Midwest Studies in Philosophy, XXIX (2005)

Freedom with a Human Face

TIMOTHY O'CONNOR

As good a definition as any of a *philosophical* conundrum is a problem all of whose possible solutions are unsatisfactory. The problem of understanding the springs of action for morally responsible agents is commonly recognized to be such a problem. The origin, nature, and explanation of freely willed actions puzzle us today as they did the ancients Greeks, and for much the same reasons. However, one can carry this "perennial-puzzle" sentiment too far. The unsatisfactory nature of philosophical theories is a more or less matter, and some of them have admitted of improvement over time. This, at any rate, is what we self-selecting metaphysicians tend to suppose, and I will pursue that high calling by suggesting a few improvements to a theory of metaphysical freedom, or freedom of the will.

Philosophers who rightly defend a robust account of freedom have tended to err in the direction of over-idealization. Concerned to protect the notion that we are *self-determining* beings from various naysayers (theological as well as philosophical), they have often suggested a picture on which our freedom of will is well-nigh absolute. We are led to entertain a capacity of will or choice that is not conditioned or structured as other terrestrial capacities are, whose only limits are the bounds of its imagination and perhaps desires. Here, for example, is the late Roderick Chisholm:

> If we are responsible, and if what I have been trying to say is true, then we have a prerogative which some would attribute only to God: each of us, when we act, is a prime mover unmoved. In doing what we do, we cause certain events to happen, and nothing—or no one—causes us to cause those events to happen. (Chisholm 2003, 34)

Chisholm himself did not defend an account of freedom as extreme as this quote might suggest, but his words are emblematic nonetheless of a tendency in

philosophical theorizing. They also nicely set up a theme I want here to develop: perfect freedom is indeed a prerogative of God alone. Human freedom, by contrast, is always limited, fragile, and variable over time and across agents. It is the sort of thing which comes in degrees, and our theorizing needs to be built around that understanding.

Failure to distinguish clearly the perfect freedom of God from the imperfect forms of freedom manifested in human actions distorts much of our theorizing. I've already noted the one extreme of exaggerating the character of human freedom, and this has concomitant distortions in theological/moral understanding. Another extreme, more popular in secular thought, is to do away with basic elements of the divine ideal and offer in their place a surrogate notion of freedom. This is done from a variety of postures. The compatibilist posture is to declare the divine ideal a fantasy that is altogether alien to our ordinary self-understanding and that has made for much mischief in philosophical theorizing about human freedom. On their view, freely willed actions are the natural products of a choice mechanism that translates beliefs about our alternatives and the stronger of our desires into plans that cause behavior aimed to realize them. Choice-making is as much a *mechanism* as any other natural capacity, albeit a rational one, responsive to reasons. Another deflationary posture holds that our ordinary concept of free action *is* roughly along the lines of the divine ideal, but joins the compatibilist in declaring the ideal an impossible fantasy, either relative to limited beings like us, or absolutely. Thus, the whole notion of human free will is a fiction, as is the notion of moral responsibility on which it rests. Finally, a more conservative cousin of the glum impossibility theorist is the revisionist, who thinks we can salvage a few strands from the impossibly tangled ideal, at the cost of having to refashion our self-understanding. In short, we need to absorb the fact that our responsibility-conferring freedom isn't nearly what it's cracked up to be, even by ordinary folk—but it ain't all bad.

I believe that all these views are variations on a *deflationary* theme—in large measure reactions to the idealized conception. Agreeing with those who apply this conception more or less intact to ourselves, the deflaters presume that the ideal does not admit a more modest, more recognizably human, version. I hope to persuade you to question this presumption. I will try to suggest how human and divine freedom can share a common metaphysical core consistent with differences that reflect the gulf between unconditioned Creator and His frail children of dust.

I. TWO COMPONENTS OF FREEDOM

The concept of metaphysical freedom in ordinary thought encompasses two quite distinct concepts, though much recent philosophical writing ignores one of these, equating freedom with the other. While analytical focus on one or another constituent concept is proper and necessary, truncating the concept of freedom in this way hinders understanding, since giving freedom a human face depends on seeing the way the two ingredients interact. The much-studied ingredient is the "openness of the future" to alternative possibilities for our actions. A person acts with freedom only if he could have chosen a different action. John Fischer likes to

employ Jorge Luis Borges's image of a labyrinth of forking paths to represent this aspect of freedom. When we act freely, our action constitutes but one of a plurality of "forking paths" that were available to us at an earlier time. The selection of any of these paths is consistent with things being as they were at that earlier time, including all the factors constraining and influencing our choosing the path we actually did. None was pre-determined, though some doubtless were more likely than others. (We may note two kinds of cases here. In one kind, the choice is itself a settling upon one of the available forks. The alternative paths remain available at all times prior to the choice. In the other kind, the alternative paths are progressively passed by prior to the action in question, such that for a period of time leading up to the action, the choosing of that action has already been guaranteed, whether the agent recognizes this or not. Choices that are undetermined until their occurrence are directly free, whereas the others are merely derivatively free, in virtue of the possibility of having avoided them by prior choices that were directly free [Clarke, 2003, 4, n.2].)

Harry Frankfurt (1969), of course, has persuaded some philosophers that alternative possibilities (at some time prior to the action) are not necessary for moral *responsibility*. He did so by describing extensions of scenarios ("R-scenarios") where a responsible agent willingly performs an action. In an extended scenario, we add the presence of an agent or device that monitors the agent's deliberation and is poised to intervene, were the agent to consider performing an unwanted action. As it happens, the agent voluntarily does what the would-be controller wants, and so the latter does nothing at all to influence or interfere with the process of deliberation and action. Surely such extensions of R-scenarios are themselves R-scenarios: they involve no change to the process of deliberation or the way in which the action is produced. And one might apply a version of Frankfurt's argument here to the concept of freedom by saying that freedom *in the sense relevant to responsibility* likewise does not require the availability of alternatives.

Frankfurt scenarios show us something, but not what Frankfurt thought. What they show us is that it is merely a contingent feature of the (ordinary) circumstances of our world that agents with metaphysical freedom are able to select from among significantly different alternatives. (We inhabit a garden of forking paths and have the means to choose from among the paths.) But freedom itself is not to be identified with the existence of those robust alternatives, which, when present, are but an outgrowth of a certain quality of the act itself, which we might term "self-determination." In Frankfurt cases, the alternatives are nipped at the buds—were the agent to contemplate choosing any but the preferred course, his control would be abrogated and he would be made to conform to a pre-selected path. To return to Borges's image, Frankfurt scenarios are analogous to a garden in which the many diverging forks along a given path are replaced by tiny side paths that almost immediately rejoin the main path, although the fact that they are so constrained is hidden to one unless one attempts to follow them. For practical purposes, it is but one path with slightly swollen "buds" along the way. Now, the significance of such alternatives (whether they are robust or mere "buds") lies in their being indicators of the self-determination manifested by one's action,

which is necessary for responsibility.[1] Self-determination is conceived even at a commonsensical level as a multi-valent potency somehow controlled by the agent. (How exactly this works is the job of philosophical theory, and I shall offer such a theory shortly.) As such, it points outside itself to a space of possibility. How broadly those possibilities differ depends on one's environment, rather than being under the direct control of the agent. In responding to Frankfurt, we insist on the necessity of at least the 'buds' of possibility because of their connection to the potency in the action itself that is essential to responsibility.

