Another disappointment in Bostock's treatment of these issues is that, at several important points, he ends up identifying a confusion as the basis of Plato's view. For example, when we come (pp. 152–153) to the discussion of what Bostock calls "Plato's principle of causation": "the cause of a thing's being P must itself be P, and cannot be the opposite to being P," we are simply told that ". . . the truth is that this principle is plainly false, and could only *seem* to be true as the result of a conceptual muddle." Bostock diagnoses the confusion; but his remarks concerning what causes must be and his consequent treatment of Plato on causes are unsatisfactorily brief and extremely restrictive, given the notorious vicissitudes of that notion.

I had hoped to see discussion of Plato's remarks on explanation and of the vexed question of how Plato understood so-called self-predication which would have revealed the positive appeal of these strange doctrines. Bostock's discussion, like many which have gone before it, does not really do this satisfactorily. But perhaps I hoped for too much. At any rate, *Plato's Phaedo* should help readers to enjoy thinking about the points it discusses, both as issues in Plato scholarship and as philosophical problems.

CONSTANCE C. MEINWALD

University of Illinois at Chicago

The Philosophical Review, Vol. XCVIII, No. 1 (January 1989)

HARM TO SELF: THE MORAL LIMITS OF THE CRIMINAL LAW. By JOEL FEINBERG. New York, N.Y., Oxford University Press, 1986. Pp. xxiii, 420.

Harm to Self is the third volume in a four-volume work by Joel Feinberg, *The Moral Limits of the Criminal Law*. The longer work addresses the cluster of conceptual issues that surround the question, "What sorts of conduct may the state rightly make criminal?" In the first volume, *Harm to Others*, Feinberg presents and discusses the "harm principle": that the need to prevent harm to persons other than the actor is always a morally relevant (although not necessarily decisive) reason in support of proposed state coercion. Feinberg defines "harm" as a wrongful setback of an interest, and he explains how the harm principle bears on particular contexts and cases by articulating and defending a number of "mediating maxims." In the second volume, *Offense to Others*, Feinberg distinguishes harm from offense; whereas harms are wrongful setbacks of interests, offenses are not setbacks of interests, although they are wrongful impositions of unpleasant and unwanted effects. Feinberg puts forward and explains the "offense principle": that offense to others is always a morally relevant (al-

though not necessarily decisive) reason in support of proposed state coercion.

In the first two volumes, Feinberg defends both the harm and offense principles. Thus, his position is different from Mill's "extreme liberalism," which accepts the harm principle but not the offense principle. In the third volume, *Harm to Self*, Feinberg discusses but rejects a third "liberty-limiting principle," "legal paternalism": that the need to prevent harm to the actor himself is always a morally relevant (although not necessarily decisive) reason in support of state coercion. In the four-volume work Feinberg undertakes to defend "liberalism" (with respect to the moral limits of the criminal law), the position that the harm and offense principles, and *only* those principles, specify good reasons for criminal prohibitions. (In the fourth volume, Feinberg will reject such principles as "legal moralism," which allows the prohibition of conduct on the ground that it is inherently immoral, apart from any consideration of harm or offense to others.)

As an example of a "mediating maxim," consider the "*volenti* maxim" (*volenti non fit inuria*): a person is not wronged by that to which he consents. On Feinberg's approach, the *volenti* maxim mediates between the harm and offense principles and particular cases, but *not* between the principle of legal paternalism and particular cases. So if one voluntarily consents to another's behavior which sets back one's interests or offends one, one is not thereby *wronged* and thus neither the harm principle nor the offense principle applies. In contrast, one's consent to behavior which will or might harm one does *not* render the principle of legal paternalism inapplicable. Indeed, it is paradigmatically in cases in which there is substantially voluntary consent to behavior which will (or might) harm one that the principle of legal paternalism becomes relevant.

