuals are morally responsible; nevertheless, this is a status they have by virtue of their participation in a system constituted by the deployment of the reactive attitudes.⁹ So, we should be taxonomically externalist about responsibility, but not locationally so. Being responsible is the same sort of property as being a predator. Since one is part of the reactive attitude system by virtue of one's behavior – in particular, by how one interacts with other people – it is appropriate to speak of the criteria of being responsible as pragmatic and public. Being a predator has the same sort of criteria: what one does to other creatures. And just as different sorts of predators realize the ability to consume their prey in radically different ways, the behavior by virtue of which one participates in a reactive attitude system could, in principle, be realized in radically different ways.

4. BEING RESPONSIBLE: THE DIFFERENCE OF STRAWSON

I take the view that the criteria of being morally responsible are not individualistically psychological but instead pragmatic as a very natural extension of Strawson's overall position that moral notions, centrally including moral responsibility, admit of no external justification. Given that many theorists who address moral responsibility locate their positions favorably in relation to Strawson, and given that his contention that moral notions stand in no need of external justification is well known, one would expect

⁹ Wilson might see this as part of the "folk psychological system," since he casts this as including social relations (2004, p. 113).

many contemporary theorists either to elaborate externalistically construed pragmatic criteria for being responsible, or to argue directly against understanding moral responsibility this way. In fact, neither of these is the case. Contemporary theorists routinely seek individualistically construed psychological criteria for responsibility even when they use Strawson as a reference point.¹⁰ Moreover, they regularly omit without argument the possibility of there being other sorts of criteria for moral responsibility that they should be arguing against.¹¹ To demonstrate the difference of Strawson, I will present the accounts of moral responsibility defended by John Martin Fischer and Mark Ravizza, R. Jay Wallace, and Philip Pettit. They all seek individualistically construed psychological conditions for being responsible. Since such criteria are “objective” in Strawson’s sense of having their natural home in discourse that is not part of the participatory perspective where moral notions are deployed, these sorts of criteria are also external to this perspective. Hence such accounts of being morally responsible depart from the Strawsonian view, and indeed from Strawson’s overall position on moral notions in general. Let’s ignore the arguments already provided against offering psychological accounts of moral responsibility. Instead, in what follows the emphasis will be on (possible) argumentation for individualistic accounts of being morally responsible.

4.1. John Martin Fischer and Mark Ravizza

John Martin Fischer explicitly casts moral responsibility in terms of being an appropriate subject of the reactive attitudes along with their closely related practices of ascribing praise and blame (Fischer, 1994). Fischer thinks this sort of matter calls for a metaphysical answer: our practices of holding people responsible are to be grounded by whatever facts about people actually make them responsible.¹² Given just these rough outlines, it should be clear that Fischer seeks psychological criteria – i.e., facts about individual psychology, regardless of context – for being responsible.

¹⁰ Perhaps this deserves defense, since the issue of individualism is not explicitly broached by most theorists. The positions that I examine all cast moral responsibility as if changing purely relational properties of individuals could not make a difference to the status of these individuals as morally responsible. Thus, we have *prima facie* reason to see these as individualistic views of moral responsibility.

¹¹ Stephen White (1990) is a notable exception. He explicitly argues for what he calls the “intrinsic property constraint” on justifications of practices of attributing responsibility. This constraint holds that such justifications must be grounded in intrinsic properties of moral agents. White thinks that dropping this constraint means one is forced to develop a kind of “direct” compatibilism that unduly violates pretheoretical ideas about responsibility (1990, pp. 418–419).

¹² Fischer, *The Metaphysics of Free Will*, p. 213.

Fischer claims that rational accessibility to the reactive attitudes and to practices of attributing praise and blame is grounded by control.¹³ Crucially, Fischer distinguishes between regulative and guidance control.¹⁴ An agent exhibits regulative control over behavior when s/he can choose amongst alternative courses of action. An agent exhibits guidance control over behavior when s/he exercises control over the course of activity even though s/he could not have chosen to do otherwise – no alternative courses of activity were available. Fischer offers the following as an example:¹⁵ suppose you are driving a car. You come to a fork where you must turn either left or right. You turn to the right, and the car goes to the right, just as you wanted. As it happens, unbeknownst to you, the car’s steering mechanism is broken. If you had turned to the left, the car would have veered to the right anyway. Thus, there was only one course of action open to you: there were no alternatives to turning right. There is a sense in which you have successfully controlled the direction of the car, but it is only in the guidance sense, not the regulative sense.