John Martin Fischer has criticized those who see significance in the fact that there will always be a bud of alternative possibility in any Frankfurt scenario—what Fischer derides as mere "flickers of freedom." He points out that the agent always fails to act *freely* in the alternatives represented by the buds (Fischer 1994, 140–47). (In a paper delivered at the 2001 Pacific Division meetings of the APA, he noted that a special case of a Frankfurt scenario is a "Fischer scenario," in which the agent immediately *dies* in all the alternatives!) But the fact that there is no freedom in the alternatives is irrelevant. Neither the scope nor the character of the alternatives bears upon the freedom of (or responsibility for) our action.[2] The question is whether I am free in, and responsible for, the one action that I in fact do, not whether I would be free and responsible in a counterfactual scenario. The buds are indicative of the nature of my agency in the actual sequence—of the fact that I am a self-determining agent, which is one side to the freedom condition on responsible action.

Now we must turn to the other, often ignored side of freedom of the will, which is self-mastery. A person acts freely to the extent that he has control of his appetites and impulses and is able reliably to direct his more significant actions towards larger aims. A self-mastered person perforce has a great deal of self-knowledge, including especially knowledge of the factors that incline him to this or that course of action. The importance of self-mastery to freedom is a familiar theme in the history of philosophical reflection on the will, while being most emphatically developed in the ancient Stoics and Spinoza. It is also the target of work in cognitive psychology, being more amenable to direct empirical study than is the self-determination condition. Prima facie, it is possible for an agent to lack self-mastery (to a significant extent) while regularly exercising a causally undetermined choice capacity. It seems, for example, that unconscious or dimly grasped desires could regularly influence such choices, or that a self-determining agent may be subject frequently to powerful impulses that, while not determining the outcome, may heavily distort the course of deliberation. In either case, the agent's freedom seems diminished thereby. (According to much Christian theology, these possibilities are more than theoretical, but in fact represent a prominent feature

1. Again, it is necessary that this quality be manifested either in the evaluated action itself, or in prior actions that shape either the disposition that produces, or the circumstance that constrains, the evaluated action.

2. Here I speak of the metaphysical condition on responsibility. My *beliefs* about the scope of my alternatives are obviously relevant as well, but this is a separate matter. In a Frankfurt scenario, the agent is unaware of the unusual circumstance surrounding his deliberation.

of the fallen human condition. The Apostle Paul speaks of our being "slaves to sin," unable to act righteously apart from divine grace, and he is of course echoed in Augustine, Aquinas, Luther, and Calvin. Some such authors may well endorse an extreme position entailing that each of our choices is fully determined, so that freedom of the will is wholly denied. Others are committed only to the view that all of our choices are "tainted" by sin insofar as they are continually informed in part by base motivations, and in that sense none of our actions is truly righteous. And Augustine famously develops the theme that our knowledge of our own motivations, particularly when making an immoral choice, is far from perfect.)

Some philosophers have become so accustomed to identifying freedom with the alternative-possibilities condition as to find the thought that self-mastery also plays a role mystifying. ("Someone who knows more about his or her motives and desires will be able to make more informed choices, just as would someone who has more knowledge about the likely outcomes of choices. Such a person would, therefore, choose more wisely. But why should we think that such a person is more *free*?")[3] And we can, if we like, simply *stipulate* that what we mean by "free" is captured by the alternative-possibilities condition. But it cannot be denied that self-mastery as I have characterized it is an important strand in ordinary thinking *surrounding* freedom. Furthermore, it is not merely an epistemic condition on (degree of) responsibility, for it has to do with how we are influenced to act and it generally is beyond our immediate control. I would, therefore, ask such philosophers whether they cannot also recognize a "thicker" concept that reflects the possession of other features whose value is related to the value of freedom narrowly conceived. (I find that the word "autonomy" tends to have the right resonance for such philosophers.) To those who can manage this, I say that *this* is the very concept I propose to analyze.

Finally, I note that exclusive preoccupation with the alternative-possibilities condition on freedom also leads one to erroneously suppose that freedom co-varies with the extent of possibilities open to one, leading to the conclusion that perfect freedom requires openness to all possibilities. Such a "free" agent would lack a nature or character—or at least a nature that is resolute in rejecting options contrary to its fundamental aims, moral or otherwise. To be open to a possibility requires, at a minimum, some degree of inclination towards that outcome. Imagine a human person who has some tendency to choose virtually any option he can conceive. His actions in the aggregate would seem whimsical, as they would not proceed from a fixed, discriminating base of general purposes and motivations. There could be, it seems, no long-term narrative that made sense (even to himself) of the basic contours of his life and actions. As an account of ideal freedom, this is absurd. A free agent knows himself well—knows his own stable purposes, desires, and beliefs—and reliably acts in a way that reflects in some way this self-understanding. To be sure, one can freely participate in the development of one's character, but such change necessarily is incremental.

3. I take this quotation from a commentary by Ed Wierenga on a version of this paper that I read at Wheaton College.

II. PERFECT FREEDOM

What does freedom look like in the limit? Hard to fathom, to be sure, but consider the component of "self mastery." In God's case, this is a misnomer, insofar as it implies hard-won attainment, but we'll stick with this term. Perfect self-mastery requires complete self-knowledge, complete knowledge of each and every wish or desire, belief or intention, that impinges on one's choices. Not only so, but the self-mastered person is steadfast in his purposes—there are no irrational impulses, nor conflict among his aims. Since God is perfectly powerful and knowledgeable as well as perfectly free, there also are no practical constraints to which He is subject. (For desirable objects that are not compossible, there is the option of creating multiple universes.)

How wide is the extent of alternative possibilities for God, who is perfectly free? Most theologians acknowledge that God's perfect goodness entails that any number of scenarios contrary to His moral nature are not genuine possibilities for Him. I would go much further and contend that there is reason to believe that if God creates anything at all, He *must* create a realm of maximal value, one which is far beyond anything that empirical observation or Christian theology gives us reason to accept. Surely His creative acts would be guided by an understanding of the intrinsic value of the possibilities.[4] Plausibly, there is no finite upper bound on possible universes ordered by their intrinsic goodness. One reason to think this is that there is no finite limit to the number of good things one can have or to the space needed to comfortably house them. (We need not assume that metaphysical goodness is "well-ordered." As one goes up the scale, possible universes might well fragment into incommensurate orderings, given the disparate characters of some distinct groupings. Still, it is presumable that each branch that is itself well ordered has no upper limit.) So which creation is God to choose? It would be exceedingly strange if an absolutely perfect being were to select just this universe, whose value is, say, 10^{100} s.g.u. (standard goodness units), while He is fully aware that there are infinitely many better than it. It is in any case implausible that a perfect being should have idiosyncratic preferences for certain kinds of universes, quite apart from their value. Of course, if there is no upper limit, the oddity would not go away were God to have chosen any other universe, no matter how much greater than this one.[5] Nor will it help for God to create two universes, or, in fact, n universes, for any finite value n. But it would help if God were to create infinitely many universes, provided there is no finite upper limit on the value of its

4. I here assume that there are degrees of objective, metaphysical goodness, a kind that encompasses, while not being limited to, moral goodness. This is in any case implied by the very concept of a perfect being.

