

Nietzsche and Murdoch on the Moral Significance of Perceptual Experience

Paul Katsafanas

Forthcoming in *European Journal of Philosophy*

Note: this is the penultimate draft, not the final version

Abstract: This paper examines a claim defended by an unlikely pair: Friedrich Nietzsche and Iris Murdoch. The claim is that perceptual experience itself—as distinct from perceptually based judgments and beliefs—can be morally significant. In particular, Nietzsche and Murdoch hold that two agents in the same circumstances attending to the same objects can have experiences with different contents, depending on the concepts that they possess and employ. Moreover, they maintain that this renders perception an object of moral concern. This paper explicates these claims, examines the way in which we might distinguish between better and worse perceptual experiences, and argues that if some version of the Murdochian/Nietzsche claim is accepted, then certain influential approaches to moral epistemology and agency must be rejected.

I want to work with an idea from Nietzsche:

All experiences are moral experiences, even in the realm of sense-perception. [*Es gibt gar keine anderen als moralische Erlebnisse, selbst nicht im Bereiche der Sinneswahrnehmung.*] (GS 114)ⁱ

That's an odd claim. What could it mean to say that perceptual experiences are moral? How could a perceptual experience have moral significance? Sabina Lovibond offers one analysis of this line from Nietzsche. She writes,

in other words, *all reports on experience* are morally significant, because our use of the particular words in which they are couched constitutes an act with a definite moral character; it exhibits our commitment to the scheme of values implicit in that vocabulary. (Lovibond 1983: 118; emphasis added)

Lovibond's claim is perfectly sensible. I have a perceptual experience: Bill, who is conducting a job interview, puts his arm around the interviewee's waist. I report this experience, either to myself or to others, as follows: Bill is flirting with the interviewee. But I might have described it differently: Bill is harassing the interviewee. If the first description is salient to me, that indicates my commitment to one set of values; if the other is salient, another set. Or, another example: I hear Ralph tell a joke involving racial stereotypes. Whether I report this experience as Ralph's being funny or Ralph's being offensive reveals something about my values. So it's easy to agree with Lovibond's claim that reports of experiences can be morally significant: the judgments that I make on the basis of perceptual

experience, including judgments about how to report the contents of my perceptions, are ethically fraught.

But Lovibond's gloss on Nietzsche's claim is inaccurate. Nietzsche doesn't say that *reports on experience* are morally significant. He says that *experience itself* is morally significant. Indeed, he emphasizes this distinction by using the word *Sinneswahrnehmung*, sense-perception. *Wahrnehmung*, usually translated as perception, could also mean cognition or conception, but *Sinneswahrnehmung* cannot. So it's clear that Nietzsche wants to say that perceptual experience itself can be morally significant. This is a much more difficult idea. We can understand how commitment to racist values might lead someone to hear a joke involving stereotypes and form a judgment that the joke is entirely appropriate. But how could experience itself be morally significant? What would that even mean?

In this paper I'll try to make sense of the idea that perceptual experience can be morally significant. Initially, that idea looks decidedly odd. It's tempting to say that two agents in the same circumstances attending to the same objects from the same perspective will have the same perceptual experiences. The important moral differences between the two agents will emerge later, in their judgments, in their emotions, in their desires, in their actions. But the perception itself will be morally neutral.

That is the view I wish to oppose. I am not alone in this: some virtue theorists claim that one's ability to perceive morally significant reasons is contingent on the possession of particular character traits (McDowell 1979).ⁱⁱ Related claims arise in metaethics, with several competing accounts of the way in which we might perceive moral facts or properties. Thus, Jonathan Dancy argues that we can perceive instantiations of moral properties (Dancy 2010); Robert Audi defends the view that we perceive moral phenomenal properties, so that we sometimes have 'a phenomenal sense of wrongdoing integrated with our perceptual

representation of the wrong-making facts' (Audi 2010: 90); Jennifer Church contends that certain perceptual experiences are infused with imagination, so that, for example, perceiving a human being *as a person* involves synchronically imagining and synthesizing certain actual and possible interactions with that person (Church 2010). Although they have different commitments, these accounts agree that perception is morally relevant insofar as it attunes us to the morally relevant features of the environment.

This essay argues that the burgeoning literature on evaluative perception has overlooked a crucial type of cognitive impact on perception. I will argue that *concepts* penetrate immediate forms of perception in morally significant ways.ⁱⁱⁱ According to the views that I'll explore, perception doesn't just attune us to important features of the environment or reveal otherwise hidden moral facts; perception *constitutes* the perceived environment in important ways.

I articulate this account by drawing on Iris Murdoch and Friedrich Nietzsche. Though they agree on little else, both of these thinkers accept the following claim: two agents in the same circumstances attending to the same objects can have experiences with different contents, depending on the *concepts* that they possess and employ. In other words, Murdoch and Nietzsche argue that perceptual experience is *cognitively penetrable* in a particular way, and they maintain that this renders perceptual experience an object of *moral concern*. Sections One and Two explain Murdoch's and Nietzsche's views, in turn. Section Three asks how we distinguish between better and worse perceptual experiences, whereas Section Four considers how we take steps toward the attainment of the better perceptual experiences. With these accounts at hand, Section Five investigates the moral significance of perceptual experience. I argue that if we accept some version of the Murdochian/Nietzschean theory of perception, then traditional accounts of moral epistemology and agency must be rethought.^{iv}

1. Murdoch's account of perception

I'll begin with a brief analysis of Murdoch's account. Murdoch endorses all of the following claims:

- (1) Two agents attending to the same circumstance can have dramatically different perceptual experiences, depending on *the concepts they possess* and *their understandings of these concepts*.
- (2) There are certain facts that are perceptually unavailable to individuals, depending on the concepts they possess.
- (3) There is a *correct* way of viewing a given situation.
- (4) The primary obstacle to correct vision is selfishness.
- (5) These differences in perceptual experience are sometimes sufficient, just on their own, to explain differences in agents' actions.

These claims aren't completely transparent in Murdoch's work; her writing tends to be highly compressed and poetic, making explication no straightforward task. I won't attempt to establish that my reading of Murdoch is beyond question; rather, I'll give some indication of why this is a plausible reading. In this section, I consider the first two claims; I address the remaining claims in Sections 3 and 4.

