THE DUALIST

Stanford's Undergraduate Journal of Philosophy

This issue is dedicated to Esbjörn Svensson.

THE DUALIST

Volume XVII ■ Spring 2012

Department of Philosophy
Stanford University

TABLE OF CONTENTS

Grounding Normativity in a Desire-Based Theory of Reasons

ARTHUR LAU
Stanford University

Abstract

In his book *Slaves of the Passions*, Mark Schroeder offers a theory of reasons providing both an analysis of what it is for an agent to have a reason to act in some way and an analysis of how the relative weights of an agent's reasons are determined. This paper investigates an important tension between these two components of Schroeder's account, one that arises because, under the theory, a given agent's reasons are grounded in that agent's own possibly idiosyncratic desires, while the weights of those reasons depend crucially on facts about other agents' psychologies. In particular, the paper considers the challenge of grounding the normative significance of the psychologies of other agents in fixing the weights of one's reasons. Two initially appealing proposals for grounding the normative force of others' psychologies, based respectively on Schroeder's ideas about primary goods and Allan Gibbard's concept of conversational demands, are then discussed and ultimately rejected. Finally, by viewing these proposals as exemplifying more general potential approaches to the problem, the conclusion suggests that the integrity of Schroeder's theory can be established only by reconciling several possibly incompatible factors: the significance of the agent's psychology, the significance of others' psychologies, and the principles governing normative discourse.

1. Introduction

A theory of reasons that is Humean in its basic conception can be constructed in various ways, yet as Mark Schroeder observes in his book *Slaves of the Passions,* every version of such a Humean theory ultimately pursues a common aim: to ground the reasons that a person has in some element of the person's psychology (1-6). In Schroeder's formulation of this view, which he calls Hypotheticalism, a reason for an agent to act in some way is, roughly, a consideration explaining why that action would increase the likelihood that one of the agent's desires is satisfied (59). His analysis of the relative weighting of reasons, however, departs strikingly from the conventional Humean model in holding that one set of reasons is weightier than another when the balance of those reasons shared by all agents engaged in deliberation favors placing more weight on the former (134-39). This paper considers a challenge to Hypotheticalism that exploits a certain tension between these two central elements of the theory by questioning how Schroeder could ground the normative significance of other agents' psychologies in determining the weight of a particular agent's reasons. Two main proposals, one utilizing Schroeder's idea of primary goods and the other based on Allan Gibbard's notion of conversational demands, will be offered to respond to this challenge. Both proposals will be developed within the model of a "proleptic mechanism" that is suggested by Bernard Williams in his article "Internal Reasons and the Obscurity of Blame" (41-42). As this paper endeavors to show, nevertheless, these proposals must in the final assessment prove inadequate, in ways related to Williams's thesis of the distinctiveness of a reasons claim, Gibbard's observation that normative judgments involve an element of endorsement, and Michael Smith's concerns about normativity and arbitrariness.

2. Williams, Smith, and Schroeder on Reasons for Action

To outline the broader context in which Schroeder presents his theory of reasons, it will be useful to consider some of its predecessors, in particular Williams's more classically Humean theory of "internal" reasons and Smith's dispositional theory of value. As Williams writes in "Internal and External Reasons," if a person has a reason to perform some action, then that action must be related to the person's "subjective motivational set" through a rational deliberative route (101-5). In this subjective motivational set would be included not only desires in the traditional sense, but also the agent's values, patterns of emotional response, loyalties, projects, and other commitments (105). Williams's argument relies on the general thesis that reasons must at some level be capable of motivating an indi-

vidual to act, and that it must therefore be possible to refer to one's reasons in an accurate explanation of one's actions (106). These internal reasons, however, bear normative and not merely explanatory significance, because they must be connected to the subjective motivational set through a process of rational deliberation that takes account of all relevant facts and does not essentially involve any false premises (102-3).

Smith's dispositional theory of value, originally articulated in his paper "Internal Reasons," incorporates an altered version of Williams's notion of a rational deliberative pathway between an action and an agent's motivational set. According to Smith, a person A has a reason to act in a certain manner under circumstances C just in case A would desire that he or she act in that way in C if A were fully rational (17-18). While accepting that rational deliberation presupposes the absence of false beliefs and the presence of all relevant true beliefs, he argues that a fully rational agent would have, in addition, a systematically justifiable set of desires (20-23). Systematic justification aims to produce a desire set that is maximally coherent and unified, one for which the agent could be confident that its elements would also be considered justifiable by other rational agents engaged in deliberation (25-27). Smith contends that the truth of reasons claims requires that the desire sets of all fully rational agents would finally converge in the process of systematic justification; if reasons exist, they must be common to all agents (33-34).

Schroeder's theory, unlike those of Williams and Smith, provides both a statement of the conditions under which a person has a reason to perform a given action, and a method for determining when the reasons in favor of an action are weightier than the reasons opposing it. As he argues, a reason r for an agent to perform some action is a proposition such that the agent has some desire with object p, and the truth of r forms part of the explanation of why the agent's acting in that way promotes p (*Slaves* 59).[1] An action promotes p, moreover, if and only if it increases the probability of p's occurring relative to the baseline condition of the agent's refraining from any action at all (113). This account of the existence of reasons follows the essentially Humean understanding of a person's reasons as dependent on that person's own, possibly idiosyncratic desires, a notion reflected in Williams's theory as well. In Schroeder's terminology, all reasons are thus fundamentally "agent-relational," or grounded in a particular agent's psychology; a reason

---

1 Schroeder uses the term 'desire' in a technical sense to mean whatever psychological state can explain how one person has a reason for some action, whereas another person does not, in typical situations involving preferences, e.g., when a person has a reason to go to a party on account of his or her desire to dance (*Slaves* 9).

is an "agent-neutral" reason shared by all possible agents just in case it is an agent-relational reason for each of those agents (18).[2] Schroeder emphasizes that Hypotheticalism utilizes a weak version of the promotion relation, in the sense that the theory does not preclude the existence of agent-neutral reasons that would apply to any agent regardless of the specific desires that the agent might have (110-11). Many different actions might promote a given desire, and consequently it is plausible that some reasons might be "massively overdetermined" by an agent's desires, almost any of which could be promoted in some way by the actions that those reasons favor (108-9). With a sufficient degree of overdetermination, it would be possible for certain reasons to be genuinely agent-neutral, since any particular agent might have at least some desires capable of grounding those reasons (108-9).

In developing the second major element of his theory, the analysis of weighting for reasons, Schroeder considers the idea that, when one set of reasons to act in a certain way is weightier than another, it is correct for the agent to place more weight on the former than the latter (129-30). He rejects the thesis, which he calls Proportionalism, that the weight of a reason for some action depends on the strength of the associated desire and the effectiveness of the action in promoting that desire (98-99). Instead, as Schroeder suggests, the correctness of doing some action A, such as weighing reasons in a certain way, consists in the fact that the set of reasons "of the right kind" for doing A is weightier than the set of reasons "of the right kind" for not doing A (134). He bases this conception of correctness on his observation that correctness is relative to an activity; for example, some some moves in chess are correct and others incorrect (134-35). As he contends, whether it is correct to make a certain move in chess is not fixed by the moral reasons that a person might have to follow the rules of chess, or by the fact that great wealth can be gained by violating those rules (135). On the contrary, the fact that a person is participating in an activity governed by certain rules should be sufficient to determine whether it is correct for that person to act in a particular way in the course of that activity (135). Schroeder concludes that the right kind of reasons to do A, the ones relevant to the correctness of doing A, are precisely those reasons common to all agents engaged in doing A, where their being engaged in doing A suffices to explain why they have these reasons (135).

---

2 More specifically, the set of agents to which a reason must apply to be agent-neutral depends on the context, often extending to all possible agents as in the case of purportedly universal moral reasons, but sometimes possibly relativized to include only a particular class of agents (*Slaves* 18; cf. 117-19). The possibility of such a relativization will be central to the discussion of the primary goods proposal below, in section 5.

Schroeder thus obtains the result that the correctness of placing more weight on set of reasons A than set of reasons B can be determined by weighing the set of reasons of the right kind for placing weight on A against the set of such reasons for placing weight on B (136-38).[3] As he explains, the activity of placing weights on reasons is simply the activity of deliberation, in which all agents are necessarily engaged, implying that the set of reasons of the right kind for weighing reasons in a certain way is precisely the set of all agent-neutral reasons for placing weight (142). Consequently, an agent's idiosyncratic reasons for placing more or less weight on a given set of reasons do not affect the weight of that set, because one does not have these idiosyncratic reasons simply in virtue of engaging in deliberation (142). In this respect, Schroeder's account of weighting diverges from the basic Humean principle that the character of an agent's reasons depends on the particular desires that the agent might possess. Although a person's desires determine what reasons the person has, the relative weights of those reasons are fixed by the set of agent-neutral reasons, not by the particular individual's psychology.

3. A Challenge Concerning Normativity and Some Remarks on Deliberation

Due to the crucial position of agent-neutral reasons in Schroeder's theory, the relative weighting of a particular agent's reasons is determined in part by facts about the reasons of other agents, which in turn depend on facts about those agents' desires. As will be explained, Hypotheticalism thus requires the attribution of normative significance to the psychologies of agents other than the one whose reasons are in question, providing for a crucial challenge to the theory. Before this challenge is presented in detail, however, some preliminary observations concerning deliberation and endorsement should be considered. One can distinguish between two basic modes of deliberation in which a single agent may be engaged, although both varieties are sometimes combined in one instance of practical reasoning. Simple practical deliberation, as one might call it, involves thinking about the desirability, usefulness, or feasibility of alternative actions; an agent deliberating in this manner notices the reasons that can be offered for and against some course of action and weighs them in relation to one another. In simple practical deliberation, the agent usually operates on tacit or habitual assumptions about what types of considerations constitute reasons

---

3 Schroeder formalizes this proposal in a recursive account postulating two ways in which set A may be weightier than set B: (i) when B is empty and A non-empty, or (ii) when the set of all reasons of the right kind to place more weight on A is weightier than the set of all reasons of the right kind to place more weight on B (*Slaves* 138).

and what relative weights different categories of reasons might have, e.g., whether moral reasons always outweigh reasons of expediency. Normative deliberation, on the other hand, consists in a concerted and reflective attempt to determine in a definitive manner what reasons one has and what principles govern the weighting of those reasons. This kind of deliberation occurs whenever one contemplates fundamental questions of value, but also whenever one seeks to revise one's previous patterns of practical reasoning and reassess one's methods for identifying reasons and weighing them against each other.

These ideas about deliberation interact in interesting ways with certain theses concerning the dynamics of normative discourse and the nature of reasons claims suggested separately by Williams and Gibbard. In his "Replies," Williams proposes that an assertion that an agent A has a reason to act in some manner exhibits a characteristic dual distinctiveness: it is a statement of a distinctive sort bearing a distinctive relation to a particular individual A (194). In other words, a reasons claim states not merely that it would be good, beneficial, or right to perform the given action, or that there is some generic reason for agents under normal circumstances to act in that manner, but that A has a reason to do so (cf. "Internal Reasons and the Obscurity of Blame" 39). This notion of distinctiveness is clearly incorporated into Williams's theory of reasons, in which any action that A has reason to perform must be firmly connected to A's subjective motivational set. As he argues in "Some Further Notes on Internal and External Reasons," it would be "bluff or browbeating" to persist in advocating a reasons claim with no conceivable connection to the agent's motivations (95). While its validity will not be defended here, the distinctiveness thesis suggests that the effectiveness of normative discourse relies on successfully engaging the evaluative standpoint or practical perspective of the agent who is thought to have a given reason.

Normative discourse involves not only an agent toward whom a reasons claim might be directed, but also a speaker or judger whose assertions are underlain by a distinctive intention or state of mind, and it is this latter subject on which Gibbard concentrates. Gibbard argues that what it means for one to call an action, belief, or feeling rational is for one to endorse it, and in particular to express acceptance of a system of norms that permits it (Wise Choices 6-7). A statement that something is rational can in turn be interpreted as claiming that it is favored by the preponderance of reasons, since identifying some consideration as a reason also implies expressing acceptance of a certain system of norms (163-64). Claims about the rationality of an action, as Gibbard writes, seem to assert that rationality is not a matter of taste, and that the action would remain

rational or irrational even if the speaker were to think otherwise (164-65). He contends that this property of claims about rationality results from the speaker's acceptance of norms that demand that one act in a certain manner even if one were to have different desires or if one were to be persuaded that some other course of action is rationally required (165). Though this paper will not necessarily concern itself with the details of Gibbard's expressivistic analysis, it will extensively employ his central observation that when a person participating in normative discourse maintains that some consideration is a reason, the person is in some way expressing a certain degree of endorsement of the action supported by that consideration.

In normative discussion over some contemplated action, when the judger offers some consideration as a reason for the action and the agent accepts the consideration as counting in its favor, the judger is endorsing a certain attitude toward the action, and the agent is acknowledging the distinctive normative significance of the consideration for himself or herself. Where this normative discourse appears most effective, moreover, the judger is asserting the existence of a certain ground for the action applying specifically to the agent, who is then accepting that he or she could possibly act for that reason. In effective normative deliberation, likewise, one both endorses an understanding of certain considerations as reasons that can be weighed relative to each other in certain ways, and commits oneself to act in a manner guided by those considerations, with their weights determined via the specified method. Consequently, normative deliberation can be likened to an internal dialogue, uniting the perspective of the judger and the perspective of the agent in a single individual. Any compelling account of normative deliberation must therefore capture the endorsement given to a set of reasons by the deliberator, along with the distinctive normative force that those reasons are understood as exerting on the deliberator.

On the basis of these observations, a challenge to Schroeder's theory arises from the question of whether it can adequately ground or explain the substantial normative significance that is accorded to the psychologies of other possible agents in demarcating the class of reasons that govern the placement of weight. An agent who engages in normative deliberation using the theory must accept that only agent-neutral reasons can be invoked in a correct and justified judgment of whether one reason is weightier than another. One can consider an agent who does not already attribute significance to other agents' psychologies in practical reasoning, e.g., by seeking the esteem of others. Suppose this agent accepts the basic framework of the theory, namely the notion that one's

reasons explain why one's acting in a certain way promotes some state of affairs, which is in fact an object of one's desires, and then proceeds to identify certain considerations as reasons with certain relative weights. Normative significance is now being granted to these considerations in various degrees, or withheld from them, and the agent is endorsing the idea that future deliberation should be guided by these results.

In this scenario, however, facts about the psychologies of other possible agents dictate whether a given method for comparing the weights of reasons is valid. Therefore, an agent who properly comprehends the theory must also recognize that normative significance is being accorded to others' psychologies, and that he or she would be endorsing different conclusions about what weights his or her own reasons have, if the set of other possible agents' psychologies were to be varied. Accordingly, the statement of the challenge is as follows: what grounds the normative significance thus assigned to the psychologies of other possible agents, in such a manner that recognition of this normative significance justifies the agent in endorsing these different results, given different facts about other agents' psychologies? Williams's distinctiveness thesis possesses great appeal precisely because it seems plausible that facts about one's reasons depend only on one's own psychology, or perhaps on general constraints of deliberative rationality. If normative force in an agent's practical reasoning is to be ascribed to any psychological features whatsoever, it seems that it should be granted to the agent's own psychology, not the psychologies of other agents.[4] What justifies the attribution of normative significance to the psychologies of other agents, even of merely possible agents with whom the individual in question would never interact?

4. Normativity and Arbitrariness: Approaches to a Solution

In his article "In Defence of The Moral Problem," Smith offers a suggestion for explaining the normative significance of convergence across rational agents that also seems appealing as a justification for ascribing normative force to other agents' psychologies. If the requirement of convergence were removed from his theory, and if the variation among desire sets could not be eliminated through systematic justification, then agents with different desire sets would also have different reasons: an individual's reasons would

---

4 Of course, if the agent antecedently assigns value to others' psychologies, then this intuitive principle no longer applies; but the theory, if it is correct, should be capable of convincing a rationally deliberating agent who does not already give normative significance to others' psychologies, as stipulated in the scenario. Similarly, if general requirements of rationality prescribe, in a Kantian fashion, that all rational agents must act with concern for others' happiness and sensitivity to facts about their psychologies, then other agents' psychologies again have normative force; yet such substantive constraints of rationality are not invoked in Schroeder's model.

depend on the idiosyncratic character of his or her desire set. As Smith writes, without convergence, normative facts about reasons and rational justification would be based on entirely arbitrary facts about what desires a given person might currently possess, but arbitrariness always undermines normativity (265). Consequently, in his theory, the existence and normativity of reasons presuppose that rational deliberators converge in the content of their desire sets. Smith's concern with arbitrariness, moreover, appears to share the same fundamental motivations as Schroeder's argument that the correctness of some action in a particular practice depends on the set of all reasons for which one's engaging in the practice suffices to explain why one has those reasons (*Slaves* 134).[5]

In Schroeder's defense, one might therefore attempt to ground the specific normative significance of agent-neutral reasons in setting the relative weights of all reasons by considering the alternative possibility that the weights of the reasons possessed by a given individual might depend on that individual's own possibly idiosyncratic set of reasons for placing weight. Then the normative fact of whether the individual's action of placing greater weight on one reason than another is correct would be determined by the fact of what reasons for placing weight the person has, which in turn would be based on the completely arbitrary fact of what desires the person presently possesses.[6] Consequently, it seems appropriate that the correct weighting of reasons should be determined instead in a non-arbitrary way by the set of agent-neutral reasons, which are the reasons that any individual engaged in placing weights would have.

Appealing to Smith's principle that arbitrariness precludes normativity would conflict, however, with several aspects of Schroeder's theory, and considerations of arbitrariness cannot, in the context of this theory, justify granting normative significance to others' psychologies. One is seeking a consideration that would explain why an agent who accepts the basic idea of grounding the existence of reasons in one's own desires would then endorse certain conclusions about the weights of reasons on the basis of recognizing the normative force accorded to the psychologies of other possible

---

5 To use Schroeder's example, if two people with different desires and thus different idiosyncratic reasons are playing chess, it does not seem that it would be correct for one person to make a certain move, but incorrect for the other person to make that move under exactly the same circumstances, merely because they have different desires. Instead, it seems that the correctness of an action can be fixed only by non-arbitrary facts, and that the normative fact of what moves in chess are correct should be determined by the reasons that any chess player would have, simply in virtue of being engaged in the activity of chess.

6 A declaration that the proper reasons for placing weight are those shared by any agent with this particular agent's system of values or any agent living at the present time would be similarly susceptible to the charge of arbitrariness.

agents. First, by believing that one has a reason for some action whenever one has a desire more likely to be realized if that action is performed, the agent is already acknowledging that a substantial degree of arbitrariness is implicated in normative facts. Reasons are essentially normative, and the theory centrally holds that the arbitrary fact of having a certain desire immediately generates a reason for performing each action that promotes the object of that desire. Any agent who accepts this foundational claim of Schroeder's theory thereby accepts the notion that arbitrariness is not necessarily inconsistent with normativity.