5. William Rowe (1993) departs from the argument in the text at just this point, seeing the makings of an argument that there cannot *be* an absolutely perfect creator. Daniel and Frances Howard-Snyder (1994) reply that, since it is inevitable that for any creative act, there will be infinitely many better choices, an arbitrary choice of a sufficiently good universe would not be deficient. You can't ask the impossible, even of a perfect being. Rowe (2004) develops his argument further. If my suggestion in the text has merit, the responses of both Rowe and his critics are hasty, proceeding from a questionable assumption concerning the available choices.

members. For example, God could simply bring about the entire, partially ordered hierarchy of possible universes. On reflection, this simple option appears unsatisfactory, since presumably there is some goodness threshold T below which God would not create.[6] More likely, then, is that God would elect to create every single universe at or above T. But notice that He could also avoid the unwanted consequence by creating every other universe, or every third one, or every *n*-th one, for all numerical values of *n*. So, even supposing God would be bound to create a realm of maximal value, He would appear to retain an infinity of adequate choices.

This simple line of reflection badly needs refinement, but it would take us too far afield to do so here.[7] The point I wish to make instead is that, even if this conclusion is correct, it does not diminish the extent of God's freedom. The conditions for freedom in the divine and human cases differ in a way that reflects the difference in ontological status between an absolutely independent Creator and a dependent, causally conditioned creature. God's choices reflect His character—and His character alone. He was not *given* a nature, nor does He act in an environment that influences the development of individualizing traits. If His character precludes His entertaining various options that are within the scope of his power, this fact cannot be attributed in the final analysis to something else (some combination of nature and nurture). Rather, their impossibility is solely and finally attributable to Him.[8] Anselm had it right: although God is certain to act with perfect goodness (and, we have added, maximal creative ambition), it is not true that God is good out of an "inevitable necessity." Instead, His perfect goodness and freedom are both attributes He has "in Himself," for all eternity.[9]

For any creature with an origin, however, there can be no freedom in actions that wholly reflect fixed features of the nature it was endowed with, or a combination of nature and nurture. For these factors are ultimately attributable elsewhere. What external forces or agents have wrought within us is not up to us. Thus, it is crucial to human freedom that we have a wide scope of options over time. They must have some attraction to us, as this is a necessary condition (given facts of human psychology) on their being actions for which there is some chance that we might choose them. When we conjoin that thought with the earlier reflection

6. Perhaps it is a vague threshold, but if the subsequent reasoning in the text is correct, this will be of no significance.

7. I have elaborated this argument in lectures given at Oxford and Beijing Universities, drawing from a chapter of my *The Necessary Shape of Contingency*, an unpublished book manuscript.

8. Of course, this allows the possibility that God's reasons might lead him to conclusively favor the creation of one specific array of universes. (We have been bold enough to suppose that God's creative choice will be shaped by certain considerations. We are not so bold as to claim to know *all* such considerations.) So, for all that we have argued, it might be that God's creative choice was both certain to occur and yet free. And this might be at odds with certain official theological positions on God's freedom in creation. As Michael Rota pointed out, for example, the Roman Catholic First Vatican General Council (1870) condemns the position that "says that God created not by an act of will free from all necessity, but with the same necessity by which He necessarily loves Himself." However, it seems that one might draw the required distinction by maintaining that God's love of Himself is a different sort of activity than would be His "necessary" choice in creation, if such it is. Plausibly, this self-love is not an intending (as is a creative choice), but is instead a kind of desire (or joint desire and affirmation).

9. *Why God Became Man*, in Davies and Evans (1998, 327).

that openness to *all* alternatives would (were it possible) clearly undermine one's freedom, we may conclude that the imperfection of freedom in human beings is inevitable.

The starting point of my reflection on human freedom was that much of our theorizing rests on a distorting idealization. That idealization is not precisely a picture on which we are "open" to choosing practically anything, despite our own native proclivities, nor one on which we enjoy the perfect self-mastery exemplified by God. Rather, it is something of a hybrid of these. We tend to represent humans as having a very wide scope of available alternatives, limited mainly by the scope of our imaginations and our desires (as we are not free to choose what we can neither contemplate nor desires to *some* degree). Further, while no one supposes that we are fully aware of the details behind every impulse moving us in this way or that, we nonetheless theorize as if when agents choose *freely*, they are always at least dimly aware of all the rational factors having a direct, appreciable influence on their deliberation—if not always of the content of such factors, then of their presence and relative degree of influence.

That the reality does not fit this idealization is evident. The scope of our freedom is always limited and varies widely from context to context, and our self-knowledge in action is often far from perfect. Our challenge is to construct a more realistic picture without deflating the very idea of freedom. On the view I shall now sketch, freedom of action comes in degrees, or at least levels, of strength, corresponding to differences in the way our reasons enter into the explanation of actions.

III. HUMAN FREEDOM

Human freedom is puzzling in part because we are embedded within a world governed by impersonal, mechanistic forces. Indeed, such impersonal forces seem to regulate not only the constant life-sustaining processes within our bodies, but also much of our behavior. A purely phenomenological perspective suggests that we exercise intermittent direct control over some actions—and often only in the broad terms of an action-plan, whereas the detailed execution is directed from we know not where. The rest of what we do seems "automated," selected and executed below the level of conscious control. (As noted above, it may be correct to say that an "automatic" action is free, but only indirectly so, resulting from tendencies fixed by directly free choices on earlier occasions.)

Taking the experience of freedom at face value, what is distinctive about episodes where we directly control what we do? Certainly, the fact that our actions in some manner are guided by reasons sets apart our activity (and that of some other animals) from the rest of the world around us. But that cannot suffice for freedom, since reasons guide even our "automated" actions. If *all* our actions were unconsciously directed without deliberation, something of great significance would be lost—something integral to our freedom. Let us suppose instead a world where agents are sometimes consciously aware of reasons for acting this way and that, deliberate on those reasons, and certain of these reasons cause them to decide on one course of action. Suppose further that the way these reasons cause their

choices is indeterministic: they had only a "statistical" or probabilistic tendency to cause the actual outcome, so that initially things might have been *exactly* as they were while yielding a different outcome. Other reasons the agent had, perhaps, might have caused a different choice and consequent action. Would all that suffice for freedom? Is it a picture consistent with how our own freedom seems to us?