Murdoch's underlying idea is that visual experience is conceptual: when we see something we see it *as* something, and what we see it *as* is influenced by our conceptual repertoire. I'll introduce Murdoch's account by considering her most famous example, which

involves a transition from seeing someone as ‘bumptious’ and ‘tiresomely juvenile’ to seeing her as ‘gay’ and ‘delightfully youthful’ (Murdoch 1997: 317). I quote the passage at length:

A mother, whom I shall call M, feels hostility to her daughter-in-law, whom I shall call D. M finds D quite a good-hearted girl, but while not exactly common yet certainly unpolished and lacking in dignity and refinement. D is inclined to be pert and familiar, insufficiently ceremonious, brusque, sometimes positively rude, always tiresomely juvenile.... Thus much for M’s first thoughts about D. Time passes, and it could be that M settles down with a hardened sense of grievance and a fixed picture of D, imprisoned (if I may use a question-begging word) by the cliché: my poor son has married a silly vulgar girl. However, the M of the example is an intelligent and well-intentioned person, capable of self-criticism, capable of giving careful and just *attention* to an object which confronts her. M tells herself: ‘I am old-fashioned and conventional. I may be prejudiced and narrow-minded. I may be snobbish. I am certainly jealous. Let me look again.’ Here, I assume that M observes D or at least reflects deliberately about D, until gradually her vision of D alters. If we take D to be now absent or dead this can make it clear that the change is not in D’s behavior but in M’s mind. D is discovered not to be vulgar but refreshingly simple, not undignified but spontaneous, not noisy but gay, not tiresomely juvenile but delightfully youthful, and so on. (Murdoch 1997: 312-3)

The situation as it would be described by external observers remains constant: it’s stipulated that D’s behavior does not change at all. However, M’s view of the situation undergoes a dramatic change: the behavior that she initially sees as pert, brusque, and rude she later sees as spontaneous, simple, and gay.

Notice that Murdoch is attributing quite sophisticated content to visual experience. She wants to say that I can have a visual experience with the content *that D is bumptious*. This is controversial. Everyone agrees that I can have a perceptually based *belief* that D is bumptious; but, some philosophers argue, I cannot have a perception with the content *that D is bumptious*. Instead, some philosophers contend that the only properties represented in visual experience are color properties, spatial properties, and perhaps a few others. Murdoch attributes a much richer range of properties to perception. I won't defend this richer view of perceptual content here; I simply want to register that Murdoch's account depends on it.^v While she doesn't explicitly distinguish between perceptions and perceptually based beliefs, she does make it clear that she is concerned with immediate visual experience rather than reports on or judgments about that experience, and she does attribute these higher-level concepts to it.

Additionally, notice that Murdoch quite deliberately focuses on certain *kinds* of concepts that figure in M's perceptual experience of D. At first, she sees D's behavior as pert, brusque, and so on; later, it is youthful and gay. Murdoch's choice of these concepts is deliberate: we can see how the same behavior might be described as tiresomely juvenile or delightfully youthful, undignified or spontaneous. These concepts pick out similar features but present them in different ways, and moreover link them to very different views of the whole person. What is at issue here is not just the picking out of some behavior and the attachment of a like or dislike, a boo or hoorah, to it; rather, the difference between seeing someone as juvenile or as youthful involves linking particular behavior to a set of assumptions about the person's character more generally. In other words, the difference between seeing someone's behavior as juvenile or youthful doesn't just involve seeing, say, the person's manner of speaking and then attaching a negative or positive evaluation to it; it

involves seeing that person's manner of speaking as attached to a more general way of life, and evaluating that way of life negatively or positively. The juvenile person is frivolous, superficial, flighty, rash; the youthful person is delightfully energetic, inquisitive, charmingly dilettantish. So the shifts that concern Murdoch aren't just changes of attitudes toward particular features, but changes in our way of viewing the whole person.

Murdoch expresses this idea rather enigmatically: she claims that 'moral differences look less like differences of choice, given the same facts, and more like differences of vision...We differ not only because we select different objects out of the same world but because we see different worlds' (1997: 82, emphasis added). What would it be to see different worlds? Elsewhere, Murdoch writes:

if the concept is withdrawn we are not left with the same situation or the same facts. In short, if moral concepts are regarded as deep moral configurations of the world, rather than as lines drawn round separable factual areas, then there will be no facts 'behind them' for them to be erroneously defined in terms of. (Murdoch 1997: 95)

The problem that Murdoch identifies is thinking of concepts as lines drawn around features of the world, rather than 'deep configurations of the world.' What might that mean? Here is a first stab. I think Murdoch is blurring two issues that it helps to hold apart: a claim about perceptions versus perceptually based judgments, and a claim about the interconnected nature of the concepts deployed in visual experience.

First, consider the perception/perceptually based judgment distinction. Suppose M at t1 and M at t2 have the same perceptual experiences, but form different perceptually based judgments or make different reports on this perceptual content at t1 and t2. Then it

would be quite right to say that concepts are lines drawn around separable factual areas: the perceptual content is constant, but the judgments formed on its basis vary at t1 and t2. M simply interprets the perceptual content differently at different times. By contrast, if the perceptual content itself differs at t1 and t2, then ‘we are not left with the same situation or the same facts.’ The experience itself differs at t1 and t2; and for that reason, it seems more appropriate to treat the concepts as ‘deep moral configurations of the world.’

Second, on a standard account of perception, it should be possible to list the features of D that M experiences at time t1 and the features she experiences at t2. Some of the features experienced at t1 will involve errors. For example, classifying D’s behavior as juvenile was a mistake; it’s really better classified as youthful. (By analogy, the vicious person might incorrectly classify empty bravado as courage.) On Murdoch’s preferred account, M doesn’t just classify the things she sees differently; she sees different things. There’s not one feature of D that used to look juvenile and now looks youthful; rather, M’s ‘total vision’ of D has shifted, so that all of D’s features fall into different relations, presenting themselves under different aspects. This is why she writes,

Communication of a new moral concept cannot necessarily be achieved by specification of factual criteria open to any observer (‘Approve of *this* area!’) but may involve the communication of a completely new, possibly far reaching and coherent, vision; and it is surely true that we cannot always *understand* other people’s moral concepts. (1997: 82)

Again, the point is that the visual experience itself changes, so that it is inaccurate to speak of seeing the same thing and categorizing it differently. Instead, Murdoch claims, the agent sees different worlds.

The image of ‘seeing different worlds’ can be misleading. After all, consider an analogous case: I can see a doughnut as a tasty treat or a disgusting mass of fried fat. The total vision of the object shifts, but it doesn’t seem right to say that I’m seeing different things, much less different worlds. Indeed, it might seem more accurate to say that my perception remains constant, but that different beliefs or emotional reactions become associated with it. This isn’t a deep problem for Murdoch, though; what her imagery is designed to evoke is simply the idea that *in some cases* perception shifts in a more profound way.

