1993

Institutional Morality

Russell Hardin

PAPERS IN POLITICAL ECONOMY

Paper No. 38

"Institutional Morality"

Russell Hardin

The UNIVERSITY *of* WESTERN ONTARIO

The Political Economy Research Group was established in the faculty of Social Science at the University of Western Ontario in 1988. Its purpose is to foster scholarship, teaching and interdisciplinary research in political economy, with a focus on:

1. the application of economic models and methods to the study of political processes and institutions,

2. the economic impact of political processes and institutions,

3. the influence of economic factors on the formation of public policy and on institutional change,

4. the politics of economic policy making,

5. the political, social, and economic effects of public policy.

Co-directors:
Ronald Wintrobe (Economics)
Robert Young (Political Science)

Board of Directors:
Peter Howitt (Economics)
B.B. Kymlicka (Political Science)
John N. McDougall (Political Science)
Peter Neary (History)
John Whalley (Economics)

Staff:
Jayne Dewar

For further information:
Political Economy Research Group,
Department of Economics,
Social Science Centre,
London, Ontario, Canada N6A 5C2
phone: (519) 661-3877
fax: (519) 661-3292

INSTITUTIONAL MORALITY

Russell Hardin

New York University
Department of Politics
715 Broadway
New York, NY 10003

Contents

INSTITUTIONAL MORALITY*

Russell Hardin
New York University

Introduction

A central problem in the design of institutions is the assignment of responsibility. Almost all moral theories are about the morality of individuals, not of supra-individual entities such as organizations or societies. Hence, in those theories, an account of institutional responsibility must be subject to unpacking down to the level of individuals. This raises questions of composition: Who is how much responsible for which part of what? It is these compositional problems that are the focus of discussion here.

In general, we may say of an institution as a whole that its morality is reasonably well defined by its purpose and its likely effects. In this sense, an institution's morality is inherently consequentialist. It might not be utilitarian, however, because its purpose could be the support of equality, justice, or autonomy, and not necessarily the support of welfare.[1] Once the institution's morality is defined, we may infer or deduce the morality of individual office holders in the institution as derivative from the purpose of the institution. Or, we may say, the rules and procedures that guide the actions of individual office holders are determined by the functional relation of their actions to the general purpose of the institution.[2] Institutional morality is structurally similar to some variants of virtue theory: Its content is functionally determined by the purpose it is to serve. And it is analogous to the problem of legal justification. We may first seek justification _of_ the law and then justification _in_ the law. The latter depends on the former for much of its content.

It also depends on social scientific facts of the way incentives and institutions can be made to work.

The consequentialist evaluation of institutions might seem to fit ill with many moral theories, which, as their proponents assert, are not consequentialist. But these theorists often evaluate institutions by how well they achieve or support autonomy, justice, fairness, and so forth. For example, one must grant that John Rawls's theory of justice is consequentialist when it focuses on institutions for achieving justice even if one supposes that, in its later Kantian variant, it is grounded in nonconsequentialist concerns. One might suppose it fine to have institutions for the achievement of various consequences, including welfare enhancement, while still insisting that nonconsequentialist considerations also play a role in a particular moral theory. In all of these cases, as well as in fully consequentialist theories such as utilitarianism, the problems addressed here are central to the moral assessment of institutional design. Indeed, it may be difficult to give a compelling anticonsequentialist justification of an institution.

Problems of composition in institutional morality take at least three forms, one of which is special to institutions. There is the problem of collective action, especially for large numbers of actors, on which there is a vast literature.[3] There are stochastic problems of probabilistic effects that are common in policy issues that are often handled by institutions.[4] And there is the problem of relating institutional to individual action, which is the focus of the present discussion. It is especially important in applied moral theory and public policy debates. It often involves the other two compositional problems.

Morality by design

When an institution fails to accomplish its purpose, we may charge the failure to one of

two causes: (1) misdesign of the institution for that purpose; and (2) failure of particular people within the institution. I am concerned with institutional design and therefore with the first class of problems. The second class of problems is of interest for design only to the extent it should be taken into account in the institutional design — for example, through creation of incentives to control individual failure or design of selection or training procedures to reduce the incidence of such failure. (The one problem of individual failure to be considered here is potential conflict between the dictates of individual and institutional morality, especially when these are ostensibly guided by the same moral vision.) But this just means the issue is again one of design. Institutional morality is in this sense a design issue: Morality must be built in.

