Prima Facie Obligations, Ceteris Paribus Laws in Moral Theory*

Paul M. Pietroski

Morty promises to meet a friend at the station by 3:00. On his way there, he sees a seriously injured child in an alley; and helping the child will make Morty late. Morty ought to help children in need, but he also ought to keep his promises. So it seems that Morty ought to help the child *and* be at the station by 3:00, even if he cannot do both. Such moral conflicts are interesting for several reasons, not least of which is that they generate now-familiar paradoxes with the following general form:

i) some moral agent, *M*, ought to do *a;*
ii) *M* ought to do *b;*
iii) if *M* ought to do *X* and *M* ought to do *Y*, then *M* ought to do *X and Y;* therefore,
iv) *M* ought to do *a and b;*
v) if *M* ought to do *X*, then *M* can do *X;* therefore,
vi) *M* can do *a and b;* but
vii) *M* cannot do *a and b*.

Letting *a* stand for (the action of) being on time and *b* for helping the child, vii is true of Morty. So the route to vi must be blocked. The inferences are valid, and the deontic principles stated in iii and v are certainly plausible. It might seem, then, that we should deny the truth of i and/or ii. Here two options present themselves: one might hold that if an agent cannot satisfy all of the obligations ascribed to him, then at least one of the sentences ascribing obligations must be false; or one might hold that, like desires, sentences ascribing obligations do not have truth conditions at all.

If we are forced to adopt this last position, then as certain authors (e.g., Bernard Williams)[1] have suggested, the existence of moral conflict

* I would like to thank David Brink, David Davies, Susan Dwyer, Jim McGilvray, and Judith Thomson for helpful comments on earlier drafts.

1. See Bernard Williams, "Ethical Consistency," in *Moral Dilemmas,* ed. C. Gowans (New York: Oxford University Press, 1987).

Ethics 103 (April 1993): 489–515

has metaethical implications. For at a minimum, being a "moral realist" commits one to the claim that sentences ascribing obligations to agents have truth conditions. The realist will, presumably, also claim that such sentences are sometimes literally true (cf. J. L. Mackie).[2] But in any case, no position meriting the label 'moral realism' could be correct, if sentences like "Morty ought to keep his promise" are not even candidates for being true. In this article, I advance a familiar diagnosis of moral conflict that is more friendly to realism; and I defend this diagnosis by embedding it in a general account of moral principles like "Agents ought to keep their promises," according to which such principles play the same role in moral theory that statements of "ceteris paribus laws" play in the special sciences. I hope thereby not only to reconcile moral conflict with moral realism but to cast doubt on the common assumption that ethical theory is different in kind from, say, evolutionary theory.

Following W. D. Ross, I want to distinguish between an agent's prima facie obligations and his actual obligations.[3] (Ross spoke of prima facie duties vs. "duties proper" or "duties sans phrase.") The English word 'obligation', I suggest, is used ambiguously. So let "*M* ought_{pf} to do *a*" mean that *M* has a prima facie obligation to do *a;* and let "*M* $\text{ought}_{\text{act}}$ to do *a*" mean that *M* has an actual obligation to do *a*. Making this distinction allows us to dissolve the paradox above by citing an equivocation in the premises: Morty ought_{pf} to help the child, ought_{pf} to be on time, and $\text{ought}_{\text{act}}$ to help the child. But, I will argue, the agglomeration principle stated in iii holds only for actual obligation; and it is not the case that Morty $\text{ought}_{\text{act}}$ to be on time. Such a diagnosis will be satisfying, however, only if coupled with some account of what prima facie obligations are and how they relate to actual obligations. Thus, I draw an extended analogy between statements ascribing prima facie obligations to agents and nonmoral statements which are true only ceteris paribus. Similarly, I draw an analogy between statements ascribing actual obligations to agents and statements which describe the phenomena we explain by citing ceteris paribus laws in nonmoral domains.

Unlike other authors, I have spoken of moral "conflicts" rather than "dilemmas." I will reserve the latter term for cases in which the act of choosing which of two (or more) prima facie obligations to violate is epistemologically difficult and/or emotionally painful. Morty's case illustrates that reflection on these hard cases is not required to bring out the apparently paradoxical features of moral conflict, since the paradox above can be formulated in spite of its being reasonably clear

2. See J. L. Mackie, *Ethics: Inventing Right and Wrong* (New York: Penguin, 1977).
3. See W. D. Ross, "Prima Facie Duties," in Gowans, ed.

that Morty ought_{act} to help the child.[4] I think we should try first to resolve "easy" cases like Morty's, and then ask—as I will below—if genuine dilemmas pose any special nonepistemological problems for moral theory.

ROSS AND THE SOURCE OF PRIMA FACIE OBLIGATIONS

It is tempting to give 'prima facie' notions an epistemological gloss. If we took only Morty's promise into account, it would seem to us that Morty ought_{act} to be at the station at 3:00. So perhaps ascriptions of prima facie obligations are just claims about how things would seem from a particular epistemically limited viewpoint. On such an account, claims of the form "*M* ought_{pf} to do *X*" would reflect more about the epistemic state of the ascriber than about the moral situation of the agent in question. But one might also introduce the notion of an "all-things-considered" (atc) obligation. Perhaps Morty "ought_{atc}" to help the child, because that is how it would seem given all the relevant nonmoral facts. This suggestion is attractive, since it speaks to the deliberative and defeasible nature of moral reasoning. It is, however, distinctly verificationist; and it raises the very difficult question: Seem to *whom?* Us? Morty? An ideal epistemic and/or moral agent?

Ross explicitly rejected the epistemic gloss. "The phrase '*prima facie* duty' must be apologized for. . . . [It] suggests that one is speaking only of an appearance which a moral situation presents at first sight, and which may turn out to be illusory; whereas what I am speaking of is an objective fact involved in the nature of the situation, or more strictly in an element of its nature, though not, as duty proper does, arising from its *whole* nature."[5] This suggests that a moral situation, *S,* may, in some sense, be "composed" of various "right-making" factors or prima facie obligations, and that what an agent ought_{act} to do in *S* may depend on all such factors. Of course, the relevant sense of "composition" is not that of physical parthood. According to Ross, an action, *a,* is prima facie obligatory in virtue of being an act of a certain kind, for example, the keeping of a promise or the giving of aid; and whether *a* is actually obligatory "depends on *all* the morally significant kinds it is an instance of."[6] Suggesting an analogy between prima facie obligations and physical forces, Ross says: "Qua subject to the force of gravitation towards some other body, each body tends to move in a particular direction with a particular velocity; but its actual movement

4. Phillipa Foot makes the same point in "Moral Realism and Moral Dilemmas" (in Gowans, ed.). We can turn the example into a dilemma by fleshing it out in unexpected ways: the child's injury is not so serious; Morty swore that he would be on time; etc. But this points to another interesting issue, namely, the defeasibility of ethical intuitions.

5. Ross, pp. 86–87.

6. Ibid., p. 86.

depends on *all* the forces to which it is subject."[7] While Ross does not develop this analogy further, I would like to.

I take as given that actions are a species of events; and I assume that, at any given time, there is typically more than one action that an agent can perform, in the morally relevant sense of 'can'—it being notoriously difficult to say just which sense this is. That said, I think we should understand talk of actions being of several moral kinds as pointing to the fact that actions, both possible and actual, may be subsumed under several moral generalizations. For it is a common view that (1) kind predicates just are the predicates that appear in counterfactual-supporting generalizations, and (2) an event, *e,* is subsumed by a generalization, *G,* in virtue of having a description under which it satisfies a kind predicate of *G*.[8] As Donald Davidson and others have pointed out, an event may have multiple descriptions; so *e* may be subsumed by a distinct generalization *G′* as well, in virtue of having a description under which it satisfies one of the kind predicates of *G′*. A particular event may be subsumed by the law of gravity, because the relevant bodies have mass, and by Coulomb's law, because the bodies have charge. Thus, both "$F = Gmm'/d^2$" and "$F = qq'/d^2$" (as well as composition principles and "$F = MA$") will be relevant to the event that is the actual motion of the body; and in such cases, we speak of the bodies being subject to both gravitational and electromagnetic forces.