This example is offered in the spirit of other scenarios devised by Harry Frankfurt (1969). Frankfurt’s examples suggest that, in contrast to common intuition, agents can be morally responsible for their activity even when they cannot do otherwise. Fischer agrees: he argues that only guidance control, and not regulative control, is necessary for moral responsibility. In their later book, Fischer and Ravizza argue that an action is under the guidance control of an agent when it issues from (1) a moderately reasons-responsive mechanism that (2) the agent has made his/her own (Fischer and Ravizza, 1998, p. 230). We shall examine these conditions in order.

A mechanism is moderately reasons-responsive when it is regularly receptive and weakly reactive to reasons.¹⁶ As the name suggests, this sort of responsiveness falls between weaker and stronger varieties. Weak reasons-responsiveness (which Fischer earlier thought was all that was necessary for moral responsibility¹⁷) posits a loose fit between there being a sufficient reason for action and the production of action. Strong reasons-responsiveness is constituted by a tight fit between action and sufficient reason. Fischer and Ravizza argue that these are respectively too little and too much to ask of agents for them to be morally responsible for their actions. Instead, a moderate fit is all that is needed.

The second condition, that the agent must have made the mechanism his/her own, is a historical condition. To make an action-yielding

¹³ Fischer, *The Metaphysics of Free Will*, pp. 3; 21.

¹⁴ Fischer, *The Metaphysics of Free Will*, p. 132.

¹⁵ Fischer, *The Metaphysics of Free Will*, pp. 132–133.

¹⁶ Further explication of what this means is unnecessary for present purposes.

¹⁷ Fischer, *The Metaphysics of Free Will*, p. 166.

mechanism one's own is something an agent does. It is accomplished by taking responsibility for acting from a certain kind of mechanism.¹⁸ This involves [a] seeing oneself as the source of behavior, and [b] seeing oneself as a fair target of the reactive attitudes. Moreover, [a] and [b] must be appropriately based on evidence. Taking responsibility for one's activity in this way can happen either through the moral education one grows up with, or through reflection. Either way, this is a historical process. Fischer and Ravizza incorporate such a process into their account of the grounds of moral responsibility because of a problem that otherwise faces attempts to ground moral responsibility in mechanisms of control. Any mechanism in itself could be installed by historical processes that compromise an agent's responsibility for the activity yielded by the mechanism.¹⁹ Would an agent really be responsible for activity that issues from a mechanism installed in the agent by a neurosurgeon? Perhaps not. If this is correct, then such processes have to be ruled out to specify the metaphysical conditions of being responsible for activity.

What should we think of this sort of position? There are a couple of issues to address. First, there is the methodological issue of defending the search for individualistic psychological criteria of moral responsibility against accounts that offer other sorts of criteria. Fischer offers nothing explicit on this matter. The closest we get is a claim that we must distinguish being responsible from being held responsible. Fischer and Ravizza worry about the close connection between being responsible and being held responsible in Strawson's position. They object that sometimes people who are not responsible are held responsible for behavior (Fischer and Ravizza, 1993, p. 18); theories of responsibility need to respect and make sense of this distinction. The distinction, they think, must be made in terms of matters external to the participant perspective. If this is true, then Strawson's position is ill-positioned to address it. However, a Strawsonian need not worry about this objection. There is lots of room in such a position to account for, e.g., being held responsible for behavior that one did not produce, or for which someone else is institutionally responsible, etc. Generally, lots of publicly observable information could serve as excuses that show that one should not be held responsible for *x* because they are not actually responsible for *x*. As for being responsible, again, if one has not demonstrated competence in the social practices that are constituted by the deployment of the reactive attitudes, then one is not responsible, but such a person could clearly be mistakenly held responsible. The individualistic approach favored by Fischer and Ravizza is not uniquely attractive as a way of dealing with differences between being held responsible and being

¹⁸ Fischer and Ravizza, *Responsibility and Control*, p. 215.