It seems not. Consistent with the above conception of agency, consider two parallel universes exactly alike up to the time of a particular choice by an agent named Al. In particular, Al and his counterpart Albert have exactly the same history and character (right down to the degree of strength of their most insignificant preferences). As their respective worlds unfold, one set of events causes another in accordance with precise probabilistic propensities. So they might have diverged at any of the innumerable "forks," or moments where more than one outcome is possible, but as it happens precisely the same paths are followed—that is, until the differing fateful choices of Al and Albert. In each universe, there are both basic physical and distinctively mental forces among the factors driving their respective outcomes. The indeterministic mental states are negligible to the grand scheme of things, of course, but crucial to the behavior of the agents inhabiting a particular planet, including Al and Albert. Now in each world, they come to occupy a podium before a large audience one October evening, alongside a fellow named George. At one point, their common irritation with what George is saying gives them an inclination to sigh disdainfully. Each also has some inclination to refrain from such a response, recognizing that it may have an unwanted effect on the mood of his audience. As it turns out, their choices diverge. In Al's case, his irritation causes him to vent his frustration through a series of sighs and groans, to his later regret. Albert's prudent assessment causes him, however, to make instead a gracious, if insincere, reply.

Did Al and Albert freely make their respective choices? Did they exercise the sort of control over their choices that we take ourselves sometimes to exercise? It's hard to accept that they did. In each case, either outcome (and perhaps others besides) was possible—they had the same likelihood for each, in fact. And, crucially, there was nothing distinctive about Al or Albert at the time of their choices—no feature or capacity over which they had control, at the time—that made the difference. If the different choices had a happy result for Albert and an unhappy one for Al, this seems entirely a matter of luck in the sense that neither outcome is properly attributed to the respective agent. Their cognitive and motivational states had identical propensities within the context, and those pointing to one course of action happened to win out in Al's case, while others won out in Albert's. Neither exercised the right sort of control over which inclination would be followed.[10]

III.1 Minimal Freedom

Thus, freedom requires more than indeterministic propensities, more than chance-like outcomes of general tendencies present within the person. Free agents actively

10. A similar conclusion is drawn by Mele (1999, 97–102) and Clarke (2003, 102–8).

determine which tendency will come to fruition on a particular occasion. We have seen that the only plausible strategy for reducing this capacity to something more basic fails. Thus, some philosophers, myself included, conclude that it is ontologically basic. Free agents deploy a basic, goal-directed ability to form intentions to act. They themselves cause their own action-triggering choices, or intentions to act. (When I speak of "intentions," I mean specifically intentions to act here and now which, in normal circumstances, immediately issue in the intended behavior. Intentions of this sort are sometimes called "executive intentions.")

Assuming this can be so, how might agent, or personal, causation fit within an unfolding mechanistic universe? Specifically, how do we understand the influence of our reasons (our beliefs, desires, and prior intentions) on the exercise of a basic agent-causal capacity to generate intentions to act? It is not merely the case that we generally do act on the basis of reasons. Freely causing an intention to act seems inherently to depend on the impetus of reasons, so that it seems correct to say that this causal capacity has a triadic structure of the form: person *S* causes intention *i* for reason *r*. Yet as we attempt to understand the "for reason *r*" part of this, two plausible claims seem to be in tension. On the one hand, reasons have some manner of causal *influence*, and to varying degrees. On the other hand, reasons cannot causally *produce* our causing our own intentions, on pain of our lacking ultimate control over what we do. For, even if the propensity of reasons is indeterministic, if they were to cause us to cause our own intentions, we would lack direct control over this prior causation—precisely the problem we saw earlier for the view that eschews agent causation and takes freedom simply to consist in the non-deterministic causal production of intentions directly by (states of having) reasons.

The resolution of the tension, I suggest, is to understand reasons as (partially) fixing the relative probabilities of particular agent-causings, while not producing them. We might say that reasons influence freely made choices by *structuring the agent-causal capacity*, giving us varying propensities towards different outcomes. On this view, our choices are embedded in the unfolding processes of nature insofar as those processes, operating through our cognitive and motivational states, make us objectively inclined to varying degrees to just a limited few alternatives. Yet the choices themselves are not a direct product of such shaping forces, but of ourselves.[11] The present view provides, I suggest, a minimalist view of (freely) acting *on* a reason: *S* acts on reason *r* in causing intention *i* just in case *r* is a non-producing probability raiser of *S's causing i*.

I call the above view "minimalist" because it entirely ignores the self-mastery side of freedom. Consider in particular the sort of self-knowledge we discussed earlier, involving knowledge of the factors motivating one to act. Consistent with the minimalist picture just described, imagine an agent unconsciously influenced by unknown reasons to cause this or that. In terms of our picture, he lacks conscious access to the content of the reasons that give him a tendency to *A*, for some action type *A*, and is aware only of having an inclination to act in a certain way. When he makes a choice under such conditions, his freedom seems quite dimin-

11. I developed an account along these lines in *Persons and Causes*, chap. 5.

ished. For the missing self-knowledge is a precondition on his being able to assess his own motivations and decide what he will do in a more reflective way. It even seems that we can imagine a worst-case scenario where his having a desire to *B* unconsciously gives him a tendency (in abnormal internal circumstance) to *A*. That is, the agent is "miswired" such that the reason's "natural" propensity—that realized in a normal cognitive condition—is "re-routed," pressed into the service of promoting a desire with an entirely different content.[12] A bizarre scenario, to be sure, but it underscores the limited variety of control (and hence of freedom) exercised by an agent who lacks good self-understanding of the psychological factors motivating him to act.

III.2 Robust Freedom

Now, self-knowledge may vary along a continuum, with limited *degrees* of awareness of *some* of the reasons being reflected in the middle range (which is perhaps most typical of human agency). I contend that freedom and moral responsibility co-vary with this feature, insofar as it is coupled with a capacity to reflect upon and potentially modify some of the aims and motivations relevant to one's deliberation. However, a still higher grade of freedom in action involves more than having some measure of self-knowledge of our own proclivities. It also reflects a difference in the way that reasons guide the action performed. A robustly free agent does not merely act *on* a reason as just described. He *also* acts *for* a reason by consciously intending to *A for the sake of G*, where *G* is the goal of a desire, *D*, that is consciously in view and that (together with the belief that *A*-ing is likely to contribute to the attaining of *G*) may or may not be the predominant antecedent probability raiser. The content of the intention one generates reflects not only the action one aims for, but also the purpose for which one pursues it. When I cause such an intention, I establish a special explanatory link between a particular motivating factor (or factors) and the action undertaken. There may be other desires that have contributed to whatever antecedent propensity I have to *A*. But these cannot explain the fact that I formed an intention with the specific content of *A-ing-for-the-sake-of-G*, which content uniquely matches the reason *for* which I act.[13]

I believe that in our most deliberate and self-aware actions, we do act on purpose-disclosing intentions of this sort. And the conscious purposiveness of these actions, it seems to me, reflects a greater realization of the ideal of freedom. In an earlier generation (circa the late 1950s), philosophers debated whether