In *Harm to Self* Feinberg points out that there are presumptive cases both for and against legal paternalism. The presumptive case against paternalism issues from our intuitive reaction of resentment and indignation when others substitute their judgments for our own: it is as if we adults are being treated as children. The presumptive case for paternalism comes from the importance of preventing harms and the intuitive reasonableness of many restrictions that appear to be paternalistic. Feinberg argues however that the decisive argument against legal paternalism is that it does not have an adequate conception of personal autonomy; that is, it does not give sufficient weight to an individual's sovereign right to govern himself. In regard to those restrictions that appear to be both reasonable and paternalistic, Feinberg argues that they are either not really reasonable or not really paternalistic. Feinberg thus opts for a position which he calls "soft paternalism": that the state has the right to prevent self-regarding harmful conduct when but only when that conduct is substantially non-voluntary, or when temporary intervention is necessary to

determine whether it is voluntary. (A "hard paternalist" would permit intervention even in self-regarding behavior that is substantially voluntary. Feinberg points out that soft paternalism is not really paternalism at all, since it does not substitute one person's choice for another's "true choice.")

According to Feinberg soft paternalism protects individual autonomy in a robust sense. An individual's choices (concerning self-regarding conduct) may be eccentric, bizarre, unfathomable, unreasonable, and so forth, but insofar as they are substantially voluntary, they must not be overridden. Feinberg emphasizes the distinction between unreasonable and irrational choices and behavior; it is only if a choice is irrational in the sense of being a product of a substantially impaired faculty of practical reasoning that it may be overridden. Clearly, Feinberg's approach requires that we make judgments about the degree of voluntariness of a person's behavior, and much of *Harm to Self* is devoted to a systematic and subtle discussion of the relationship between the voluntariness of behavior and the permissibility of intervention.

We might consider whether it would be permissible to intervene coercively in another's self-regarding behavior in a great variety of different contexts. One approach to the issue of intervention would be context-invariant: it would specify some level of voluntariness and then claim that any behavior that is at least *that* voluntary must be protected, no matter what the context. In contrast, Feinberg argues cogently for a context-relative approach. On this sort of approach, the level of voluntariness required for the protection of choice and behavior from intervention varies with the context. For example, in the context of a choice of a person to end his own life, which is both extremely important and irrevocable, a very high standard of voluntariness is appropriate. In contrast, in the context of a decision about gambling a small amount of money a rather lower standard of voluntariness is appropriate. In general, the more risky and irrevocable the conduct, the greater the degree of voluntariness required if the conduct is to be permitted.

Feinberg distinguishes between "single-party" and "two-party" cases. In single-party cases, we consider a person's own self-regarding behavior. In two-party cases, we consider the effects of the behavior of one person on another who has consented to the behavior. In both kinds of cases, there are three kinds of "voluntariness-reducing" factors: compulsion and coercion, ignorance, and incapacity. In one-party cases, these factors bear directly on the voluntariness of the choice and behavior of the agent, and in two-party cases, these factors bear directly on the consent of the one person to the behavior of the other. Feinberg's discussion of the three voluntariness-reducing factors is remarkably comprehensive and always illuminating.

I find Feinberg's discussion of compulsion and coercion particularly fas-

cinating. In a case of "compulsion proper," some force makes it literally impossible for the agent to do otherwise: a person is sent reeling by a hurricane wind or an explosion, and so forth. In contrast, coercion does not destroy an alternative, but it "destroys its appeal by increasing its cost" (p. 191). Coercive pressure exists insofar as the alternative to the action is rendered undesirable to some extent. "Coercion proper" exists when the alternative is rendered *too costly* (although not literally impossible). To apply Feinberg's distinction to a slightly different set of issues, a person who is compelled to do what he does would not (in general) be morally responsible for his behavior, whereas a person who is coerced is morally responsible for what he does, although (perhaps) not blameworthy for so acting. A person who is compelled is not even a rational candidate for moral praise or blame; in contrast, a person who is coerced is accountable for his behavior, although it may be unreasonable to praise or blame him. Feinberg usefully develops different ways of measuring the degree of coercive pressure on an agent (pp. 199–210). He distinguishes threats from offers (pp. 216–228), and he argues that both threats and offers can be coercive (pp. 229–268).