¹⁹ Fischer, *The Metaphysics of Free Will*, p. 208.

responsible. More explicit argument is needed to defend the anti-Strawsonian approach.

Second, we should hesitate to follow Fischer and Ravizza in their argument for a historical condition for being morally responsible. This is the claim that the reasons-responsive mechanism must be made the agent's own. One makes a mechanism one's own by taking responsibility for it; this adds a historical dimension to this account of moral responsibility. Fischer and Ravizza include a historical dimension because, "[Any mechanism] could *also* be produced in a responsibility-undermining way – say, by direct electronic stimulation of the brain, subliminal advertising, and so forth." [original emphasis].²⁰ Hence Fischer and Ravizza specify a specific sort of historical process which purportedly helps constitute agents as morally responsible. They have, however, over-stepped what is really needed here. Let's call the responsibility-undermining ways of installing mechanisms for behavior "bad historical processes." The claim is that bad historical processes suffice to undermine moral responsibility. Schematically:

If Bad historical processes **then** no moral responsibility.

Fischer wants moral responsibility. So, working *modus tollens*, the inference should be:

If Moral responsibility **then** no bad historical processes.

That is, if certain ways of installing mechanisms for producing behavior undermine responsibility, then all that is necessary for moral responsibility is the absence of such processes. The absence of "bad historical processes" is not equivalent to the presence of specific responsibility-enabling processes. Fischer and Ravizza strive for more than they need. Consequently, their position should be modified to claim that moral responsibility is grounded by moderate reasons-responsiveness and the absence of ways of installing such an ability that would undermine moral responsibility.²¹

This argumentative move is more closely related than one might think to the search for individualistic psychological criteria for moral responsibility. Having identified a set of conditions that can undermine responsibility, Fischer seeks something specific about individuals to pre-empt the problem. However, mere logic calls not for this, but for just the absence of the problem. The absence of such processes is not a fact about an individual,

²⁰ Fischer, *The Metaphysics of Free Will*, p. 208.

²¹ It is natural to expect accounts of moral education, including learning about moral responsibility, to make central reference to history, but this is a distinct issue from the present topic from Fischer and Ravizza.

but rather about broader states of affairs. Let me speculate that this strategy is rooted in an assumption that moral responsibility is to be grounded in facts about individuals and individual psychology. To assume this is to beg the question against the sort of approach to moral responsibility implicit in Strawson's work.

4.2. R. Jay Wallace

R. Jay Wallace's account of moral responsibility is more subtle than the work of Fischer and Ravizza with regard to its relation to Strawson's work. Wallace identifies two traditional ways of interpreting the issue of being responsible (Wallace, 1994, p. 85):

- (A) The *metaphysical* interpretation holds that there is a realm of facts absolutely independent of moral norms and practices and to which these norms and practices answer. Some of these facts constitute being responsible. The work we have seen from Fischer and Ravizza falls into this category. Wallace rejects this approach simply because he finds the positing of such a domain of facts that can serve this purpose implausible. We can add this sort of complaint to the problems already seen.
- (B) Wallace identifies a *pragmatic* interpretation of the issue of being responsible, but he provides merely a negative characterization of it: "Extreme pragmatist interpretations, on the other hand, abandon the idea that there is any fact of the matter about what it is to be responsible. . ." (Wallace, 1994, p. 85). He dismisses this as overly pessimistic, to be adopted only as a last resort (Wallace, 1994, p. 88).