12. Randy Clarke presses this sort of case (Clarke 2003, 140–41).

13. Here, I depart in subtle respects from a picture I defended in *Persons and Causes*. There I held that the content of such intentions refers not to the goal of one's action but to the antecedent desire itself (as that which one aims to satisfy). Reflection has persuaded me that this is mistaken. The present view is similar to that of Hugh McCann (1998), chap. 8—except that McCann denies that agents cause such intentions. Instead, they are undetermined events that possess an *intrinsic* property of "activeness" or "spontaneity." For my part, I cannot help but think of this as an inexplicit suggestion of agent causation. But see also Thomas Pink (2004 and forthcoming) for a defense of the claim that metaphysical freedom is conceptually and ontologically basic, and quite distinct from the idea of causation.

reasons explanations of actions were necessarily non-causal in form. Donald Davidson (1980) convincingly argued that there is no conceptual bar to reasons also being causes, and he noted that thinking of them in this way often has great explanatory power. Suppose Ed wants to inherit his rich uncle's fortune and also wants his uncle's longstanding suffering to cease. For which reason did Ed pull the plug on the respirator? Well, imagine that Ed is prone to deceive himself about his baser motives. In that case, were we able directly to examine the causal propensity of each of the two *available* reasons at the time, we might correctly conclude that he did it solely for the money, despite what he sincerely says when caught in the act. But cases of this sort, involving varying degrees of self-deception, should not blind us to other cases where we have our aims fully in view. In the latter, our having particular reasons non-causally explain what we do, owing to the direct reference to the content of our reasons in the content of the intentions that is at the core of our action. This is a distinctive form of explanation, based solely in the reason's content, as opposed to its causal propensity, and it corresponds to a heightened measure of freedom.

Some are inclined to understand this notion of acting for a reason as implausibly maintaining that an agent can declare by fiat what is the motivation for his action.[14] But the objection rests on a misunderstanding. The agent actively *forms* the intention, and this is not a mere declaration *accompanying* the action, possibly mistaken, but the initiating core of what he freely does, and its content guides the completion of the action. (We are not asserting a completely general thesis that the goal-directed contents of any agent's executive intentions necessarily explain the action that ensues. Were the action-triggering intention formed via an external manipulation of some kind, the fact that the agent had a desire whose content matched that of the intention would not be explanatory. It is essential to this form of explanation that the agent himself directly produced the intention. I have argued elsewhere [O'Connor 2000, 52–53] that an agent's causing an intention cannot itself be caused to occur by any independent causal factor.)

But now imagine a case in which a desire D1 gives me a very strong antecedent tendency to cause an intention to A—95%, say—and yet I intend to A for the purpose of achieving G, where this content matches not D1, but some relatively weak prior influence, D2. On my view, I am acting *on* both D1 and D2, but I am acting *for* D2 alone, and this last fact accords a heightened measure of freedom. So suppose I accept all this and, rather portentously, develop the habit of forming intentions to act that reflect in this special way my most *noble* aims, regardless of their actual pull on me when I deliberate. Isn't this a kind of sham, an attempt to dress up what is really going on?[15]

My reply is that while the imagined procedure is indeed farcical, its possibility is not implied by my account. First, I emphasize that the fact that a particular reason gave me a strong propensity to act *does* help to explain both my considering that possibility while deliberating and my carrying it out—assuming I

14. See, for example, Randy Clarke's (2003, 138–43) criticisms of an earlier formulation of the present proposal. Some material from this paragraph is drawn from O'Connor (2002).

15. Susan Brower-Toland objected along these lines in conversation.

continue to have the relevant desire at the time of my action. But these prior factors do not produce the action itself. Given all the prior causal influence on my deliberation, it remains that for the action to occur, I must exercise a causal power to intend to A or not. And it is possible for me, in at least some such cases, to reflect on my own motivations and, as a result, to deem the goal of a particular one of them as being the one I ultimately pursue. Now, doing so involves a change in perspective, something that, as a matter of human psychology, *takes time*. It implies that a major motivation diminishes in importance prior to the act itself. The scenario described above draws its absurdity from imagining instead that I can, so to speak, "package" my intention willy-nilly, apart from any change in the relative strength of motivations. But this is either to misconceive the nature of an intention to act or (more likely) mistakenly to think of the goal-directed content as part of a separate state that merely accompanies the intention.

III.3 Freedom and History

The medieval philosopher, Robert Grosseteste, in one place invites us to imagine God's creating an angel that exists for a single instant only.[16] In that instant, we are to imagine, the angel exercises his freedom by an instantaneous act of the will. In order to circumvent quite sensible worries about the coherence of Grosseteste's thought experiment, let us stretch out the angel's life to a few seconds. We will also suppose that the angel springs into existence with a fully developed psychology (of the typical angelic sort), complete with a bunch of pseudo-memories of a long history. And, as in Grosseteste's telling, he comes with a disposition to decide some matter straightaway. Finally, let us imagine that he is deciding between a plurality of alternatives, each having some attraction, is fully aware of his own motivations, and has the capacity to determine himself to do any of them. In short, he has a very high degree of the robust freedom I described a moment ago. Here's my question: in making his one and only choice, is our angel (whom we may call Andrew) just as free as a counterpart (Angela) who differs from Andrew only in that she really does have a history, one filled with many prior choices that have partly shaped her present inclinations and intentions?

Andrew and Angela do share the same choice-making capacity and exercise it to a similar "extent" (being fully and consciously self-aware and causing an intention that identifies the reason for which the action is pursued). But for Angela, unlike Andrew, the very factors that shape her choice were to some extent of her own making. Like Andrew, she began with a set of psychological and behavioral dispositions that were merely "given." But over time, as she habitually made certain choices, her psychological makeup reflected less and less this "givenness" and more and more something that is her own free creation.

And so, too, for ourselves. We come into the world with powerful tendencies that are refined by the particular circumstances in which we develop. All of these

16. See his *On the Freedom of the Will* (translated into English by Lewis [1991]). Grosseteste's views influenced (via Henry of Ghent) the arch-champion of freedom of the will, John Duns Scotus.

facts are for us merely "given." They determine which choices we have to make and which options we will consider (and how seriously) as we arrive at a more reflective age. However, presuming that we are fortunate enough not to be impacted by traumatic events that will forever limit what is psychologically possible for us, and, on the positive side, that we are exposed to a suitably rich form of horizon-expanding opportunities, the structure of our choices increasingly reflects our own prior choices. In this way, our freedom *grows* over time.

The common failure to appreciate this point again reflects a myopic view of freedom, one that is fixated on the basic metaphysical requirements of freedom to the exclusion of features reflected in the idea of self-mastery. For a further reason to think this is right, consider a scenario involving an agent much as we take ourselves to be—except that his psychology is regularly manipulated, altering some of his preferences and the "weights" of others. Owing to the marvelous wireless neural-intervention technology of the late 21st century, all this occurs while he remains wholly oblivious. Even if his capacity to choose remains robust, it seems clear that we must judge his freedom, his autonomy, to be diminished.[17] The integrity of the self-formation process is a further component of freedom, or of freedom of the most valuable sort. This conclusion is reflected in corresponding judgments about moral responsibility. If Andrew and Angela each contemplate a morally significant matter, are inclined to a degree to both a virtuous and a vicious action, and choose the virtuous one, Angela seems the more praiseworthy. The action is *hers* to a greater degree.