2. Nietzsche’s account of perception

I’ve given some indication of why Murdoch wants to say that different agents see different worlds. But the idea remains a bit hazy. In this section, I’ll try to spell out the view in more detail by looking at the way in which Nietzsche develops a version of it. Surprising as it may be to turn to Nietzsche for a precisification of a claim, I think that in this case it succeeds.

By way of orientation, let me start with a story about Kant. According to Kant, the way in which we cognize things depends on the ‘pure concepts of the understanding,’ or ‘Categories,’ which are—at its simplest—a specification of our most fundamental concepts and the relations among them. Kant takes these concepts and conceptual relations to be uniform and ineluctable for all rational agents (*Critique of Pure Reason* A80/B106 ff.; see also

Critique of Practical Reason 5:67ff.). In the nineteenth century, there was a widespread reaction to this story, which went something like this: while Kant thought that he could show that the Categories govern all self-conscious thought for all rational creatures, he was wrong. We can accept the basic claim that our minds impose a conceptual structure on thought, so that what we experience is determined not just by the nature of the world, but also by the nature of our minds. Yet we might reject the second component of Kant's account: the claim that this conceptual structure is fixed and ineluctable, uniform for all rational agents at all times. For example, we might adopt a more Hegelian view, according to which conceptual schemes vary historically, structuring the way that cognition unfolds at certain moments but later being supplanted by different schemes. So, the conceptual scheme governing Antigone's thought in Greek antiquity leads her to experience her choices in one way; the conceptual scheme governing Kant's thought in the nineteenth century leads him to experience his choices in a very different way. (Suppose, for example, that there is nothing inevitable about the experience of choice as undetermined by desire.)

With this Kantian backdrop in place, I think we can already grasp a way in which agents could, as Murdoch puts it, see different worlds. Suppose it's possible for sufficiently dissimilar agents to have their conscious experiences partly constituted by different sets of concepts; then the agents would, indeed, see different worlds. For example, one of Kant's pure concepts of the understanding is *causality and dependence*, and one of his pure concepts of practical reason is *perfect and imperfect duty*. Imagine a creature that cognizes things without seeing them as causally conditioned; or, more intelligibly, imagine an agent that reasons practically while lacking any concept of perfect and imperfect duty. These agents would have experiences sufficiently dissimilar to us that it would make sense to speak of them as seeing different worlds.

Much has been written on Nietzsche *perspectivism*. I'd argue that it is best understood along the lines sketched above: there is no one set of concepts and conceptual relations that we necessarily impose on experience. More precisely, the view can be reconstructed as follows:

- (A) Our conscious experience is structured by the systems of concepts that we employ. Two agents in the same circumstances attending to the same objects would have experiences with different contents, depending on the concepts that they possess.
- (B) These conceptual systems include not only classificatory concepts, but also evaluative ones. Even some of our most fundamental concepts, such as *matter*, are included here. So, too, are standards governing interconceptual relations and permissible patterns of inference.
- (C) These conceptual systems change over time.
- (D) Particular conceptual systems can be assessed as better and worse than one another. However, there is no way of assessing one as best.

I'll illustrate these claims with an example. I won't *defend* the attribution of these claims to Nietzsche, for two reasons. First, the above reading is relatively uncontroversial among Nietzsche scholars. Second, Nietzsche's texts resist quick exegesis: given his style of writing, decisively establishing that he endorses any particular claim requires a great deal of textual work. That would take us too far afield.

I introduced Murdoch's account by focusing on her example of the individual first perceived as bumptious and then perceived as delightfully youthful. The examples that

concern Nietzsche are far more complex—so complex that they are not initially recognizable as accounts of perception. Consider the conceptual shifts that Nietzsche describes in works such as the *On the Genealogy of Morality*. As is well known, Nietzsche argues that ancient and modern moralities differ in their evaluations: what the ancients labeled good (strength, conquest, power, health, beauty, and so on) the early Judeo-Christian system labels evil; what the ancients labeled bad (commonness, ordinariness, humility, weakness, and so on) the early Judeo-Christian system labels good. However, it's important to note that while the ancient and modern moral systems employ different values, they differ not only in these affective and evaluative respects, but also in the ways in which they conceptualize agency, the self, freedom, and responsibility. These are complex points, so I'll linger for a bit on just one aspect.

Nietzsche is at pains to emphasize the way in which one conception of agency supplanted another:

Just as the common people separates lightning from its flash and takes the latter to be a *deed*, something performed by a subject, which is called lightning, popular morality separates strength from the manifestations of strength, as though there were an indifferent substratum behind the strong person which had the *freedom* to manifest strength or not. But there is no such substratum; there is no 'being' behind the deed, its effect and what becomes of it; 'the doer' is invented as an afterthought,—the doing is everything. (GM I:13)

Here, Nietzsche draws attention to the way in which our modern conception of agency treats the agent as distinct from what he does. At its most extreme, this is the libertarian

conception of free will, which holds that the agent is to be identified with a characterless and utterly undetermined capacity for choice. The ancient conception, Nietzsche suggests, does not recognize this distinction: the agent's character is given by the nature of the agent's action, by what the agent actually does. The agent of antiquity does not take himself to be distinct from his past, his social relationships, and his community, nor does he see his actions as things that might or might not express his character; put simply, he *is* his deeds and relationships. (These accounts of agency obviously require far more explication than I can provide here.)^{vi}

Nietzsche's concern is not whether any particular philosopher (or non-philosophical person, for that matter) explicitly adopts just these thoughts about agency. Rather, he is concerned with the way that a person who internalizes something like this sense of agency will experience his action. To the extent that I tacitly adopt the modern conception of agency, I will tend to experience my choices as undetermined by and unreflective of my character. I will tend to see punctual moments of conscious choice as of overriding importance. I will tend to see myself as self-defining and isolated from my environment.

So Nietzsche's point is not just that our detached, abstract reflections on agency vary with changes in conceptual systems. It's true, of course, that Aristotle's philosophical reflections on agency differ from Kant's; that's hardly a surprise. But what is surprising is that these changes in conceptual systems, which at first seem to be merely a matter of detached reflection, in fact influence the agent's perceptual experiences. In other words, Nietzsche's claim is that Aristotle perceived persons and their actions differently than did Kant.

The way in which I experience my action depends upon the conception of agency that I tacitly adopt. For example, we moderns are perennially tempted to say, with Kant,

that our experience of deliberation commits us to taking ourselves to be free: we cannot engage in genuine deliberation without presupposing that we are free to determine our forthcoming action.^{vii} But Nietzsche's point is that just this sort of experience is historically contingent. It is not a datum, not a starting point, but a link in a complex story.