However, we have good reason to be suspicious of Wallace's quick rejection of pragmatic interpretations of being responsible. I have characterized a Strawsonian approach as pragmatic, yet I have provided a positive characterization: to be responsible is to demonstrate competence in the social practices that are inherently characterized by practices of holding people responsible. There clearly is a fact of the matter about whether any given person has in fact demonstrated such competence and earned the social status of being responsible. Only if one thinks that there must be an *individualistically* construed fact of the matter could one claim that this pragmatic interpretation of the issue gives up on there being truths in this domain. For, as we have seen, when characterized this way, being responsible has contextualist, non-individualistic criteria. One implication of this is that what counts as a demonstration of social competence in one domain might not be adequate in another domain. A person could be responsible in one social context, but not in another. This possibility, of course, does

not entail that there is no fact of the matter: there clearly are facts about how the person has lived up to the social criteria in the two contexts. Interestingly, Wallace does not have the resources to reject this idea on the grounds that it ties responsibility too closely to our practices of holding people responsible, for he has already embraced such a close connection in his rejection of the metaphysical interpretation.

Since he has rejected both traditional approaches to the issue of being responsible, Wallace provides a new, third option:

- (C) The issue is a *normative* one. The issue should be seen as one of specifying the conditions under which it is *fair* to hold someone responsible. As a normative issue, the matter of being responsible arises within social practices characterized by moral norms. The issue no longer calls for handling in terms external to these practices. Wallace offers the following schema for thinking about the conditions of being responsible: (Wallace, 1994, p. 91).

(N) *S* is morally responsible (for action *x*) if and only if it would be appropriate to hold *S* morally responsible (for action *x*).

The task now for Wallace is how to explicate this schema. Wallace proceeds, from a perspective internal to our participatory stance, metaphysically. He claims that the possession of rational powers, ". . . the power to grasp and apply moral reasons, and the power to control one's behavior by the light of such reasons. . ." (Wallace, 1994, p. 7), makes it appropriate to hold the possessors morally responsible. These are similar to the ideas offered by Fischer and Ravizza. We can now see why Wallace's position is subtle: he follows Strawson very far, yet still caters to the metaphysical intuitions that drive much contemporary investigation of this issue.

There are several points to note here. One is that Wallace's normative schema is virtually identical to the way I characterized the natural understanding of being responsible from a Strawsonian perspective. I then pursued a pragmatic interpretation: demonstrating social competence and earning a certain social status makes it appropriate to hold one morally responsible. This suggests that the normative strategy is not so distinct from the pragmatic strategy as Wallace thinks. Secondly, and crucially, it indicates that even from a stance internal to the participatory perspective, the pragmatic route is available. Wallace rejected metaphysical and pragmatic approaches *prior* to the choice of a normative approach; now, within the normative approach, Wallace pursues, by his standards, a metaphysical interpretation. A pragmatic interpretation *within* the normative approach is equally within reach. So, despite Wallace's rejection of these ways of approaching the issue, it seems that the choice between them reappears after Wallace's reinterpretation. The question is how to choose between them.

Wallace does not address this issue directly. He does not return to consider the merits of the pragmatic interpretation (understandably, since he had not characterized it very charitably to begin with). Instead, he claims that it is “natural” to think that what makes one morally responsible is some sort of ability (Wallace, 1994, p. 6). That is, very early in his examination of being responsible Wallace assumes an individualistic approach. We have seen the same sort of assumption about this issue from Fischer and Ravizza.

The clearest defenses of individualism presented by Wallace start from considerations of what might undermine a person’s responsibility. Wallace casts his position against incompatibilism. He offers the following as normative interpretations of the incompatibilist’s complaints about determinism and responsibility:

- (1) “The incompatibilist wishes to say that people do not deserve to be blamed or sanctioned for the individual actions they perform, if determinism is true.” (Wallace, 1994, p. 107).
- (2) “The thought is that it would be unreasonable to hold people accountable if they lack freedom of the will. . . .” (Wallace, 1994, pp. 108–109).

Wallace comments that these sorts of ideas are “most naturally interpreted” as indicating that a certain sort of ability is a condition of the appropriateness of holding someone responsible. This is an individualistic approach.