IV. IS THE CONSCIOUS WILL AN ILLUSION? (*OR*, WHY ARE THEY TRYING TO GET RID OF ME?)

Cognitive psychology and neuroscience have begun to study human experience of conscious willing—the conditions under which we have such experience and the factors that appear to influence its occurrence. The Harvard psychologist Daniel Wegner has recently drawn together much of this work in *The Illusion of Conscious Will*. As the title suggests, his conclusions are startling. If Wegner is right, my foregoing effort to clarify the nature of human freedom has been futile. Philosophers may continue to take an interest in the concept of freedom, but it can have no practical application. The notion that we are free rests on an illusory feeling of purposeful control, a feeling that masks the true, subconscious mechanisms that actually generate our actions. However, a careful scrutiny of Wegner's main evidence reveals that it does not support his thesis. It is indeed an empirical question whether anything like our account of human freedom actually applies to some of our actions, but evidence to the contrary has yet to be adduced.

IV.1 Wegner's Case against the Efficacy of Conscious Will

Here is how Wegner describes the experience of willed agency:

17. Al Mele (1995, chap. 9) makes this point and Clarke (2003, 16 n.4, 77) concurs. For a dissenting view, see Daniel Dennett (2003, 281–87).

> We each have a profound sense that we consciously will much of what we do, and we experience ourselves willing our actions many times a day. . . . We feel that *we cause ourselves to behave*. (2)

Will, as Wegner thinks of it, is *a feeling of doing*. The question, then, is how this feeling of doing varies with doings themselves, that is, voluntary actions. As Wegner notes, there are four possible scenarios. In the two predictable scenarios, they are jointly present or absent. When I am in fact acting voluntarily, I experience myself as willing my action. When I am not so acting, I correspondingly lack any experience as of willing. Wegner marshals a range of data suggesting there are also abnormal cases corresponding to the remaining two scenarios of doing without willing and willing without doing. It is cases of these latter two sorts that provide the basis for his ultimate conclusion that the experience of willing is quite incidental to the genesis of voluntary action *even in normal scenarios*.

Consider first *automatisms*, voluntary actions unaccompanied by any feeling that one is doing something.[18] We may subdivide these in turn into individual actions and group actions. An example of an individual automatism is "alien hand syndrome," where a person experiences his hand move as uncontrolled by himself, although the movements can display a complex purposiveness. Other automatisms involve the phenomenon of automatic writing (where one scribbles randomly until words begin to appear seemingly unbidden) and hypnosis. As a group-level automatism, Wegner cites the table turning episodes within the nineteenth-century spiritualist movement. People seat themselves around a table with finger tips lightly touching the edge. As they await a message from a dead person, they begin ever so slightly to rotate the table—while quite sincerely believing that they are entirely passive observers of this event.[19]

Wegner then turns to cases involving an *illusion of control*, where there is a feeling of doing which does not correspond in the right way to any actual voluntary action. In one type of case, the feeling occurs without *any* voluntary action. Here, Wegner cites the familiar scenario of one's "playing" a video game in an arcade that turns out to have been a replay of a previous game, shown for promotional effect. This false feeling of activity can also be induced as a result of brute social influence, as when a pretend instructor persuades the hapless learner that he pressed a key that destroyed a valuable computer file.

Experimenters have been able to induce a more subtle kind of illusion in which one's feeling of will allegedly *misidentifies the true cause* of the action. Wilder Penfield's "open-head" studies, in which he electrically stimulated motor-control parts of the brain, are justly famous. These were followed by similar studies by Delgado (1969). Like Penfield, Delgado was able to induce coordinated actions from his conscious patients. His patients, however, also cited reasons for their actions. Wegner suggests that these are likely to be fabricated, after-the-fact justi-

18. These cases are discussed in chapters 2 and 4 of *The Illusion of Conscious Will*.

19. A similar phenomenon occurs in the popular Ouija board game. And Wegner and a colleague devised an experiment permitting more careful documenting of this sort of phenomena in their "I Spy" study (Wegner and Wheatley 1999).

fications for what are nevertheless experienced as consciously willed (Wegner 2002, 46–47). More recently, Brasil-Neto et al. (1992) asked subjects to freely choose to move either their left or right index finger when signaled. While they awaited the signal, a magnet was moved across the motor area of either the left or right side of the brain. It turns out that subjects tended to move the finger contralateral to the side that was stimulated, although they "continued to perceive that they were voluntarily choosing which finger to move" (Wegner 2002, 48).

However, the most famous cases purportedly involving misidentification of the true cause of one's action arise in Benjamin Libet's (1985) studies on the timing of the experience of willing. Libet devised a study in which people are asked to wiggle their finger within a short interval of time (thirty seconds or so). The experimenter instructs them to do so whenever they wish—though spontaneously, not by deciding the moment in advance. They are then to watch a special clock and note the precise moment at which they felt the urge or wish to move the finger. During the experiment, a device measures electrical activity on the agent's scalp. Libet discovered that a steady increase in this activity (dubbed the "readiness potential," or RP) consistently preceded the time the agents cited as when they experienced the will to move. By averaging results over hundreds of experiments, Libet determined that the RP preceded the "experience of will" by some 500 milliseconds, a significant interval in the context of neural activity. Libet and others conclude that conscious will is not the initiator of voluntary acting but instead a *consequence* of an unconscious physical process that triggers the action.

From this wide range of data involving automatisms and illusory or misleading experiences of willing, Wegner constructs a general picture of what is going on with our conscious willings. His picture is summarized in the following passages:

> The feeling of consciously willing our actions . . . is not a direct readout of . . . will power. Rather, it is the result of a mental system whereby each of us *estimates* moment-to-moment the role that our minds play in our actions. (p. 15)
>
> We can't possibly know (let alone keep track of) the tremendous number of mechanical influences on our behavior because we inhabit an extraordinarily complicated machine. So we develop a shorthand, a belief in the causal efficacy of our conscious thoughts. We believe in the magic of our own causal agency. (pp. 27–8)
>
> [My] theory of apparent mental causation, then, is this:
>
> *People experience conscious will when they interpret their own thought as the cause of their action.* (p. 64)
>
> The real and apparent causal sequences relating thought and action probably do tend to correspond with each other some proportion of the time. After all, people are pretty good information processors when given access to the right information. The occurrence of conscious intention prior to action is often wonderful information because it provides a fine clue as to how things that are on the person's mind might pertain to what the person does . . .