So some of the conceptual shifts that interest Nietzsche involve changes in valuations. But others involve concepts of agency, subjectivity, and selfhood (BGE 16). Others are deeper still: Nietzsche complains that our grammar misleads us into positing superfluous subjects, reifying causes and effects, and so on (BGE 20-22).

In short, Nietzsche is questioning basic components of our conceptual scheme: concepts such as *matter, agency, will, self, responsibility, good, bad*, and so forth. Note that modifying these any one of these basic concepts would force us to modify a cluster of related concepts. If *will* is modified, so too will be *deliberation, agency, reflection, person, responsibility*, and so on.

This explains why Nietzsche tends to engage in genealogies rather than analyses of particular concepts. Precisely because concepts are interrelated in these ways, we need to look at clusters or systems of concepts—in Nietzsche's terminology, *perspectives*—and critique them. Notice that I said 'clusters or systems.' We could think of perspectives as encompassing *some* or *all* of the concepts that the agent possesses. Though Nietzsche never makes it entirely clear how extensive perspectives are supposed to be, he seems to treat them as clusters of relatively basic, interrelated concepts. For example, there's a cluster of concepts concerning agency: agent, action, person, self, responsible, free, intentional, duty, obligation, and so on. These concepts are interrelated in the sense that a change to one will affect the others; they are relatively fundamental in the sense that, whereas concepts such as *bashful* or *humble* also concern agency, they seem to require reference to these other concepts.

3. Distinguishing perceptual improvements from regressions

So far, I've mentioned that both Murdoch and Nietzsche think we can evaluate the concepts deployed in perceptual experience, showing some as defective. But how, exactly? What makes one case of perception better than another, morally speaking? Here, an important divergence between the two thinkers emerges. Murdoch claims that we can decide between these perspectives solely on *epistemic* and *ethical* grounds. Nietzsche argues that we cannot: there is a way in which epistemic and—especially—ethical considerations are internal to particular perspectives, and thus cannot give us a neutral criterion of assessment for perspectives. Accordingly, Nietzsche wants to argue that *practical* considerations must play a role in the evaluation of perspectives. Let me explain.

Murdoch maintains that the varying conceptual schemes adopted by agents aspire to conformity with a putatively correct set of concepts (or, as she'll sometimes put it, a putatively correct *understanding* of concepts). There is a *correct* way of viewing a situation, which we can move toward. One of Murdoch's favorite claims is that the attentive individual experiences 'a deepening process, or at any rate an altering and complicating process' in her grasp of concepts (Murdoch 1997: 322). For example, 'we have a different image of courage at forty from that which we had at twenty' (1997: 322); what at first appeared to be 'daring of the spirit' or 'self-assertive ferocity' is later understood as 'a particular operation of wisdom and love' which 'would enable a man coolly to choose the labour camp rather than the easy compromise with the tyrant' (1997: 378). As our grasp of concepts deepens, we depart from the often superficial ways that those concepts are

employed in everyday discourse: our knowledge of the concept 'is something to be understood, as it were, in depth, and not in terms of switching on to some given impersonal network' (1997: 322).

This raises a complication: as we acquire 'a specialized personal *use* of a concept' (1997: 319), our private understanding of concepts may increasingly diverge from the public understanding. Thus, Murdoch writes, 'Each of us lives and chooses within a partly private, partly fabricated world' (1997: 199). The world is partly private because our visual experiences are partly constituted by the concepts we employ, and deepening understandings of concepts takes us away from the (typically superficial) understandings prevalent in our communities. Thus, the 'movement of understanding is onward into increasing privacy' (1997: 320-22).

In sum: for Murdoch, the agent improves her visual experience as she departs from the ordinary, communal interpretations of concepts and acquires a deeper insight into 'what is real.' Shifting her understandings of concepts, progressively clearing away distortions, she opens herself to the correct, undistorted view of the world. Thus, Murdoch tells us that goodness is 'a refined and honest perception of what is really the case, a patient and just discernment and exploration of what confronts one, which is the result not simply of opening one's eyes but of a certain perfectly familiar kind of moral discipline' (1997: 330).

We can see, then, that Murdoch ranks perspectives in epistemic terms. There are also suggestions that she ranks them by appeal to evaluative criteria: she frequently claims that *just* and *loving* perceptions are superior to selfish ones. For example, she writes, 'What M is *ex hypothesi* attempting to do is not just to see D accurately but to see her justly and lovingly' (1997: 317). There need not be an *error* in seeing D as bumptious; it may simply be

more just or loving to see her as youthful. I won't try to defend these ideas here; I merely want to register that Murdoch appeals to evaluative criteria in addition to epistemic ones.

Nietzsche dispenses with the idea that there is a single best or correct perspective. He has at least two reasons for this claim. First, Nietzsche believes that there is no unique set of basic concepts that is forced upon us by the nature of the world. As Lanier Anderson describes this point, 'the world does not determine answers to basic ontological questions independently of our variable conceptual assumptions' (1998: 1). For example, if physical theories take *force* as their basic concept, this ramifies throughout higher-level concepts and generates one possible perspective; if they take *matter* as basic, we get a different perspective.

Second, to determine what would count as a best perspective, we would have to appeal to some normative standard. But Nietzsche thinks that, when it comes to evaluating perspectives, there are multiple opposed normative standards that cannot be jointly realized. Consider a simpler case: biographies can be assessed as better or worse, but none can be unqualifiedly best. They can be more and less comprehensive, they can focus on different aspects of the same life, they can put aspects of that life into different relations, they can draw attention to new ways of understanding events within a life. They can be inadequate, superficial, sentimental, distorting. One might be recognized as best so far. But none is best, full stop. And, although we have no difficulty with the idea that we can compare the respective merits of two biographies, we have no independently specifiable criterion that we can employ to determine what it would even be for a biography to be best. (The most entertaining? The most comprehensive? The most illuminating? The simplest?) Nietzsche is claiming that the same point applies to our perspectives on the world.

This raises complicated questions about the way in which we adjudicate between competing perspectives, and I don't want to focus on that point here. Suffice it to say that we can sometimes see, from within a new conceptual scheme, that former schemes are inadequate either because they harbor internal tensions which the new scheme resolves, or because the new scheme alerts us to features of the world which the earlier scheme ignored or obscured. So some shifts between perspectives can be explained on epistemic grounds.