This pattern of inference deserves scrutiny. It seems cogent, but in fact it is enthymematic, and the missing step is just the assumption of individualism. Generally, the lack (or malfunction) of something *X* can prevent an event or item from working or from being of a certain kind without the presence of *X* being criterial of successful functioning or the kind in question. *X* could, for instance, be part of a system whose overall constitution realizes the kind in question. Imagine that one lives in a political jurisdiction where voting is done by marking an “x” in a box. If one does not mark an “x” in a box, then one does not vote. This, of course, does not mean that marking an “x” in a box constitutes an event as a vote. Other very important conditions have to be met: one has to be registered to vote, one has to go to the appropriate place, receive an official ballot, etc. As another example, consider submitting assignments as part of taking a course. Suppose a student does not pass the course because s/he did not pass in assignment *X*. This, of course, does not mean that submission of *X* constitutes success in the course. This student might have done adequately well on assignments *U*, *V*, & *W*, but *X*, worth 40% of the final mark, was needed to pass. Submitting *X* alone would not constitute success in the course.²² The lesson

²² Reasoning from absence or malfunction to the conditions that constitute good functioning is common in the sciences. It has been particularly important in neuroscience: a

here is that explicit positive argumentation is needed to establish the *positive* criteria of moral responsibility.²³ The factors that defeat ascriptions of moral responsibility do not necessarily correspond in a simple one-to-one fashion with the factors that warrant it. Crucially, such objective, individualistic defeaters are perfectly compatible with pragmatic criteria of being responsible. Strawson himself discusses psychological problems as particular reasons to give up the participatory perspective in particular cases. Conclusions about the criteria of some status cannot be simply read off, in a one-to-one fashion, from premises about conditions that defeat that status. Criterial and defeating conditions *might* have such a one-to-one correspondence, but this must be independently demonstrated. To assume that they line up in this way is not warranted.

On the matters of critical and defeating conditions of desert and reasonableness – the normative notions deployed in Wallace’s two characterizations of incompatibilism – consider the following cases:

- (1) *S* is a professional hockey player. One day while practicing, *S* injures his ankles. As a result of his injury, *S* will not be able to play for the rest of the season. Most would agree that it is not *fair* for the team’s General Manager to fire *S* because of his injury. That is, *S* does not *deserve* to be fired. However, it is not natural to think that having healthy ankles is a condition of *S*’s deserving the job. And, even if one thinks that having healthy ankles is some sort of condition of deserving the job, it clearly is not a central or important one.
- (2) Suppose that you want to read a long book that you have been meaning to read for some time. As it happens, you are in the middle of moving house, and you have no place to sit other than the hard floor. So you sit down and start your book. Your spouse comes in and tells you that it’s not *reasonable* to attempt such a long read without a good place to sit. However, it is not natural to think that having a good place to sit is a condition of the reasonableness of reading a long book. And even if one is tempted to think that it is at least some sort of condition of the

brain malfunction that correlates with some sort of disability calls out for further investigation. But, of course, the inference that successful functioning of *that* part of the brain is responsible for successful performance of a certain kind of ability is too hasty an inference. Malfunction of a part of the brain can interfere with abilities in lots of ways. For a *good* example of this sort of reasoning, see Pinker (1994, p. 298).

²³ Something is defined in terms of negative criteria when it is defined in terms of the absence of things that cause problems for it. As an example, consider a definition of health solely in terms of the absence of disease and infirmity. A definition in terms of positive criteria specifies conditions the *presence* of which purportedly realizes the phenomenon in question. The WHO definition of health equates it with a complete state of physical, mental, and social well-being; these are its purported positive criteria.

reasonableness of such an undertaking, it is clearly neither a central nor an important one.

The lesson of [1] and [2] is that it is not, in fact, natural to infer directly from a defeating condition to a criterial condition. Moreover, even if the defeating condition does in fact indicate some sort of criterial condition, it need not be a centrally important one. The criterial condition indicated by the defeating condition could well be relatively superficial and hence relatively uninformative about the nature of the phenomenon at issue. Direct argumentation is needed both to demonstrate that a defeating condition is a good indicator of a criterial condition and that this criterial condition is in fact an important one. Individualistically oriented accounts of being morally responsible routinely fail to provide such direct argumentation.²⁴