In moral theory, many (though perhaps not all) generalizations will be of the form: if initial conditions of kind *K* are met, then *M* ought_{pf} (not) to perform an action of kind *K′*. For example: if *M* promises to make some state of affairs, *s,* the case, then *M* ought_{pf} to perform an action, *a,* that makes *s* the case. For the moment, let us make the simplifying assumption that whenever initial conditions of kind *K* are met, exactly one possible action, *a,* of *M* is of kind *K′*. Then *a* will be prima facie obligatory (or prohibited) in virtue of satisfying the consequent of an obligation-ascribing generalization whose antecedent has been satisfied. An action will be of several moral kinds, if it satisfies the consequents of several such generalizations. For example, *a* may be the keeping of Mildred's promise *and* the murdering of a person; it may also be the dethroning of a tyrant and, hence, the promotion of freedom, etc. Here *a* is both prima facie obligatory and prima facie prohibited for Mildred; whether *a* is actually obligatory for her depends on all the moral generalizations under which *a* falls.

Mildred's case is a clear example of moral conflict, since she cannot perform the action *a and not a.* But this is only a special case, in which

7. Ibid., p. 94.

8. Here I follow Donald Davidson (see "Mental Events," in his *Essays on Actions and Events* [Oxford: Clarendon, 1980]) and Jerry Fodor (see "Special Sciences" in his *RePresentations* [Cambridge, Mass.: MIT Press, 1981]).

(1) it is logically impossible to perform the action in question and (2) one of the relevant "sub-actions" is a "negative action." In general, conflict will arise whenever a is prima facie obligatory in virtue being of kind K', b is prima facie obligatory in virtue being of kind K'', but the agent cannot (in the morally relevant sense, henceforth, ITMRS) perform *a and b*. There is no reason to expect that, whenever two possible actions of an agent fall under moral generalizations as described above, there will be a single course of action open to the agent which satisfies the consequent of each obligation-ascribing principle. Indeed, this is often not the case. Hence, an agent may be prima facie obligated to perform each of a set of actions but unable to perform all such actions. This is Morty's situation; and, I claim, his situation no more raises a paradox than does the fact that an event may be subsumed both by the law of gravity and by Coulomb's law when the initial conditions for both laws are satisfied.

The proposal, then, is that an agent acquires a prima facie obligation when the initial conditions of an appropriate moral generalization are satisfied. Or from the other side, M ought_{pf} (not) to do a if a is an action which satisfies the kind predicate in the consequent of an obligation-ascribing generalization whose antecedent has been satisfied. This is, I believe, in the spirit of Ross's remarks. It allows us to make sense of "morally relevant kind" and "right making factors" talk. For intuitively, promise keeping and aid giving are two such factors relevant to the question of what Morty $\text{ought}_{\text{act}}$ to do. We can explain this by noting that the antecedents of two distinct generalizations are satisfied; that is, two kinds of morally relevant initial conditions are met. Moreover, the proposed account is distinctly nonepistemological. Mildred, and everyone else for that matter, may be unaware that her flipping of a switch is the launching of a missile and thereby a killing of millions. But Mildred's action would still be (at least) prima facie impermissible, since an action's being of a certain kind does not depend on anyone knowing that such is the case. Thus, ignorance turns out to be irrelevant with respect to what prima facie obligations an agent has, although ignorance may be relevant when assessing blame. I think this is the right result.

Let me conclude this section with some brief remarks on how to complicate the simplifying assumption made above, namely, that there will be exactly one possible action of M which satisfies the consequent of the relevant moral generalization. Often, agents have several morally acceptable options available to them. There may be any number of particular actions that Manny could perform which would count as keeping his promise to "help out with the charity ball." He might donate ten dollars or decorate the hall; he might donate by signing a check or handing over a note, and he might decorate by stringing streamers or blowing balloons, etc. Here I think we want to say that

what Manny is prima facie obligated to perform is not any particular one of these actions but at least one of the many possible actions that would count as "helping out."[9] While this feature of obligation can complicate matters in particular cases, I do not think it poses any special theoretical difficulty.

On the other hand, there seem to be cases in which no possible action satisfies the consequent of the relevant moral generalization. Suppose the school bus has broken down. Manny announces, "I promise to fix the bus tomorrow"; but not knowing a carburetor from a carbohydrate, he is unable to do the work. Here the antecedent of a moral generalization is satisfied—Manny promised to make a state of affairs, *s*, the case. But by hypothesis, there is no possible (ITMRS) action, *a*, such that Manny's performing *a* will make *s* the case. It might be suggested that Manny fails to make a genuine promise here, precisely because a promiser must be able (ITMRS) to keep his promise. But I cannot believe that one can get "off the moral hook" that easily. Manny may have any number of "residual" (at least prima facie) obligations, ranging from apologizing to paying for the repairs, depending on the circumstances; and I do not see how he could be "on the moral hook" if he did not really make a promise.[10] If Manny did promise to fix the bus, however, the relevant initial conditions have been met; and so on the proposed account, Manny ought_{pf} to fix the bus—even though doing so is not possible (ITMRS) for him. I think this too is the right result, although it is potentially a problematic result. At a

9. More formally, let '$O(M,s)$' mean that M ought (in the specified sense) to make s the case. '$O(M,$ M performs an action, x, of type T)' is ambiguous as between (1) $\exists x[Tx \,\&\, O(M,$ M performs $x)]$, and (2) $O(M, \exists x[Tx \,\&\, M$ performs $x])$. The latter reading is the one relevant to Manny's helping out.

10. One might suggest the "promising-sentence principle" (PSP): M ought_{pf} not to say "I promise to do X," if M cannot (ITMRS) do X. I accept PSP; but only because I accept the "promise-making principle" (PMP): M makes a promise by uttering a promising sentence in a suitably serious context. The PSP is of interest here, only if coupled with a rejection of PMP; and the idea that M can fail to make a promise, yet still (somehow) violate the "institution of promising," strikes me as an unnecessary theoretical epicycle. Moreover, why do we think *particular* parties are wronged in these cases, if not because they are the parties to whom the promise was made? An epistemic variant on PSP whose antecdent is 'M believes that M cannot (ITMRS) do X' may have prudential value as a "rule of thumb" maxim for those with reliable beliefs. But I see no reason to adopt such a maxim as an ethical principle in its own right, unless we are convinced that PMP is false; and I see no reason to be so convinced, given that (*a*) promise making guarantees only prima facie obligations, (*b*) actual obligation is limited, as we shall see, by the "ought implies can" doctrine, and (*c*) we can distinguish failing to meet an obligation from being blameworthy. To be sure, deciding whether a context is "suitably serious" can be hard. Suppose you say to a stranger, "I promise to kill myself"; or to your captors, "I promise to return with a ransom if you free me." But no theory is likely to avoid these difficulties.

minimum, it conflicts with the "tidy" claim that M ought_{pf} to do a if and only if there is a possible action, a, of M which is prima facie obligatory for M. I return to these issues below. But perhaps for now, we can put aside questions concerning particular prima facie obligations which cannot be satisfied. For in the cases at hand, for example, Morty's, the agent can satisfy each, but not all, of his prima facie obligations.

ACTUAL OBLIGATIONS AND CETERIS PARIBUS LAWS

I have claimed that moral generalizations ascribe prima facie, not actual, obligations. But we need to see why this is so, and we need to see how the two notions of obligation are related. The short answer to both questions lies in the familiar fact that it is (almost?) impossible to find an exceptionless ethical principle. Universal generalizations like "If M promises to make some state of affairs, s, the case, then M $\text{ought}_{\text{act}}$ to perform an action, a, which makes s the case" are easily refuted by familiar kinds of counterexample. In Davidson's terms, such generalizations are not "strict." But ethical generalizations are hardly unique in this regard. Consider, for example, the following principle of "billiard ball mechanics" which subsumes many particular cases:

> If a force, f, impinges on a ball at rest, the ball will move in the direction of f at a velocity, v, where v is a function of the magnitude of f and the mass of the ball.

The truth of the antecedent does not guarantee the truth of the consequent. For there may be multiple forces impinging on the ball; and actual trajectories and velocities are typically the products of many interacting forces. But other things being equal, a ball will indeed move in the direction of a force impinging on it. The generalization above is thus best understood as implicitly containing an "other things equal" or "ceteris paribus" (henceforth, CP) clause.