But this isn't always the case: changes in conceptual schemes needn't arise merely because of perceived errors or inadequacies in the earlier schemes. As the *Genealogy* sets out to show, conceptual schemes sometimes get a grip for practical reasons. The early Judeo-Christian perspective spreads, Nietzsche claims, primarily because it makes sense of the emotional inadequacies of agents. The slaves are attracted to this perspective not because it appeals to their refined rational sensibilities and philosophical proclivities, but because it enables them to interpret their own weaknesses and inadequacies as strengths: the slave's weakness,

thanks to the counterfeiting and self-deception of powerlessness, clothed itself in the finery of self-denying, quiet, patient virtue; as though the weakness of the weak were itself—I mean its *essence*, its effect, its whole unique, unavoidable, irredeemable reality—a voluntary achievement, something wanted, chosen, a *deed*, an *accomplishment*... [This] facilitated that sublime self-deception whereby the majority of mortals, the weak and the oppressed of every kind could construe weakness itself as freedom, and their particular mode of existence as an *accomplishment*. (GM I:13)

So Nietzsche wants to be realistic about the reasons for which conceptual changes occur. Often, it isn't in response to rational considerations, but to affects, desires, social circumstances, and so on. So, Nietzsche argues, it's a mistake to think that conceptual shifts have occurred primarily for epistemic reasons. And, he adds, it's a mistake to think that they *should* occur solely in response to epistemic considerations: a particular moral perspective might be superior not because it more accurately captures the facts, but because it makes possible a more attractive or more exemplary way of life. Or because it makes possible the affirmation rather than the rejection of life; this is one of Nietzsche's points in the *Gay Science*.

In other words, Nietzsche argues that there is no particular set of concepts that is rationally mandated. Some sets will contain internal contradictions, some will present impoverished views of their objects, but even when we clear away the egregiously conflictual and inadequate ones, we will be left with many others. So epistemic grounds alone won't provide a single answer. And this opens us to non-epistemic considerations involving, for example, the conceptual schemes that enable a form of self-affirmation or optimism rather than pessimism, or—to use Nietzsche's favored criterion—the conceptual schemes that enable us to maximally will power.^{viii}

So this, too, is a departure from Murdoch. Increased attention to the world might lead us to perceive inadequacies in our current conceptualizations. But it might also lead us to revise our conceptual repertoire. And it might do so not solely for epistemic reasons, but also for practical ones.

4. Pursuing improvements in vision

So we have competing ideas about whether there's a single best set of concepts and whether the grounds for assessing conceptual systems are purely epistemic. Let's now turn to a different question: how shifts in concepts are to be attained.

Murdoch focuses on the plight of the solitary individual; she presents *attentive looking*, self-scrutiny and honest introspection as the paths to correct vision. The M/D example illustrates this. 'M *looks* at D, she attends to D, she focuses her attention. M is engaged in an internal struggle' (1997: 317). As M looks carefully and clears away the distortions induced by jealousy and so forth, M's perception improves, and 'she sees D as she really is' (1997: 329). And all of this is a personal, private struggle.

Part of the reason for this is that Murdoch thinks the primary threat to adequate vision is simply egoism. She writes, 'Most of the time we fail to see the big wide real world at all because we are blinded by obsession, anxiety, envy, resentment, fear. We make a small personal world in which we remain enclosed' (1997: 14). To counteract this, 'to contemplate and delineate nature with a clear eye,' we must 'silence and expel self' (1997: 352). The 'intellectual ability to perceive what is true ... is automatically at the same time a suppression of self' (1997: 353). For the self is prone to 'fantasy, the proliferation of blinding self-centered aims and images' (1997: 354).

If that were true—if what cut us off from adequate vision were selfish motives—then we could, perhaps, counteract these motives by focusing on the world more clearly, as M does in the famous example. But is this an adequate account?

Nietzsche writes, 'whatever they may think and say about their "egoism", the great majority nonetheless do nothing for their ego their whole life long' (D 105). Think, here, of familiar facets of contemporary culture. It's almost a truism that many agents, in an attempt to pursue self-interest, end up pursuing exactly the opposite. Duped by the ways in which

issues are described and framed in their community, the individual who would most benefit from a governmental program such as universal health care or increased labor rights is precisely the one who most vehemently opposes it. Or, to use a more Nietzschean example: if it were true that conformity to contemporary morality undermines flourishing, while presenting itself as doing just the opposite, then we'd have another example of the way in which putatively egoistic pursuits end up harming the agent.

For Nietzsche, the idea that we are in a position to distinguish egoistic and unegoistic motives is facile. Our motives are presented and shaped under the aegis of concepts whose meanings and connections largely elude us. The true threats to adequate vision are the concepts that surreptitiously structure our experiences.

So there's a deep distinction between Murdoch and Nietzsche on moral significance, here. Though Murdoch often inveighs against the Kantian picture of morality as solely concerned with acts of individual will, she is still in the grip of the idea that morality concerns individual action and thought. Nietzsche, though, is much less interested in this. The social proliferation and reinforcement of inadequate conceptual schemes is the primary problem; egoism is something that arises only *within* these conceptual schemes, and is thus of secondary interest.

One intriguing feature of Nietzsche's account—shared, of course, with Hegel and others—is that investigation of experience's conceptual structure takes a historical form (see GM Preface). To see why, it helps to return, once again, to Kant. In Kant, the uncovering of conceptual structure takes a transcendental form: we ask what would have to be the case for conscious experience to be as it is, and we are thereby led to the Categories. But transcendental arguments would only reveal the necessary conditions of experience *as it is presently configured*. If Nietzsche were correct in arguing that experience changes (in

sufficiently profound ways) over time, then transcendental arguments would yield different results at different times.

The turn to history can reveal that apparently necessary features of experience are, in fact, local and contingent. Accordingly, history can be useful in helping us to understand whether the features that purport to be present in all conscious experience were, in fact the conditions of a local and contingent configuration of conscious experience. To answer that question, we would need to ask whether experience *could be* otherwise; and one of the easiest ways to do that is to ask whether it *has been* otherwise.

Why is this? Why can't I simply consider changing various concepts, asking what would be the case if I began to operate with (say) the archaic conception of agency instead of the modern one? The problem is that concepts are interconnected: as I mentioned above, a change in the concept of agency would resound throughout a host of other concepts. In trying to envision such a change, it's highly unlikely that I'd be able to anticipate, just in a moment of armchair reflection, what would be entailed by it.

This difficulty can generate illusions of fixity. Features of experience can appear to be fixed and necessary because of the interconnected nature of concepts. A change in a concept such as *agency* or *matter* would cascade throughout a series of related concepts. If these related concepts are held fixed, the target concept will also appear to be fixed. (For an example of this, consider the Strawsonian arguments about reactive attitudes [Strawson 1974]. Briefly, Strawson's argument is that concepts such as *responsibility* are non-optional for us because if we changed them we would be forced to abandon our distinction between, for example, acknowledging someone as a person and responding to them as a mere force of nature. Nietzsche's point would be that this presumes a fixed conception of what it is to respond to someone as a person. We might jointly shift the notions of personhood and

responsibility, as well as related concepts, and thereby come to a different conceptual scheme.)