4.3. Philip Pettit

Philip Pettit's position has the most in common with the present Strawsonian account of being responsible. Pettit addresses freedom in terms of fitness to be held responsible. He explains this in terms of what he calls discursive control. Discursive control has two aspects, one individualistically individuated, the other contextually individuated: "People enjoy freedom as discursive control so far as they have the ratiocinative capacity to enter discourse and so far as they have the relational capacity that goes with having only discourse-friendly linkages with others." (Pettit, 2001, p. 103). Since, for Pettit, being free, and hence responsible, is ineliminably relational, it is not amenable to a purely individualistic interpretation. However, Pettit curtails the contextuality of this position in dramatic ways, and it is worth wondering whether they are merited. For example, Pettit claims that it is the *concept* of freedom that is perspective-dependent. The *property* of being free – i.e., the phenomenon picked out by our concept – and hence of being responsible, is not perspective-dependent (Pettit, 2001, p. 28). If "perspective" is taken, as I have taken Strawson's remarks about the participant perspective, as equivalent to "context", then, on Pettit's view, being responsible is not contextually individuated. By contrast, the position defended here clearly makes being responsible itself context-dependent.

²⁴ Wallace comes close to providing the necessary argumentation in Sections 5.2 and 6.1 of his book. These sections are on various sorts of psychological phenomena and their role in either excuses or justifications of behavior. However, these sections are inconclusive on the present matter because they are not directed explicitly at it. Examining these sections would take us too far from the present topic, so I will not pursue them.

It is not entirely clear to me why Pettit curtails the externalism of his position, but I suspect it is because he mainly assumes the relevance of metaphysical individualism to the understanding of moral responsibility. This is illustrated in Pettit's discussion of a "conundrum" that supposedly comes with the conceptualization of freedom as fitness to be held responsible. The conundrum is generated by the purported recursive nature of responsibility:

Suppose that I am responsible for an action. Presumably this will be so because that action is under the control of some other factor in me: say, my particular beliefs and desires. The recursive character of responsibility appears in the fact that, by ordinary intuitions, this means that I must be responsible in turn for those beliefs and desires. . . . But if I am responsible for the beliefs and desires that are in control of the action, then presumably I am responsible for them in virtue of their being under the control of some further factor still in my make-up: say, my habits of forming and revising beliefs and desires. . . . And so on, it appears, indefinitely. (Pettit, 2001, pp. 10–11).

Let's note two things about this passage. The first is the assumption of individualism. To assume that one is responsible because of the ability to control something is to recast a pragmatically cast competence in a one-to-one fashion with an individualistically individuated ability. But we have seen reason to think that competences do not necessarily entail abilities in this fashion.

Second, the purported conundrum is independently dubious. Consider examples where responsibility and control come apart:

- (A) Consider an institutional example. A CEO is responsible for the activities of his/her corporation. However, these activities might well be produced through processes over which the CEO has, in some sense at least, no control. For instance, these activities might be produced through processes that are themselves generated by decisions that institutional subordinates have the authority to make themselves. Such decisions are out of the CEO's control, yet s/he is responsible for the activities of the corporation.
- (B) Consider a biological example. The neural (and other) processes that, in some sense, control, e.g., the trajectory of my hand as I plunge a knife into my victim are out of my control, in that I play no direct role in these biological workings. Yet the suggestion that this might relieve me of responsibility for killing my victim is not even remotely tempting.

In short, Pettit's argument for an individualistic rendering of the pragmatically cast competences that constitute being responsible is not compelling. This position should be resisted along with the others examined here.

5. CONCLUDING REFLECTIONS: THE DIFFERENCE STRAWSON SHOULD MAKE

To recap: at the beginning of inquiry, both externalistic and individualistic approaches are available for explaining what it is to be morally responsible. We have some *prima facie* reason to think of it as a competence_E, and Strawson's well-known account of attributing moral responsibility gives us more principled reason to pursue the externalistic approach. Despite this, contemporary work on moral responsibility is overwhelmingly individualistic. In fact, the individualistic approach is largely assumed: it is very difficult to find explicit arguments in its favour in the work of contemporary theorists on this topic. Besides being problematic in itself, this assumption leads to both gaps and errors in argumentation. Overall, both the subtle power of a Strawsonian account of being responsible and the corresponding methodological burden on those who would offer a different sort of account of responsibility should be clear. I will conclude with reflections on some differences the acceptance of a Strawsonian approach to moral responsibility should make.