Second, for every reason that a person possesses, there exists a reason, perhaps an idiosyncratic one, for that individual to place weight on that reason, because every reason must be connected to some desire that the person has. Placing weight on the original reason increases the probability that the person will act in the manner prescribed by the reason, and undertaking the specified action would, as the reason guarantees, promote the object of one of the person's desires. Therefore, for every reason that one has, one has a corresponding reason, that is, a normatively significant consideration in favor, for one to accept arbitrariness in any instance of deliberation and to place great weight each time, quite arbitrarily, on the original reason. Even more strikingly, for any desire in an agent's motivational set, the presence of that desire induces a reason to pursue the object of that desire and a reason to place weight on the original reason. As this observation shows, any given agent will have numerous reasons and strong motivations to weigh all his or her reasons using a scheme in which the weight of reasons depends on the intensity of the associated desires, rather than the balance of agent-neutral reasons.

Third, and most importantly, Schroeder maintains that normativity in general should be analyzed in terms of reasons, as stated in his thesis of "reason-basicness" (81), and the principle that arbitrariness undermines normativity is itself a manifestly normative assertion. If this principle were to ground the ascription of normative significance to other agents' psychologies, one would thus expect that any agent would have a reason to avoid arbitrariness in normative deliberation, because the definition of correctness in placing weights is applicable to any agent. Suppose, then, that an agent with only a few, strong, but highly idiosyncratic desires engages in normative deliberation, guided by the premise that the presence of desires generates various reasons to act in ways that promote those desires. Adopting a general policy of avoiding arbitrariness in the relative weighting of reasons would almost certainly lead the agent to place greater weight on the reasons for actions with only a minimal capacity to increase the probability that the agent's desires

are satisfied, and less weight on the reasons for actions with a more robust such capacity. Actions with direct relevance to realizing the objects of this agent's exceptionally unusual desires would surely be unlikely to promote the desires of most other agents, and the agent in question would hesitate to weigh reasons using an arbitrary standard that could not be universalized to all other agents. Arguably, therefore, acknowledging an inescapable incompatibility between normativity and arbitrariness would decrease the likelihood of satisfaction for every one of the agent's desires. Indeed, the agent would have absolutely no reason to observe a principle of uniform non-arbitrariness in deliberation, and a multitude of reasons to weigh reasons in a manner that emphasizes attaining the objects of the agent's desires. One may conclude that such a principle of non-arbitrariness, on its own, cannot justify every agent who antecedently accepts Schroeder's account of what reasons exist in admitting moreover that facts about other agents' psychologies determine how those reasons should be weighed.[7]

Schroeder's thesis of reason-basicness suggests, however, that it might be more fruitful to identify some broader reason not merely to avoid arbitrariness but to respect certain aspects of other agents' psychologies, a reason common to all possible agents or at least to all agents capable of social interaction with others. Here, 'respect' might refer quite generally to many different forms of action, ranging from altruistic behavior to simply considering the opinions of others in one's own practical reasoning, but each case crucially involves some sort of evaluative attitude that presupposes the normative significance of facts about other agents' psychologies. One might seek to construct an argument in this direction on the model of Williams's notion, articulated in his paper "Internal Reasons and the Obscurity of Blame," that what he describes as "focussed blame" can operate through the "proleptic invocation of a reason" which arises as a result of the expression of the blame (41). According to Williams, when one is blamed by others for acting in a certain way, the blame is meant to convey the thought that one had at that time a reason for acting in a different way, but failed to act on that reason (41). Even if one does not actually possess any direct motivation to act in that manner, one's subjective motivational set may nonetheless contain dispositions to avoid the disapproval of others, such as the desire to be esteemed reciprocally by those whom one

---

7 Considerations of arbitrariness alone cannot ground the ascription of normative significance to other agents' psychologies, when the alternative is not attributing such normative significance, except to the extent that acting in accordance with one's own psychology already involves doing so. Nevertheless, the inclination against introducing arbitrariness remains strong, and the principle of non-arbitrariness can in fact successfully undermine other proposals for grounding the normative significance of others' psychologies in a particular manner, when the alternative is accepting this normative significance in some other manner, as will be discussed later.

esteems (41). When, through the expression of blame, one learns that others expect one to behave in a different way, one gains a reason to alter one's behavior that was not necessarily present before the blame was expressed (41-42).

One could consequently seek to locate generally applicable grounds for granting normative significance to the psychologies of other possible agents in an agent-neutral reason to respond in a certain manner to the expectations that others have for one's own behavior, or would have if one were to come into contact with them. One might then have a further reason to act in some particular way, such as a reason to adopt a certain weighting scheme or method of deliberation, that is ultimately connected to one's own desires by way of facts about others' expectations. More a research program than a specific solution for the problem of normativity, this strategy resembles Williams' idea of prolepsis in ascribing reasons to agents who would not ordinarily have these reasons except through the intervention of this mechanism involving the expectations of other agents. At this point it is appropriate to examine two specific formulations of this schema for responding to the problem of grounding the normative significance of others' psychologies before finally returning to assess the broader viability of the proleptic strategy.

5. The Primary Goods Proposal

For the first proposal arising from this approximately proleptic model, one might argue that there exist certain "primary goods" that can serve as universal means in attaining the object of any, or nearly any, conceivable desire that an agent could possess, as Schroeder suggests in his discussion of massively overdetermined reasons (*Slaves* 109 n.10). Then, responding to various criticisms in his "Reply to Shafer-Landau, Mcpherson, and Dancy," Schroeder offers a more detailed account of how agent-neutral reasons, and specifically moral reasons, could be shown to exist (8). If moral requirements are such that other agents have a reason to expect one to conform with them, then the universality of moral reasons might be explained by a reason, applicable to all moral agents, to fulfill the legitimate expectations of others (8).[8] One could combine the features of both approaches by suggesting, in a Hobbesian or Rawlsian fashion, that any agent is more likely to enjoy a number of primary goods, including safety, goodwill from others, and access to resources, if the agent's behavior is sensitive to the legitimate expectations of other agents. Since these primary goods would be useful to the agent in satisfying almost any desire, the agent would

---

8 For a given agent, this reason could be derived from any desire that would be promoted by fulfilling others' expectations in this way, such as a desire to be respected or even not to be despised, or another essentially social desire (*Slaves* 8).

thus acquire a reason to respect the expectations that others would have for the agent's actions if they were to come into contact with the agent. This action of granting respect to others' expectations could then be generalized in some way to generate a reason for attributing normative significance to others' psychologies more broadly, providing grounds for considering their psychologies in the determination of the correct relative weights of reasons.

While this mechanism could prove effective in establishing the universality of moral reasons, it encounters an important obstacle in grounding a sufficiently comprehensive ascription of normative significance to the psychologies of other possible agents in normative deliberation. Merely possible agents who do not actually exist, as well as historical agents and agents who might be living in the distant future, do not exert any influence on one's success in acquiring primary goods, and one's respect for their expectations would not furnish any benefits in achieving one's ends. Consequently, because the reason to fulfill the expectations of other agents must be connected to an increased likelihood of satisfying one's own desires, this proposal does not provide adequate grounds for according normative force to every possible agent's psychology in setting the weights of one's reasons. It might be thought that this mechanism could be extended to consider the possible effects of one's actions in promoting one's desires, in a counterfactual sense that would allow merely possible agents to be included.[9] Given the nature of the promotion relation in Schroeder's theory, however, normative significance cannot be extended to possible agents' psychologies in this way. One has a reason for an action, such as modifying one's behavior out of sensitivity to others' psychologies, only if acting in that manner now, in the actual world, increases the probability, measured at the present time, of satisfying some desire that one has. Nonetheless, every agent is supposed to have a reason, at all times, to place weights on reasons in accordance with a theory that attributes normative force to the psychology of any possible agent simply in virtue of that agent's engagement in deliberation.

Perhaps the proposal can be adjusted through a 'relativization' of the domain of agents who are considered in determining which reasons are agent-neutral: the weights presently associated with the reasons of a given agent would then depend on the set of reasons common to all agents now existing, in the actual world.[10] Then the notion that every agent has a reason to act in a way that responds

9 For example, if some merely possible agent were to exist in actuality, then one's respect for that agent's expectations would have some possibility of affecting that agent's actions in a manner that would eventually produce benefits in terms of the satisfaction of one's desires.

10 Correspondingly, the definition of correctness could be revised to refer to all reasons that any presently existing agent would have simply in virtue of engaging in the relevant activity.

to the legitimate expectations of other presently existing agents would indeed justify the agent in privileging the set of reasons that are agent-neutral with respect to the relativized domain. With these modifications, nevertheless, the primary goods proposal can no longer fully capture the aspect of endorsement implicated when an agent acknowledges the normative significance of other agents' psychologies in fixing the weights of his or her own reasons. In normative deliberation, an agent endeavors to determine, through rigorous reasoning from clear premises, what considerations can be regarded as reasons and what methods are appropriate for weighing reasons against each other. When this process is completed, one can endorse one's conclusions as justified, and commit oneself to decide upon one's later actions through practical reasoning in which the conclusions from previous normative deliberation may simply be presupposed as true. As the theory stipulates, considerations of agent-neutrality contribute to determining the weight of every conceivable reason, and even agent-neutral reasons are weighed in accordance with higher-level agent-neutral reasons favoring or opposing the associated actions. Endorsing the results produced by the theory therefore constitutes an exceptionally strong commitment to the normative significance of other agents' psychologies.

Using Gibbard's terminology, one might say that an agent who recognizes this normative significance and utilizes Schroeder's weighting scheme in deliberation seems to be expressing acceptance of quite generally applicable norms fundamental to the entire complex of the agent's normative judgments. In particular, these norms are most naturally understood as applying even in the counterfactual situation where the psychologies of the other agents included in the relativization are different. With the model as it stands, however, only the exact present character of the psychologies of other presently existing agents receives normative significance; one has no reason to respect the expectations that some other agent would have were that agent to display the slightest variation in his or her psychology. Only the actual psychology of that agent can exercise any influence on one's attaining the objects of one's desires, and alternative possibilities for the character of others' psychologies cannot affect the assessment of reasons and their weights.[11] Suppose that the whole complex of agent-neutral reasons identified using the relativized domain at the present time were to be reconceptualized as a system of norms governing the weighing of reasons. Clearly, this system of norms could be sensibly employed in deliberation only to determine the relative weights of reasons for the

---

11 By the term 'actual psychology' is also meant those psychological features from which one can infer reasonable predictions about near-term future changes in the agent's psychology, but any variation not projected from these features is excluded, and this result is all that is needed for the argument.

actual world and the present time, because in some other hypo-
thetical situation the psychologies of other existing agents would
change, and the relevant system of norms would change along with
them. Consequently, weighing one's reasons at any given time in
accordance with the relativized primary goods proposal seems ar-
bitrary, in a way that interferes with the necessary endorsement of
the normative significance attributed to other agents' psychologies.
This relativization cannot explain why an agent would judge that
certain orderings of reasons in terms of relative weight are correct,
because the concept of correctness does not accommodate this form
of arbitrariness.

At this point, one might object that, as was demonstrated earlier,
an agent who already acknowledges the basic idea of deriving rea-
sons from his or her own desires would not be justified in accept-
ing any general principle that arbitrariness necessarily precludes
normativity. Nevertheless, whereas in the previous case the arbi-
trariness was found in the notion of weighing reasons in a man-
ner dictated by one's own desires, here the arbitrariness afflicts the
decision to grant normative significance to the actual psychologies
of other presently existing agents, rather than the psychologies of
all possible agents. When one already accepts the normative force
of one's desires in generating reasons, endorsing their further nor-
mative significance in fixing the relative weights of reasons does
not seem unjustified. When, on the other hand, one is considering
whether to grant normative significance either to all other possible
agents or to a select few, on apparently arbitrary grounds, it does
seem unjustified to choose the second option. Even for an ideologi-
cally Humean agent completely persuaded of the normative sig-
nificance of his or her idiosyncratic psychology, there would be no
inclination to endorse the idiosyncrasies of other agents' desires in
determining the weights of the agent's own reasons. Thus, either
the primary goods proposal will provide a reason to respect the
psychologies of all possible agents, in which case it is inconsistent
with the promotion relation, or it must be relativized, in which case
it cannot ground the required endorsement of the normative force
of others' psychologies.

6. The Conversational Demand Proposal

Despite the difficulties attending the previous proposal, the general
strategy of locating the source of the normative significance of oth-
ers' psychologies in their expectations for one's own actions has
not been exhausted of its potential. Another proposal can be con-
structed in this larger, proleptic model on almost entirely separate
foundations, centering on the concept of "conversational demands"
developed by Gibbard to explain the apparent claim to objective

validity involved in a person's assertion that something is rational. In normative discussion, as Gibbard contends, a "conversational demand" is imposed on the audience to attribute a degree of "fundamental authority" to the speaker, authority that is not founded on a set of norms shared by both parties or on the consequences of norms that the audience already accepts (Wise Choices 174-76). This fundamental authority is grounded in one's trust in one's own normative judgments, because one's present normative thinking has been formed under the extensive and ineliminable influence of other individuals' normative views (179-80). By having confidence in one's capabilities as a normative judge, therefore, one implicitly ascribes legitimacy to past influences from others, and it would be unjustified to withhold all authority from the future influence of others' judgments merely because it has not already been exerted (180). One must consequently accord some fundamental legitimacy to the normative judgments of any other individual who can be regarded as an equally qualified judge of normative matters (180-83).

If another person were to hold that a certain consideration constitutes a reason of a certain weight, then, in accordance with the conversational demand, one would be obliged to accept that the person's opinion bears some degree of authority in one's own deliberations. In this manner, Gibbard's idea might be profitably employed as the basis for the normative significance granted to others' psychologies in determining the weights of reasons.[12] Importantly, unlike the previous proposal, which concentrated on the tangible benefits for achieving one's own desires that would result from being respected by others, this proposal can ascribe normative significance to the psychologies of merely possible agents. In analogy to the argument concerning the past and future influence of others, it would be unjustified to refrain from acknowledging the legitimacy of merely possible but equally qualified individuals' judgments, when one would have accorded them authority if they had contributed to the development of one's own normative views.

To assess the prospects of the conversational demand proposal, it will be useful to examine what function the purported fundamental authority of equally qualified normative judges might hold in interpersonal normative discussion, before extending the conclusions to the normative deliberation of individual agents. For simplicity, one can consider discourse between an agent and a judger who claims that the agent has a conclusive, rather than pro tanto, reason to perform some action. While the judger seeks ultimately to induce the agent to act in the specified manner, the endorsement

---

12 Various methods could be devised to show that one has a reason to respect the conversational demands issued by others' normative judgments; perhaps a form of self-trust that can withstand doubts about one's own normative authority could be considered a primary good.

that the judger gives to the reasons claim must be understood in a particular, distinctive sense: the judger would not be satisfied if the agent should seemingly accept the claim, proceed to perform the action, and later explain that he or she merely wished to placate the judger, or was impressed by the judger's rhetoric. Instead, the judger appears to be asserting that the agent would act in a certain way if the agent were considering the matter rationally on the basis of accurate information, or possibly if the agent were to have had the appropriate upbringing or experiences. What the judger is endorsing must therefore be not merely the agent's performing the action specified in the reason, but the agent's acknowledging the objective validity of the reasons claim through rational means and then acting for the reason in performing the action. Accepting the judger's claim in its fullest sense means coming to think about the subject under discussion, at least in this particular respect, in the same way as the judger, that is, coming to regard the action as objectively rational given the circumstances.

One must consequently ascertain how the judger conceives of his or her own justification for asserting that the agent has conclusive reason to undertake the action, in order to determine what sorts of grounds or bases are characteristically invoked in effective normative discourse. In particular, it seems evident that the judger does not view the reasons claim as being grounded in the judger's trust in his or her own qualifications as a normative adviser, and is not purposefully endeavoring to impose a conversational demand on the agent. From the judger's standpoint, the considerations supporting the specified action are fixed prior to the judger's conclusion in favor of the action, regardless of whether the judger considers the reason for the action as deriving from the agent's subjective motivational set or from substantive constraints of rationality. No appeal to the judger's fundamental authority is needed, precisely because the judger is claiming that the consideration offered as a reason bears a specific normative significance for the agent in question. In the judger's thinking, if such fundamental authority were indeed to provide the basis for a genuine reason for action, then the reason could not apply distinctively to this particular agent, because any individual to whom the judger were to express the reasons claim would thereby acquire the reason as well. Likewise, the reason could not follow from universal requirements of rationality, because in that case the agent could discover the reason on his or her own simply by deliberating in accordance with the appropriate rational constraints. Clearly, the judger would not be confident that the agent has actually been persuaded of the reasons claim if the agent's acceptance of it essentially involves invoking the judger's fundamental authority.

Simply observing that the judger would not fully endorse the agent's appeal to the judger's fundamental authority does not, of course, demonstrate that in effective normative discourse the agent is never justified in accepting a claim on the strength of the judger's authority as a qualified normative adviser. One may appreciate that the judger possesses a separate point of view, informed by different experiences, values, and knowledge, that may provide the judger with some greater insight on certain matters; a second opinion on some normative concern can furnish a valuable understanding of different possible ways to approach the issue. For most claims about reasons, however, one takes oneself to be discussing a definite subject matter, a body of normative facts that are true independently of a person's normative judgments. Due to one's own limited information and experience, one accords fundamental authority to the normative judgments of others because such judgments provide another point of access, so to speak, on these normative facts. These ideas have important ramifications for the proposal of applying Gibbard's notion of the conversational demand to the normative deliberation of an individual agent, since deliberation also involves recognizing certain considerations as normative facts and endorsing a commitment to be guided by those considerations in one's actions.

According to Schroeder's theory, certain normative facts, namely those concerning the relative weights of one's reasons, are analyzed in terms of facts about other possible agents' psychologies, and from this observation there emerges a difficulty that ultimately undermines the conversational demand proposal. It seemed, initially, that the normative significance of other agents' psychologies as factors in determining the weights of one's reasons could be derived from the fundamental authority granted to others' normative judgments about one's reasons. This ascription of fundamental authority, nevertheless, presupposes that those judgments can be understood as being about the content of an antecedently fixed body of normative facts, and that the correctness of placing more weight on one reason or another already holds normative force for the agent. If the question of whether one should have weightier reasons to perform some action than to refrain from that action were merely a matter of opinion, or of taste, then it would not be unjustified to privilege one's own normative judgments over those of other equally qualified individuals. This principle shows why in effective normative discourse the judger must present the claim asserted by his or her judgment as an independently grounded fact, which would under ideal circumstances be accessible through rational inquiry to the agent as well.