This is not the place to defend a theory of truth conditions for lawlike statements containing CP clauses in any detail. But at a minimum, "exceptions" to a true CP law must be explicable as the result of some kind of "interference." Billiard ball mechanics is not falsified by the mere fact that some particular ball fails to move north in spite of being pushed north. We are, however, committed to there being an explanation for why such a ball fails to move north. Thus, we might say that a CP clause offers a "promissory note" to the effect that there is a correct explanation for each exception to the lawlike statement it modifies. The explanation will typically take the form of citing a further factor, for example, another force impinging on the ball from another di-

rection.[11] Proffered explanations will typically carry many empirical commitments—and can thus be independently tested—since the presence of an interfering factor can manifest itself in many ways besides that of producing the exception to the CP generalization in question. A famous example will make the point. Original observations suggested that Uranus's orbit did not conform to Newtonian mechanics. The explanation offered was that an as yet unknown planet was exerting a gravitational influence on Uranus. The empirical commitments of this explanation were clear and subsequently confirmed: the interfering factor, that is, Neptune, manifested itself to astronomers who looked in the right place.

This is hardly all that can be said about the theoretical role of CP clauses. But I hope it is enough for present purposes. Perhaps even more to the point, however, is the fact that, regardless of whether or not philosophers have a satisfactory account of how CP laws work, the generalizations discovered in the special sciences (evolutionary biology, economics, psychology, etc.) are not exceptionless; and these generalizations are not rejected on those grounds.[12] The methodological moral I draw is that (1) we should be unsurprised if generalizations like "Agents ought$_{act}$ to keep their promises" are also CP generalizations and (2) we should not reject such generalizations on the grounds that they are not exceptionless. That is, it is no more problematic to appeal

11. More formally, suppose a CP law has the form CP[$(\forall x)(Fx \text{—>>} (\exists y)Gy)$], where x and y range over objects in some domain; F and G are kind predicates of the science in question; and —>> represents the appropriate connective for nomic conditionals. Then the proposal in the text would be:

'CP[$(\forall x)(Fx \text{—>>} (\exists y)Gy)$]' is true, if and only if
$\forall x\{Fx \text{—>>} \{(\exists y)Gy \vee (\exists H)(\exists z)[(H \neq F) \ \&\ ([Hz] \text{ explains } [\neg(\exists y)Gy])]\}\}$.

That is, if initial conditions for the CP law are satisfied, then either the consequent condition is satisfied, or there is some independent factor, H, such that Hz (i.e., the fact that z has property H) explains why the consequent condition was not met. Note that CP laws would be far from tautologous on this account, since such laws would carry empirical commitments whenever their initial conditions were satisfied. The empirical commitment would not be simply that the consequent condition was satisfied, but this is just what we should expect given CP laws. For an elaboration and defense of this account, see Paul Pietroski and Georges Rey, "All Other Things Being Equal" (McGill University, 1992).

12. This is one of the morals of Fodor's "Special Sciences." Imre Lakatos makes similar suggestions in "The Methodology of Scientific Research Programmes," in *The Methodology of Scientific Research Programmes,* ed. John Morrall and Gregory Currie (Cambridge: Cambridge University Press, 1978). Perhaps, as N. Cartwright argues in *How the Laws of Physics Lie* (Oxford: Clarendon, 1983), the fundamental laws of physics are not strict; but resolution of this question depends in part on the possibility of unified field theory—a topic best left alone here.

to CP clauses in ethics than in, say, evolutionary biology. I propose, then, the following schema:

'M ought$_{pf}$ to do X' is true iff CP, M ought$_{act}$ to do X.

If what I have said so far is correct, we can read the English sentence "Agents ought to keep their promises" as (*a*) a claim which ascribes prima facie obligations or (*b*) a claim which ascribes actual obligations but which also contains an implicit CP clause. It is, perhaps, worth noting here that the F in "$F = Gmm'/d^2$" can also be read in two ways. If read as synonomous with the F in "$F = MA$," then the law of gravity must be read as a CP generalization if it is to be true at all. The magnitude, MA, will not be equal to Gmm'/d^2, for example, if the bodies in question have charge. If on the other hand—and this is the more common reading—the F in "$F = Gmm'/d^2$" is taken to mean the force due to gravity (a prima facie force?), we do not have to modify the law of gravity with a CP clause. But in that case, we are talking about a factor which contributes to the "resultant force" (which is equal to MA). Or, as Ross would put it, we are talking about "an objective fact involved in the nature of the situation, or more strictly in an element of its nature, though not arising from its whole nature."[13] I suspect, though I cannot argue here, that force talk is interchangeable with generalization talk; and again, at least in the special sciences, the relevant generalizations hold ceteris paribus. But this last claim is the crucial one. So I turn now to its defense by focusing on one of the best candidates for a special science law, Darwin's principle of evolution by natural selection. For we can easily see that a "CP-less" version of Darwin's principle would simply be false.[14]

A BIOLOGICAL DIGRESSION

Darwin's great insight can be formulated as follows: If (i) organisms possessing a property, P, are better able to survive and reproduce than organisms possessing an alternative property, P', and (ii) P and P' are heritable, then (iii) the proportion of organisms in the population having P will increase. Given Mendelian genetics, ii becomes ii*: P and P' are controlled by—or in a limiting case just are—gene complexes g and g', respectively. If i and ii*, then (iii*) the proportion of organisms in the population having g, and hence P, will increase. Given the common characterization of evolution as a change in gene frequencies in a population, iii* implies that (iv) evolution will occur "in the direction of" g, and hence P. Thus: if i and ii*, then iv.

13. Ross, p. 87.

14. In what follows I draw heavily on Elliott Sober, *The Nature of Selection* (Cambridge, Mass.: MIT Press, 1984).

As Elliott Sober points out, however, "[A] *ceteris paribus* clause needs to be added here: heritable variation in fitness will result in evolution only on the assumption that no counteracting forces cancel its effects."[15] Evolutionary theory itself tells us why this is so. If (1) certain organisms have a gene, g, which controls a slightly useful phenotypic trait, T, and (2) the competitor of g at its locus controls a selectively inert trait, then we expect the proportion of organisms having g and T to increase in the population. But this will not always happen. When a given gene is replicated, sometimes a nearby gene is copied as well. This is one of several ways in which the presence of genes controlling distinct phenotypic properties can be "linked." Suppose now that (3) g is linked to some other gene g'; (4) g' controls a phenotypic trait T', which is very deleterious; and (5) the competitor of g' at its locus is selectively inert. The proportion of organisms in the population having g and T will actually decrease.[16] The natural thing to say, I think, is that Darwin's principle is a CP law and that ceteris is not paribus if 3–5 hold. This should be unsurprising, since 3–5 explicitly introduce interfering factors into the example, and these explain why the proportion of organisms with T fails to increase.

Genetic drift, the biological analogue of sampling error in statistics, would be another source of counterexample to a "CP-less" principle of natural selection. If ten balls are drawn from a box containing fifty red and fifty blue balls, there is a chance that more than five (and perhaps all ten) of the balls drawn will be red. Similarly, the distribution of genes that come to reside in zygotes (the would-be children of the next generation) may differ from the distribution of genes in the sexually reproducing population at large. For each zygote will be composed of one gamete from each parent, each gamete having only one of the pair of parental genes for each locus. Because of "sampling error," the distribution of genes in zygotes is almost certain to differ somewhat from the distribution in the parental gene pool from which the zygotes were chosen. If organisms with gene g are somewhat better able to survive and reproduce than organisms with a competitor gene, g', but there is strong drift in favor of g', the result may be evolution

15. Ibid., p. 27.

16. Of course, this claim holds only CP as well. The gene g might also be linked to g'', which controls a very useful trait. (Compare n. 4 above.) I should note that when Sober speaks of organisms with property P being "better able" to survive and reproduce than those with some alternative property, he means that the former organisms are made "fitter" by virtue of having property P rather than the alternative to P. Suppose the swifter antelope have longer legs, but suppose they also happen to suffer from a serious disease—say, because of the kind of genetic "linkage" discussed in the text. We can still say that swiftness (i.e., long legs) make antelope fitter, even though the actual long-legged antelope do less well at surviving and reproducing (because of their other deleterious properties) than their short-legged and slower (but healthy) relations. See Sober, chap. 2, and also the discussion (on pp. 285–86) of Simpson's paradox.

in the direction of g'. Or again, if gene g mutates into g' more often than vice versa, selective pressure in favor of g may be nullified by mutation, or a combination of mutation and drift. Or to take a final mundane example, organisms with gene g may just happen to get struck by lightning or swallowed up by earthquakes more often.