Given the interrelated nature of concepts, in undertaking conceptual change we need to assess clusters or even whole systems of concepts. In theory, we could assess interconnected sets of concepts without history—we could consider hypothetical situations, being careful to let changes in one concept appropriately cascade through others. The danger, though, is obvious: in stepping back from history (or, more generally, from some kind of extended narrative), we will inadvertently hold fixed elements of our conceptual repertoire that in fact need to vary for the focal concept to vary.^{ix}

Aside from the interconnectedness of concepts, there is a second difficulty that plagues ahistorical investigations of perspectives. It's difficult to observe our perspectives directly because they are that *through which* we observe. For example, it's difficult to observe my perspective on agency by reflecting on my intuitions about agency, viewing particular cases of agency, thinking about my associations with it, and so on, precisely because these operations are all structured by my perspective on agency. I need to step back from that structuring.^x Consider, by analogy, the way that emotions influence perception. If I want to know how my love for someone is affecting my perception of her, I don't keep viewing her through it. I try to imagine how others view her; I try to distance myself from my typical emotional reactions to her; I try to think about how I viewed her before I loved her.

Or, for a more complex example of great ethical relevance: consider Weber's reflections on the way in which an ideal of *efficiency* gradually comes to pervade modern thinking (Weber 2002). It's become increasingly difficult for modern individuals to appreciate the idea of something's having non-instrumental value. Witness the modern debates about and defenses of the humanities (or the university more generally): again and

again, discussions are structured around the idea that we can defend the value of an education only by showing that it leads to some distinct, valued end (where this is typically taken to be a well-paying career). It can be extraordinarily difficult to entertain the increasingly alien idea that education might be valued for its own sake, that an enriching and broadening of the mind might not need to *produce* anything in order for it to be valued. It's easy to *state* these alternative ideas, of course; but the point I wish to make is that it is a struggle for many modern individuals to *think* in this way. To put it in Nietzschean terms: our current perspective, with its commitments to an ideal of efficiency, continually structures our ways of viewing the world and our habits of thinking, such that reflections on the possibility of non-instrumental value can be, for most individuals, only difficult reminders that are not put into everyday practice.

This, then, is the value that Nietzsche sees in genealogy: it can wrest us outside of our current perspective, forcing us to inhabit a different one. It can cultivate new affective reactions, present new classifications, foster new ways of organizing experience, show us new distinctions.

5. The moral significance of perception

The previous sections outlined Murdoch's and Nietzsche's accounts of perceptual experience. Let me summarize the argument. Both Murdoch and Nietzsche argue that the meanings of concepts are holistically determined, so that a change in one concept resounds throughout a whole system of concepts. In light of this, Nietzsche emphasizes that a given concept can seem ineluctable precisely because changing it would require changing a group or even a whole system of concepts. Moreover, because concepts structure thought and

experience, standing back from them is difficult. Murdoch and Nietzsche have different solutions to this difficulty: Murdoch thinks attentive looking and the purging of self-interest will be sufficient, whereas Nietzsche counsels a turn to history. However, it's plain that these approaches can be complementary: each may be effective in certain domains.

With these accounts at hand, we can now assess the idea that perceptual experience can be conceptually penetrated in morally significant ways. Let's begin by clarifying what would it be for perceptual experiences to have a moral character at all. Put simply, it would be for two agents in the same circumstances attending to the same objects from the same perspective to exhibit morally significant differences, based solely on the perceptual experiences that they are having.^{xi} So if we hold all else constant—situation, character traits—but vary the perception, this should make a moral difference.

As I mentioned in the introduction, standard accounts of evaluative perception recognize only two ways in which this could be so. First, there is an Aristotelian view according to which certain morally significant facts can be perceptually unavailable to individuals who lack the right character traits. Second, Audi and others have argued that we can directly perceive moral facts or properties, such as the wrongness of various acts.

However, the Murdochian/Nietzschean account yields two additional ways in which perception can have moral import. First, there is an epistemological point: moral theory is not independent from experience in the way that we might expect. Second, there is a point about agency: the target of evaluation and the strategies for moral improvement are to some extent transformed if we accept the Murdochian/Nietzschean view. I'll explain these points in turn.

5.1 Moral epistemology

Moral theory is standardly taken to have a particular form of independence from perception: data from perception is assumed to be unaffected by the moral theory and moral concepts with which one operates. Though this assumption is rarely made explicit, it informs a number of accounts. This is most vivid in accounts that attempt to move from descriptions of situations to moral judgments. Let me give just two examples. Frances Kamm claims that we should do moral philosophy by ‘present[ing] hypothetical cases for consideration and seek[ing] judgments about what may and may not be done in them’ (Kamm 1993: 7). We then attempt ‘to construct more general principles from these data’ (1993: 8). Though operating with a quite different picture of moral theory, Michael Smith makes an analogous point, writing, ‘It is agreed on nearly all sides that moral knowledge is relatively a priori, at least in the following sense: if you equip people with a full description of the circumstances in which someone acts, then they can figure out whether the person acted rightly or wrongly just by thinking about the case at hand’ (Smith 2004: 203). Both Kamm and Smith present the work of ethics as the movement from description to prescription: we are given some description of a scenario, and the work for moral theory consists in forming (and justifying) judgments about what is to be done.

That’s a tempting view. However, the Murdochian/Nietzschean account complicates it. An analogy may be helpful. In the early twentieth century, philosophers discussing science tended to assume that we could specify some set of theory-neutral observational terms and evaluate competing scientific theories in light of them. It’s now a familiar point that this won’t work: observational terms are theory laden. When we recognize the theory-ladenness of observational terms in science, we thereby recognize that we need a more complex epistemology. Just so in ethics. If perception were morally neutral,

we could appeal to it to develop moral theories and we could evaluate competing theories in light of it. But if the Murdochian/Nietzschean view is correct, this won't work. Let me illustrate this with two different examples.

First, consider Murdoch's M and D. In the quotations above, Kamm and Smith recommend that we begin with a full description of the case and then ask what moral conclusions follow. But Murdoch's point is that much of the work of morality has already been done at the stage that Kamm and Smith are taking for granted: the formulation of the 'correct description.' The way that the agent sees and describes the situation will be influenced by the concepts with which she operates. At an early stage, M's perceptual experience will present D as bumptious; at a later stage, as youthful. Different descriptions of and reactions to the situation will seem appropriate to her, depending on which perceptual experience she has. In short, we're not entitled to assume that there are theory- or concept-neutral descriptions of the situations that many moral philosophers take as starting points.