With this proposal, therefore, one attempts to explain the norma-

tive force of other agents' psychologies on the basis of the normative significance of their judgments about one's reasons, but their judgments hold this significance only insofar as they relate to certain antecedently normative facts, which in this case are actually constituted by facts about others' psychologies. Now the normative force under dispute is being transmitted in a circle, and this reasoning can be admitted as legitimate only if it should be possible for facts about the judgments of others concerning some normative question to constitute the very principles which determine the answer to that question. Perhaps, the advocate of the conversational demand might respond, another person's judgment that one has more reason to perform some action can make it true, ipso facto, that one has weightier reasons in favor of that action, even if there would otherwise be no fact of the matter about relative weights. This response might be defensible on other grounds, but the aim of the proposal was to explain how the features of others' psychologies could hold normative significance specifically for a particular agent, such that the agent would endorse the appeal to others' psychologies in normative deliberation. It is simply an empirical datum that in effective normative discussion neither the agent nor the judger has recourse to such a claim about the constitutive significance of a person's normative judgments as the objector here is offering.

Even if some exceptional instances could be found, the proposal cannot succeed without explaining how every agent who understands the theory would be justified in endorsing different results for the weights of his or her reasons on the basis of recognizing the normative force of others' psychologies. As remarked in a previous section, normative deliberation is analogous to normative discourse, because this form of deliberation unites the perspective of the agent with the perspective of the judger. Accordingly, it is reasonable to suppose that the standards for what sorts of appeals to normative authority are accepted would be similar in both cases. Thus, because the assertion that a speaker's judgment about some normative fact carries an independent, constitutive normative significance is generally or even universally inadmissible in normative discourse, it is highly doubtful that this assertion could ground the necessary endorsement in normative deliberation. Consequently, the conversational demand proposal does not seem plausible as a solution for the problem of how the theory might provide a compelling explanation for the attribution of normative significance to other possible agents' psychologies

7. Conclusions

Having observed that both the primary goods and conversational

demand proposals fail to provide adequate grounding for the normative force granted to other agents' psychologies, one can now turn to assess the prospects of the broader proleptic model from which these proposals arise. Unfortunately for Schroeder, the challenge of grounding his weighting scheme seems possibly insuperable, because solving this problem requires reconciling three critical and perhaps incompatible factors: the significance of the agent's psychology, the significance of others' psychologies, and the principles governing normative discourse. To understand the tension among these factors, one must recognize that the primary goods and conversational demand proposals are actually instances of two more general methods for applying the proleptic model.

Many of the ways of constructing a specific solution from the framework of Williams's ideas about proleptic mechanisms can, it appears, be classified as following what might be called the pragmatic strategy or what might be called the rationalistic strategy.[13] With the pragmatic strategy, of which the primary goods proposal is a specimen, one appeals to the potential of others' actions to affect the satisfaction of the desires that a particular agent possesses. Almost any instance of this strategy, however, seems susceptible to difficulties regarding whether the promotion relation actually licenses the claim that the agent has a reason to attribute normative significance to the psychology of any possible agent. Thus this strategy is hindered by a lingering conflict between the normative force already conceded in the theory to one's own psychology and the normative force granted to others' psychologies. On the other hand, with the rationalistic strategy, exemplified by the conversational demand proposal, one attempts to derive the normative significance of others' psychologies directly from the nature of justification in normative deliberation. Again, however, many different instances of this strategy would arguably be shown to be problematic once one considers the principles underlying the interaction of the judger and the agent in effective normative discourse. When accepting a reasons claim, the participants in normative discussion view the claim as reflecting some aspect of a fixed body of normative facts, in such a manner that a given consideration holds specific normative significance for a particular agent. This normative significance is generally explicated either in terms of a distinctive relation to the agent's subjective motivational set, or in terms of a universally applicable requirement of rationality. In neither

---

13 The terms 'pragmatic' and 'rationalistic' as descriptions of these strategies are taken from Kieran Setiya's paper "Intention, Plans, and Ethical Rationalism," to which there is an interesting parallel. Setiya considers two methods of grounding the normative force of the requirement of means-end coherence in Michael Bratman's planning theory of intention, the first being a pragmatic account that appeals to the practical advantages of coordinating one's intentions, and the second being a rationalistic account that derives the coherence requirement from considerations about the nature of agency (5-15).

case is an independent normative force credited to other agents' psychologies, casting suspicion on the purported capability of the rationalistic strategy to explain the necessary endorsement of the significance of others' psychologies in normative discussion or, correspondingly, in normative deliberation.

Consequently, what a compelling defense of the normative force of others' psychologies in the system for weighting reasons would need is an approach along the lines of the rationalistic strategy that would be based not on some form of roughly proleptic mechanism but instead on considerations of arbitrariness. Only by invoking some variety of the notion that arbitrariness is inconsistent with normativity, a notion that seems to generate the intuitions ultimately supporting Schroeder's conception of the correctness of weighing one reason over another, can the conflicts surrounding the significance of others' psychologies be resolved. With considerations of arbitrariness, one might be able to ground the normative force of others' psychologies in fundamental principles connected to the ideas of rationality and objectivity, rendering this normative force consistent with the nature of normative discourse. It suggests a profound tension within Schroeder's theory that such an appeal to constraints of non-arbitrariness would once again encounter difficulties with the normative significance of one's own psychology, a significance that marks Hypotheticalism's inescapably Humean character.

References

Gibbard, Allan. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, MA: Harvard UP, 1990.

Schroeder, Mark. "Reply to Shafer-Landau, Mcpherson, and Dancy." *Philosophical Studies*. Forthcoming [30 Oct. 2010]: 1-12.

Schroeder, Mark. *Slaves of the Passions*. Oxford: Oxford UP, 2007.

Setiya, Kieran. "Intention, Plans, and Ethical Rationalism." 19 Nov. 2011. 28 Nov. 2011 <http://www.pitt.edu/~kis23/Plans.pdf>.

Smith, Michael. "In Defence of The Moral Problem: A Reply to Brink, Copp, and Sayre-McCord." *Ethics and the A Priori*. Cambridge, UK: Cambridge UP, 2004.

Smith, Michael. "Internal Reasons." *Ethics and the A Priori*. Cambridge, UK: Cambridge UP, 2004.

Williams, Bernard. "Internal and External Reasons." *Moral Luck*. Cambridge, UK: Cambridge UP, 1981.

Williams, Bernard. "Internal Reasons and the Obscurity of Blame." *Making Sense of Humanity*. Cambridge, UK: Cambridge UP,

1995.

Williams, Bernard. "Replies." *World, Mind, and Ethics*. Ed. J. E. J. Altham and Ross Harrison. Cambridge, UK: Cambridge UP, 1995.

Williams, Bernard. "Some Further Notes on Internal and External Reasons." *Varieties of Practical Reasoning*. Ed. Elijah Millgram. Cambridge, MA: Massachusetts Inst. of Technology P, 2001.

On the Conceivability of Zombies

KAREN DUFFY
College of Saint Benedict

Abstract

A zombie, as employed in philosophical thought experiments, is a creature that is physically and functionally identical to humans, but lacks consciousness. The possibility of zombies has been used to argue that consciousness cannot be explained in terms of the physical, that is, that physicalism is false. Using zombies in this way is contingent on two premises: first, that zombies are conceivable, and second, that conceivability entails possibility. In this paper I focus on the first premise, the conceivability of zombies. I begin by considering arguments in favor of the conceivability of zombies, and then evaluate whether these arguments are convincing. I argue that the pro-conceivability arguments are not convincing because the "physically identical" portion of the zombie hypothesis means anything that is causally necessary in reference to our physical facts will be part of the concept of a zombie world as well. This makes the zombie hypothesis vulnerable to the possibility of a causally necessary entailment from our physical facts to consciousness that would render zombies inconceivable. I conclude with an argument for the existence of such a causally necessary entailment.

Introduction

Most contemporary philosophers seem to believe that some form of physicalism – the doctrine that everything can be explained by the physical facts – is likely to be true. However, it seems that most non-philosophers, at least in the western world, believe in some form of dualism – that the mind is ontologically different from the body, and thus that the functions of our mind, our consciousness, cannot be explained by the physical facts. So despite the trend in

the philosophical community, dualism certainly is not dead. Of course, simply the fact that many people believe in dualism is not enough to imply that the doctrine has any validity. It is extremely possible that the majority of people are simply mistaken. However, there do exist prominent contemporary philosophers who argue for dualistic views. One such philosopher is David Chalmers.

In this paper, I will be looking at a particular thought experiment – the "Zombie" thought experiment – which Chalmers uses to argue for the conclusion that physicalism is false (and thus that some form of dualism must be true). In order for Chalmers's argument to succeed, these "zombie" creatures must be logically possible. Chalmers employs what is known as the "argument from conceivability"[1] to this end. The argument is this: (1) Zombies are conceivable, (2) Whatever is conceivable is possible, (3) therefore Zombies are possible.[2] Now, each of the two premises of this argument can be called into question. I will focus on the first premise – whether or not zombies are conceivable. If it turns out that zombies are not conceivable, then the argument from conceivability is not sound, and this would call into question any argument that physicalism is false based on the possibility of zombies.

My examination of the conceivability of zombies will proceed as follows: in Section I, I will explain the zombie thought experiment and the argument that the possibility of zombies implies that physicalism is false. In Section II, I will give arguments for the conceivability of zombies. Section III seeks to determine whether these arguments are convincing or not. I argue that the pro-conceivability arguments as they stand are not convincing. Specifically, Chalmers's nonstandard realization argument applies only to the question of a logically necessary relationship between physical facts and consciousness. The argument is vulnerable to the possibility of a causally necessary entailment from the physical facts of our world to consciousness, which would render zombies inconceivable. I will end in Section IV with an argument given by Sydney Shoemaker that, if successful, I believe would establish the existence of such a causally necessary entailment.

Section I

For most readers, the word "zombie" might conjure up images of decaying, brain-eating monsters, emerging from graves to wreak havoc on the world. These are the zombies of Hollywood. Unfortunately for anyone seeking a bit of a thrill, these are not the type of zombies that are relevant to arguments against physical-

---

[1] Robert Kirk, *Zombies and Consciousness*, (Oxford: Clarendon Press, 2005), 27.

[2] Ibid.

ism. From this point on, "zombie" will be taken to mean a creature that is physically and functionally identical to a human being, but which lacks consciousness.[3] By "functionally identical," it is meant that zombies have the same sensory inputs and manifest the same outputs as a normal human – externally, humans and zombies are completely indistinguishable. When a zombie stubs its toe, it winces and says, "Ouch!" just like a human would. Confronted with a human and its zombie replicate, there is no test one could employ to identify which is the zombie – the two creatures will act identically. The difference between a zombie and a human is an internal difference. In the toe stubbing case, there is something "it is like" for a human to stub its toe – it *hurts*, the sensation of pain is *felt*. For a zombie, however, there is nothing "it is like" to stub its toe – *nothing* is felt.[4] Zombies have no point of view, no awareness, and experience no phenomenal properties ("raw feels," also known as "qualia").

Chalmers's argument that the possibility of zombies implies that physicalism is false runs as follows: (1) In our world, there are conscious experiences, (2) There is a logically possible world physically identical to ours in which the positive facts about consciousness in our world do not hold [a world inhabited only by zombies], (3) Therefore, facts about consciousness are further facts about our world, over and above the physical facts, (4) So materialism is false.[5]

The idea is that in a world that lacks consciousness, the physical facts about that world cannot entail consciousness (there is no consciousness to be entailed). So, for instance, the physical facts about a zombie world cannot entail consciousness. But the physical facts about the zombie world are the physical facts about *our* world. And if the physical facts of the zombie world cannot entail consciousness, then the physical facts of our world cannot entail consciousness, because they are the same physical facts. So if the zombie world is possible, then the conclusion is that our physical facts cannot entail consciousness – the physical facts and consciousness are independent. Therefore, physicalism, which would hold that consciousness must be entailed by the physical facts, must be false.

Notice that the argument against physicalism appeals to the *possibility* of zombies and recall that Chalmers argues for the possibility of zombies based on two premises: (1) Zombies are conceivable, and (2) Whatever is conceivable is possible. If either (1) or (2) is

---

[3] Robert Kirk, "Zombies," *Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta (Spring 2011): <http://plato.stanford.edu/archives/spr2011/entries/zombies/>.

[4] David Chalmers, *The Conscious Mind*, (New York: Oxford University Press, 1996), 95.

[5] Ibid., 123.

false, the zombie-based argument against physicalism falls apart. So we now come to the main content of this paper: examining the conceivability of zombies. The next section gives arguments in favor of conceivability.

Section II

Before we progress, I would like to make clear that for the remainder of this paper, I will understand "conceivable" to mean "involving no conceptual contradiction." Thus something that is inconceivable is something that does involve a conceptual contradiction in its description. Whether this definition of "conceivable" implies or is identical to "logically possible" is beyond the scope of this paper.

Now for the actual arguments. In first addressing the question of how to argue for the conceivability of zombies, Chalmers admits that creating such an argument is not entirely straightforward. For example, how does one argue that a mile-high unicycle is conceivable? Chalmers says that it just seems obvious.[6] Similarly, the conceivability of zombies seems equally obvious, at least to Chalmers. We can describe a zombie – a physical and functional replica of a human that lacks consciousness – without any apparent contradiction or incoherence. Chalmers notes that arguments for conceivability come down in part to intuition. It *seems* to us that zombies are conceivable. We have no reason to doubt this intuition unless someone, an anti-zombist, provides a compelling reason why our intuition in favor of conceivability is mistaken. But it is up to the anti-zombist to do this. The burden of proof, according to Chalmers, is on the anti-zombist.[7]

However, Chalmers does make actual arguments for the conceivability of zombies as well. Chalmers notes that there are indirect ways to argue for conceivability – start with an obviously conceivable situation and argue that if the first situation is conceivable, so is the situation you have in mind. For instance, we can argue indirectly for the conceivability of a mile-high unicycle by asking readers to first conceive of a normal unicycle going down the street and then to imagine that everything in the scene expands by a factor of one mile.[8] There does not seem to be any difference in conceivability between the first situation and the second situation – if the first situation (the normal unicycle) is conceivable, then it seems like the second situation (the mile-high unicycle) should be as well.

We can employ a similar strategy regarding zombies. Chalm-

---

[6] Ibid., 96.

[7] Ibid.

[8] Ibid., 97.

ers does this by asking us to imagine "nonstandard realizations" of our functional organization.[9] For example, suppose a hoard of extremely tiny people infiltrate my brain and disconnect it from my body.[10] The people then arrange themselves isomorphically to my brain – one person for each neuron, and the people use cell phones to communicate correspondingly to the firing of my synapses. My brain and the population of tiny people are organized and function exactly the same – the difference is only in the material by which the organization is realized. A more straightforward example is to imagine a silicon isomorph to the functional organization of my brain, in which silicon chips correspond to each of my neurons.[11]

The question Chalmers poses is whether a nonstandard realization of my functional organization would become conscious.[12] Would my silicon isomorph become conscious? Chalmers notes that the point of the question is not whether consciousness would arise *in fact*, but whether it is a meaningful possibility that the nonstandard system would not be conscious.[13] This is because Chalmers's goal with this argument is establish that there is not a logically necessary relationship of entailment from the physical facts to consciousness. Logical relationships are not about what the actual facts are, but rather about what facts are *possible*. If the counterfactual situation is possible, there is no logical relationship. So if it is a meaningful possibility that the silicon system lacks consciousness, Chalmers takes this to establish that there is no logical entailment between the functional organization of my brain and consciousness.[14]

Chalmers has the intuition that it *is* a meaningful possibility that the silicon system lacks consciousness. This is meant to show that there is no logical relationship between my functional organization and consciousness. Now, if we substitute the organic material of our brains back in for the nonstandard material, we should be left with the conclusion that it is conceivable that the organic system also lacks consciousness. For why should the material of realization make a difference to a relationship of *logical entailment* from organization to consciousness?[15] If it is conceivable that the organization realized in a nonstandard material lacks consciousness, then it should be conceivable that the same organization realized in the

---

[9] Ibid.

[10] Kirk also advances this argument in *Zombies and Consciousness*.

[11] Ibid.

[12] Ibid.

[13] Ibid.

[14] Ibid.

[15] Ibid.

standard material lacks consciousness.   So zombies are conceivable.

A second argument for the conceivability of zombies is put forth both by Robert Kirk and Torin Alter. Kirk asks us to imagine a person, Dan, who one day experiences a peculiar change in himself.[16] On this day, Dan stubs his toe. Dan winces, and exclaims, "Ouch!" as usual in toe-stubbing situations. However, Dan is surprised to find that he *felt* no pain when he stubbed his toe. Dan is surprised at his wince and exclamation, since before he only winced or exclaimed "Ouch!" in response to a pain sensation. However, now Dan finds these behaviors obtaining in the same way his leg jerks when a doctor checks his reflexes.[17] Slowly over time, Dan loses his other phenomenal properties as well. One day he can no longer hear any sounds, although he still responds to all auditory stimuli just as he did before. The next day, Dan cannot taste anything, although he still finds himself saying things like "Yuck!" or "This soup is delicious!"

The argument comes down to this: Kirk and Alter seem to maintain that it is hard to deny the conceivability of Dan losing the phenomenal feel of pain but continuing to act the same. Likewise with the conceivability of Dan losing his auditory phenomenal properties and his gustatory phenomenal properties. But then, according to Kirk and Alter, we should also be able to conceive of Dan losing *all* his phenomenal properties, rendering him a zombie.[18] If an anti-zombist were to say that we cannot conceive of Dan losing all his phenomenal properties, where in the sequence does the inconceivability come in? If we can conceive of Dan losing one property, two properties, three properties, and so on, then the anti-zombist would have to claim that at some number of lost phenomenal properties, the scenario suddenly becomes inconceivable.[19] But this seems arbitrary. So zombies are conceivable.

We now turn to examining whether or not these arguments are convincing.

Section III

I would like to start by making some statements about the Dan argument. Recall that indirect arguments for conceivability are supposed to take us from an obviously conceivable situation to the conceivability of zombies by showing that if the first situation

---

[16] Kirk, *Zombies and Consciousness*, 25.

[17] Ibid.

[18] Ibid. and in Torin Alter, "Imagining Subjective Absence: Marcus on Zombies," *Disputatio* 2, no. 22 (May 2007): 95.

[19] Ibid.

is conceivable, then so are zombies. The first, "obviously conceivable," situation in the Dan argument is that we can conceive of Dan losing one phenomenal property, pain, while remaining physically and functionally identical to how he was before. However, I am not convinced that this is conceivable at all. If we look at instances of people who actually do lack certain phenomenal properties, such as people who are deaf or blind or do not feel pain, they differ both physically and functionally from humans who retain these phenomenal properties (obviously, deaf people do not respond to auditory stimuli, blind people to visual stimuli, and people who lack pain sensations do not respond to bodily damage).