Yet even with so many possible and interacting causes of evolution—I have mentioned only a few here—coherent evolutionary theory is possible. The biologist lays down a set of generalizations governing the phenomenon of evolution, each of which tells us something about a single factor that can contribute to evolution (e.g., natural selection, drift, or mutation). Exceptions to these "singleton laws," as Sober calls them, are to be explained, we hope, by citing other factors. Exceptions to Darwin's principle are to be explained, for example, by citing drift, mutation, lightning strikes, etc. But we know in advance that singleton laws will describe phenomena only CP. For by hypothesis, they take account of only one of the factors that may have contributed to a particular instance of the phenomenon in question.

This pattern is the norm in the special sciences. We can, if we like, retain the term 'law' for strict generalizations. The important point is that increases in the supply of X's are not always followed by decreases in the price of X's in market economies; but economists still retain the law of supply and demand. The AND gates in a computer "normally" register a 1 given two 1's as input, but not always; glitches occur. Psychological generalizations, especially those which make claims about behavior, are notoriously nonstrict. But such facts should come as no surprise, since there are a vast number of influences on phenomena like price and behavior. This does not preclude the possibility of true counterfactual-supporting generalizations in these domains. It is just that any such generalizations are likely to hold only CP.

BACK TO ETHICS

I take obligation-ascribing principles to be the CP generalizations of ethical theory. As such, they are the analogues of singleton laws in evolutionary theory. This is not to say that the obligation-ascribing principles we currently accept are true or unrevisable, or even that familiar examples of ethical principles have the right "theoretical depth."[17] It is, rather, to make a claim about the explanatory role of such principles. Generalizations like "Fitter traits increase their percentage shares in populations" hold CP, and they play important roles

17. Perhaps the obligation-ascribing principles we currently accept (or modifications of them) can be reduced to more basic ethical principles, e.g., "never treat persons as means." But I make no such claims here. Nor do I make any claims about the sources (e.g., rationality or utility maximizing) in virtue of which basic ethical generalizations are true.

in the explanation of phenomena in their domains. Similarly, obligation-ascribing principles tell us what agents $\text{ought}_{\text{act}}$ to do, CP, and these CP generalizations play an important role in explaining why agents have the actual obligations they do have. I take the fact that Morty $\text{ought}_{\text{act}}$ to help the child (and similar facts) to be among the phenomena we seek to explain by citing, inter alia, obligation-ascribing principles. As such, these moral facts are the analogues of biological facts of the form, "Population *P* evolved in manner *m*," which we seek to explain by citing, inter alia, the CP generalizations of evolutionary biology.

Of course, the facts explained in ethics and biology are different. In particular, appeal to evolutionary forces like natural selection can explain why populations "behave" (i.e., evolve) as they do, in the same way that appeal to Newtonian forces can explain why bodies behave (i.e., move) as they do. But appeal to prima facie obligations will not explain why an agent behaves as he does, except perhaps (*pace* Gilbert Harman) when the agent performs some action *because* he was obliged to do so.[18] The primary phenomenon we seek to explain by appealing to prima facie obligations in a given case is the fact that some particular action (as opposed to relevant alternatives) is actually obligatory, whereas by appealing to biological or Newtonian forces, we seek to explain why some particular trajectory of evolution or motion (as opposed to relevant alternatives) was actual. It is important to remember, though, that the behavior of objects is not the only kind of phenomenon one can legitimately seek to explain. For example, on the basis of many thermometer readings in controlled environments, a scientist may conclude that the melting point of lead is 327° C. Having made this inference, she will be far more interested in explaining this rather abstract fact about lead—that it melts at a certain temperature (as opposed to relevant alternatives)—than in explaining the behavior of any particular thermometer or lead-cum-thermometer system.[19]

One might, of course, advance antirealist arguments to the effect that there is no such phenomenon as actual obligation and, hence, nothing to explain by appealing to prima facie obligations. But I am trying to defend an account of moral conflict friendly to realism *given*, at least provisionally, the assumption that agents have actual obligations. If one assumes that statements ascribing actual obligations are never true, there is no point in discussing whether moral conflict, or anything else, poses a problem for moral realism. On the other hand, if one

18. See Gilbert Harman, *The Nature of Morality* (New York: Oxford University Press, 1977).

19. The example is from J. Bogen and J. Woodward (see their "Saving the Phenomena," *Philosophical Review* 97 [1988]: 303–52), who argue persuasively that the normal case in science is to use observations as evidence for the existence of some underlying phenomenon, and then explain that (typically nonobservable and nonbehavioral) phenomenon.

grants at least provisionally that (actual) obligation-ascribing statements are sometimes literally true, then there are phenomena—that is, instances of agents having actual obligations—we can seek to explain by citing prima facie obligations.

Let me sum up to this point. "Agents ought to keep their promises" does not mean that every promise made $\text{ought}_{\text{act}}$ to be kept, anymore than the corresponding biological generalization means that fitter traits always increase their percentage share in populations. We can say that, CP, a promise made $\text{ought}_{\text{act}}$ to be kept and, CP, fitter traits increase their percentage share. Or we can say that a promise made ought_{pf} to be kept and fitter traits tend to increase their percentage share in populations. But regardless of how we talk, ethical generalizations are on all fours with biological and other special science generalizations in this respect.

CONFLICTING OBLIGATIONS AND CONFLICTING FORCES

I have said that moral conflict arises when two or more distinct possible actions are subsumed by moral generalizations, with the result that each action is prima facie obligatory, but the agent is capable of performing only a proper subset of the actions in question. Thus, each action is such that, ceteris paribus, it $\text{ought}_{\text{act}}$ to be performed, but it cannot be the case that all such actions will be performed. This account, I now want to argue, is analogous to the account of "conflicting forces" we give in other domains over which we state CP generalizations.

Like physics, evolutionary biology is what Sober calls a "theory of forces." Theories of this form state singleton laws, which we have already discussed. We also want to be able to say something about the many cases in which more than one force is at work. Thus, a theory of forces will also state principles governing the various possible combinations of forces. But each theory must work out its own "combination principles." Newtonian vector addition is quantitative; and each vector figures in the calculation of the actual trajectory. Equal pushes to north and east will (CP) send a ball northeast. In an intuitive sense, actual motion is the result of "splitting the difference" between component vectors. But force theories do not have to follow this model. Compositional principles might "weight" particular forces more or less heavily (or perhaps not at all), depending on the circumstances. One can imagine a world in which a ball given equal pushes north and east moves north-northeast or even due north. In the latter case, we might say that a northward push "trumps" an eastward push. (Compare Ronald Dworkin's claim that rights trump utilities.)[20]

20. See Ronald Dworkin, "Rights as Trumps" (in *Theories of Rights*, ed. Jeremy Waldron [Oxford: Oxford University Press, 1984]).

Moreover, we sometimes aim for qualitative claims, and qualitative claims can have surprising compositional consequences. Drinking acid will (CP) make you sick; so will drinking a base. Yet drinking both will not make you twice as sick. Evolutionary biologists are often more concerned with qualitative claims of the form "There will be (some) evolution in the direction of trait *T*" than with quantitative claims which make exact predictions about the change in ratios of certain traits in a population. And we often settle for qualitative claims when that is all we are likely to get. Discovering principles that determine what happens to populations when selection, drift, and mutation are all at work is extraordinarily difficult. We can expect matters to be even more complicated in ethics. But it is worth noting that at least some portions of normative ethics (e.g., the literature on saving vs. not killing) certainly look like attempts to work out what happens when conflicting moral generalizations apply to a single case.

This is not yet to say that composition principles for ethics are waiting to be found. Perhaps the best we can do is to say that "intuition must decide" how to weight the various CP generalizations in various cases. But perhaps not; and I do not think we have to take a stand on this (hard) question in advance of the investigation. For I know of no good argument—as opposed to an assertion of metaphysical faith—for the claim that singleton CP generalizations can be true only if there are compositional principles governing all possible cases of interaction. Moreover, an objection to the present account based on (the lack of) composition principles in ethics would have to establish an asymmetry between moral theory and the special sciences in this regard. But given the paucity of composition principles discovered so far in the special sciences, establishing such an asymmetry is going to be difficult.