To oversimplify an intricate and nuanced discussion, when a person makes an offer he gives another the opportunity to do better than what he could normally expect. When a person makes a threat, he warns another that if he doesn't cooperate, he will do worse than what he could normally expect. Of course, it is important to disambiguate the crucial phrase, "what he could normally expect." Feinberg distinguishes four interpretations, and he defends a version of the "statistical" interpretation (pp. 219–228). On Feinberg's approach, an offer can coerce even though it is "freedom-enhancing" (that is, the offer increases the options available to an agent). Such an offer is coercive insofar as it forces an agent to take a course of action by making the alternatives unduly costly.

Consider Feinberg's example of the lecherous millionaire who offers to pay for the expensive surgery that alone can save a woman's child provided that she becomes for a period his mistress (p. 229). The millionaire's proposal is an offer because it "does not threaten any harm beyond what would happen anyway without his gratuitous intervention" (p. 230). Indeed, it proposes a consequence—that the child is saved—that is preferable to what the mother can "normally expect," given her child's illness. It is coercive in part because the alternative to cooperation is unreasonably bad. (In this respect a coercive offer is different from a merely enticing offer: pp. 233–235). Thus, the lecherous millionaire's offer is coercive even though it enhances the options available to the mother.

This book is extremely wide-ranging and rich, and it is impossible to do justice to it in a short review. I shall be content to raise a few questions about Feinberg's arguments for liberalism. I believe that the most trou-

bling challenge to Feinberg's anti-paternalism comes from the view that one can sometimes override substantially voluntary choices in circumstances in which the choices concern trivial or unimportant things and the risks of harm are great. A proponent of this sort of view is Gerald Dworkin.[1] Dworkin argues that the state may require such things as the use of safety belts and helmets and the wearing of certain clothes when hunting on the grounds that these are only "trivial interferences with autonomy" and that they protect individuals against great harms. On this approach, an individual's right to make and act upon a very wide range of decisions concerning important matters is preserved, but protection of the individual against great harm is weighed more heavily than his autonomy in certain contexts.

In response to this sort of view, Feinberg emphasizes the distinction between *de facto* and *de jure* autonomy. *De facto* autonomy is the actual condition of self-government, whereas *de jure* autonomy is the *moral right* to self-government. On Feinberg's view, the *de jure* autonomy of a person is analogous to the sovereignty of a nation. Feinberg argues that national sovereignty (and thus personal sovereignty) is an "all or nothing affair": "a nation's sovereignty is equally infringed by a single foreign fishing boat in its territorial waters as by a squadron of jet fighters flying over its capital" (p. 55). Sovereignty implies "absolute control over whatever is within one's domain however trivial it may be" (p. 55). Thus, for Feinberg, whereas *de facto* autonomy is a matter of degree, there cannot be a "trivial interference" with *de jure* autonomy, and Dworkin's view (at least as expressed by Dworkin) is incoherent (p. 94).

Feinberg considers the possibility that a proponent of Dworkin's view will simply re-define the boundaries of the *domain* of personal sovereignty. On this view, the domain of personal sovereignty would include the vital life-decisions, and the individual would have absolute control over this domain. But Feinberg claims that there is no plausible alternative to employing the notion of self-regarding behavior as the boundary of the domain of personal sovereignty (p. 57). Feinberg appears to believe that a proponent of narrowing the boundaries of personal sovereignty so as not to include trivial issues is either completely without any intuitive conception of autonomy, or lacks a conception which is sufficiently robust (insofar as the conception of autonomy seems to collapse into concern for the agent's good).

I am unconvinced by Feinberg's arguments, even as I am attracted to his

[1]Gerald Dworkin, "Paternalism: Some Second Thoughts," *Paternalism*, ed. Rolf Sartorius (Minneapolis, Minn.: University of Minnesota Press, 1983). See also: Gerald Dworkin, "Paternalism," *Morality and the Law*, ed. Richard A. Wasserstrom (Belmont, Calif.: Wadsworth, 1971).

anti-paternalistic conclusion. When one distinguishes between *de facto* and *de jure* autonomy, I do not see why one needs to say that the only (or the relevant) notion of *de jure* personal autonomy is analogous to the "all-or-nothing" notion of national sovereignty. It seems to me that it is coherent to have a *de jure* notion of personal autonomy that is not a form of (or analogous to) sovereignty, in which case Dworkin's claim that there can be relatively trivial interferences with autonomy is not incoherent. I see no argument in Feinberg that the all-or-nothing form of personal autonomy is the *only* coherent model of *de jure* autonomy.