On my own speculative analysis, it seems plausible that this is the case because a physical difference causes a lack of phenomenal property, and the lack of phenomenal property causes the functional difference (the lack of response to certain stimuli). This suggests that phenomenal properties play a causal role in our functional organism. Therefore, if Dan really is to lose a phenomenal property while remaining functionally the same, it seems there would have to be a change in causal laws so that the causal role usually filled by the experience of that phenomenal property can be filled. But this would change the physical facts of the situation, which contradicts the hypothesis that Dan also remains physically the same. Then it would not be conceivable that Dan loses even one phenomenal property while remaining physically and functionally the same.

The analysis just offered above indicates that there might be a causally necessary relationship of entailment from our physical facts to consciousness. This will ultimately be the crux of my argument that the conceivability arguments are not convincing. To begin, let us distinguish between logical and causal necessity. I will functionally define "logically necessary" as "cannot conceivably be false" and "causally necessary" as "cannot be false in reference to a specified set of causal laws."

We need to take a careful look at what actually needs to be shown to establish the conceivability of zombies. At first glance, it may seem that zombies are conceivable if it is not *logically* necessary that physical facts entail consciousness – that is, it is conceivable that physical facts hold without consciousness (the entailment is conceivably false). I would grant that it is not logically necessary that physical facts in a general sense entail consciousness. We can conceive of *some set of physical facts* that does not entail consciousness. This lack of logical necessity in general, however, is not enough to establish the conceivability of zombies. A zombie world is not a world with just *any* set of physical facts, it is a world with *our* physical facts. So we have to modify our previous conjecture about what suffices for the conceivability of zombies: zombies are conceivable

if it is not necessary in reference to the physical facts of the actual world that the physical facts entail consciousness (it is not causally necessary that the physical facts of our world entail consciousness).

I claim that the nonstandard realization argument can only establish a lack of logical necessity in general, and is thus insufficient for establishing the conceivability of zombies. The argument banks on an intuition that any perceived causal relationship between the functional organization of the brain and consciousness does not hold for the same organization realized in some other material. This is meant to show that we can conceivably bring the brain matter back in while leaving perceived relationship between organization and consciousness behind, since this relationship does not *have* to hold. We are then supposedly left with the meaningful possibility that our brains could be physically and functionally identical to how they are now, but lacking consciousness.

However, I think this is an invalid intuition. It does seem that we tend to perceive a causal connection between the functioning of our brain and our consciousness (e.g., the cases discussed above where physiological differences are paired with phenomenal differences, or the observation that when people get hit really hard on the head, they tend to lose consciousness for a while). If this perceived causal connection is correct, since it is between the brain's *functioning* and consciousness, it will have to hold even in the nonstandard realization case since the brain's full functional organization is preserved. Then in order for it to be a meaningful possibility that the functional organization lacks consciousness, a new set of laws will have to be adopted, which changes the physical facts. This case can therefore only establish a lack of logical necessity between the physical in general and consciousness (we can conceive of *some* set of physical facts that does not entail consciousness).

If we are wrong about the perceived causal connection between the functioning of the brain and consciousness, then I think our only option is to say that we do not know the nature of the relationship between the brain's functional organization and consciousness. To conclude that there is *no* causal relationship between the brain's functional organization and consciousness from the fact that the perceived connection is incorrect would be to beg the question in favor of the conceivability of zombies. If we do not know the nature of the relationship between the brain's functional organization and consciousness, then while I admit we are able to conceive of *some* set of physical facts that does not entail consciousness, we would have no way of knowing whether these facts are the facts of the actual world or not. Again, the only conclusion that is warranted from this is a lack of logical entailment from the physical in general to consciousness.

As stated above, this is insufficient for establishing the conceivability of zombies since the conceivability of zombies depends on a lack of necessity from *our* physical facts to consciousness. Keeping this in mind, we can see that if it is *causally* necessary that the physical facts of our world entail consciousness, then zombies are inconceivable. For if it is causally necessary that our physical facts entail consciousness, then it is causally necessary that the physical facts entail consciousness in zombie world (the zombie world is physically identical to ours, so it is governed by the same causal laws). This is a contradiction, since a zombie world also lacks consciousness by hypothesis.

Both the nonstandard realization argument and the Dan argument are vulnerable to this possibility. In both cases, a causally necessary entailment from the physical facts to consciousness would make the initial scenario inconceivable. In the nonstandard realization argument, fully replicating all the causal functions of the brain, even in a different material, would necessitate consciousness, based on the causal laws. In the Dan argument, as outlined above, a lack of some conscious property would have to be accompanied by a change in the physical facts – either Dan's physiology or the causal laws governing him. If the starting situation in a conceivability argument is not actually conceivable, then there is no hope that the argument can establish the conceivability of the target situation (zombies), since the conceivability of the latter was supposed to be based on the conceivability of the former.

A pro-zombist would not be able to escape the existence of such an entailment by dropping the physical identity portion of the zombie hypothesis. The physical identity stipulation is essential if a zombie world is to serve as an effective counterexample to physicalism. The question of physicalism is whether *our* physical facts entail consciousness. The existence of a possible world in which the physical facts do not entail consciousness has no bearing on this question unless the physical facts of the possible world are the physical facts of our world. Therefore the physical identity portion of the zombie hypothesis is essential in its use against physicalism.

The conclusion of this section is the conditional: *if* there is a causally necessary entailment from our physical facts to consciousness, *then* zombies would be inconceivable. I think this conditional conclusion gives enough reason to conclude that the pro-conceivability arguments as they stand are unconvincing, since they do not preclude the possibility of an entailment that would render zombies inconceivable. However, we may still feel dissatisfied with this, since this conclusion does not actually provide an answer to the question "are zombies conceivable?" Therefore, in an attempt to bolster the case that the answer to this question is in fact "no," I

will devote the final section of this paper to an argument given by Sydney Shoemaker that, if successful, I believe would establish the existence of a causally necessary entailment from our physical facts to consciousness.

Section IV

A certain objection to functionalism holds that a functional analysis of qualitative experience (what I have been calling "consciousness") cannot be given because it would have to consist of functional analyses of individual qualia. These analyses cannot be given, the argument goes, because of the possibility of "inverted qualia," which refers to the possibility that two functionally indistinguishable people may experience opposing qualia, such as the same objects appearing as green to one person and as (what we would normally refer to as) red to the other.[20] In this case, the two states are functionally the same but qualitatively different, which seems to show that a functionalist account of individual qualia would be inadequate.[21]

Shoemaker, however, suggests that we can give a functional account of qualitative experience that does not involve functional analysis of individual qualia. Shoemaker's suggestion is that we can give a functional definition of the relationships of qualitative similarity and difference that hold between individual qualitative states.[22] We do this by defining a relationship between individual qualitative experiences as one of "qualitative similarity" if it tends to produce beliefs that there is a corresponding objective similarity between the objects/events that the individual qualitative experiences are about.[23] Likewise for the difference relationship – a relationship between qualitative experiences is one of "qualitative difference" if it tends to produce beliefs that there is a corresponding objective difference between the objects/events that the individual qualitative experiences are about. These beliefs about objective similarities/differences can themselves be functionally defined in terms of tendencies to assent to or dissent from sentences about the objective similarities and differences of the objects or events in question.

Consider a simple example. I have two qualitative experiences of color – one of the color of a cube and one of the color of a sphere.

---

[20] Ned Block, "Are Absent Qualia Impossible?" *The Philosophical Review* 89, no. 2 (April 1980): 257.

[21] Sydney Shoemaker, "Absent Qualia are Impossible," *The Philosophical Review* 90, no. 4 (Oct 1981): 597.

[22] Sydney Shoemaker, "Functionalism and Qualia," *Philosophical Studies* 21, no. 5 (May 1975): 305.

[23] Ibid.

The relationship between these two individual qualitative experiences will be one of qualitative similarity if it produces in me a belief (a tendency to assent to sentences) that there exists an objective similarity in the color of the cube and the sphere. The relationship will be one of difference if it produces in me a belief (a tendency to assent to sentences) that there exists an objective difference between the color of the cube and the sphere.

Providing such a functional analysis of the relationships of qualitative similarity and difference would ultimately fix all the qualitative facts in reference to the physical facts of our world. For if such a functional analysis of these relationships can be given, then the relationships of qualitative similarity and difference must be physically realizable (because functional analysis defines things in terms of causal inputs and outputs,[24] and causal laws are part of the physical world).[25] If the relationships between qualitative experiences are physically realizable, then they will be fixed in reference to any given physical realization. That is to say, since these relationships would be realizable by our physical facts, the relationships between qualitative experiences would be fixed in reference to our physical facts.

What would it mean to fix the relationships between qualitative experiences? Fixing all the similarity/difference relationships between qualitative experiences would require that the qualitative character of each individual experience be fixed as well (if the qualitative character of an individual experience was allowed to fluctuate, then its relationships of similarity and difference to other qualitative experiences would fluctuate as well).[26]

Finally, since the "qualitative character" of a given experience is simply whatever particular quale is being instantiated by the experience, fixing the qualitative character of each experience entails that all individual qualia that are experienced are fixed as well. Therefore, if the functional analysis of the similarity/difference relationships between qualitative experiences is successful, all qualia will be entailed by our physical facts. That is to say, there would be a causally necessary entailment from our physical facts to consciousness (all the qualia).

As noted above, the success of this argument would establish the inconceivability of zombies based on the arguments given in this paper. Ned Block calls the success of the argument into question,

---

[24] Janet Levin, "Functionalism," *Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta (Summer 2010): <http://plato.stanford.edu/archives/sum2010/entries/functionalism/>.

[25] Sydney Shoemaker, "On David Chalmers's *The Conscious Mind*," *Philosophy and Phenomenological Research* (June 1999): 440).

[26] Ibid.

posing an objection to the first portion of the argument – that a functional analysis can be given of qualitative similarity and difference.[27] Shoemaker in turn replies to Block's objection.[28] Their discussion centers on the possibility of a situation extremely similar to the zombie case, which they term the "absent qualia" case. The absent qualia hypothesis, for Block and Shoemaker, is the possibility that two people may be functionally identical while one has states with qualitative character and the other has states that lack qualitative character.[29]

Where zombies are used to argue against physicalism, cases of absent qualia are employed analogously against functionalism. To examine the exchange between Block and Shoemaker on this matter with any amount of care would involve a discussion of functionalism at large, and is thus beyond what I am able to accomplish in this paper. What I hope to have shown with this final section is that there is a plausible case to be made for the existence of a causally necessary entailment from our physical facts to consciousness, thereby adding support to the position that zombies are in fact inconceivable.

References

Alter, Torin. "Imagining Subjective Absence: Marcus on Zombies." *Disputatio* 2, no. 22 (May 2007): 91-100.

Block, Ned. "Are Absent Qualia Impossible?" *The Philosophical Review* 89, no. 2 (April 1980): 257-274.

Chalmers, David. *The Conscious Mind.* New York: Oxford University Press, 1996.

Hill, Christopher and Brian Mclaughlin. "There are Fewer Things in Reality Than are Dreamt of in David Chalmers's Philosophy." *Philosophy and Phenomenological Research* 59, no. 2 (June 1999): 445-454.

Kirk, Robert. "The Inconceivability of Zombies." *Philosophical Studies* 139, no. 1 (May 2008): 73-89.

Kirk, Robert. "Zombies." *Stanford Encyclopedia of Philosophy.* Edited by Edward N. Zalta. Spring 2011. http://plato.stanford.edu/archives/spr2011/entries/zombies.

Kirk, Robert. *Zombies and Consciousness.* Oxford: Clarendon Press, 2005.

Shoemaker, Sydney. "On David Chalmers's *The Conscious Mind.*"

---

[27] See Block, "Are Absent Qualia Impossible?"

[28] See Shoemaker, "Absent Qualia are Impossible"

[29] Block, "Are Absent Qualia Impossible?" 258.

Philosophy and Phenomenological Research 59, no. 2 (June 1999): 439-444.

Marcus, Eric. "Why Zombies are Inconceivable." *Australasian Journal of Philosophy* 82, no. 3 (2004): 477-490.

Shoemaker, Sydney. "Functionalism and Qualia." *Philosophical Studies* 27, no. 5 (May 1975): 291-315.

Shoemaker, Sydney. "Absent Qualia are Impossible – A Reply to Block." *The Philosophical Review* 90, no. 4 (Oct 1981): 581-599.

Levin, Janet. "Functionalism." *Stanford Encyclopedia of Philosophy.* Edited by Edward N. Zalta. Summer 2010. http://plato.stanford.edu/archives/sum2010/entries/functionalism.

Killing and Letting Die: An Ineffective Distinction

AMANDA GORMAN
University of Mary Washington

The distinction between causing and allowing harm, particularly between killing and letting die, has undergone great philosophical scrutiny in recent years, primarily because of its bearing on the issues of end-of-life care and euthanasia. However, the status quo in the public sphere is to acknowledge a distinction between the two. In fact, the American Medical Association's position is completely in accordance with the position that so called "active" euthanasia is morally unacceptable, whereas "passive" euthanasia is allowed to be carried out for patients by physicians and is even obligatory for physicians in certain cases of patient demand (Decisions). In practice, under this demarcation, many cases have surely been appropriately judged, but the theoretical foundation of the active/passive distinction demands a closer look. Making a distinction between active and passive euthanasia requires the framework of a theory of action that can incorporate the acts and omissions doctrine.

Many theorists have tried to defend this doctrine in regards to the distinction between killing and willfully letting die with a number of different rationales, but none have as of yet been accepted as conclusive. If such a theory were to exist, it should not only capture a difference in the way we use the terms "kill" and "let die" but also provide a normative moral basis for such a distinction. In this

paper I will show that the current arguments for a distinction have each failed to meet one or both of these requirements. I will then propose that no such distinction exists, and explain what accounts for the gap that needs to be bridged between common-sense intuitions about the distinction and the lack of reasoned justification for it. I will primarily analyze these matters outside the context of the euthanasia debates, but will end with a brief discussion of how accepting the falsity of the acts and omissions doctrine might affect the direction of the ongoing conversation about euthanasia.

Perhaps the reason many have so fervently defended the distinction between killing and letting die is because the equivocation of the two terms would presumably make us accept responsibility for deaths of those who live millions of miles away that could have been prevented if aid was given. This runs counter to the common-sense viewpoint that tends to regard the needy as not having a positive right to assistance. As Philipa Foot states, "we are not inclined to think that it would be no worse to murder to get money for some comfort such as a nice winter coat than it is to keep the money back before sending a donation to Oxfam or Care. We do not think that we might just as well be called murderers for one as for the other" (816). However, a refusal to uphold the distinction between killing and letting die does not necessarily mean that watching someone die before your eyes and not saving them, and allowing someone to die in an impoverished country are equivalent. Other factors such as difference of intent or physical distance from the situation may, in fact, affect the determination of moral responsibility.

The moral relevance of physical distance to a situation may seem arbitrary, but upon further reflection, physical distance is almost always a factor in our moral judgments. Failure to tend to the needs of people in distant countries tends to stem from practical concerns such as lack of awareness, and uncertainty of outcome. Furthermore, as Peter Singer points out, there is a lack of malicious intent present when failing to give aid, which would ordinarily be present in a situation of watching a person die and failing to intervene (162). A more thorough examination of the relevant moral factors related to distance is a topic for another paper, but the mere presence of these other factors shows that it is not necessarily the distinction between killing and letting die that is solely responsible for the idea that people are not murderers for failing to provide aid.

Others have tried to defend the distinction by saying that a man would be responsible for the death resulting from a killing, but if he has let someone die he has merely let nature take its course. This argument faces two serious challenges. One is the problem with defining human action outside the scope of "nature". The other is that it remains to be proven that there is some sort of inherent value

in letting nature take its course; this is the normative aspect that would need to be justified for the distinction to be a morally important one. Yet it is too often that this argument goes unchecked and unquestioned. As John Stuart Mill writes, "[the word "nature"] has thus become entangled in so many foreign associations, mostly of a very powerful and tenacious character, that [it] ha[s] to excite, and to be the [symbol] of, feelings which [its] original meaning will by no means justify, and which ha[s] made [it] one of the most copious sources of false taste, false philosophy, false morality, and even bad law" (8). It seems one of the two ways "nature" can be defined is as the sum of all things that take place, in which case it seems peculiar to discount human action from falling within the scope of this definition. If nature is defined in a way that does exclude human agency, on the other hand, it seems then that everything we do is outside of "letting nature take its course".

Imagine Bob has been fatally poisoned by Ed and can only be saved with an antidote that Joe has. It is hard to imagine what might be meant by giving a definition of nature that includes the results of Ed's poisoning Bob while simultaneously excluding Joe's actions of administering an antidote. Even if such a definition of nature could be found, it would remain to be proven that acts that fall within the scope of nature are inherently more valuable than acts that fall outside the scope of nature. This kind of justification seems unlikely when no grounds for distinction have even been drawn.

Perhaps, then, as some have thought, the distinction lies in whether or not one has caused the death of another. On the surface it seems as though there may be a clear correlation between the term "killing" and a sense of causing a death, whereas we can find another cause in cases of "letting die". To test the validity of this claim, it seems essential to inquire a bit into nature of causation in general. What do we mean when we make a statement about a single event causing another event? Nothing is caused by one thing alone, but rather, there are always multiple factors involved. Even the most seemingly simple cause and effect relationships in the physical world do not have single causes. For example, a ball rolls when pushed not only because the right amount of pressure was exerted onto it, but also because someone decided to push it, because the ball was round, because the ground was not sticky, etc. If any of these were not the case, the ball would not roll.

There is no inherent hierarchy of these factors in determining whether or not the ball will roll. We pick out the amount of force as the relevant cause because it is relevant to something; perhaps we are trying to explain a principle of physics. Similarly we pick out one particular cause of a death because of our moral and practical

intuitions about the relevancy of the various causes leading up to a death. It is not from some indicator of the one most important cause that we form our moral ideas, but rather it is from the fact that we have brought our preexisting moral ideas into our interpretation of causation. Once we acknowledge the presence of multiple factors in all causal relationships it seems strange to exclude a person's "inaction" of not saving someone when they had the choice to as one of the causes of the death.

In cases of killing it seems like an agent is calling the other's life to be called into question, compared to cases of letting die where the life is already called into question and he must merely decide if they should die or not. In reality, though, life is called into question every second of every day. Each time someone decides to stay on the road whilst driving, each time he decides not to step off a balcony, he is making the choice to continue his life. People similarly make decisions not to push their friends off of balconies and to call 911 when someone is injured, and so their own lives and the lives of those around them are constantly at the risk of the decisions that they make –the constant choices of action and so-called "inaction." The realization of this fact also seems to make moot a point of distinction raised by Jonathan Bennett. He attributes the difference between killing and letting die to a difference between cases in which most of the ways someone present could have acted would have lead to the person's death, and cases in which most of the ways someone present could have acted would not have lead to the person's death (48).