Returning to our original example, Morty ought_{pf} to be at the station on time, and Morty ought_{pf} to help the child in need. This is analogous to a situation in which selection in favor of trait *T* and drift in favor of a competitor trait, *T'*, are both at work in a population. Given the selection, CP, there will be evolution in the direction of *T;* given the drift, CP, there will be evolution in the direction of *T'*. Similarly, CP, Morty $\text{ought}_{\text{act}}$ to be at the station on time, and CP, Morty $\text{ought}_{\text{act}}$ to help the child. In the ethical case, the question is what Morty $\text{ought}_{\text{act}}$ to do, given his two prima facie obligations. In the biological case, the question is how the population will actually evolve, given the two evolutionary forces. But the former question strikes me as no more puzzling, at least in principle, than the latter. For the intuitive answer is surely the right answer in each case: it depends on the relative strength of the relevant forces. And here is the crucial point with respect to moral conflict.

Saying that Morty $\text{ought}_{\text{act}}$ to go to the station *and* that Morty $\text{ought}_{\text{act}}$ to help the child makes no more sense than saying that the population will actually evolve in the direction of T *and* that it will actually evolve in the direction of T'. The population will evolve in the direction of T, or it will evolve in the direction of T', or perhaps it will not evolve at all. (The 'or' here is exclusive; I return to the important third disjunct below.) But the population cannot evolve in both directions. Nor can a ball simultaneously pushed north and pushed south move north and move south. Similarly, Morty $\text{ought}_{\text{act}}$ to be at the station on time, or (exclusive) he $\text{ought}_{\text{act}}$ to help the child. But it is not the case that he $\text{ought}_{\text{act}}$ to do both. For he cannot do both. Thus, we can describe Morty's case as follows:

i) M ought_{pf} to do a (be at the station on time);
ii) M ought_{pf} to do b (help the child);
iii) M cannot do a *and* b;
iv) if M $\text{ought}_{\text{act}}$ to do X, then M can do X; therefore,
v) it is not the case that M $\text{ought}_{\text{act}}$ to do a *and* b;
vi) if M $\text{ought}_{\text{act}}$ to do X and M $\text{ought}_{\text{act}}$ to do Y, then M $\text{ought}_{\text{act}}$ to do X *and* Y; therefore,
vii) it is not the case that Morty $\text{ought}_{\text{act}}$ to do a, *or* (inclusive) it is not the case that he $\text{ought}_{\text{act}}$ to do b.
viii) In fact, Morty $\text{ought}_{\text{act}}$ to do b, and it is not the case that he $\text{ought}_{\text{act}}$ to do a.

This diagnosis maintains both the voluntarist principle and agglomeration for actual obligation, in iv and vi, respectively.

I think there is a great deal to be said in favor of agglomeration for actual obligation and very little to be said against it. The consequent of the principle can look odd, since we are not used to thinking about actions like being-on-time-and-helping-the-child. But many ordinary actions are also "compound." Traveling to Europe involves packing, loading the car, going to the airport, getting on the plane, etc. We often allow for temporally "disjoint" actions as well, and these can be the objects of obligation. It may be that what I $\text{ought}_{\text{act}}$ to do is bake the cake today and then give you the cake tomorrow. If we countenance compound events, we should allow that (at least in nondilemmatic cases) if M $\text{ought}_{\text{act}}$ to do X and M $\text{ought}_{\text{act}}$ to do Y, and then M $\text{ought}_{\text{act}}$ to do X *and* Y. For as long as we ascribe truth conditions at all to the consequent of this claim, the consequent is clearly true when both conjuncts in the antecedent are true.

To say that nondilemmatic cases conform to agglomeration is not yet to justify the deontic principle, since one might take cases of moral conflict to be sufficient reason for rejecting agglomeration. But this

would be a mistake, I think, for the fact that (at least) nondilemmatic cases so conform cries out for explanation. Moreover, if Morty were able to both keep his promise and help the child, it would be the case that (*a*) Morty $\text{ought}_{\text{act}}$ to be on time, and it would also be the case that (*b*) he $\text{ought}_{\text{act}}$ to help the child. From *a* and *b*, it would seem plausible to infer that (*c*) Morty $\text{ought}_{\text{act}}$ to help the child and be at the station on time; and I know of no principle other than agglomeration that can explain the permissibility of this and many similar inferences. To reject agglomeration is therefore to render a large class of intuitively compelling inferences mysterious; and other things being equal, this is bad methodology. So, given an account of moral conflict that allows us to do so, we should retain agglomeration for actual obligation.

It does not, however, follow from i and ii above that Morty ought_{pf} to be at the station on time and help the child. For ceteris paribus reasoning does not license such agglomerative inferences. Given natural selection in favor of gene *g*, there will be evolution in the direction of *g*, CP; and given counteracting drift in favor of a competitor gene, *g*′, there will be evolution in the direction of *g*′, CP. But it does not follow that given the selection and the drift there will be evolution in both directions, CP. Similarly, given forces impinging on the ball from the south and north, it may be that, CP, the ball will move north and, CP, the ball will move south. But it does not follow that, CP, the ball will move north and south. Such claims are absurd. We have reason to reject agglomeration for prima facie obligation, then, since prima facie obligations have been characterized in terms of CP generalizations. To reiterate: if *p* will actually obtain and *q* will actually obtain, then, trivially, *p* and *q* will actually obtain; but if *p* will obtain, CP, and *q* will obtain, CP, it does not follow that *p* and *q* will obtain CP.

Similar considerations may render less surprising the fact that the "tidy" claim mentioned above—*M* ought_{pf} to do *a* if and only if there is a possible action, *a*, of *M* which is prima facie obligatory for *M*—is false. The generalization "CP, bodies will fall at 16 feet/second^2" does not entail that bodies affected by friction fall at that rate, CP or otherwise. Indeed, given friction, it is a matter of physical law that bodies will not fall at exactly 16 feet/second^2.[21] For the same reasons, we cannot conclude that for every falling body, *b*, there is a physically

21. In short, ceteris paribus reasoning is nonmonotonic; i.e., the inference from "CP (if *F*, then *G*)" to "CP(if *F* & *H*, then *G*)" is not licensed. Compare "If you promised to return Fred's gun . . ." and "If you promised to return Fred's gun and discover that Fred is a homicidal maniac. . . ." This raises interesting questions about the proper degree of specificity for the antecedents of CP generalizations. Should we say that CP(if *F*, then *G*), counting *H* as a condition in which other things are not equal; or should we say that CP(if *F* & $\neg H$, then *G*)? Does one have a prima facie obligation, violated relatively often, to keep promises; or does one have a prima facie obligation, violated relatively rarely, to keep promises in certain circumstances? I cannot resolve these substantive questions here. But note that "CP(if *F*, then *G*)" and "CP(if *F* & *H*, then

possible event, *e*, such that *e* is a falling of *b* at 16 feet/second^2, since it may well be physically impossible for friction (and other impinging factors) to be absent. Or again, the ideal gas law abstracts away from the molecular attractions governed by Coulomb's law. Yet it is physically impossible for these attractions to be absent. But this just shows, I think, that the truth of a CP generalization, *G*, in the physical domain does not require that it be physically possible for ceteris to be paribus with respect to every event that *G* subsumes. If this is correct, then we should allow that "CP, *M* ought$_{act}$ to do *a*" can be true, even though given certain physical/psychological/etc. facts about *M*, it is not morally possible (i.e., not possible in the morally relevant sense) for *M* to perform *a*.