For a rather different example, consider Nietzsche's claim that perceptions of our own agency are conceptually penetrated. If this were right—if the experiences of agency vary depending on the concepts of agency that we embrace—then doubts would arise for theories that attempt to move from facts about our experience of agency to substantive ethical conclusions. Thus, the Kantian claim that autonomy is presupposed by deliberation from the first-person standpoint, as well as Darwall's claim that autonomy is presupposed by deliberation from the second-person standpoint (Darwall 2006, Chapters 9-10), operate by taking an experience of deliberation as fixed and analyzing its presuppositions. If these experiences of deliberation are rationally optional—if, for example, they are dependent on historically local conceptions of agency—then the Kantian and Darwallian arguments are

conditional: they show what follows if we operate with a certain concept of agency, but do not establish that there is any rational fault in operating with different concepts. (Obviously, these are difficult points and a number of responses are available. I am not claiming that the Murdochian/Nietzschean view constitutes a decisive indictment of the Kantian or Darwallian projects, but merely that it raises important concerns.)

While much more could be said, my hope is that these examples illustrate the way in which the Murdochian/Nietzschean account raises new concerns for moral philosophy. While the standard accounts of evaluative perception focus on the reliability of perception in detecting morally relevant features (à la Aristotle), or defend claims concerning the ability to perceive moral facts (à la Audi), the Nietzschean/Murdochian view is in some ways the opposite: it brings to light the untrustworthiness of perception and the way in which any perception of moral features is theoretically entangled. This necessitates a more complex account of moral epistemology.

5.2 Agency

Let's now turn to the point about action. Morality's goal is to influence action, prompting us to perform some actions and refrain from others. If perception can constrain or influence choice to a significant extent, then moral philosophy ought to be concerned with the agent's perceptual experiences.

By way of illustration, consider Murdoch. She claims that her theory diminishes the importance of punctual moments of choice. Take the situations that moral philosophers usually focus upon: to use some of Kant's examples, I need to determine whether I should make a lying promise in order to secure some cash, or I need to decide whether I should

cheat my customers in order to maximize profits. In these situations, the facts of the case are taken as given. But, as we saw above, this is problematic: what the agent takes the facts of the case to be is influenced by the concepts with which she operates. Murdoch argues that the way in which the agent sees the situation exerts a decisive influence on her actions. After all, 'I can only choose within the world I can *see*' (Murdoch 1997: 329); the way in which I conceptualize a situation restricts the range of possible choices. But, even stronger, the way in which I see the world sometimes *determines* my choice. As Murdoch puts it, 'at crucial moments of choice most of the business of choice is already over'. For 'one is often compelled almost automatically by what one can see' (1997: 35-6). When M sees D as brusque and rude, certain forms of behavior (shunning her, avoiding her, discouraging her son from continuing his relationship with her) seem rationally mandated. When M sees her as delightful and youthful, different forms of behavior seem mandated.

So Murdoch embraces an exceptionally strong claim about the relationship between vision and choice. She writes, 'If I attend properly, I will have no choices and this is the ultimate condition to be aimed at' (1997: 331). In other words, Murdoch wants to say that the way in which I conceive a situation often determines my decision (in the same way that arriving at a proper understanding of a geometric problem can compel the conclusion). This depends on two further theses: (1) that properly conceptualized situations of choice entail only one correct outcome; (2) that conceptualizations determine, rather than merely dispose, choice.

Although Murdoch defends these strong claims,^{xii} we can preserve much of the force of her argument while weakening the claims. We could argue that proper conceptualization is compatible with multiple correct outcomes, and that conceptualizations merely dispose choices. How damaging would these claims be for traditional ethical theories? The Kantian

focus on moments of choice, as well as the recent proliferation of work that focuses on reasons for action, tend to miss these more basic aspects of our relation to the world. If thinkers including Murdoch and Nietzsche are correct, the emphasis on first-personal experiences of choice, perceptions of reason, and so on is, for all its import, a focus on surfaces—surfaces determined by subtler aspects of our mental economies. An ethic aspiring to completeness, an ethic aspiring to show us how to *deepen* and *improve* the conceptual scheme through which we view the world, would take a very different form. It would enjoin an attentive look at our concepts, a skepticism about present experience, an eye toward history for enlightenment.

The proponent of this view can maintain that choices are important. Indeed, one way in which they are important is by shifting the clusters of concepts through which we view the world. Consider Murdoch's wonderful line: 'at crucial moments of choice most of the business of choice is already over'. We are now in a position to see why she distinguishes 'the business of choice' from the 'moment of choice': the extended engagement with elements of one's conceptual repertoire, the decision to look again at the way one has been perceiving the world, the efforts to achieve a more adequate vision of the world—all of this involves choice, all of this involves deliberation. What it does not involve is merely surveying a set of possible actions and picking one. It involves a retreat from the moment of choice and a greater engagement with the precursors for these moments.

Nothing precludes Kantians, utilitarians, and others from doing this. But their concerns are to some extent transformed if we accept the Murdochian or Nietzschean picture. After all, if something like the Murdochian or Nietzschean view is correct, then choice is largely determined by the agent's vision of the world, and the agent's vision of the world is largely determined by the set of concepts that she employs. An important way of

influencing agents' behavior, then, is by changing the set of concepts with which they operate.

So this, I think, is a second area in which perception has moral significance. To generalize, modern ethical theory tends to begin with the deliberating agent who considers reasons for and against various actions. But another view of ethics sees deliberation and judgment as downstream from a more immediate manifestation of the agent's ethical stance: if concepts penetrate perceptual experience, then ethics needs to address questions concerning which concepts we ought to operate with. For it is only in light of these concepts that deliberation begins.

6. Conclusion

I began this paper by claiming that, despite the recent proliferation of work on evaluative perception, an important form of cognitive impact on perceptual experience has been overlooked: the way in which concepts structure perceptual experience. I've discussed two different accounts of the way in which this might be so: Murdoch's quasi-Platonic account and Nietzsche's historical account. Murdoch discusses the way in which our perceptions are suffused by concepts, our grasp of which can be deepened and transformed through attentive reflection. Nietzsche offers a more complex version of this picture: first, rather than individual concepts we're concerned with sets or clusters of concepts, which interact with one another in unexpected ways; second, these clusters of concepts both contain and in some cases constitute evaluations; third, the concepts themselves are historically fluid; fourth, there's no one best set of concepts; fifth, practical concerns, not just epistemic ones, force these shifts in perspective.