This seems hard to judge without factoring in some sort of boundary to the possibilities of the action of the person present. If someone is choking on a pretzel it seems like the choices of another person present are between dislodging the pretzel, and doing other things like standing around, or walking away. But if all possibilities are to be accounted for, the possibility that he will drop a piano on his head, or shoot an arrow through his heart should also be included. Making this kind of speculation of all possible actions the person present can take necessary to evaluate the morality of someone's role in causing a death is clearly absurd.

It seems then as though we might be able to group together all actions that result in the death of a person as being equal. In fact, many who have argued against the distinction between killing and letting die have come at it from the standpoint of consequentialism. James Rachels notes how naturally this position sits within this kind of framework: "killing and letting die have the same consequences--the victim ends up dead--and so, if rightness is determined by consequences, killing and letting die should be morally equivalent" (750). Yet this viewpoint seems to be an oversimplification of the

matter at hand. This is part of the similarity between killing and letting die, but a stronger claim may also be made about the similarity of intentions in many of these cases. In situations where it is easy to help, to prevent a death, a person must have a lot invested either in something else or in the death itself for him to allow it to happen. In cases where he is not invested in the death itself but in doing something instead that he has determined to be more important than saving a life, his morals should probably still be called into question, for there should probably be nothing that he should find more morally important than stopping someone from dying. This is no different from the moral judgment most people pass on paid assassins who are more invested in the money they will receive than in the actual death of the person. The commitment to the outcome of another's death remains, regardless of the reason.

Another argument for a moral distinction has been advanced by Richard Louis Trammell. The difference between killing and letting die that he points out, which he calls "optionality", is that when someone merely lets another die, he is allowing for the logical possibility that someone else might come along and save the dying man. As Richard Brook rightly points out in his rebuttal to Trammell, whereas the logical possibility for survival may still exist, in many cases no physical possibility exists (198). Perhaps there is no one else in the room, or no one else who knows how to save him. This distinction, though initially attractive because of its descriptive accuracy, runs into the problem of moral relevance. There is no reason why the mere fact that it is not a logical truth that the person will die should be a matter of moral importance.

Furthermore, Trammell's distinction may not be as much in line with common moral thought as it first seems. If Trammell's distinction holds, it follows that we are less responsible for saving lives if there are other people present who could also save that life, who, as Brook importantly notes, are equally unaware of what the others will do (198). That our duties diminish when we are in a crowd seems neither intuitively true nor rationally true. This particular case may be hard to imagine because in real life we can often predict the behavior of others to some degree.

If there was no knowledge of what the others might or might not do to save a dying person there would be nothing by which to alter the perceived duty of a man in a crowd to save the dying person. For Trammell's distinction to be meaningful, though, one should also be able to imagine a case in which it is physically impossible for the man to be saved by anyone else and agree that its logical possibility is what significantly differentiates it from cases of killing. It seems Trammell may have pointed to a small descriptive difference in the way that we use the terms letting die and killing,

but not one that accounts for the grossly different reactions we have morally to hearing about cases in which someone has killed and those in which someone has let another die.

What then, does account for our common-sense intuitions that seem to point to an important distinction between killing and letting die? Despite the lack of an adequate rational and moral distinction between the two, if someone did in fact want the death of a man, and had the choice between shooting him (a case which would be referred to in our society as "killing"), and having him choke on a pretzel and not intervening (a case which would be referred to in our society as "letting die") I would still most likely choose the latter. His motive in the choice would be reliant, however, upon the fact of other people in his society believing that there is an important distinction between the two options.

If he lived in a society that did not have equivalents of the phrases "killing" and "letting die" but instead lived in a society which commonly used the phrase "let a life end" and considered the details as to how the person went about letting the life end a secondary matter of clarification, he would not have particularly favored either the choking option or the shooting option. To drive the point home even stronger, if his society were to first draw the line on a perhaps more morally relevant ground such as identifying these acts as "maliciously letting a life end" and "non-maliciously letting a life end" and then identifying his part in the matter as an action or inaction, he would feel that he would be in equal trouble for the allowance of pretzel-choking as for the killing because he would have maliciously let a life end in either circumstance. Because his society, though, has used the terms killing and letting die and assumed importance in the distinction, there now has become an actual slight difference in his choice to "let die" rather than "kill", all other factors being equal. This is not morally relevant as his choice to let die is merely a choice to evade responsibility for an act that he know is just as heinous. (It is not out of the question that he should be punished even more harshly for trying to get away with maliciously letting a life end in that manner.)

This reasoning that actually affects choices between "killing" and "letting die" is what tends to get glossed over as "common-sense moral intuition", but is really better explained as a willingness to buy into the false authority of the distinctions of ordinary language in determining moral value. The consequences resulting from awareness of the perceived distinction end up accounting for what is perceived of as the distinction itself. Since there is no perceptible secure normative basis for morally distinguishing between killing and letting die, people should not be swayed by the fact that such a distinction is in fact used in our common language.

There are important consequences to not supporting a moral distinction between killing and letting die. Taking away the active/passive distinction changes the conversation about end of life decisions, forcing medical ethicists to reexamine old cases. Of particular interest are cases free of malicious intent in which it is not really the death of the person that is wanted, but rather, there is something else wanted more than their continuing living. Euthanasia cases of stopping or not beginning treatment will oftentimes fall into this category. If it is no longer the fall-back to justify these cases as "passive" or as cases of "letting nature take its course", the real ethical value claims at stake can be more clearly evaluated. This may bring into the foreground vital questions including whether or not there should be a monetary limit for keeping a person alive with the aid of expensive treatment, or whether there is a stronger duty to alleviate suffering or to value the sanctity of life.

Examination of this sort may result in awareness that in some cases, there is too much at stake in saving a person's life. On the other hand, it may lead to the resolve that the duty to preserve life outweighs in these circumstances. Either way, the proper method of analysis must start with a disavowal of fruitless distinctions, and a refusal to evade the reality that these cases all meaningfully deal with ending life in just as serious a way as taking someone's life with a bullet or with poison.

References

Bennett, Jonathan. "Morality and Consequences." *The Tanner Lectures on Human Values*. II. S. McMurrin (ed) Salt Lake City: University of Utah Press, 1981.

Brook, Richard. "Dischargeability, Optionality, and the Duty to Save Lives" *Philosophy and Public Affairs*, Vol. 8, No. 2 (Winter, 1979), pp. 194-200.

"Decisions Near the End of Life." CEJA Report B – A-91. American Medical Association. April 18. <http://www.ama-assn.org/ama1/pub/upload/mm/code-medical-ethics/220a.pdf>

Foot, Philippa. "Killing and Letting Die". *Ethics: History, Theory, and Contemporary Issues*. Cahn, Steven M. and Peter Markie. NYC: Oxford University Press, 2009. pp. 815-821.

Mill, John Stewart. "On Nature." Lancaster E-text 1996. <http://lancs.ac.uk/users/philosophy/texts/mill on.htm>

Rachels, James. "Killing and Letting Die." *Encyclopedia of Ethics*, 2nd edition, ed. Lawrence Becker and Charlotte Becker (New York: Routledge, 2001), vol. 2, pp. 947-50.

Singer, Peter. *Practical Ethics*. 2nd edition, Cambridge, 1993.

Perceptual Parameters and Intentional Content

Brian Tracz
Boston College

Introduction

A classic problem of perception (PoP) goes something like: how does perception, an observer-dependent mental state capable of justifying epistemic states, take as its object non-epistemic, observer-independent physical things? This paper does not dissolve this problem but presents a third way aside from the recent dichotomy, that between conceptual and non-conceptual content, upon which recent solutions have relied. A few Interesting Considerations bring the PoP into focus:

(IC 1) *Perception is world-presenting and has unique phenomenal qualities.* Objects are manifest in perception, and there is something it is like to perceive *x*.

(IC 2) *Perception is necessarily an immediate, conscious experience.* Unlike believing *x*, perceiving *x* is necessarily a conscious event. Much of the epistemic weight of perception is derived from this immediacy.

(IC 3) *Perception can be transitive or intransitive.* One can perceive something of which one is not directly aware, whereas it is impossible to judge "that *x*" or to imagine "*x*" if one is not also aware of *x*. Even in attending to *x*, there are "fringe" percepts *y* and *z* that can be recalled at a later time.

(IC 4) *Perception is not meaning-neutral.* Photographs are meaning-neutral in that they do not bias one part of the image (minus the effects of focusing the camera lens). Perception is biased

toward one thing, namely, the object *x,* and usually toward only one aspect of *x* at that.

I will not further pursue a direct argument for these considerations, but I hope that they are somewhat intuitive. Sections I and II describe sense data theory at length, which claims, among other things, that perception essentially involves non-conceptual contents; sense data theory is meant to be the "negative" position which a type of conceptualism is specifically designed to avoid. I define this conceptualist position that denies non-conceptual content, critique it, and then adopt a picture of intentional acts that avoids the shortcomings of conceptualism and sense data theory. I argue for a different type of perceptual content—Parametric Content—based on an analysis of Husserl. I show this by developing a fifth Interesting Consideration:

(IC 5) *Perception depends on a passive, associative time-consciousness.*

Some table-setting remarks: Intentionality is a technical term to delineate the content of certain experiences. A denial of perceptual content amounts to a denial of its intentionality under this definition. Content is usually thought of as representational in some sense, though it is not necessary to argue for or against this here. To say that perception is conceptual has recently been taken to mean that the content of perception is of the same sort as the content of other intentional states, "that it is the sort of thing one can also, for instance, judge."[1] This is usually coupled to an argument denying that passively received, non-conceptual perceptual content exists: "receptivity does not make an even notionally separable contribution to the co-operation" that results in perception.[2] These remarks are oft quoted and, I believe, represent the present state of discussion.

Section I

Sense data theory (SDT) broadly denies any conceptual or propositional perceptual contents. SDT responds to the PoP by claiming that the *objects* of perception are the same as the very *contents* of visual experience themselves, which, in turn, are reliably linked, perhaps causally, to things and states of affairs.[3] SDT maintains that sense information is an *immediate* deliverance to the mind, whereas add-ons to this data, such as theories, inferences, judgments, and the like, are formed quite long after the fact from this non-conceptual perceptual content. So it amounts to three major claims regarding perception:

---

[1] McDowell, *Mind and world*, 9.

[2] Ibid.

[3] For a complete discussion, see Searle, *Intentionality*, 53–61.

(a) The contents of perception are sense data of which we are *immediately* aware.

(b) Sense data indirectly represent *properties* of things or states of affairs to the perceiver.

(c) Perception is reliable. This justifies *a posteriori* knowledge.[4]

For instance, John Locke maintains a paradigmatic SDT position. In his view, perception (thought/understanding) and volition (willing) are the only active powers of the mind, and these active powers never precede the reception of ideas from the senses.[5] These ideas are the immediate contents of perception, satisfying (a). Furthermore, perception (the understanding) and volition cleanly deal with ideas, whereas external objects instantiate "primary" qualities. Primary qualities, as Locke puts it, are "utterly inseparable from" external objects, whereas secondary qualities, unlike primary qualities, are merely powers of an object to excite a certain sense.[6] Ideas of primary qualities resemble external things, whereas ideas of secondary qualities do not bear this resemblance.[7] In practice, though, secondary qualities reliably represent physical objects, satisfying (b) and (c). The notion of secondary qualities as separable from active spontaneity is, in effect, to deny that perception is conceptual.

Though we need not burden ourselves with Locke's formulation, it suffices to illustrate that the goal for SDT is to derive non-inferential knowledge of objects from separate foundational building blocks—sensations or "ideas"—that the senses render as properties. The separability of these building blocks from subjective spontaneity is taken to vindicate observer independent objectivity: the object of perception simply *is* this set of separable sensation-properties. As a solution to the PoP, SDT seems at first glance to be an appealing option. The trick in addressing the PoP is to avoid asserting, *tout court*, that epistemic states are composed of simpler non-epistemic states. SDT claims to avoid this critical issue by positing an intermediate, that of sense data, which already has all of the trappings of epistemic currency. Before considering this further, here are the SDT accounts of each Interesting Consideration.

(IC 1) SDT has trouble here. We are assured by claim (c) that sense data indirectly, though accurately, represent the world. However, it is not clear how the world itself can be an epistemic basis of justification of those sense data. How, for instance, can we suppose that the same primary, public qualities of the world will be represented with the same secondary qualities to different observers? This rests

---

[4] For an extended treatment and history of sense data, see Huemer, "Sense-Data."

[5] Locke, *An Essay Concerning Human Understanding*, II.vi.2; II.i.20.

[6] Ibid., II.viii.9.

[7] Ibid., II, viii.15. Locke averts skeptic counterarguments by arguing for a reliabilist epistemology, appealing to the fact that, in practice, secondary qualities represent *objects* as they really are.

upon a reliabilist epistemology; with a defeat of this epistemology goes the world-directed nature of SDT.

(IC 2) Not a problem for SDT. Perceptions are, quite simply, the sensing of data. Sense-data motivate any perceptual experience.

(IC 3) SDT does an admirable job with non-transitive perception. There seems to be a framework for dealing with data that are, so to speak, never processed or attended to. Transitive, object-directed perception, as we will see, is more problematic. This will shed doubt on SDT's ultimate capacity to deal with non-transitive perception.

(IC 4) SDT fails here. SDT does not explain how certain objects of perception, and thus certain data, are biased over others. All ideas (data) from the senses are conveyed to the observer, as Locke notes, before any active biasing by the observer occurs. This seems, at best, a phenomenologically thin analysis.

Section II

Despite its initial appeal in dissolving the PoP, sense data theory *does* claim that epistemic facts are derived from non-epistemic building blocks.[8] SDTs accept perception as an autonomous receptive capacity that collects simple sense-data. These data, furthermore, need not be manipulated or translated, for the information (or, for Locke, the idea) already represents, in a reliable way, external objects. In his *Empiricism and the Philosophy of Mind*, Wilfrid Sellars provides a trenchant criticism of sense data theory. SDTs, Sellars argues, equate having sense contents with being conscious.[9] In other words, another SDT hypothesis is, according to Sellars:

(d) Reception of immediate sense data of x is *sufficient* for consciousness of x.[10]

In other words, SDT allows that passive reception of qualities is sufficient for an intentional mental episode in normal waking life, a characterization consistent with Locke's theory that thought does not precede the reception of ideas from sensation. The objects of perception, in a sense, are "prepackaged" with an organization that makes them epistemically accessible; no activity on the part of the perceiving subject is required. According to Sellars, much of the heavy-lifting that is really performed by the conceptual capacities of the subject is mistaken by SDT as being performed by an extra-subjective entity, like the object itself. Sellars famously labels this problem the 'Myth of the Given.'

---

[8]   See Sellars, *Empiricism and the Philosophy of Mind*, § 5.

[9]   Sellars, § 6.

[10]   This might not seem quite fair to Locke, since he holds that recollections are also sources of simple ideas. However, one might take Sellars' stronger interpretation of Locke's position (i.e. that ideas received in perception are *sufficient* for conscious episodes) if one notes that no ideas are innate and that they are all ultimately derived from sensation.

I find this criticism persuasive: the Myth of the Given under-mines the SDT answer to the PoP. It points out, first, that being conscious of *x* does indeed involve more than sense reception from *x*. Thesis (d) is false. For example (and to illustrate IC 3 and 4), I might walk around a truck, inspecting it from many sides, in order to determine its overall build, its durability, and other things about it. In this type of experience, I am *transitively conscious* of the truck: I want to know more about the truck, so I attend to an object—the truck—in order to determine some of its properties. Transitive con-sciousness is "consciousness of" something, and it is something in which I, the perceiver, take interest. Indeed, it is often (and usually harmlessly) equated with *interest*, *attention*, or *awareness*. Interest, however, cannot be accounted for purely in terms of sense data. Of course, sense data will *cause* attention to be swayed, but this does not account for the part-active aspect of my *interest*, given a person-al history and background, in viewing certain items like trucks. In failing to account for IC 3 and 4, SDT does not provide an adequate dissolution of the PoP. It is this active or spontaneous aspect of the perceptive event that needs further development.

Section III

To that end, the remainder of the essay is a partial attempt to come to terms with Kant's dictum: "Thoughts without content are empty, intuitions without concepts are blind…. The understanding cannot intuit anything, the senses cannot think anything."[11] This, it seems, summarizes Sellars' diagnosis of the problem with SDT. In light of this, Sellars reconstructs sensation within a Kantian transcendental framework by defending a deeply *intentional* account. I will review this as a segue to John McDowell's treatment of the PoP. Though I am much indebted to McDowell's analysis of Sellars in his Wood-bridge Lectures on this analysis, my argument ends up in a quite different destination from his.

In his Woodbridge Lectures, McDowell illustrates Sellars' overall scheme well with an "above the line"/"below the line" image. In short, "above the line are placings in the logical space of reasons, and below it are characterizations that do no do that."[12] Sellars, in the Myth of the Given, essentially warns against a philosophical attempt to make things *below* the line—natural states of affairs—do the work of things *above* the line—epistemic states of affairs. In this illustrative apparatus, the area above the line corresponds to con-ceptual capacities. McDowell thus restates the central thesis of the Myth of the Given: "The space of reasons, the space of justifications

---

[11]   Kant, *Critique of Pure Reason*, A51/B76.

[12]   McDowell, "Lecture I," 433.

or warrants, extends more widely than the conceptual sphere."[13] This observation holds good for our analysis: SDT seems to resolve the PoP because certain non-inferential, non-epistemic data serve as objects for perception. As objects of perception, sense data serve to justify beliefs. This results in an enlargement of the space of reasons beyond the conceptual sphere, and a descent into the Myth.

It is precisely this cross-the-line contamination that Sellars and McDowell want to avoid. However, the PoP is precisely a *problem*, so avoiding the Myth is not sufficient; Sellars and McDowell need to present a view that more coherently dissolves the PoP itself. Let us consider Sellars' view first, as it shapes McDowell's subsequent view. As Sellars argues, to describe an experience as a "seeing that x" is, indeed, to do more than merely describe an appearance.[14] In referring to a state such as "seeing that x," one is characterizing such a state as making a *claim.* All assertoric statements both describe a state of affairs and *endorse* that statement as being true to fact. The force of Sellars' argument comes in the claim that, from the perceiving subject's point of view, these belief-claims are the content of perceptual experience. Tim Crane calls this the *Propositional Attitude thesis*.[15]

Sellars thus wants to show that perceptual contents are "propositional claims."[16] Sellars proceeds to note that (1) "x looks green to Jones" and (2) "Jones sees that x is green" are different descriptions of Jones' experience. Description (1) extracts a proposition from Jones' experience suggesting that Jones sees a "green something" while remaining agnostic on the status of object x. Description (2) extracts a claim from Jones' experience illustrating Jones' endorsement of a particular claim about object x—"that it is green." According to McDowell's reading, Sellars denies that one can separate one's description of some *object-neutral appearance* (1) from one's endorsement of some *object-centered proposition* (2). If propositions merely *express* the placement of objects under concepts, then Sellars is restating that intuitions without concepts are blind: there is no mere appearance of "greenness" without some concept—"amorphous blob" or "triangle"—being associated with it.