This brings us to the role of the voluntarist principle in cases of moral conflict. For it is this doctrine, together with the empirical facts about Morty's situation, that explains why it is not the case that Morty ought$_{act}$ to do that which he ought$_{pf}$ to do. The voluntarist principle strikes me as incontrovertible for actual obligation, since I have trouble even making sense of the claim that an agent ought$_{act}$ to do that which he cannot do. But for reasons already discussed, I do not think the voluntarist principle holds for prima facie obligation. Manny's being unable to fix the bus can explain why it is not the case that he ought$_{act}$ to do so. His inability does not, however, keep him from having a prima facie obligation to fix the bus. But precisely because I am making the voluntarist principle for actual obligation do theoretical work here, my claim of an asymmetry between prima facie and actual obligation in this regard needs independent motivation. Luckily, some has just been given. For we might express the point of the previous paragraph this way: if it is always possible (in *some* sense of 'possible') for other things to be equal with respect to physical generalizations, the relevant sense of possibility is wider than that of physical possibility. Analogously, a sense of possibility in which all prima facie obligatory actions are possible actions would be wider than "moral possibility," that is, wider than the morally relevant sense of possibility we have been appealing to all along. Given the plausible assumption that the sense of possibility relevant to the voluntarist principle just is this morally relevant sense, actions can be prima facie obligatory yet impossible relative to the voluntarist principle.

Two concerns facing the diagnosis of moral conflict sketched here must now be addressed. The first involves comparison of very different kinds of prima facie obligations, for example, Sartre's case of the young man whose country has been invaded by the enemy and whose dying mother desperately wants him by her side. The worry here is not

---

$\neg G$)" can both be true. While (CP) falling bodies accelerate at a rate of 16 feet/second^2, bodies affected by friction do not.

merely that we lack precise compositional principles which could handle this case but that such prima facie obligations are so unlike that neither can be said to be stronger than the other. We must also take up the related question earlier put aside: do genuine moral dilemmas raise any special nonepistemological difficulties for the moral realist? Second, prima facie obligations often seem to override one another "without remainder." A billiard ball pushed north and east can actually move northeast. But it is not the case that Morty $\text{ought}_{\text{act}}$ to somehow divide his energies, trying a little bit to get to the station and expending the rest of his energy helping the child. I have denied that Morty $\text{ought}_{\text{act}}$ to help the child *and* keep his promise, in favor of the claim that Morty $\text{ought}_{\text{act}}$ to help the child (period). There seems to be "moral residue," however, which this last claim does not speak to. For example, if Morty has done the right thing by helping the child, and if his prima facie obligation to be at the station on time was just simply overridden, regret about the moral "road not taken" would seem to be inappropriate. As a start to addressing these issues, it will be useful to press another aspect of the analogy between ethical theory and theories of forces.

ETHICAL FORCES AND MORAL CHOICES

Crucial to theories of forces is what Sober calls a "zero force law." Intuitively, such a generalization tells us what happens in the absence of forces. Put another way, it tells us what counts as "not changing" for the purposes of a theory that, inter alia, explains how and why change occurs. Newtonian theory provides the most famous example of a zero force law: a body in motion tends to stay in motion; a body at rest tends to stay at rest. In evolutionary biology, the Hardy-Weinberg law of population genetics tells us that, in the absence of evolutionary forces, the distribution of genes in a population will remain constant. The introduction of evolutionary forces, for example, selection, will (CP) bring about evolution. The zero force law thus provides a background against which singleton laws—the CP generalizations that tell us about individual forces—are introduced.[22] A plausible candidate for an ethical analogue that says what an agent's (actual) obligations are in the absence of "ethical forces" is the following:

> In the absence of prima facie obligations, any action is permissible.

I assume, standardly, that X is permissible for M if and only if it is not the case that M $\text{ought}_{\text{act}}$ not to perform X. (I will not be using the notion of 'prima facie permissibility'.)

22. Zero force laws can figure in explanations. If a moving body continues to move in the absence of any force, we can cite Newton's first law as the explanation. Changing

Suppose I must choose between chocolate chip mint and peanut butter fudge ice cream. As long as there are no prima facie obligations relevant to my choice, either option is morally permissible. If for some baroque reason choosing peanut butter fudge would cause millions to die, then I ought_{pf} not to choose peanut butter fudge.[23] But at least often, a moral agent, M, must choose among options $O_1, O_2, \ldots, O_n$, without it being the case that M ought_{pf} (not) to choose O_j for any j. There is no reason to think that, given any choice whatsoever, there must be a moral difference among the options; and if there is no j such that M ought_{pf} (not) to choose O_j, then any option, O_j, is morally permissible.

This proposal for an ethical zero force law has an interesting consequence for the diagnosis of genuine moral dilemmas. For consider what might be called a "net-force zero state" for theories of forces. If two equal and opposite forces impinge on a body at rest, the body remains at rest. Over the course of a number of generations, the combination of selective pressure in favor of gene g, drift in favor of g', and mutation in favor of g' may result in nó evolution at all. In an intuitive sense, the forces at work "cancel out" in such cases. This canceling need not be representable mathematically as in vector addition, especially when qualitative results are in question. Drinking an acid and a base can cancel out with respect to making a person sick. An obvious and important feature of these net-force zero states is that they "look" just like the corresponding zero force states; and because "[v]arious forces may have been present and canceled out each other. . . . [C]hange implies a force, but the absence of change does not imply the absence of force."[24]

We expect descriptions of zero force states to serve also as descriptions of the corresponding net-force zero states. So the proposed account of prima facie obligations leads us to expect that when an agent's prima facie obligations "cancel out," any action is permissible. I believe this to be case. Suppose that Aggie is forced by the fates to kill one of her twin daughters, Zoe and Zelda; and suppose that Zoe and Zelda are equally talented, personable, important, blameless, etc. Aggie has a very strong prima facie obligation not to kill Zoe, but similarly for Zelda; and by hypothesis, these prima facie obligations

---

zero force laws thus amounts to a change in what calls for explanation by citing forces. Explaining why moving bodies keep moving was difficult for Aristotle. For Newton, it is the "default" case which calls for no further explanation than citing the zero force law.

23. We can flesh out the notion of 'relevance' by saying: if M ought_{pf} (not) to do X, and if M would do X by doing Y, then M ought_{pf} (not) do Y. See Judith Thomson, *The Realm of Rights* (Cambridge, Mass.: Harvard University Press, 1990), for a defense of this principle.

24. Sober, pp. 34–35.

have equal moral weight. If the only two options available to Aggie are to kill Zoe or to kill Zelda, then my intuition is that either course of action is morally permissible. Either course of action will be terribly painful. But if those are her only options, I do not think Aggie acts wrongly, whichever option she chooses.[25]

It is sometimes thought that such cases present a difficulty for moral realists, who are supposed to (1) think there is a right answer to every moral question, and therefore (2) be embarrassed by the possibility of "moral ties." The former claim is dubious, but the inference to the latter is a non sequitur. For the right answer to Aggie's dilemma can be that either course of action is permissible. The realist is no more committed to (*a*) the claim that either Aggie $\text{ought}_{\text{act}}$ to kill Zoe or Aggie $\text{ought}_{\text{act}}$ to kill Zelda than to (*b*) the parallel claim concerning the choice between chocolate chip mint and peanut butter fudge ice cream. An important difference between these cases is that in the latter, no prima facie obligations are at work, while in the former, two very strong prima facie obligations happen to cancel out. As a result, it is not the case that, CP, I $\text{ought}_{\text{act}}$ to choose chocolate chip mint; but it is the case that, CP, Aggie $\text{ought}_{\text{act}}$ not to kill Zoe. Moreover, Aggie is not permitted to kill both daughters (since, presumably, that would be an unnecessary violation of a second prima facie obligation), whereas I may (CP, of course) resolve the ice cream dilemma by having both. Nonetheless, just as it is morally permissible for me to choose either flavor, it is morally permissible for Aggie to kill either daughter.

The prima facie obligation not to kill Zoe is, intuitively, very similar to the prima facie obligation not to kill Zelda. Moral dilemmas in which the agent's prima facie obligations are intuitively very dissimilar are, by hypothesis, more difficult to assess. But one might think that such cases present more than merely epistemological difficulties. For one might think there are prima facie obligations so dissimilar that, in no intuitive sense of 'strength', is either stronger than the other. But the moral realist can offer the same kind of diagnosis here: if two (or more) prima facie obligations really are so dissimilar that neither can properly be said to be stronger than the other, then any course of action among the set of relevant alternatives is permissible. If the prima facie obligation to aid the resistance is that dissimilar to the prima facie obligation to be with one's dying mother, then my intuition is that one may do either. Again, this can be the right moral answer to a difficult moral question.[26]

25. Of course, if Aggie could have done something earlier to avoid this predicament, she should have. Note also that Aggie could be obliged to kill one of her daughters without being obliged to kill either. (See n. 9 above.)