5.1. General – The Role of Moral Psychology

Although I have characterized contemporary work on moral responsibility as variously individualistic, the overwhelmingly obvious thing about these positions is that they seek *psychological* conditions of being morally responsible. On these sorts of view, the issue of moral responsibility is a moral psychological one. In particular, these theorists conceive of their task as to provide general positive moral psychological conditions of moral responsibility.²⁵

The Strawsonian account of responsibility demands a rethinking of the role of moral psychology. Instead of relatively *a priori* theorizing about the universal psychological conditions of moral responsibility, an *a posteriori* and locally contingent approach to psychology is appropriate. What is needed is what might be called *psychological anthropology*: descriptions should be made of culturally, geographically, and historically specific practices of deploying the reactive attitudes. Psychological theorizing about how the abilities that constitute moral responsibility in a particular setting can follow such descriptive work. Such work will exemplify Wilson's *integrative synthesis*, discussed in Section 4.2 above. It is likely that such theorizing will be largely explicitly local. However, depending on the psychological patterns revealed by such work, it might be possible

²⁵ Paul Russell makes this claim explicitly, with regard to how properly to evaluate the Pessimist's metaphysical commitments (1992, pp. 300–301).

to tentatively propose universal psychological attributes of moral responsibility. Such universals would be framed on this empirically informed, *a posteriori* foundation, not posited right away at the beginning of principled inquiry into this phenomenon.

There is another implication for moral psychology. Besides recasting the task of providing a positive account of the psychology of moral responsibility, the Strawsonian externalist view of moral responsibility makes the task of cataloguing and explaining phenomena and conditions that *defeat* moral responsibility more important than it currently is. Indeed, this is arguably where extant accounts of moral responsibility have made their lasting impression. The reason for this is that it is possible to identify such defeating conditions accurately even if one is mistaken about the general nature of moral responsibility, or about the implications to draw from the defeating conditions for the task of providing a positive account of being responsible. As transformative of contemporary work on moral responsibility as this might be, the investigation of defeating conditions is already part of post-Strawsonian work on responsibility – we saw this in all of the positions considered here – so the adoption of a Strawsonian approach does not entail the rejection of all work in recent approaches.

5.2. Specific – An Example

As an example of how anti-individualism both transforms thought about moral responsibility yet preserves important insights at the same time, consider the following. Wallace notes approvingly that Susan Wolf has remarked that there is a certain depth to our ascriptions of moral responsibility that goes beyond both our identification of a person as beginning a causal chain that resulted in the behavior in question and our assessment of the moral worth of the behavior (Wallace, 1994, p. 52; Wolf, 1990, p. 4). Wallace attempts to account for such depth in terms of assessment and appreciation of agents' rational powers of reflection and action in accordance with reflection. Pettit and Fischer and Ravizza offer similar individualistic features to ground moral responsibility, so it is fair to see their positions as explicating such depth in similar ways. By contrast, the Strawsonian position developed here implicitly suggests another way to account for this aspect of attributions of moral responsibility. Besides causal contribution and the worth of the behavior, ascriptions of moral responsibility either implicitly or explicitly relate a person to his/her social context. They serve to highlight a status an agent has by virtue of taking part in certain interpersonal practices. The metaphor of "depth" suggests that we should look more closely into the person, underneath the skin as it were, but we should always be careful about how literally we interpret metaphors. Alternative metaphors of "range" or "focus" or even "weight" might dissuade theorists

from hastily assuming an individualist interpretation. Overall, Strawson's account of moral responsibility should make us rethink the connection between individuals, social contexts, and what is "deeply" important about our practices of holding each other morally responsible and the concepts characteristic of these practices.