This is worth further explanation. Suppose Jones says, "x looks green to me now." Sellars parses Jones' statement as reporting "the fact that [Jones'] experience is … intrinsically … *as an experience*, indistinguishable from a veridical one of seeing that x is green."[17] In

---

[13]   McDowell, *Mind and world*, 7.

[14]   Sellars, § 16.

[15]   Crane, "Is Perception a Propositional Attitude?".

[16]   Sellars, 39.

[17]   Sellars, 41.

other words, experience has one epistemic mode—that of validating a certain claim about objects. Thus, Sellars writes, "x is red iff x looks red to standard observers in standard conditions."[18]

This equation equates an epistemic state (on the left) with a state that was previously thought to be a neutral description of a non-epistemic appearance (on the right). For Sellars, 'epistemic' tends to connote something like 'concept involving.'[19] Put differently, Sellars does not only note that "intuitions without concepts are blind," but also claims that conceptual content is propositional—the Propositional Attitude thesis. It is the latter thesis that I will reject below.

In analyzing Sellars' position, McDowell remarks that Sellars implies that *ostensible seeings* (episodes of "seeing that x") contain their claims "as ostensibly *necessitated* by an object ostensibly seen."[20] This provides a way of disambiguating, at the level of conceptualization, perceptual episodes from other episodes, such as judgments, that are *conceptual* but nevertheless *free* or *spontaneous* exercises of conceptual capacities. For instance, one could judge that x is fat, where x is Barack Obama, purely spontaneously and incorrectly. However, one would never be able to extract the proposition "Barack Obama is fat" from any sort of veridical perception of the man himself.[21] But, in standard conditions, the point is that although *concepts*, like 'fat,' are a 'common currency' between perceptions and judgments, a *particular* conceptualization occurs with a certain necessity in perception and sensation, a necessity lacking in other episodes like judgment.

This distinguishes perception from other modes experience, so what's left to be seen is how perceptive experiences themselves are exhaustively described in epistemic terms that are "above the line," that are conceptual. In his reading of Sellars, McDowell claims that

> Conceptual episodes of the relevant kind are already, as the conceptual episodes they are, cases of *being under the visual impression* that such-and-such is the case. It is not that as conceptual episodes they are phenomenologically colorless, so that they need to be associated with visual sensations in order that some complex composed of these conceptual episodes and the associated visual sensations can be recognizably visual.[22]

With Willem deVries,[23] I take McDowell's reading here to be that

---

[18] Sellars, 43.

[19] I follow McDowell, "Lecture I," 433.

[20] McDowell (1998), 440.

[21] Understood here: in standard conditions by a normal observer with everyday concepts. The words in a proposition, again, are symptoms of the concepts from which they derive their meaning.

[22] McDowell, "Lecture I," 442

[23] DeVries, "McDowell, Sellars, and Sense Impressions," 182–201.

such visual experiences are intrinsically qualitative and phenomenal because of and through their conceptuality, accounting for IC 1. This position indicates that perceptive consciousness is not only *shaped* by concepts but is, indeed, *constituted* by conceptualization. At this point, I will simply note that this claim is, at least, ambiguous, and I will revisit it.

Section IV

We can give a preliminary evaluation of this view of perception in light of the Interesting Considerations. I will freely call the view elucidated above, originated by Sellars but developed and presently defended by McDowell, "conceptualism" for convenience (except, of course, where the two differ). This will avoid a critique of the accuracy of McDowell's reading of Sellars.[24] To start, an account of perception needs to disambiguate perceptive states from other states, such as belief, imagination, and emotion. One is left with the question of why perceptions are *immediate* and *epistemically primary* encounters with objects in a way that beliefs or judgments are not. This is the intuition behind IC 2.

It is unclear, however, how one can do this simply by altering the *modality* of the propositional content of conceptual experience. Ostensibly seeings, recall, have a particular conceptual content by *necessity*, whereas imaginings and judgings have their particular conceptual contents *contingently*. One can "decide" to imagine or judge as one likes in a way that one cannot "decide" what to perceive. Even if it is granted that perception involves asserting a proposition, which I contest below, this idea is still problematic. If the contents of perception have the mode of necessity, then how do we deal with such curiosities as "duck-rabbits"? One can, in essence, observe the same duck-rabbit and, based on a whim, change what one sees. If I observe a duck-rabbit, the propositional content of my perception would be either (a) "I see that this is a picture of a rabbit" or (b) "I see that this is a picture of a duck." I choose which propositional content I get, and I can literally "change my mind" as to what I see. Another interpretation might be: I see *necessarily* the same object, but it appears differently. This, however, is a forbidden interpretation, for it smacks of SDT; Sellars explicitly argued for the equivalence of "looks as if x" and "sees that x." If the duck-rabbit example is too peculiar, then there are also examples of subconscious priming which cause one to perceive certain images in certain ways based on a contingent experimental manipulation. If a

---

[24] Cf. DeVries, "McDowell, Sellars, and Sense Impressions." For instance, Sellars, in Empiricism and the Philosophy of Mind, actually seems to argue that conceptual states are separable from phenomenal states. See Sellars, Empiricism and the Philosophy of Mind, Part V and VI. As deVries points out, Sellars makes the distinction between ideas and impressions explicit.

subject is shown a many pictures of ducks, then one is more likely to see a certain image *as* a duck. In addition to these phenomenological problems, it is also unclear how and why the modality of perceptual contents are "flipped" from necessity to contingency as the conceptual contents are "detached away" as epistemic justifications for beliefs. At any rate, these considerations weaken the argument that sheer *modal* considerations of a proposition as sufficient to disambiguate perceptions from other states.

In a move that McDowell believes to be mistaken, Sellars endorses a type of "sheer receptivity" which, though itself non-conceptual, serves to guide the conceptualization of intuitions.[25] The view goes that "our entitlement to see elements in the conceptual order as intentionally directed toward elements in the real order has to be transcendentally secured from outside … the conceptual order."[26] McDowell, on the other hand, argues that we can "spell out a transcendental role for sensibility in terms of the immediate presence of objects to intuitionally structured consciousness."[27] Take, for example, a relation "that x is y." Therefore:

> In statements of meaning and aboutness, we relate the conceptual order to the real order, mentioning elements of the real order [e.g. x] by making ordinary uses of the words [e.g. y] on the right-hand side of these statements. But we affirm these relations without moving outside the conceptual order—without doing more than employing our conceptual capacities.[28]

So, in relating some real thing x to some descriptor y, we need not remove ourselves from our conceptual understanding in order to make this assertion.

This is insightful but altogether misses the issue at hand in the PoP. The form of assertoric statements might indeed have exactly this format. Concepts themselves might, too.[29] But, according to IC 2, we need to account not only for the *content* but for its *immediacy* in perception. Sellars tries to ensure this by arguing that the mind→world fit in perception is disambiguated from other states both conceptually (via a change in modality) and non-conceptually (via 'sheer receptivity'). McDowell claims that this mind-world fit is conditioned by nothing outside of the conceptual realm and that, therefore, we can expect perceptive events to be completely constituted by their conceptual, propositional content.

---

[25] McDowell, "Lecture II," 453.

[26] Ibid., 488.

[27] Ibid.

[28] Ibid., 489.

[29] McDowell notes in the Woodbridge lectures that this is essentially a Tarskian semantics applied to a Kantian transcendental method.

As Walter Hopp notes, this implies a 'Detachability Thesis':

> C is a conceptual content if and only if it can serve as the content of a mental state M in which the relevant objects, properties, and/or states of affairs that C is about are not perceptually given or present in M.[30]

If perception has conceptual content, which runs deeply from the constitution of objects (x) and properties (y) to the very relating of objects and properties (x is y), and if this conceptual content is detachable in whole from an initial perceptual event, then we have no way of saying, *merely considering the conceptual content*, that perceptions are epistemically prior to beliefs or imaginings. The Detachability Thesis prohibits this by drawing a deep equivalence between a concept across different experiential modes. We already noted the difficulties in modal distinctions at the conceptual level. Any appeal to sheer receptivity, remember, is not allowed for McDowell, because it is *not* detachable from the original perceptive event. On this point, McDowell's account fails to account for IC 2 and the aspect of the PoP associated with it. If conceptualism is to maintain the epistemic priority of intuitions, it is forced to abandon the Detachability Thesis, that is, it must acknowledge that there are special conceptual contents only had in perceptual experiences.

Conceptualist views are in tension with IC 3 and IC 4 as well. If perception is couched in terms of object-centered beliefs, then where does intransitive perception belong (IC 3)? Even though conceptualism accounts for IC 4 directly, there is still a lingering question of *why* certain objects and properties in one's visual field are biased over others. Consider Fred Dretske's claim that our sensory experience is informationally rich in a way that our cognitive, conceptual utilization of it is not.[31] For Dretske, perception is a channel for the analogue ("noisy" and "diffuse") reception of information, while cognition is the digital ("clear" and "discrete") extraction of relevant information from this analogue channel in order to form concepts. At first glance, this seems to complicate the PoP—inviting non-epistemic perceptions to feed epistemic cognitions sounds like the Myth of the Given—but Dretske's point here is worth consideration for a different reason. For example, suppose an observer sees some object (say, a paint chip sample) containing an assortment of shades of orange, ordered from brightest to darkest. Let's say there are thirty shades of orange, such that the transition from dark to bright occurs somewhat slowly as one observes the object. First question: does the observer *see* all thirty shades of orange? The answer, obviously by my estimation, is yes. The observer can see

---

[30] Hopp, "Conceptualism and the Myth of the Given," 370

[31] Dretske, Knowledge and the Flow of Information, 150.

the border of each shade to the next one. If there were 3,000 shades, it would certainly be more difficult or impossible to see the borders, but disambiguating thirty shades is not difficult. Second question: does the observer *conceptualize* all of the shades of orange?

The answer to this is more complicated. In one sense, all of the shades are conceptualized as (1) orange and (2) bright/dark. In another sense, how could orange$_{18}$ be *conceptually different* from orange$_{19}$? When an observer sees orange (1-30), she seems to be *cognitively limited* in the amount of conceptual "fine-grain-ness" that she can afford.[32] Phenomenologically, the observer sees "that object with a large spectrum of bright and dark shades of orange." The problem is that, in comparison to the apparent richness in perceptual phenomenology, the conceptual resources of the perceptual experience overall are impoverished and, rather quickly, exhausted.

As a response to this sort of objection, McDowell notes that it is possible to acquire the concept of orange$_{18}$:

> In the throes of an experience of the kind that putatively transcends one's conceptual powers—an experience that *ex hypothesi* affords a suitable sample—one can give linguistic expression to a concept that is exactly as fine-grained as the experience, by uttering a phrase like "that shade", in which the demonstrative exploits the presence of the sample.[33]

I urge here that McDowell implies a flawed account of visual phenomenology, which I will simply label *Retrospective Analysis*. It goes like this: We can, *in principle*, go back to a certain time-point in one's stream of visual phenomenology and cash it out in terms of conceptualized content, which for McDowell means in terms of propositional content. The problem is that perception has non-transitive contents at time *t* that one cannot simply *post facto* interpret as having been propositional at time *t*. As Crane observes, "describing the content and *being* the content are not the same thing."[34] Even if one perceives something and only later conceptualizes it—say, the ambient surroundings or temperature that one "subconsciously," so to speak, senses—this does not imply that this *post facto* judgment was derived from any *prior* conceptualization. We, the analytic audience of a hypothetical "recording" of a series of lived perceptual states, could extract a propositional content from those states.

I find no prima facie problem with the idea that one can, *in principle*, assign propositional content to any sensing or experience, though the fanciful notion of a frame-by-frame recording of phe-

---

[32] Both Dretske (in the same work) and Gareth Evans make this point. See Evans, The Varieties of Reference.

[33] McDowell, Mind and world, 56–7.

[34] Crane, "Is Perception a Propositional Attitude?".

nomenology is unappealingly atomistic and too 'photographic'. Nevertheless, I could even adopt a principle:

For any perceptual experience *E*, there is a sentence which gives the content of *E*.[35]

At any rate, one often does *post facto* conceptualize a stretch of his past experience when prompted in certain terms. Suppose that Jones is driving a car through the Pennsylvania turnpike in winter, attending to skillfully navigating the serpentine and icy roads. If Jones is asked *post facto*, "What color were the trees?" it is perfectly plausible that he could answer the question without having ever conceptualized such content at the same time he perceived it—he was too busy with the roadway. He could even answer, correctly, with such a peculiar turn of phrase as, "I *don't even remember* there being leaves on the trees at all!" or "I *seem to remember* the leaves being reddish." We are not concerned with such exercises of reflective judgment, such retrospective claims, but about Jones the *perceiver*. Hopp points this out in a thought exercise: "Think, for a moment, of all the objects and properties and relations you currently perceive, and how few 'that is this' (or 'that and that are related thus and so') style thoughts you are currently entertaining."[36] It seems that there is a large residue left over after the 'that is this' objects are subtracted, even though some of the residue is available for conceptualization at a later point. This illustrates how Retrospective Analysis does not work as a *method* to derive a thesis regarding conceptuality based merely on propositional content. Therefore, if propositions are symptoms of conceptualization, and if conceptual processes are methodically derived from propositions via Retrospective Analysis, then perception is not necessarily conceptual.

Here is a summary of the conceptualist report card as it stands:

(IC 1) Mixed success. On the one hand, the conceptualist view avoids the Myth of the Given and is seemingly tailored to put epistemic states in contact with real objects. Issues with the other Considerations complicate this.

(IC 2) Conceptualism fails here due to the Detachability Thesis and the subsequent lack of a way to vindicate it without appeal to modality of ostensible seeings. The epistemic priority ("immediacy") of perception is left hanging.

(IC 3) Conceptualism again fails. According the Retrospective Analysis, McDowell implies that any putative intransitive content could be cashed out as conceptualized, transitive content. Whereas I maintain that subsequent judgments are *primordial* propositional claims, McDowell maintains that subsequent judgments are merely

---

[35] Following Crane's formulation.

[36] Hopp, "Conceptualism and the Myth of the Given," 374.

recollections of a *prior* propositional claims. If conceptual content is simply a propositional endorsement, then I hold that conceptualization occurs after perceptive experience.

(IC 4) Conceptualism is, at best, weak here. It is not clear how there is an epistemological gradient in which a meaningful object in one's perceptive field is steadily and increasingly "better known" over a stretch of perceptions.

Section V

In the previous sections, I shed doubt on both conceptualism and SDT as answers to the PoP given a small list of Considerations. With SDT, I agree that perception has special features that enable an interface with the world. Its numerous problems have been outlined. With conceptualism, I agree that transitive perception requires conceptual guidance that doesn't succumb to the Myth of the Given. But this is not much of a concession. Against conceptualism, I have argued (1) that Propositional Attitudes are not perceptual content, and (2) that the Detachability Thesis does not disambiguate perception from other states or grant perception the epistemic primacy needed to dissolve the PoP, and (3) that Retrospective Analysis is a specious method by which to derive conceptuality. I have also indicated doubt that conceptualism can account for phenomenal content.

Even if the McDowell's account were descriptively adequate, in any case it could not be *sufficient* for a solution. If all of the above conceptualist strictures on perception were met, nevertheless this would not guarantee any perceptual experience to be epistemically foundational as a justification. I will argue that an associative time consciousness is an essential structure of perception (IC 5), and that this leads to a crucial distinction between determinability and active determination of content. This distinction hones the way conceptual content is characterized in the first place: I claim that perception has *Parametric Content*. Though this analysis suggests a way of understanding normalized perception, this is not the place to launch into a discussion of disjunctivism, and critical treatments of McDowell's flavor of disjunctivism can be found elsewhere.[37]

Rather, the goal here is to show that IC 5 is a key part of the dissolution of the PoP, which I take to be primarily a problem regarding intuitional content and determinability.

To motivate this discussion, it will be helpful to distill Edmund Husserl's understanding of time consciousness, upon which association is based.[38] Intentional acts rely on association and time-con-

---

[37] See Sturgeon, "Visual Experience."

[38] A compact account appears in Husserl, Experience and Judgment, § 23. A broader summary appears in Zahavi, Husserl's Phenomenology, 80.

sciousness. Husserl does not consider this a "theory" but, rather, transcendental conditions "without which subjectivity could not be."[39] Husserl, in short, argues that single, momentary experiences are actually very temporally robust events. The "lived present," he urges, is more than just a "now point"; it is also a constantly evolving series of passive expectations and passive recollections. To make this clear, he calls the passive expectation aspect of the present *protention*, in order to disambiguate it from everyday spontaneous expectations, intentions, and desires. He calls the passive recollection aspect of the present *retention*, in order to disambiguate it from everyday spontaneous memories, recollections, and imaginings. Protentions and retentions issue from a single *now-point*, which, in the context of our discussion, singles out a single timepoint in the overall constitution of a perceptual experience. This can be represented as in the diagram below.

As this diagram shows, it is a mistake to think of retention or protention as "acts" that a conscious subject can commission. One can initiate a recollection of something, whereas one cannot initiate a retention of anything. Retention and protention are passive background operators of consciousness; even an act of recollection itself has a retentional and protentional component since recollection is, itself, an intentional act.

All of these aspects of the present—retention, now-point, and protention—are at work in apprehending any object.[40] One cannot isolate the now-point, for instance, and attempt to characterize a single moment of apprehension in terms of a simple proposition, since the now-point is itself simply one of a series of *fulfilled* pro-

---

[39] Husserl, *Analyses concerning passive and active synthesis*, § 26. Husserl notes elsewhere: "Consciousness is inconceivable without retentional and protentional horizons, without a co-consciousness (althout a necessarily non-intuitive one) of the past of consciousness and an anticipation of an approaching consciousness (no matter how indeterminate it may be)." See Ibid., sec. 4. "Static and Genetic Phenomenological Method."

[40] Husserl often refers to apprehension as "interpretation."

('Lived') Present

Past Events   Retention   Protention   **Future Events**

Now Point

tentions.[41] For Husserl, the axis of fulfillment runs from *empty* acts (or "signative acts," like believing and propositional judgments) to *intuitive* acts (like perceiving), both of which are fundamentally different intentional acts.[42] Furthermore, all intentional acts have a matter and quality, and these can span between different modes of fulfillment. *Matter* denotes the aspect of a state responsible for its intentional directedness. This includes both (1) an 'objective reference' and (2) an 'interpretive sense' (or 'apprehensive sense'). *Quality* denotes the overall character—judgment, emotion, perception, and so forth. In summary:

> Matter: Sense, reference
> Quality: Presentation, emotion, judgment, ...
> Fulfillment: Empty, Intuitive

The distinction between intuitive and non-intuitive acts essentially appears twice—both under quality and under fulfillment. I will discuss its placement below.