More work remains to be done in determining whether these accounts are ultimately defensible and assessing the extent to which they are in conflict with more traditional forms of ethical theory. However, my central goal in this paper has been to articulate these accounts and give some indication of the way in which they complicate moral epistemology and accounts of agency.^{xiii}

Paul Katsafanas

Boston University

745 Commonwealth Avenue, Room 534

Boston, MA 02215

USA

pkatsa@bu.edu

References

- Anderson, R. L. (1998), 'Truth and Objectivity in Perspectivism', *Synthese*, 115: 1-32.
- Audi, R. (2010), 'Moral Perception and Moral Knowledge', *Proceedings of the Aristotelian Society Supplementary Volume*, LXXXIV: 79-97.
- Church, J. (2013), *Possibilities of Perception*. Oxford: Oxford University Press.
- Dancy, J. (2010), 'Moral Perception', *Aristotelian Society Supplementary Volume*, 84(1): 99-117.
- Darwall, S. (2006), *The Second Person Standpoint*. Cambridge: Harvard University Press.

- Kamm, F. (1993), *Morality, Mortality, Volume I: Death and Whom to Save From It*. Oxford University Press, New York.
- Kant, I. (1781 and 1787/1998), *Critique of Pure Reason*, Paul Guyer and Allen Wood (eds.). Cambridge: Cambridge University Press.
- (1788/1999), *Critique of Practical Reason*, Paul Guyer and Allen Wood (eds.). New York: Cambridge University Press.
- (1797/1996), *The Metaphysics of Morals*, Mary Gregor (ed.). New York: Cambridge University Press.
- Katsafanas, P. (2015), 'Fugitive Pleasure and the Meaningful Life: Nietzsche on Overcoming Nihilism?', *Journal of the American Philosophical Association*, 1(3): 396-416.
- (2016), *The Nietzschean Self: Agency, Moral Psychology, and the Unconscious*. Oxford: Oxford University Press.
- Lovibond, S. (1983), *Realism and Imagination in Ethics*. Minneapolis: University of Minnesota Press.
- McDowell, J. (1979), 'Virtue and Reason', *The Monist*, 62: 331-50.
- Moran, R. (2012), 'Iris Murdoch and Existentialism', in J. Broackes (ed.) *Iris Murdoch, Philosopher*. Oxford: Oxford University Press, 181-96.
- Murdoch, I. (1997), *Existentialists and Mystics: Writings on Philosophy and Literature*. New York: Penguin.
- Nietzsche, F. (1881/1982). *Daybreak*, trans. R.J. Hollingdale. Cambridge: Cambridge University Press.
- (1882 and 1887/1974). *The Gay Science*, trans. W. Kaufmann. New York: Random House.
- (1886/1967). *Beyond Good and Evil*, trans. W. Kaufmann. New York: Random House.

- (1887/1967). *On the Genealogy of Morals*, trans. W. Kaufmann and R.J. Hollingdale. New York: Random House.
- Setiya, K. (2013), 'Murdoch on the Sovereignty of Good', *Philosopher's Imprint*, 13(9): 1-21.
- Siegel, S. (2010), *The Contents of Visual Experience*. Oxford: Oxford University Press.
- Smith, M. (2004), 'Moral Realism', in *Ethics and the A Priori*. Cambridge: Cambridge University Press, 181-207.
- Strawson, P. (1974), 'Freedom and Resentment', in *Freedom and Resentment and Other Essays*. London: Methuen.
- Väyrynen, P. (2008), 'Some Good and Bad News for Ethical Intuitionism', *Philosophical Quarterly*, 58(232): 489–511.
- Weber, M. (2002), *The Protestant Ethic and the Spirit of Capitalism*. New York: Penguin.

ⁱ I cite Nietzsche's texts using the standard English abbreviations of their titles: BGE is *Beyond Good and Evil*, D is *Daybreak*, GM is *On the Genealogy of Morality*, and GS is *The Gay Science*.

ⁱⁱ One complication with these views is that it is often unclear whether virtue theorists mean to say that we *literally perceive* reasons, or whether they simply use the language of 'seeing' and 'perceiving' loosely, to emphasize the immediacy and uncodifiability of an agent's grasp of reasons.

ⁱⁱⁱ Some authors, including Dancy (2010) and Väyrynen (2008), discuss a version of this possibility in passing, but do not offer sustained analyses of its import.

^{iv} Although I focus on Nietzsche and Murdoch, other philosophers explore related ideas. Weil, Sartre, Heidegger, and Husserl spring to mind. However, Nietzsche and Murdoch are an interesting pair because, as I'll explain below, their disagreements on topics in ethics and metaphysics lead to competing accounts of a shared point: that concepts impact perceptual experience in morally significant ways. Seeing how Murdoch and Nietzsche articulate this point, despite their substantial disagreements in ethics and metaphysics, helps to clarify the conceptual space.

^v For some discussion of the view that perception can represent the sorts of properties that concern Murdoch, see Siegel (2010), who calls this the 'Rich Content View.'

^{vi} For more extended analyses of these matters, see Katsafanas (2016).

^{vii} Kant claims that human choice 'can indeed be *affected* but not *determined* by impulses... *Freedom* of choice is this independence from being determined by sensible impulses' (*Metaphysics of Morals* 6:213-214). In the second critique, Kant provides a 'table of the categories of freedom'. These are the most fundamental concepts that are employed in

agency; they are the pure concepts of practical reason. See the second chapter of the *Analytic of Pure Practical Reason*.

^{viii} For further discussion of these points, see Katsafanas (2015) and Katsafanas (2016).

^{ix} So far, I've emphasized the difficulty with imagining changes to foundational concepts such as *agency* and *matter*. The same points apply to less fundamental concepts. For example, Nietzsche argues that *compassion* is connected to particular understandings of *suffering*, *selfhood*, *achievement*, and *flourishing*.

^x As Nietzsche is keen to emphasize, stepping back from one perspective just is coming to inhabit a new perspective. We should not imagine that I step back from my understanding of agency by achieving some detached, impartial view; instead, we do so by governing our thoughts with a different model of agency.

^{xi} Here I'm using 'perspective' in the ordinary, non-Nietzschean sense.

^{xii} See Setiya (2013) and Moran (2012) for helpful discussions of this point.

^{xiii} For comments on drafts of this paper, I am indebted to audiences at Boston University, University of Toronto, University of Pittsburgh, University of Chicago, and Harvard University. Thanks especially to Kathryn Lindeman for commenting on this paper at Pittsburgh.