26. I am taking pains to avoid the term 'incommensurable' here to avoid a hornets' nest of issues. But one can treat incommensurable obligations—if such there be—as special cases of dissimilar obligations. Of course, we may discover that certain moral values are not so dissimilar after all.

We are now in a position to take up the question of regret. Here Williams's example of Agamemnon's choice between killing his daughter and carrying out his mission will be useful. Let us suppose that Agamemnon's prima facie obligation to carry out his mission is, in the relevant sense, stronger than his prima facie obligation not to kill his daughter. On the account developed here, he $\text{ought}_{\text{act}}$ to carry out his mission, and it is not the case that he $\text{ought}_{\text{act}}$ not to kill his daughter. That is, it is permissible (indeed, obligatory) for Agamemnon to kill his daughter. No doubt, Agamemnon will in fact feel regret about killing his daughter. But we may ask the justificatory question: should Agamemnon feel this way about having done what was, in the circumstance, the right thing? And if so, why is regret appropriate? The response that one would feel "natural regret" at having lost a daughter is singularly unsatisfying. For we think that Agamemnon's regret is morally appropriate, and in general, we want to find a place in our moral theory for a distinctively moral regret concerning the "road not taken" when an agent faces moral conflict. We also want our theory to explain, if possible, why regret is such a common (and perhaps necessary?) feature of moral experience. This looks to be a problem for the realist who says that Agamemnon did what he $\text{ought}_{\text{act}}$ to have done and, moreover, there was no other thing he $\text{ought}_{\text{act}}$ to have done.

It is important to remember that 'prima facie obligation' is a metaphysical notion. It may be that M ought_{pf} to do a without it being the case that M $\text{ought}_{\text{act}}$ to do a. But it does not follow that M's prima facie obligation "disappears," or never existed, if it is not the case that M $\text{ought}_{\text{act}}$ to do a. If a combination of drift and mutation outweigh selective pressure in favor of some gene, g, there will not be evolution in the direction of g. But it hardly follows that there is no selective pressure in favor of g. Nor does it follow that prior claims of the form "CP, there will be evolution in the direction of g" were false. Or again, if the force from the south is equal to the force from the north, the ball will remain in its place. But that hardly shows that no forces are being exerted on the ball; and these forces may affect the ball in many ways, without actually making it move. The unmoved ball may become mishapen, or even crushed. Similarly, just because M's prima facie obligation to do a does not manifest itself in its being the case that M $\text{ought}_{\text{act}}$ to do a, it does not follow that the prima facie obligation has no moral effect on M at all.

Regret, I claim, is morally appropriate whenever a prima facie obligation has been violated, regardless of whether it was morally permissible to violate it. An agent who has faced a dilemma may well find consolation in the fact that she did the right thing. But this does not change the fact that she violated a prima facie obligation in the course of doing so. Speaking metaphorically, we might say that a prima facie obligation to do a can "leave its mark" on a moral agent, M, even

if it is not the case that M $\text{ought}_{\text{act}}$ to do a; and one such mark is morally appropriate regret. Prima facie obligations may also manifest themselves as residual obligations. After the child has been put in the ambulance, Morty ought_{pf} to let his friend know that he will be late. The beginning of an explanation of this fact is that Morty's prima facie obligation to keep his promise did not "go away," simply because it was not the case that Morty $\text{ought}_{\text{act}}$ to keep that promise. This is not the place to develop a full account of residual obligations, even if I could do so. But it seems to me that appropriate regret and residual obligation are alike in being by-products of prima facie obligations that have been violated.[27]

Let me conclude this section with one last point about regret. We worry about the appropriateness of regret, I suspect, because we think there is at least something to the Aristotelian idea that having a certain kind of moral character (e.g., the kind that would feel regret in situations like Agamemnon's) is both morally appropriate and desirable. For even if one thought that acting in a morally appropriate manner was all there was to being a good moral agent, one would still have room for the importance of moral character, because a good, and perhaps the best, strategy for acting rightly in general is to make oneself into the kind of person who acts rightly in particular cases. If the picture of obligation I have sketched is accurate, then what one $\text{ought}_{\text{act}}$ to do is a function of one's prima facie obligations. It does not strictly follow that forming reliable beliefs about what one $\text{ought}_{\text{act}}$ to do in particular cases requires reliable beliefs about one's prima facie obligations. But I do not know how else one would go about forming reliable beliefs about what one $\text{ought}_{\text{act}}$ to do. If, CP, one $\text{ought}_{\text{act}}$ to keep one's promises, then one will have a better chance of acting rightly if one is sensitive to the moral fact that promisers have certain prima facie obligations, and similarly with respect to other prima facie obligations.

Such moral sensitivity—that is, being reliable with respect to noticing relevant prima facie obligations—has a personal cost. For when one prima facie obligation outweighs another, or when two cancel out, the agent will recognize that each of two competing prima facie obligations is relevant to the moral situation at hand. Choosing one or another course of action will not extinguish this sensitivity. Regret, I suggest, is thus the expression of a sensitive moral character in a complex moral situation. So we might well take the absence of regret to reflect badly on an agent's character; and we might well take the lack of

27. The claim that violations of prima facie obligations yield residual obligations, however, may not be true in general. See Thomson, who argues that the notion of violating a right is crucial here.

regret as evidence that the agent is not sensitive to her obligations and, hence, less likely to act appropriately.

TAKING THE ANALOGY (AND A DISANALOGY) SERIOUSLY

There are, to be sure, many respects in which moral and nonmoral theories differ. But my suggestion has been that there are important similarities and that for present purposes, the differences are irrelevant. Of course, no analogy is perfect. But one might think that (i) causal construals of biological (and other special science) generalizations are appropriate, (ii) causal construals of moral generalizations are inappropriate, and (iii) this disanalogy undermines the major claims of this article. I argue below that, even granting i and ii, the important aspects of the analogy drawn here between moral and nonmoral generalizations survive. But I think that ii may well be false; and while defending the contrary claim is a task for another day, I want to suggest that the matter is anything but obvious. My point, put succinctly, is this: the question of whether moral generalizations are causal depends largely on what the right theory of causation is. An adequate discussion would require both a survey of the plausible candidates for a theory of causation and a discussion of their implications for ethics. But since my present claim is only that moral generalizations are not obviously noncausal, it may suffice to consider briefly the "leading ideas" of several approaches to causation.[28]

Suppose we said that events of type S cause events of type T if S events raise the probability of T events. (Other, more plausible, probabilistic theories of causation—e.g., those of E. Eells and E. Sober or Brian Skyrms—are significantly more complicated; but such complications are likely to be irrelevant for present purposes.)[29] Now at least intuitively, it is more likely that Mildred $\text{ought}_{\text{act}}$ to give you ten dollars if she has promised to do so than if she has not. I certainly do not take such considerations as decisive. Perhaps an intuitive notion of probability is inadequate for these purposes. But it is far from clear that probabilistic theories of causation will count 'CP, if an agent promises to do X, then the agent $\text{ought}_{\text{act}}$ to do X', as noncausal.

28. Again, an objection along these lines must reveal an asymmetry between moral and nonmoral generalizations; i.e., the best theory of causation must also count the generalizations of physics and the special sciences as causal. Perhaps this is an adequacy constraint on a theory of causation, but perhaps not. In any case, the point is relevant to Davidson's account, which I mention only briefly here. If causal laws are strict laws, current biological generalizations are not causal laws; and the claim that such generalizations will be strict "when corrected" is an article of faith.

29. For Eells and Sober's theory, see E. Eells and E. Sober, "Probabalistic Causality and the Question of Transitivity," *Philosophy of Science* 50 (1983): 35–57; for that of Skyrms, see Brian Skyrms, *Causal Necessity* (New Haven, Conn.: Yale University Press, 1980).

Or again, suppose that *S* events cause *T* events if (*a*) the conditional claim 'Given an *S* event, a *T* event occurs' supports counterfactuals and (*b*) appeal to *S* events helps to (correctly) explain the occurrence of *T* events. Such a "pragmatic" account holds that the notion of causation is derivative on that of explanation, in the sense that we attribute causation wherever we offer explanation. The slogan here might be: all causation is "becausal." I have argued at length that obligation-ascribing principles figure in explanations; and claims like "Had John promised to help out, then (CP) he would have been obliged_{act} to help out" certainly seem to be true.[30] For similar reasons, it is unlikely that a counterfactual account of causation (e.g., David Lewis will count ethical generalizations as noncausal.[31] My tentative conclusion is that theories of causation do not support and, indeed, may help to revise the belief that ethical generalizations are noncausal. There are, however, independent worries about construing ethical generalizations causally. I think these can be diffused; though again, I will just offer some brief remarks here.