REFERENCES

- Berofsky, B., *Liberation From Self: A Theory of Personal Autonomy*. Cambridge: Cambridge University Press, 1995.
- Brooks, R., Intelligence Without Representation. *Artificial Intelligence* 47 (1991), pp. 139–159.
- Clancey, W., *Situated Cognition*. Cambridge: Cambridge University Press, 1997a.
- Clancey, W., The Conceptual Nature of Knowledge, Situations, and Activity, in Feltovich et al. 1997b.
- Clark, A. and Chalmers, D., The Extended Mind. *Analysis* 58.1 (January, 1998), pp. 7–19.
- Fischer, J.M., *The Metaphysics of Free Will: An Essay on Control*. Oxford: Blackwell Publishers Ltd., 1994.
- Fischer, J.M., Recent Work on Moral Responsibility. *Ethics* 110 (October, 1999), pp. 93–139.
- Fischer, J.M. and Ravizza, M. (eds.), *Perspectives on Moral Responsibility*. Ithaca: Cornell University Press, 1993.
- Fischer, J.M. and Ravizza, M., *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge: Cambridge University Press, 1998.
- Frankfurt, H., Alternate Possibilities and Moral Responsibility. *Journal of Philosophy* 68 (1969), pp. 5–20.
- Hurley, S.L., *Consciousness in Action*. Cambridge, Massachusetts: Harvard University Press, 1998.
- Hutchins, E., *Cognition in the Wild*. Cambridge, MA: MIT Press, 1995.
- LaFollette, H. and Shanks, N., *Brute Science: Dilemmas of Animal Experimentation*. London: Routledge, 1996.
- Paprzycka, K., False Consciousness of Intentional Psychology. *Philosophical Psychology* 15(3) (September, 2002), pp. 271–295.
- Pettit, P.A., *Theory of Freedom: From the Psychology to the Politics of Agency*. Oxford: Oxford University Press, 2001.
- Pinker, S., *The Language Instinct*. New York: Harper Perennial, 1994.
- Rowlands, M., *The Body in Mind: Understanding Cognitive Processes*. Cambridge: Cambridge University Press, 1999.
- Russell, P., Strawson's Way of Naturalizing Responsibility. *Ethics* 102(2) (January 1992).
- Scheman, N., Feeling Our Way Toward Moral Objectivity, in L. May, M. Friedman, and A. Clark (eds.), *Mind and Morals: Essays on Ethics and Cognitive Science*. Cambridge, Massachusetts: The MIT Press, 1996.
- Sneddon, A., Towards Externalist Psychopathology. *Philosophical Psychology* 15(3) (September, 2002), pp. 297–316.
- Strawson, P.F., Freedom and Resentment, in *Freedom and Resentment and Other Essays*. New York: Methuen, 1974.
- Strawson, Way of Naturalizing Responsibility. *Ethics* 102(2) (January, 1992), pp. 300–301.

- Wallace, R.J., *Responsibility and the Moral Sentiments*. Cambridge, Massachusetts: Harvard University Press, 1994.
- Watson, G., Reasons and Responsibility. *Ethics* 111 (January, 2001), pp. 374–394.
- White, S.L., Rationality, Responsibility, and Pathological Indifference, in O. Flanagan and A.O. Rorty (eds.), *Identity, Character, and Morality*. Cambridge, Massachusetts: The MIT Press, 1990.
- Wilson, R., Wide Computationalism. *Mind* 103 (1994), pp. 351–372.
- Wilson, R., *Cartesian Psychology and Physical Minds: Individualism and the Sciences of the Mind*. New York: Cambridge University Press, 1995.
- Wilson, R., Two Views of Realization. *Philosophical Studies* 104(1) (May, 2001), pp. 1–31.
- Wilson, R., *Boundaries of the Mind: The Individual in the Fragile Sciences – Cognition*. New York: Cambridge University Press, 2004.
- Wolf, S., *Freedom Within Reason*. New York: Oxford University Press, 1990.

Department of Philosophy
University of Ottawa
70 Laurier East Street
P.O. Box 450, Stn. A
Ottawa, Ontario, Canada K1N 6N5
E-mail: asneddon@uottawa.ca