In respect to the categorization of intuition under "fulfillment," how exactly is 'empty' distinguished from 'intuitive'? Husserl clarifies:

> The "horizons" of perceptions are another name for *empty intentions* (intentions in our clarified and precise sense) that are integrally cohesive and that are actualized in the progression of perception in and through different orientations. All of these intentions must accord with one another should the perception be normal.[43]

I will suggest that horizons of perceptions and empty intentions can be thought of as "parameters for further determinability" that consciousness passively sets. This grants perception itself with a type of Parametric Content. This, in my estimation, upholds the Husserlian insight that perceptual experience comes with "potential knowledge of something that has not yet come into view," a knowledge gained from future experiential orientations.[44]

These parameters await intuitive content (i.e. confirmation or disappointment), but they already have *intentional* content; that is, protentions are passively posturing the subject, creating "conditions of satisfaction," that can be realized in a subsequent, actual perception. I intend this to resonate with John Searle's theory of in-

---

[41] Husserl also puts this into epistemic language: "Belief is not only directed toward the present, but also toward the anticipated future and toward the memorial past; manifold memory-beliefs and expectant-beliefs emerge that can be verified or rejected." See Husserl, *Analyses concerning passive and active synthesis*, § 23.

[42] See Hopp, "Husserl on Sensation, Perception, and Interpretation." Husserl describes this in investigations 5 and 6 in Husserl, Logical Investigations, Vol. 1.

[43] Husserl, *Analyses concerning passive and active synthesis*, § 22.

[44] Husserl, *Experience and Judgment*, § 8.

tentionality and his idea of "conditions of satisfaction," though his theory deals more with a static description of certain qualitative intentional states like emotion, hope, and love.[45] The present account provides a framework in which to understand the genesis of such intentional conditions of satisfaction. Once intuitively fulfilled (or disappointed), the intentional content is more clearly determined, and each of the parameters that were set in advance are fixed with a higher (or lower) epistemic 'value.' This resonance between *requirement* (fixed by a protentional parameter) and *thing required* (a fulfilling intuition) provides a continuity over a series of now points that constitutes an object.[46] This anticipation of an object that is fulfilled results in a primordial form of *evidence*.[47]

Consider the modified diagram below.

The grey area indicates empty/signitive acts, aspects of the conscious present that constitute an object *over time*. The white box in the middle represents the area of intuitive fulfillment. Heuristically, the diagram shows that apprehension of an object is actually distributed over areas that are not themselves intuitively filled. The mistake that conceptualism makes is to suppose, in the terms of this framework, that spontaneity is describable as a "black box," namely, a simple propositional claim or endorsement (diagram below, on left). The mistake that SDT makes is to grant all of this associative work to non-subjective entities (diagram below, on right). Both result in a type of temporal atomism of varying degrees and types.

My suggestion is that a conscious form of evidence, the product of passively set parameters meeting certain conditions for satisfaction, provides the epistemic heft *at the conceptual level* to begin addressing the PoP. Since the conditions of satisfaction for *perceptive* states are intuitive fulfillment conditions, and since intuition itself fulfills empty intentions only insofar as these intuitions meet certain conditions of satisfaction that are *specified by the empty intention itself*, we have no problem of the Myth of the Given because only

---

[45] Searle, *Intentionality*, 10–36.

[46] Ibid., 13.

[47] Husserl, *Analyses concerning passive and active synthesis*, § 23.

Temporal Black Box

? Temporal and Conceptual Black Box

Object

Preformatted data ⟵ Object

"above the line" properties—like retention and protention—confine intuitive receptivity to specific *determinable* parameters. In perception, the deliverance of intuitive receptivity to these conditions is a *determination*. The point of an analysis of time-consciousness is to elucidate more clearly *where* in a perceptive act one could say that an "object" is manifest *at all*. Husserl more clearly defines the "present" and, in so doing, cashes out IC 2: perception has an immediate quality because it fulfills certain protentions and past retentions, both of which are *empty*—with intentional content but without a determinate object—before an intuitive event fulfills them. The above account can be summarized, therefore, as a conceptual account in that it is "above the line." It, however, rejects any type of Proposition Attitude thesis.

The language of "fulfillment" and "disappointment" will become clearer after an example.[48] Suppose Jones is attentively listening to a melodic tune of which he has no previous knowledge. He is naturally surprised at certain turns in the melody, and at other times he finds the melody predictable. Some parts of the melody are particularly interesting and novel, while certain other parts are less interesting and novel. Ryle lucidly describes the situation, worth extended quotation:

> To describe [the listener] as knowing the tune is at the least to say that he is capable of recognizing it, when he hears it; [that is,] if after hearing a bar or two, he expects those bars to follow which do follow; if he does not erroneously expect the previous bars to be repeated [etc.] …When we speak of him expecting the notes which are due to follow and not expecting notes or bars which are not due to follow, we do not require that he be actually thinking ahead. Given that he is surprised, scornful or amused, if the due notes and bars do not come at their due times, then it is true to say that he was expecting them, even though it is false to say that he went through any processes of anticipating them. In short, he is now recognizing or following the tune, if, knowing how it goes, he is now using that knowledge; and *he uses that knowledge not just by hearing the tune, but by hearing it in a special frame of mind, the frame of mind of being ready to hear both what he is now hearing and what he will hear, or would be about to hear*, if the pianist continues playing it and is playing it correctly [emphasis added].[49]

Recognizing a tune, gaining a certain epistemic affordance of it, re-

---

[48] This example, which Ryle gives without attribution, is famously given in the phenomenological literature of Husserl. Cf. Husserl, *Ideas Pertaining to a Pure Phenomenology and to a Phenomenological Philosophy*, 24.

[49] Ryle, *The Concept of Mind*, 226–7.

quires a series of expectations and upset expectations. The frame of apprehension is temporally broad, so object-constitution is also distributed over both empty and intuitive modes of fulfillment.

It might seem that this time structure has lost its "object directedness," but this is far from the truth. As Hopp argues, in transitive perception, a particular object (a *referent*) can be fruitfully considered to persist over a span of perspectives that constantly change (*senses*).[50] The idea is that this quasi-Fregean sense and reference schema can properly "pick out" the object of a perception over a span of experiences: "The *way of thinking of an object* to which the general Fregean conception of sense directs us is, in the case of a dynamic Fregean thought, a *way of keeping track of an object*."[51] Taking the more abstract example of a tune, the referent remains fixed—the (musically informed) listener is consistently interested in the overarching melodic line. The sense, however, is in flux—certain upward movements or unexpected melodic intervals change *how* the overarching melody line is understood.

Let us expand on this tracking idea, since the object-presenting nature (IC 1) of this account depends on it. In Husserl's language, it is the matter of an intentional act, with its interpretive sense and objective reference, that is modulated in tracking an object over a stretch of associative apprehensions. In Husserlian language, perception is *self-giving*: "Every perception within itself is not only, in general, a consciousness of its object, but [also] gives its object to consciousness in a distinctive manner…. It is not grasped mediately as if the object were merely indicated by signs or appearing in a reproduced copy."[52] This follows from the fact that perception is both intuitive *and* signitive—i.e. not purely propositional, indicative, or linguistic. Objects of perception are *expressible* symbolically or linguistically though they are not *disclosable* in symbols, contra Propositional Analysis. This is the key insight of fulfillment. Objects in perceptive acts are disclosed in ever-changing perspectives in a way that symbolic representations are not. Intentionality, the capacity for a mental event to focus itself on one object of interest in a relational manner, has a referential capacity to 'pick out' a particular object that persists over a stretch of time. But it is a mistake to take this referential capacity and run with it. In Husserlian terms, the Detachibility Thesis is mistaken in isolating perceptual content as a simple question of *matter*—sense and reference. Rather, perceptual content is also attendant on certain *partially* constituted objects and their different senses over time. In other words, perception has

---

[50] Hopp, "Conceptualism and the Myth of the Given," 377.

[51] Evans, *The Varieties of Reference*, 196.

[52] Husserl, *Analyses concerning passive and active synthesis*, 140. I will cite Husserl based on the English translation pagination.

Parametric Content that is unique to perceptual experience itself. This satisfies IC 2—perception has its own structure of immediacy in virtue of its Parametric Content.

With this framework, we can turn conceptualism on its head. Whereas conceptualism regards conceptual experience as the sole justification for any epistemic state, it is unclear, as mentioned above, how concepts account for phenomenal properties (IC 1). On the other hand, the sense/reference framework allows that the *phenomenal* character of perception can serve as a basic avenue of epistemic justification. It makes sense of phenomenal content, though it is not meant to describe *how* it arises.[53] This is the first step in accounting for IC 1. Imagine, for instance, a fire truck driving down a shady street. Initially, the fire truck's color is muted: the fire truck drives through a patch of shade that does not allow its color to be disclosed. The expectation is that it will, naturally, be intensely red. To ensure that the fire truck is not an abnormal color, you maintain it in your gaze. The truck, indeed, becomes more and more intensely red. Sensations differ fundamentally in degrees of intensity, but not in degrees of clarity.[54] Reasoning can be done well or poorly, more or less clearly, and it can yield true or false conclusions; but reasoning has no gradient of intensity. Sensation cannot be done well or poorly, clearly or unclearly; there are no "vague" sensations. But sensation can be more intense, and an epistemic affordance "that the fire truck is red" becomes more justified as the sensation of redness becomes more intense, and as the truck drives through a variety of different lighting settings. Phenomenal aspects of a given property—say "red"—*do* seem to play an important epistemic justification role, and the example above shows that this increase in sensational intensity requires a protention with conditions of satisfaction, namely, "that the redness increase in intensity." But the protentional contents, and the conditions of "redness" that satisfy it, are non-propositional: this is a case of *describing* intentional content propositionally without that proposition *being* its content. IC 1 is well accounted for without Retrospective Analysis.

Section VI

I will conclude with a final comment on McDowell in light of a theory of time-consciousness and then say something about intuitive content.

Auditory examples and increases in phenomenal intensity highlight the fact that active, 'momentary' judgments require pas-

---

[53] The argument here that the phenomenal character of perception is "accounted for" is not intended as any sort of solution to, say, the "hard problem of consciousness."

[54] For an interesting discussion, see Bennett and Hacker, Philosophical foundations of neuroscience, 124–125.

sive connection over time, that is, *association*. This is necessary for transitive consciousness: "It is precisely the analysis of associative phenomena that draws our attention to the fact that consciousness necessarily must not be a consciousness of a single object for itself," that is, consciousness of an object in a single, temporally atomized manner.[55] The single protention/now-point/retention cluster can be thought of as one slice from a loaf of bread, with the thickness of the loaf corresponding to a stretch of time. Object-constitution is distributed not only laterally over a single slice but also longitudinally over a series of slices. Association bonds these slices together; a persistent interest in an objective referent can be traced over a series of slices. Time-consciousness keeps each slice together.

This said, it is important to realize that a single slice—a single present moment—has both a *determined* and *determinable* side of both retention and protention. The axis, therefore, is not a conceptual/non-conceptual axis but a determined/determinable axis. The rich phenomenology of perception, indicated in IC 3 as both transitive and non-transitive perception, translates to a high degree of *determinable* contents but a comparatively low degree of *determined* or *fulfilled* contents. This imbalance of determined/determinable contents is related to the issue of "conceptual exhaustion" discussed above with the shades of orange example. Regarding this, McDowell makes an interesting remark in his response to J.M. Bernstein:

> But contrary to what Bernstein suggests, this is no problem for my claim that the whole content of experience is actualization of conceptual capacities, because I do not connect actualization of conceptual capacities with bringing things *into focus*…. Why should we stipulate that conceptual capacities are operative only where there is "conceptual grasping" *in that sense*? As far as I can see, it is only that stipulation that makes Bernstein think that my conception cannot accommodate the richness of experience. [emphasis added][56]

McDowell thus claims that his notion of conceptualization is not equivalent to bringing things "into focus," or, in our terminology, it is not equivalent to active determination. McDowell would certainly oppose any attempt to segregate the "experiencing of *x*" and the "conceptualization of *x*": there could be no separation, even "notionally," between receptivity and spontaneity in such experiences. But so would Husserl, assuming McDowell were to make some modifications to his account, like forgoing Retrospective Analysis method of deriving perceptual content, the Propositional Attitude thesis, and the Detability Thesis that constrains intuitive

---

[55] Husserl, *Analyses concerning passive and active synthesis*, § 26.

[56] McDowell, "Response to J.M. Bernstein," in Smith, *Reading McDowell*, 299–300.

content to matter, only one part of a three-part table. Husserl can actually be seen as providing a more explicit account of how objects and concepts are brought into focus genetically, "in real time." Over a stretch of present moments, an observer maintains interest in a single referent and, over an association of these moments, increasingly determines or fulfills certain passive, signitive protentions. The key is that conceptual synthesis—Husserl's *passive synthesis*—is involved from the very beginning of an "experiencing of *x*." From the outset of the experience, there is a horizon of determinability for an object of perception—both transitive and intransitive. What changes *over time* is how *determined* this empty horizon of *determinability* becomes. In a lot of ways, this account of perceptual content resembles that of Gareth Evans, in which he claims that perceptual experience is, fundamentally, an "input to a *thinking, concept-applying, and reasoning system*."[57] This paper provided a framework for placing such a robust account of perceptual input "above the line."

Though in historically dissimilar ways, McDowell and Husserl both acknowledge a great debt to Kant. They are both explicit in the fact that they are guided by the Kantian insight that "the senses cannot think anything." The question is what kind of decorations to append to this insight, and one will decorate based on how seriously one takes the initial distinction between *receptivity/sensibility* and *spontaneity*. Indeed, one might wonder whether such a distinction at the outset violates McDowell's idea that "receptivity does not make an even notionally separable contribution to the cooperation."[58] For these distinctions between receptivity and spontaneity seem to be, at least, *notional* distinctions. Husserl's account has the merit of describing the purpose of this initial distinction: it allows one to accommodate empty content and intuitive content. With Husserl, there is a method—the method of genetic analysis—that "follows the history, the necessary history of … objectivation and thereby the history of the object itself as the object of possible knowledge."[59]

Perhaps the key decoration lacking here is an account of intuition, though this is not the place to fully develop a theory of intuition. Husserl's theory of *hyle* as the substrate for intuition is still in need of an extended treatment. Some commentators are more optimistic on this project than others,[60] but I will shelve this for now. The systematic issue with Husserl's account of intuition is not just a problematic theory of *hyle*. Rather, as noted above, intuition itself

[57] Evans, *The Varieties of Reference*, 158.

[58] I owe this observation to Barber, "Holism and Horizon."

[59] Husserl, *Analyses concerning passive and active synthesis*, sec. 4. "Static and Genetic Phenomenological Method."

[60] For an excellent attempt to reconstruct *hyle*, see Gallagher, "Hyletic experience and the lived body."

appears twice under the divisions that comprise intentional acts. The major oddity is that intuitive sensation is both a unique *quality* and a unique *mode of fulfillment,* namely, intuitive. This asymmetry in terms of its intentional character is, perhaps, to be expected if intuition is to be *the* world-presentifying mode, satisfying IC 1. Like Sellars, I am content to indicate that a final analysis of intuitive passivity is necessary in an account of perception, and for an ultimate resolution of the PoP.

This paper presented the Myth of the Given as a substantive objection to a certain move, that of positing pre-formed sense data, to dissolve a classic problem of perception. An alternative view, conceptualism, was shown to imply the improper method of Retrospective Analysis, the failure of a Detachability Thesis, and a phenomenologically inadequate Propositional Attitude thesis. An alternative, I urged, was to accept the formative aspects of conceptualism while rejecting propositional perceptual content. The new Husserlian framework deals more specifically with conceptualization in time-consciousness and addresses the PoP via an "above the line" theory of fulfillment. I concluded by suggesting that Parametric Content allows for phenomenally driven epistemic states (IC 1), disambiguates perception from other states (IC 2), accounts for content that is both determinable and intransitive (IC 3), and makes it intelligible how subjectivity is characterized by meaning-relevant, interest-driven perception (IC 4).

References

Barber, Michael D. "Holism and Horizon: Husserl and McDowell on Non-conceptual Content." *Husserl Studies* 24, no. 2 (January 23, 2008): 79-97.

Bennett, M. R., and P.M.S. Hacker. *Philosophical foundations of neuroscience*. Malden, MA: Blackwell, 2003.

Crane, Tim. "Is Perception a Propositional Attitude?" *The Philosophical Quarterly* 59, no. 236 (July 1, 2009): 452-469.

DeVries, Willem A. "McDowell, Sellars, and Sense Impressions." *European Journal of Philosophy* 14, no. 2 (2006): 182-201.

Dretske, Fred I. *Knowledge and the Flow of Information*. Cambridge: MIT Press, 1983.

Evans, Gareth. *The Varieties of Reference*. Oxford: Clarendon Press, 1982.

Gallagher, Shaun. "Hyletic experience and the lived body." *Husserl Studies* 3, no. 2 (1986): 131-166.

Hopp, Walter. "Conceptualism and the Myth of the Given." *European Journal of Philosophy* 17, no. 3 (September 1, 2009): 363-385.

Hopp, Walter. "Husserl on Sensation, Perception, and Interpretation." *Canadian Journal of Philosophy* 38, no. 2 (2008): 219-245.

Huemer, Michael. "Sense-Data." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. 2011st ed., 2011. http://plato.stanford.edu/archives/spr2011/entries/sense-data/.

Husserl, Edmund. *Analyses Concerning Passive and Active Synthesis: Lectures on Transcendental Logic*. Dordrecht: Kluwer, 2001.

Husserl, Edmund. *Experience and Judgment*. Translated by James Spencer Churchill and Karl Ameriks. Evanston: Northwestern University Press, 1975.

Husserl, Edmund. *Ideas Pertaining to a Pure Phenomenology and to a Phenomenological Philosophy: Second Book*. Dordrecht: Kluwer, 1990.

Husserl, Edmund. *Logical Investigations, Vol. 1*. Translated by J.N. Findlay. New York: Routledge, 2001.

Kant, Immanuel. *Critique of Pure Reason*. Translated by Marcus Weigelt. Revised. New York: Penguin Classics, 2008.

Locke, John. *An Essay Concerning Human Understanding*. Indianapolis: Hackett Publishing, 1996.

McDowell, John. "Lecture I: Sellars on Perceptual Experience." *The Journal of Philosophy* 95, no. 9 (1998): 431-450.

McDowell, John. "Lecture II: The Logical Form of an Intuition." *The Journal of Philosophy* 95, no. 9 (1998): 451-470.

McDowell, John. *Mind and World*. Cambridge: Harvard University Press, 1996.

Ryle, Gilbert. *The Concept of Mind*. Chicago: University of Chicago Press, 1949.

Searle, John R. *Intentionality*. Cambridge: Cambridge University Press, 1983.

Sellars, Wilfrid. *Empiricism and the Philosophy of Mind*. Edited by Richard Rorty and Robert Brandom. Cambridge: Harvard University Press, 1997.

Smith, Nicholas H., ed. *Reading McDowell: On Mind and World*. New York: Routledge, 2002.