In the extreme case of misdesign, we may even suppose that a particular purpose cannot be served by any institution. If we take seriously the dictum that 'ought' implies 'can', we may wish to give up on that purpose. For example, Lon Fuller argued against the attempt to create adjudicative regulatory agencies to handle allocative decisions, such as setting fares and allocating routes for the airlines. Such agencies, Fuller thought, cannot be designed to do the job well enough to justify government interference.[5] In a similar move, John Rawls has recognized the possibility that we must give up on the achievement of equality of income and wealth in order to achieve mutual advantage. The reason for doubting income equality is that actual people may well be differentially motivated in part by the rewards they get for producing goods and benefits. Hence, unequal rewards may be necessary for general productivity.

In the extreme case of apparent irresponsibility within an institution we may not be able to assign specific responsibility to anyone. The standard problems of collective responsibility arise. If the institution's failure cannot be adequately regulated through the assignment of responsibilities, the institution may fall under the censure of misdesign rather than internal irresponsi-

bility.

If an institution's morality is defined by its purpose or likely effects, then its morality might be judged harshly from a particular moral perspective if its purpose is one that is judged immoral from that perspective. For example, from almost every western moral perspective the institutions of Nazi genocide had a morally reprehensible purpose. But the functional relation between the actions of various office holders in those institutions to the purposes of the institutions might have fit very closely. It might seem to be only a matter of definitional convention whether we say that the functionally fit actions were moral in those institutions or whether we say, by inference from the larger morality of the institutions' purpose, that they were immoral simpliciter. The choice, however, may not be trivial. We cannot say without confusion that the competence of a superb secretary is immoral simpliciter. Yet, the actions of such a secretary may be necessary for the success of an organization with a heinous purpose.

Much of the debate over the inherent morality of the law seems to have been a debate at cross purposes from these two quite different linguistic conventions. Fuller, assuming that law is to have some function or purpose, then argued that it has an internal morality. Much of his argument was muddled, but it makes relatively clear sense if it is seen as an argument from functional fit within the law rather than as an argument for specific moral content in the law. He blurred these two categories somewhat when he further assumed bits of human psychology to argue that certain purposes could not readily be achieved by law and were therefore generally not consistent with the morality of law, and when he assumed specific content for the external morality of law, which he thought must be to facilitate human interaction.

Institutions typically have the structure of law roughly as Fuller saw it. Again, this means merely that the rightness within an institution of a particular action by one of its office holders is

determined contingently by the functional fit of that action with the institution's purposes. This

need not allow an Eichmann defense against personal responsibility for one's actions within an

institution, because one's personal moral responsibility may be to oppose an institution's purpose

and the means it uses to achieve that purpose. The discussion here will focus on the internal

problem and will simply assume that the purposes of the institutions under discussion are exter-

nally justifiable.

Note that the internal morality of an institution is, as Hume called duties of justice, artifi-

cial. That is, it depends for at least part of its content on what humans have first created. In this

sense, the morality of a profession is artificial and what Fuller roughly called the internal morali-

ty of law is artificial. What it is right for a professional or an officer of the law to do depends on

how the profession's services or the law's institutions are organized. Hence, the argument for in-

stitutions more generally has its more specialized counterparts in professional ethics and the in-

ternal morality of law.[6]

Collective responsibility

No one would want to live in a world in which just anyone could have significant effect

on major policy issues. For example, if, during the recent superpower standoff, just any individu-

al American or Soviet citizen could have had major impact on whether nuclear weapons would

be put to use, all of us would have been at much more grievous risk even than what one might

suppose we did face from those weapons. There is even a logical problem in supposing that each

of us could have a major impact. What could it mean to suppose that you and I could each indi-

vidually have great impact on whether the weapons were to be used this week if it is not logically

precluded that my impact would be to cause their use and yours would be to prevent their use?

Yet, despite the obvious sense in which it is absurd to suppose many of us can have substantial impact on such policies, many of us were inclined to think we were somehow partly responsible for those weapons. If 'ought' implies 'can', as we generally believe it does, then our sense of responsibility in such cases is typically wrong.

In common usage and in much of our moral concern with it, responsibility is a causal notion. This is not the only sense of the word 'responsibility'. Dictionaries typically make causal notions the first sense of 'responsible'. One or something is responsible that is the primary cause, motive, or agent, the cause or explanation, for example, of an accident. Dictionaries then include a second sense, which generally applies only to persons or collective persons. One is responsible if one is able to answer for one's conduct and obligations, able to choose for oneself between right and wrong. Hence, by taking on ambivalent meaning, responsibility plays a role in both consequentialist and deontological moral theories.