First, we apparently know the truth of claims like "Promises ought to be kept" a priori. Observation seems to be irrelevant to their confirmation, and these maxims seem almost trivial. Thus, it has been suggested (e.g., by John Searle) that such moral generalizations reflect conceptual and, hence, not causal, truths.[32] But even if (*pace* Quine) we grant that there are conceptual truths, determining whether a given sentence expresses a conceptual or empirical truth is notoriously hard. Moreover, Saul Kripke has argued persuasively that a priori-ness does not entail analytical-ness.[33] If we know the truth of certain moral generalizations in an a priori fashion, that tells us something about our epistemology. It does not, by itself anyway, tell us anything about the nature of the propositions known. The claim that ethical generalizations express conceptual truths should also trigger a sense of déjà vu. It has been (mistakenly) thought that Newton's laws and Darwin's principle were analytic truths; and the connection between mental states and behavior was, not long ago, taken to be conceptual and, hence, not causal. We should be wary of drawing similar conclusions about ethical generalizations too hastily. Certain moral generalizations

30. On the "pragmatic" account, see Hilary Putnam, "Why There Isn't a Ready-made World," in his *Realism and Reason*, vol. 3 (Cambridge: Cambridge University Press, 1983). Fodor makes similar suggestions. Note that in these counterfactual claims, the subscript *act* functions like an indexical, picking out the possible world(s) in which the promise in question is made. It need not pick out our actual world.

31. See David Lewis, "Causation," *Journal of Philosophy* 70 (1973): 556–57.

32. See John Searle, "How to Derive 'Ought' from 'Is'," *Philosophical Review* 73 (1964): 43–58.

33. See Saul Kripke, *Naming and Necessity* (Cambridge, Mass.: Harvard University Press, 1972).

may appear obvious. But being obvious does not even entail being true; a fortiori, it does not entail being conceptually true.

Second, the familiar temporal aspect of causal claims is missing from ethical generalizations. It is not that a promise is made and then an obligation is acquired. This suggests that the relation between promise making and obligation incurring is, if not conceptual, a relation of constitution or some other noncausal relation. There are, however, clear counterexamples to the claim that causes always precede their effects. Examples like Kant's lead ball on a cushion are perhaps open to the objection that elasticity ensures a slight causal precedence. But M. Brandt points out that John's going down on a seesaw synchronically causes the nonidentical event of the seesaw going down;[34] and Richard Taylor notes that the motion of a pencil is caused by the motion of a hand holding it, and the wind causes a leaf to flutter.[35] There are, of course, asymmetries between causes and effects. The wind is, in some sense, "responsible" for the motion of the leaf. Typically, the motion of the hand is what explains the motion of the pencil. Holding "background conditions" fixed, the same cause will have the same effect; but same effects need not have same causes. But similarly, an agent's prima facie obligations are, in some sense, "responsible" for his actual obligations. The direction of explanation will typically be from prima facie to actual obligations. Two agents cannot be alike with respect to all their prima facie obligations yet be unlike with respect to their actual obligations (holding background conditions—including other sources of actual obligation and ability to carry out obligations—fixed); but agents can be alike with respect to their actual obligations yet differ with respect to their prima facie obligations. If anything, this parallel of asymmetries speaks in favor of construing ethical generalizations causally.

Third, one might claim that normative properties just do not have causal powers. But I do not think this claim adds anything to the debate except a certain tone of voice. Since ethics is obviously concerned with the normative, insisting that normative properties cannot be causal is to beg the very question at hand. Nonetheless, let me now grant for purposes of argument that the relation between, say, promise making and obligation incurring is not one of cause and effect; and let us return to Morty.

Morty $\text{ought}_{\text{act}}$ to help the child; he $\text{ought}_{\text{act}}$ not to go to the station; he will feel some regret about having broken his promise; and he has a residual prima facie obligation to let his friend know that he will be late. But consider what Morty's situation would have been if

34. See M. Brandt, ed., *The Nature of Causation* (Urbana: University of Illinois Press, 1976).

35. See Richard Taylor, "Causation," in ibid.

circumstances were slightly different. If he had not come across the child, then it would have been the case that he $\text{ought}_{\text{act}}$ to be at the station on time (as long as no new interfering factors were introduced). On the other hand, had Morty not made the promise, he would feel no regret later; nor would he have had the residual obligation to get in touch with his friend. These points are painfully obvious, but they bear attention, for facts about obligation seem to be embedded in a counterfactual structure, just as facts about the evolution of populations and the motions of bodies are embedded in such a structure.

If selection in favor of gene *g* is outweighed by mutation, we expect the following to be true (CP): had it not been for the mutation, there would have been evolution in the direction of *g*. Similarly, if equal and opposite forces impinge on a ball from the south and north, the ball will not move. But had the force from the north been absent, the ball would have moved north. One might insist that the ethical examples are still disanalogous, on the grounds that (1) all counterfactuals must be grounded in causation and (2) causal construals of ethical generalizations are inappropriate. But this is now a double article of faith. Moreover, and crucially, one cannot consistently retain both dogmas, since manifestly, ethical generalizations support counterfactuals. In the physical, biological, and moral cases, the counterfactual facts suggest that actual situations are often complex, that is, they have underlying structure. Alter some part of that underlying structure, and the surface phenomenon changes. Remove a Newtonian force, and the ball may move in a different direction. Change the evolutionary forces, and evolution may take a different path. Alter the agent's prima facie obligations, and what the agent ought to do may change.

This kind of situation is familiar enough in the sciences. Theory often illuminates structure which underlies the phenomenon in question; and the phenomenon is thereby shown to be the product of underlying forces. We should be unsurprised, then, if ethical theory reveals the phenomenon of (actual) obligation to have underlying structure. Whether or not we call prima facie obligations 'forces' is irrelevant. If we wish to reserve the term 'force' for causal factors and refuse to admit ethical generalizations as causal, so be it. The important point is that prima facie obligations seem to play the same theoretical role that forces play in Newtonian physics and evolutionary biology. Claims about obligation are embedded in a counterfactual structure just as claims about motion and evolution are; and holding that something grounds these moral counterfactuals is simply to reject the miracle theory of moral counterfactuals. So if you do not like calling prima facie obligations 'ethical forces', call them 'those things (whatever they are) that play the relevant role in moral explanation and which ground the relevant moral counterfactuals'. I prefer the shorter term. Moreover, when push comes to shove, force talk may amount to nothing more

than explanatory-role/counterfactual-grounding talk. So in the absence of argument to the contrary, I think we should treat realistic construals of physical/biological forces and similar construals of prima facie obligations as on a par.

Perhaps we can even kick away the ladder of appeal to CP laws. Suppose that, contrary to what I believe, biology and other special sciences do not appeal to CP clauses, and that appeal to such clauses would be inappropriate in ethics as well. We know that special science "laws"—for example, Darwin's principle of natural selection—are false, unless modified by something like a CP clause. So anyone who eschews such clauses owes us an account of how the special sciences manage to provide good explanations with false laws.[36] Given such an account, we could still draw our analogy between, say, a generalization concerning the obligation of promise keeping and the principle of natural selection. Both would be false, construed as exceptionless generalizations. But this would not undermine the analogy developed here in terms of theoretical and explanatory role. Perhaps eschewing CP clauses would lead us to a kind of antirealism in physics, biology, and ethics as well. Perhaps this is some reason for not eschewing CP clauses. But in any case, such antirealism would not be peculiar to ethics. For we could still be as realist about ethics as we could be about biology. This is, I think, not a bad position for the moral realist to be in at the end of the day.

36. Cartwright tries to provide such an account for physics. But cf. R. Laymon, "Cartwright and the Lying Laws of Physics," *Journal of Philosophy* 86 (1989): 353–72; and Pietroski and Rey. Moreover, it is not clear how to extend Cartwright's strategy to the special sciences.