> However, if conscious will is an experience that arises from the interpretation of cues to cognitive causality, then apparent mental causation is generated by an interpretive process that is fundamentally separate from the mechanistic process of real mental causation. The experience of will can be an indication that mind is causing action, especially if the person is a good self-interpreter, but it is not conclusive.
>
> The experience of will, then, is the way our minds portray their operations to us, not their actual operation. We come to think of . . . prior thoughts as intentions, and we develop the sense that the intentions have causal force even though they are actually just previews of what we may do. (p. 96)

IV.2 Where Wegner Went Wrong

To see where Wegner goes wrong, we must first revisit the target of Wegner's analysis. Let us make a few simple distinctions:

1. *Minimally voluntary action*: an action that accords with one or more of one's desires or intentions, conscious or unconscious, and that lacks external or internal compulsion.
2. *Feeling of wanting to perform an action*: what we might refer to with the terms "urge," "desire," and "wish."
3. *Feeling of will*: the experience of purposely initiating and/or directing one's action while it occurs.
4. *Forming a future-directed intention*: deciding to act in the future at the appropriate time or when one detects the appropriate stimulus.
5. *Forming a present-directed intention*: deciding to act here and now, which immediately issues in the action itself.

None of these phenomena, as we ordinarily conceive them, are identical to any of the others. The proper target for Wegner, given his conclusion, is the class of consciously formed, present-directed intentions. These we may also call "willings." An example is my intending *that I move my arm here and now*, which I experience as produced by me (for this or that reason) and which I believe to be efficacious in producing the movement. The target is *not* after-the-fact *reports* of one's intendings/willings, nor is it concurrent or after-the-fact *beliefs* about one's intendings/willings.

Wegner is not at all careful in keeping these items distinct over the course of his review of case studies. Perhaps this is because he attributes to ordinary folk a view of agency even more exaggerated than the philosopher's conception I criticized at the outset:

> Pointing to will as a force in a person that causes the person's action is the same kind of explanation as saying that God has caused an event. This is a stopper that trumps any other explanation but that still seems not to explain anything at all in a predictive sense. Just as we can't tell what God is going to do, we can't predict what the will is likely to do. (Wegner 2002, 13)

Now, common sense thinking on conceptual matters is to a degree inchoate. But, at a minimum, it is widely appreciated that much of what we do voluntarily is not consciously willed at all, but is simply a more or less unconscious outworking of previously settled intentions to act. Only some of the time do we consciously deliberate about what to do—and these are the actions we think of ourselves as most directly responsible for. And we easily recognize on reflection that when we deliberate, our self-knowledge is less than complete. The most we can claim is that we are *usually* aware of at least the *central* factors that are influencing our deliberation. Though the buck does indeed ultimately stop with us when it comes to initiating our actions, numerous factors influence our activity, making us likely to a degree to act in this way or that. Finally, we should observe that, because our ordinary self-conception is not highly detailed, it is *malleable*, capable of incorporating surprising data from the cognitive sciences.

Viewed from this relaxed perspective and with our simple distinctions in hand, much of the evidence that Wegner adduces underwhelms. Recall first the cases of *automatism*, such as alien hand movements or corporate table turnings at seances. Here, agents act voluntarily, in the minimal sense that their actions are clearly guided "from within" by certain of their own beliefs and intentions, although they do not experience the acts as something they are consciously controlling. Contrary to Wegner, however, it is no threat to our self-conception to allow that some voluntary actions are not consciously willed at all (as in the cases just mentioned) or are not consciously willed as to certain of their determining features (as perhaps with automatic writing or the group Ouija-board action).

Consider next alleged cases involving an *illusion of control*: one feels that one is consciously willing an action but is simply mistaken. Here, Wegner cites the unwitting "playing" of a video game demo and instances where one is pressured to acknowledge that one pressed the key that damaged the computer, even though one did no such thing. This is the crucial category for Wegner, yet it is woefully supported. The illusion of control in the video (and other machine-interaction) examples are merely cases of *belief*, based in external clues, that one's actions are determining certain outcomes beyond one's action. Clearly, these cases lend no support to thesis that the agent's *experience of willing* his own action is illusory. And the "you-wrecked-the-file" case is obviously a case of *after-the-fact* revision in the agent's *beliefs* about his prior behavior. That we can be induced to re-write our recollections of past actions is hardly a new discovery. The scientist observing this phenomenon should not be tempted to be partly complicit in it by agreeing that there was, after all, an experience of willing that fits the fabricated version of events.

The final set of data are the cases that, on Wegner's telling, involve a feeling of willing that misidentifies the true cause of one's action. A brain surgeon can directly induce an action that the unwitting patient claims was done for some purpose of his own. And the less intrusive procedure of subjecting part of a person's brain motor area to magnetic stimulation can significantly affect the likelihood of his performing one of two requested actions. The capacity to manipulate behavior in both of these ways is certainly striking. But neither needs to be interpreted in the manner Wegner suggests. As Wegner (following the original experi-

menter) notes, it is impossible to determine whether the brain-induced action is experienced as being consciously willed or is merely rationalized after the fact by the agent in an attempt to make sense of his own behavior. Furthermore, there is no reason at all to think that the agent falsely experiences himself as willing the action *for the reason he later cites*. In the case of magnetic stimulation affecting the outcome of a forced choice, we are presented with a surprising instance of a familiar phenomenon—a powerful (albeit obscure) unconscious factor affecting one's tendency to choose to act in certain ways. Such a case surprises only because of the blunt, external manner of the influence. But as we've noted, there are degrees of freedom, or autonomy, and the lesser varieties are consistent with powerful unconscious influences and even a measure of self-deception about one's purposes.

Taken cumulatively, the cases considered thus far do not justify Wegner's extraordinary conclusion that, quite generally, our experience of freedom is illusory—rather than being reflective of an actual self-determination of one's own decisions, the experience is epiphenomenal, caused by processes quite separate from the actual mechanisms of agency. But we have yet to consider Wegner's ace-in-the-hole, the finding by Libet that, under specific circumstances, a spiking of electrical activity in the motor cortex precedes an agent's experience of deciding to perform a specified action. Here is Wegner's summation of its significance:

> "The timing of will, finally, seems to seal the fate of that elusive lightbulb. The detailed analytical studies of the timing of action indicate that conscious will does not precede brain events leading to spontaneous voluntary action but rather follows them." (Wegner 2002, 60)

To assess the significance of Libet's studies, we need to interpret them in the light of our foregoing distinctions. Unfortunately, Libet shares Wegner's penchant for lumping several items together. In the experiments, a subject agrees at the outset to perform a specific action within a short interval of time. All that is left to be determined is the precise time of its occurrence. Though there is more that needs remarking on, we should not join Libet and Wegner in ignoring a very obvious fact here: in agreeing to cooperate with the experiment as described, *the agent has already decided to perform a specific action*. In our terms, she forms the intention to Ø, for a specific action-type Ø (say, flicking her wrist). No evidence is adduced that there is a slow build up towards a readiness potential before *this* decision. Libet and Wegner, however, will say that this is inconsequential. For the agent goes on to form another, ostensibly free, intention to perform the action here and now, and his choice is shown to have a neural antecedent associated with the movement itself.