Sturgeon, Scott. "Visual Experience." *Proceedings of the Aristotelian Society* 98. New Series (January 1, 1998): 179-200.

Zahavi, Dan. *Husserl's Phenomenology*. Stanford: Stanford University Press, 2003.

An interview with David Chalmers

September, 2011
Australian National University

Each year The Dualist *includes an interview with a contempoary philosopher chosen by the staff. This year, we are very pleased to have David Chalmers, author of* The Conscious Mind, *answer questions posed by* The Dualist *and by the Stanford Philosophy Department.*

The Dualist:

Through your website (http://consc.net) and now your PhilPapers initiative, you have been a pioneer in exploring the internet's potential to spread ('democratize'?) philosophical discussions. How do you believe this project has progressed? How do you envision the future of philosophical uses of technology?

David Chalmers:

I've been an Internet addict since the early days. As a graduate student in an AI laboratory in the late 1980s and early 1990s, I was involved in a lot of philosophical discussion over Internet newsgroups. I also compiled an annotated bibliography of the philosophy of mind and made it available over the Internet -- initially by email and ftp, later via the web. I'd compiled this bibliography for my own purposes: I'd come into philosophy from the outside with a lot of views about issues in the philosophy of mind, and I decided that to do philosophy properly I'd need to learn what real philosophers had said, and how they go about saying it. So I went to the library, sat down with back issues of journals, and tried to read as much as I could. Only later did I discover that most philosophers read only a small fraction of the journal literature, so I ended up being pretty well read in that area. (Fortunately, I read quickly.)

I circulated the bibliography privately at first, and various people suggested I make it available publically, so I did. Part of the hope was that others like me who were interested in philosophical issues but who didn't know anything much about the philosophical literature could begin to get a sense of what it's all about and eventually make contributions themselves.

Since then the project morphed into the website MindPapers and eventually produced PhilPapers. The key development here was my collaboration with David Bourget, who first contacted me when he was a graduate student in Toronto interested in adding search mechanisms and other sorts of automation to the bibliography. He did that and MindPapers was the result. Before long he came to ANU to do a Ph.D., writing a marvelous thesis on consciousness and intentionality, part of which has been published in major journals. Along the way we developed PhilPapers, meant as an analog of MindPapers for all of philosophy, with all online and offline material to be covered. It has now exploded into a huge project with 300,000+ entries, 25,000 registered users, and 300 active editors. The British government has put major funding into the project (through the Institute for Philosophy in London, where David is now working), and many further developments are in progress.

I hope that this has had the effect of making philosophy more accessible to nonphilosophers -- though I know that it can also have the effect of overwhelming them. We are moving toward providing more guidance on PhilPapers about introductory articles, key papers in the literature, and the like. I don't know to what extent nonphilosophers have been led to contribute to philosophical discussions as a result of this. Doing professional-level philosophy is hard, and like all academic fields doing it credibly requires serious training in the practices of the field. But I hope that it's at least improved things around the margin. I hope that we've also provided a service to professional philosophers in helping to circulate philosophy among professionals and students in quick and easy ways.

It remains to be seen whether all this makes any substantial difference to the progress of philosophy. Perhaps it won't. Much as I'd like to, I can't honestly say that I've seen a major wave of new philosophical insights that can be attributed to these Internet tools. Still, if it at least increases the openness and interactivity of the field, I'll be happy.

The Dualist:

Do you have any thoughts about connections between the philosophy of mind and ethics? For example, you have famously discussed the possibility of philosophical zombies -- creatures who are just like you and me

but who are not conscious. Are there interesting questions about, say, whether it would be wrong to torture a philosophical zombie? Or do you think the relationship between philosophy of mind and ethics is not so straightforward?

David Chalmers:

I've never written anything on ethics beyond a paragraph or two here and there, but I've once or twice given talks on the topics. The most relevant is an after-dinner talk I once gave at a consciousness conference on the role of consciousness in ethics. The centerpiece was a number of audience polls on "zombie trolley problems": who do you kill, the one conscious being or the five zombies? Most people vote to kill the zombies and save the conscious being. Then one can make things more extreme: e.g. who do you kill, the one conscious chicken or a whole planet of humanoid zombies. In that case I remember just one philosopher voting for the zombies, saying "They're just meat". I take the upshot of the poll to be that people have a strong intuition that consciousness matters for moral worth, while at the same time it isn't the only factor.

I suppose that I'm inclined toward a similar view. Sometimes I am inclined toward the stronger view that consciousness is the ground of all value in life, and also the ground of all meaning in life. But perhaps that is just the stance of a philosopher who has a hammer (consciousness) and sees everything as a nail. Certainly I would be very reluctant to torture a zombie, even if I knew it was a zombie. There is a decent case that zombies have intentional states including desires, and in torturing them I'd certainly be acting contrary to their desires, producing bad memories and various other sorts of functional trauma. Still, I'd torture a zombie before torturing a conscious human.

There are also questions about the role of consciousness in moral agency. E.g. is Hitler's zombie twin, a physical and behavioral twin who we can stipulate sends millions of conscious beings to their death, as morally responsible as Hitler? Presumably we could make a good case for punishing zombies at least as a deterrent, but what about as retribution? Again I find myself inclined toward the view that consciousness plays an essential role in being a fully responsible agent, without having a strong argument for that view.

The Dualist:

You are well known for your confidence in the usefulness of a priori methods in philosophy. How much philosophical progress do you think is ultimately attainable through a priori methods? Is the scope of this progress (more or less) tightly circumscribed, or is philosophy more like mathematics, where the progress seems to be continual?

David Chalmers:

I am a pluralist about philosophical methods. I think that empirical methods play a vital role and that a priori methods do too. I think that one gets much further with a combination of the two than with either alone. A priori methods alone can be sterile, while empirical methods alone can be shallow. Empirical results very often provide a marvelous springboard for interesting philosophical conclusions-- but at the same time the springboard almost always needs a big a priori premise to function properly. It's usually an interesting exercise to dig around for this premise in putatively empirical work.

I think one can get a long way with a priori methods, although of course there are limits when it comes to clearly a posteriori matters, such as the contingent fundamental character of our world, or even the functioning of human minds. I think the existence of consciousness itself is something we know empirically, albeit through introspection. But again, a priori methods can bootstrap empirical results to yield more interesting conclusions than either alone.

The question of progress in philosophy is a tricky one. I think there certainly has been a good deal of progress over the years. At the same time, any honest philosopher must admit that progress has been less than in the sciences, or at least that it takes a very different form. The big questions in philosophy don't get decisively settled -- rather, we come to understand them better, we get a sense of the best arguments on each side and the best versions of each view, but nothing ends up being black and white. Of course some areas end up splitting off, such as logic and psychology. These areas and other sciences have certainly resolved some philosophical questions, but they tend to leave the deepest philosophical questions unanswered. The deepest questions seem only to get answered to a few philosophers' satisfaction at any one time, and never attain consensus. Even the answers that we have tend to involve a complicated mix of a priori reasoning and of science. I doubt that this will change any time soon.

The Dualist:

Your work on consciousness and cognition has often emphasized that we have a special sort of awareness of, or acquaintance with, our own conscious experiences. What kinds of epistemic features and standards do you take this special awareness to have? That is, what does it take to be acquainted with a thing, and what kinds of knowledge do we get by being acquainted with our own conscious experiences?

David Chalmers:

I find myself increasingly sympathetic with Russell's views about acquaintance. I think there is a form of acquaintance with an entity that puts us in a position to know what that entity is in a strong sense. The sense is strong enough that Frege puzzles are ruled out: that is, if "a" and "b" are both expressions under which one is acquainted with an entity in the relevant strong sense, then one can know "a=b" a priori. Of course it follows from this that most ordinary proper names are not expressions involving acquaintance. In fact, like Russell I think one can't be acquainted in this way with concrete external objects. But I do think one can stand in this relation to some abstract objects such as numbers and also to some properties, not least the properties of consciousness.

I think the acquaintance we have with consciousness enables us to form special concepts of conscious states, the "what-it-is-like" concepts that pick out experiences as they are in themselves. This acquaintance also is responsible for the special knowledge that one gains upon having color experience for the first time. I wrote about acquaintance with consciousness initially in my paper "The Content and Epistemology of Phenomenal Belief". More recently, acquaintance has been playing an increasingly significant role in my thinking about epistemology and the philosophy of mind, especially in my forthcoming book "Constructing the World".

The Dualist:

In recent work you defend an anti-realist view about metametaphysics — about what we talk about when we talk about metaphysics and ontology. You have compared metametaphysics' relationship with ontology to metaethics' relationship with ethics: each meta-discipline inquires into the status of the claims of the discipline. Do you think that debates about metametaphysics (and your anti-realist position, in particular) have important consequences for how we do metaphysics, and philosophy more generally? Or do you take the attitude of many metaethicists, who hold that their inquiries should leave undisturbed the claims whose status is their subject?

David Chalmers:

These are interesting questions. I suppose I think that accepting a deflationary view in metametaphysics should affect one's first-order metaphysical practice. In particular it suggests that certain sorts of "deep metaphysical questions", for example about whether certain sorts of mereological sums exist, should be seen either as unanswerable (because they have no determinate answer) or as shallow questions to be answered by a sort of folk ontology. So

the sort of first-order metaphysics that allows that there are deep revisionary answers here (saying e.g. that tables don't exist) looks less well-motivated.

I think that there's a disanalogy with the ethical case, in that strongly revisionary views are much less common in ethics then in metaphysics. I don't have settled views in meta-ethics, but I'm inclined toward some sort of moral anti-realism, on which our moral views are grounded in certain basic evaluative attitudes that aren't fundamentally truth-evaluable. That view might be developed as a form of expressivism, subjectivism, or relativism -- I don't have a strong view about which. Given that perspective, I suppose a revisionary view that holds that "folk ethics" is fundamentally wrong -- i.e. that our fundamental moral commitments should be rejected for other moral commitments looks poorly-motivated. But I don't see many normative ethicists putting forward views like that. Rather, many take themselves to be exploring the most coherent way of making sense of folk ethics and of our fundamental moral commitments. I don't have any objection to a parallel program in metaphysics, although perhaps there's a case that folk ontology is much less central to our lives than folk ethics.

The Dualist:

You have recently written about "the singularity" -- the possible future advent of machines that are more intelligent than humans. Do you think the singularity raises any distinctively philosophical questions?

David Chalmers:

I think that the singularity raises many interesting questions that philosophers would do well to think about. For a start, there is a prima facie compelling argument (due to I.J. Good) that once there is greater-than-human intelligence, there will shortly be superintelligence. Machines more intelligent than us will be better than us at designing machines, so they will be able to design machines more intelligent than the most intelligent machines we can design, so they will be able to design machines more intelligent than themselves. Repeat that process and we have a rapid spiral to superintelligence. Perhaps there's something wrong with that argument, but if so it's not obvious what. It's just the sort of argument that deserves philosophical analysis, which I've tried to give in a recent article.

There are also important questions about ethics and values, and about consciousness and identity. On the former, we need to think about what sort of values we program into machines, if they are to produce a world that we value. After all, if the machines of the future control the world, then their values will largely dictate its

shape. We also need to think about whether we ourselves might be those machines. The best way for us to control that future is to be part of it. But the limitations of biology may require that to do that, we will need to upload our cognitive architecture onto new hardware that will be enormously faster and upgradable. That raises the question of whether uploaded versions of ourselves will be conscious, and of whether they will be ourselves. These are versions of old philosophical chestnuts that suddenly take on enormous practical relevance. I wouldn't be surprised if our descendants have fierce philosophical arguments along the lines of Derek Parfit's discussion of personal identity in "Reasons and Persons".

In any case I think there's a lot for philosophers to chew on here. There will be a forthcoming symposium with responses to my paper on the singularity, and a number of philosophers as well as AI theorists, cognitive scientists, and science fiction authors are supposed to be contributing. I'll be very interested to see if a serious philosophical literature might come to be developed in this area.

The Dualist:

By present lights, what do you think is the falsest thing you argued in The Conscious Mind*?*

David Chalmers:

Tricky question. My views have evolved a bit since The Conscious Mind, but often in areas where my views in that book were qualified in any case. For example, I'm less sympathetic to epiphenomenalism now than I was then, though I was officially neutral then. I think I dismissed interactionism overly quickly in the book. I'm also somewhat more sympathetic to representationalism and to close links between consciousness and intentionality. I think the information-based theory in chapter 8 of the book probably doesn't work, but even in the book I said it would probably turn out to be wrong.

Perhaps the biggest flat-out mistake is in the final chapter on quantum mechanics, where I say falsely that in a quantum superposition of two system's states, any computation implemented by either system will also be implemented by the superposition. For a version of that principle to hold, one needs to add the further constraint that it is a "decoherent superposition" in which the two branches don't interact with each other, or at least interact minimally. Something like that is what I meant to say all along, and is enough to make the points I wanted to make about the Everett interpretation of quantum mechanics. But I certainly said it wrongly in the book.

Resources for Philosophy Undergraduates

This section includes listings of journals, contests, and conferences — all of which are available to undergraduates in philosophy. If you have comments, suggestions, or questions, or if you would like to be listed here in the next issue, please contact us and we will gladly accommodate your request.

JOURNALS:

There are numerous journals, published both in print and online. The information is as recent as possible, but contact the specific journal to ensure accurate information.

> **Aporia**: Brigham Young University. Submissions due early fall. Papers not to exceed 5,000 words. Send submissions to: Aporia, Department of Philosophy, JKHB 3196, Brigham Young University, Provo, UT 84602. Visit: http://aporia.byu.edu/

> **The Bertrand Russell Society Quarterly:** Edinboro University of Pennsylvania. Visit: http://www.lehman.edu/deanhum/ philosophy/BRSQ/

> **The Dialectic**: University of New Hampshire. Submissions due in April. Essays (15-20 pages), short critical articles, book reviews, artwork. Send submissions to: The Dialectic, c/o Department of Philosophy, University of New Hampshire, Hamilton Smith 23, Durham, NH 03824. Visit: http://www.unh.edu/ philosophy/Programs/dialectic.htm

> **Dialogue**: Phi Sigma Tau (international society for philosophy). Published twice yearly. Accepts undergraduate and graduate submissions. Contact a local chapter of Phi Sigma Tau for details or write to Thomas L. Predergast, Editor, Dialogue, Department of Philosophy, Marquette University, Milwaukee WI 53233-2289. Visit: http://www.achsnatl.org/society.asp?society=pst

> **The Dualist**: Stanford University. Submissions due early 2013. 10-30 page submissions. For more information, see http:// www.stanford.edu/group/dualist/ or contact the.dualist@ gmail.com. Check website for information on submitting a paper and updates on the submission deadline.

> **Ephemeris**: Union College. For more information, write: The Editors, Ephemeris, Department of Philosophy, Union College, Schenectady, NY 12308. Visit: http://www.vu.union.edu/ ~ephemeris/.

> **Episteme**: Denison University. Due November 14. Maximum

4,000 words. Contact: The Editor, *Episteme*, Department of Philosophy, Denison University, Granville, Ohio 43023. Visit: http://www.denison.edu/philosophy/episteme.html

Interlocutor: University of the South, Sewanee. Direct questions to Professor James Peterman at jpeterma@sewanee. edu. Send submissions to Professor James Peterman, Philosophy Department, 735 University Avenue, Sewanee, TN 37383-1000. Visit: http://www.sewanee.edu/Philosophy/Journal/2006/current.html

Janua Sophia: Edinboro University of Pennsylvania. Submissions and inquiries sent to Janua Sophia, c/o Dr. Corbin Fowler, Philosophy Department, Edinboro University of Pennsylvania, Edinboro, PA 16444. Visit: http://www.edinboro.edu/cwis/philos/januasophia.html

Princeton Journal of Bioethics: Princeton University. Visit http://www.princeton.edu/~bioethic/journal/.

Prolegomena: University of British Columbia. Visit http://www.philosophy.ubc.ca/prolegom/ or write prolegom@hotmail.com or Prolegomena, Department of Philosophy, 1866 Main Mall, Buchanan E370, University of British Columbia, Vancouver, B.C., Canada. V6T 1Z1.

Prometheus: Johns Hopkins University. *Prometheus* strives to promote both undergraduate education and research, and looks for submissions that originate from any scholarly field, as long as those submissions clearly demonstrate their applicability to philosophy. Visit http://www.jhu.edu/prometheus/. Write prometheusjhu@hotmail.com or Prometheus, c/o Philosophy Dept., 347 Gilman Hall, Johns Hopkins University, Baltimore, MD 21218.

Stoa: Santa Barbara City College. For more information, write The Center for Philosophical Education, Santa Barbara City College, Department of Philosophy, 721 Cliff Drive, Santa Barbara, CA 93109-2394. Visit: http://www.sbcc.edu/philosophy/website/CPE.html.

The Yale Philosophy Revew: Submissions due February 14. Visit: http://www.yale.edu/ypr/submission_guidelines.htm

CONFERENCES:

There are many undergraduate conferences, so contacting the philosophy departments of a few major schools in a particular area or researching on the web can be quite effective. The conferences below are by no means an exhaustive list.

American Philosophical Association: The APA website, http://www.apa.udel.edu/apa/opportunities/conferences/, contains

an extensive list of conferences.

Butler Undergraduate Research Conference: Butler University. Conference in mid-April. See http://www.butler.edu/urc/index.html for details.

National Undergraduate Bioethics Conference: Notre Dame. Visit http://ethicscenter.nd.edu/events/nubec.shtml or write bioethic@nd.edu.

Pacific University Undergraduate Philosophy Conference: Pacific University. Conference in early April. Visit http://www.pacificu.edu/as/philosophy/conference/index.cfm for details.

Rocky Mountain Philosophy Conference: University of Colorado at Boulder. Visit: http://www.colorado.edu/philosophy/rmpc/rmpc.html

ESSAY CONTESTS:

The essay contest listed below aims at a broad range of undergraduates, but there are many other contests open to students enrolled at specific universities or interested in particular organizations.

Elie Wiesel Essay Contest. Open to undergraduate juniors/seniors with faculty sponsor. Questions focus on current ethical issues. Due in late January. Top prize $5,000. For more information, visit: http://www.eliewieselfoundation.org/EthicsPrize/index.html

THE DUALIST is a publication dedicated to recognizing valuable undergraduate contributions in philosophy and to providing a medium for undergraduate discourse on topics of philosophical interest. It was created by students at Stanford University in 1992 and has since featured submissions from undergraduates across North America. If you would like to receive an issue of *THE DUALIST* or to submit a paper, please contact us at the address below. We prefer that submissions be formated acccording to the *Chicago Manual of Style* guidelines. See http://www.chicagomanualofstyle.org/ Papers should be submitted in electronic form only.

Visit our website for submission information:
        http://www.stanford.edu/group/dualist

Please email us with any inquiries:
        the.dualist@gmail.com

Or write to:
        The Dualist
        Philosophy Department
        Stanford University
        Stanford, CA 94305