Further, I do not see any reason to suppose that one who employs either a model of *de jure* autonomy that is not a form of sovereignty or a model of *de jure* autonomy that is a form of sovereignty with narrower boundaries than self-regarding behavior cannot have an intuitive conception of autonomy that is attractive and robust. Such a person holds that individuals have the right to control vital and important decisions; indeed, it is only a small class of relatively unimportant decisions that one may override. A proponent of the Dworkin-type strategy will not intervene in an individual's choice of career, friends, spouse, religion, and so forth in order to protect him from even very significant harm. This sort of view does not appear to imply a collapse of autonomy into protection of an individual's good.

Now I certainly am not suggesting that Dworkin-type paternalism is obviously acceptable. I am simply pointing out that there does not appear to be a convincing argument in Feinberg that rules it out. The analogy with national sovereignty is interesting, but it is not decisive in the absence of an argument that personal sovereignty is the *only* coherent model of *de jure* autonomy or that a narrower way of fixing the domain of personal sovereignty must abandon any sufficiently attractive conception of self-government. Perhaps Feinberg will rest his case on the difficulty of identifying the vital life-decisions (p. 93), and there is no question that this is very difficult. But it is not clear to me that it is substantially more difficult to identify the vital life-decisions than the "self-regarding" ones. Surely we can say that whether a hunter is to wear a red or green jacket is *not* among the vital choices, whereas one's choice of a career is, and it is not evident that we are unwilling to override a particular hunter's desires in this matter in order to protect him against significant harm.

Harm to Self is an outstanding book. It is lucid, comprehensive, and extremely insightful throughout. The book contains just the right mix of philosophical theorizing and consideration of particular examples—both hypothetical and actual. Indeed, Feinberg fruitfully applies his ideas to a wide range of current political and legal issues, thus exhibiting the power of a good philosophical theory. I believe that *Harm to Self* is the best systematic discussion of the problem of paternalism that is available. This

book, together with the previous two volumes, constitute the most complete and illuminating discussion of the group of issues pertaining to individual liberty and the limits of state coercion with which I am familiar. I am certain that the four-volume work will become a classic in its field.

JOHN MARTIN FISCHER

University of California, Riverside

The Philosophical Review, Vol. XCVIII, No. 1 (January 1989)

THE PHILOSOPHY OF RIGHT AND WRONG. By BERNARD MAYO. London, England, Routledge and Kegan Paul, 1986. Pp. 176.

This book is an introduction to moral philosophy written for non-specialists and intended to be suitable for use as a textbook.

The greater portion of the book is devoted to a consideration of questions about the meaning of moral judgments. Mayo develops five tests for a satisfactory theory of meaning: 1. the theory must not postulate the existence of occult properties; 2. the theory must account for the motivational force and action-guiding character of morality, and 3. the theory must give a satisfactory account of disagreements about morality; 4. the theory must be able to explain the universalizability of moral judgments, and 5. it must be able to make sense of our "moral autonomy." These conditions (at least the first four) are plausible as necessary conditions for the adequacy of a theory of meaning, but Mayo provides no reasons for thinking that they are sufficient. Mayo rejects intuitionism and naturalism on the grounds that they fail all or most of the tests. He regards emotivism as a more plausible theory, but argues that it cannot satisfy the third condition.

Mayo defends a version of prescriptivism. He argues plausibly that the theory meets all of his tests; however, he does not show or attempt to show that it is the only theory capable of meeting these tests. Prescriptivism is very briefly characterized as the view that morality is analogous to the law. Mayo then begins a long and interesting comparison of morality and the law. Unfortunately, he never reformulates his version of prescriptivism in light of this discussion.

According to Hare's version of prescriptivism, there are no substantive restrictions on the content of moral judgments. Any prescription that we are willing to universalize and treat as overriding constitutes a moral judgment. Mayo devotes two separate chapters to a consideration of substantive restrictions on the content of moral judgments. (These are apparently intended as restrictions on the content of *reasonable* moral judgments, not as restrictions on the content of what can be classified as moral as opposed