But is this really so? In describing the instructions that the experimenter gives to the subject, as Mele (1997) points out, Libet (1985) uses several terms interchangeably: "urge," "desire," "wish," and "intention." The subject is asked to note the precise timing of such an urge/desire/wish/intention for each of several movements that will be performed in a single setting. But this set up and set of instructions invites interpretations differing sharply from Libet's own. First, the

action is hardly a "spontaneous" one. Even though the subjects are instructed not to pre-plan the timing of their actions, and instead wait for the urge/desire to do so, the action-type itself is pre-planned, and even its timing is to a significant degree, as they are instructed to make the movement within a thirty-second or so interval of time. Secondly, by asking the subjects to carefully introspect to pinpoint the timing of the impulse to move, the experimenter is inviting the subject to adopt the role of observer in relation to his conscious experience, and specifically to wait for an unplanned *urge* to occur. This certainly encourages a passive posture. Having decided that one will move, one looks for the *urge* to do so in order to act upon it. The result of a pre-formed intention to act upon the right internal "cue" for which one looks, I suggest, is that an unconscious process is triggered that promotes the occurrence (or perhaps *evolution*) of a conscious appetitive state that is not actively formed. In context, the state's default is to trigger the pre-planned activity, absent a last-second "veto" by the agent. In another study, Libet confirms the possibility of such vetoes subsequent to both the RP and the urge experience. In the confines of the original scenario, of course, the agent has no reason to do so. (Quite the contrary.)

I conclude that the design of Libet's studies renders them incapable of yielding the anti-libertarian goods. Taken together with our previous assessment of Wegner's other data, our verdict on the empirical case against freedom is negative: case unproved.[20]

V. CONCLUSION

We are created in the image of God—but it is closer to the truth to think of this as the form of a rough analogy than that of a small-scale replica. Part of the *Imago Dei* is our responsibility-conferring freedom. There is a univocal, common core to the concept of freedom applicable to God and those at various times applicable to ourselves. But God's perfect freedom is mysterious, whereas human freedom is merely a hard problem. If my reflections are on the right track, the major difficulties concerning human freedom stem not from the general level on which there is commonality with God's, but from that which is specific and different.[21]

20. Despite this limitation, Libet's pioneering studies are helpful in pointing to the need for greater attention to the precise phenomenology of different instances of willing or desiring to act. One wonders, for example, whether some or all of the subjects in Libet's experiments experience a growing anticipation that they are about to act, one that is not experienced vividly at first. Clearly, conscious awareness of our intentional states (as with our perceptual states) comes in degrees. Moreover, these states have phenomenal aspects that we cannot readily articulate. It may be that a better taxonomy of these features will yield another dimension to the self-mastery side of freedom.

21. Various ideas broached above differ from views I expressed in earlier work, including *Persons and Causes*. I am indebted to criticism of that work by Hendrickson (2002), Fischer (2001), Clarke (2003), Bishop (2003), Buckareff and Feldman (2003), and Mele (in correspondence). I read a shortened version of the present paper at the October 2004 *Divine and Human Freedom* conference at Wheaton College. I thank my commentator on that occasion (Ed Wierenga) and the audience for helpful comments and criticisms. I also thank an audience at St. Louis University, especially Michael Rota, Susan Brower-Toland, Scott Ragland, and Garth Hallett.

REFERENCES

Bishop, John. 2003. "Prospects for a Naturalist Libertarianism: O'Connor's *Persons and Causes*." *Philosophy and Phenomenological Research* 66:228–43.

Brasil-Neto, et al. 1992. "Focal Transcranial Magnetic Stimulation and Response Bias in a Forced Choice Task." *Journal of Neurology, Neurosurgery, and Psychiatry* 55:964–6.

Buckareff, Andrei, and Feldman, Richard. 2003. "Reasons Explanations and Pure Agency." *Philosophical Studies* 112:135–45.

Chisholm, Roderick. 2003. "Human Freedom and the Self." Reprinted in *Free Will*, 2nd ed., ed. Gary Watson, 26–37. New York: Oxford University Press.

Clarke, Randolph. 2003. *Libertarian Accounts of Free Will*. New York: Oxford University Press.

Davidson, Donald. 1980. "Actions, Reasons, and Causes," In *Essays on Actions and Events*, ed. Donald Davidson, 3–19. Oxford: Oxford University Press.

Davies, Brian, and Evans, Gillian R. 1998. *Anselm of Canterbury: The Major Works*. Oxford: Oxford University Press.

Delgado, Jose M. R. 1969. *Physical Control of the Mind: Toward a Psychocivilized Society*. New York: Harper and Row.

Dennett, Daniel. 2003. *Freedom Evolves*. New York: Viking.

Fischer, John. 1994. *The Metaphysics of Free Will*. Oxford: Blackwell.

Fischer, John. 2001. "Review of Timothy O'Connor, *Persons and Causes*." *Mind* 110:526–31.

Frankfurt, Harry. 1969. "Alternative Possibilities and Moral Responsibility." *Journal of Philosophy* 66:829–39.

Hendrickson, Noel. 2002. "Against an Agent-Causal Theory of Action." *The Southern Journal of Philosophy* 40:41–58.

Howard-Snyder, Daniel, and Howard-Snyder, Frances. 1994. "How an Unsurpassable Being Can Create a Surpassable World." *Faith and Philosophy* 11:260–8.

Lewis, Neil. 1991. "The First Recension of Robert Grosseteste's *De libero arbitrio*," *Mediaeval Studies* 53:1–88.

Libet, Benjamin. 1985. "Unconscious Cerebral Initiative and the Role of Conscious Will in Voluntary Action." *The Behavioral and Brain Sciences* 8:529–66.

McCann, Hugh. 1998. *The Works of Agency*. Ithaca: Cornell University Press.

Mele, Al. 1995. *Autonomous Agents*. New York: Oxford University Press.

———. 1997. "Strength of Motivation and Being in Control: Learning from Libet." *American Philosophical Quarterly* 34:319–32.

———. 1999. "Kane, Luck, and the Significance of Free Will." *Philosophical Explorations* 2:96–104.

O'Connor, Timothy. 2000. *Persons and Causes: The Metaphysics of Free Will*. New York: Oxford University Press.

———. 2002. "Dualist and Agent-Causal Theories." In *Oxford Handbook on Free Will*, ed. Robert Kane, 337–55. New York: Oxford University Press.

Pink, Thomas. 2004. *Free Will: A Very Short Introduction*. Oxford University Press.

———. forthcoming. *The Ethics of Action: Action and Self-Determination*. New York: Oxford University Press.

Rowe, William. 1993. "The Problem of Divine Perfection and Freedom." In *Reasoned Faith,* ed. Eleonore Stump. Ithaca: Cornell University Press.

———. 2004. *Can God Be Free?* New York: Oxford University Press.

Wegner, Daniel. 2002. *The Illusion of Conscious Will*. Cambridge, MA: MIT Press.

Wegner, Daniel, and Wheatley, T. 1999. "Apparent Mental Causation: Sources of the Experience of Will." *American Psychologist* 54:480–91.