In a utilitarian moral theory 'responsibility' is inherently a causal notion. To say, on a utilitarian account, that I am responsible for result x is to say that I caused or helped to cause x. It follows that a proper assignment of responsibility requires a relevant assessment of causal effects. But causal assessments are notoriously difficult in many contexts, perhaps most especially in social contexts. Difficulties in making causal claims may be epistemological or informational in large part but the more important difficulties are conceptual in the following sense. If my action is one of many contributions to a result, there may be no compelling way to resolve the full cause of that result into its partial causes. The collection of causes may violate Mill's principle of the composition of causes.[7]

It is often supposed that the possibility of doing experiments has been the major reason for the relative success of certain of the hard sciences as compared to the social and behavioral

sciences. A far more important consideration may be differences in the causal structures of the phenomena of interest in these sciences. For many physical laws, especially in classical physics, the principle of the composition of causes is valid or approximately so. For example, a classical electron may be accelerated by the combination of gravitational and electrical forces, each of which can be partialed out from the sum of the two forces. It is because the effects of gravitational and electromagnetic forces are straightforwardly separable that the Millikan oil-drop experiment for determining the charge of the electron was possible.[8]

For much, perhaps most, of what interests us in the social sciences and, hence, in moral theory, the principle of composition of causes is clearly not valid. There have been theories, such as Arthur Bentley's group theory of politics,[9] that presumed such simple relationships. But the once commonplace phrase, "the vector sum of forces," has rightly fallen out of vogue in the social sciences. We are now too accustomed to the importance and pervasiveness of threshold and other complex effects in social causation to suppose lightly that the result of a large number of actions by many people is somehow merely a vector addition of the results of the actions taken singly.

Suppose my institution is badly designed and it does harm (or it fails to do the good it could do) or that it is well designed and it does good. And suppose I fulfill my role in the institution as I am supposed to do. What is my responsibility for the harm or the good that my institution now does? This question is merely an instance of the common question of how causal weights should be assigned in cases of plural causation. To keep the general question clear, let us consider varied contexts in which it might arise, contexts that display strategic as well as substantive variety.

Some examples. In an election, many of us may vote. Some of us vote for the losing side.

Do we share responsibility for the result? The votes of many of those who vote for the winning side may be otiose — victory does not depend on them, given the other votes cast. Many of us may not vote. If the victor wins by, say, a million votes, does any particular nonvoter bear any responsibility for letting the loser down? Do a million and one nonvoters bear any responsibility? During the Watergate scandal that followed Richard Nixon's landslide reelection as President in 1972, there was a widely appreciated bumper sticker that said, "Don't blame me, I'm from Massachusetts." While every other state went for Nixon in his reelection, 54 percent of voters in Massachusetts voted against Nixon.[10] Was someone in Massachusetts who voted _for_ Nixon less responsible than someone in Illinois who voted with a state majority for Nixon? Less responsible than someone in Illinois who voted, forlornly, for George McGovern? One might have said, with Bentley, that the Nixon and McGovern causes were like two forces in opposition and that the greater force prevailed, as though politics were merely social physics in its classical variant. But the question of the prior composition of each of those forces is much messier.

Between crossing the Niemen River into great Russia in June 1812 and leaving Moscow in despair some months later, Napoleon's army declined from 422,000 to merely 100,000; by the time they recrossed the Niemen, they were down to 10,000.[11] These numbers trumpet defeat. We may suppose that many of the losses were from desertion; many were from disease and exposure in an unusually early and harsh winter and some were even from battlefield casualties. Suppose at Smolensk — with more than half the army already gone, winter coming on, and the main Russian army yet to be met — some soldier figured (rightly, one might suppose) that defeat was assured by the conditions and the continual losses. Hence, that soldier's continuation on the march could not bring about victory but might well contribute to the casualty lists. Arguably, that soldier's desertion was better for France than his continuing on the march. If he deserted, did he

then bear part of the causal responsibility for the eventual French defeat?

A manufacturing company may use several hundreds or thousands of workers to design, produce, and market its products. Most of them might individually be replaced without harm to the firm's output. Indeed, many of them might miss work without affecting the firm's output. Yet, if any great number of them failed to contribute, the firm would be less productive and might fail. What is the causal responsibility of any one of the workers for the firm's success (or failure)?

The problems of voters and of Napoleon's soldiers on the march to Moscow are variants of the more general class of collective action problems. In rational choice theory, the logic of collective action is about the conflict between a group's interest in providing x and the individual group member's incentives or interests. The concern is with incentives and their misfit, rather than with responsibility as in the voting example above. But all of these are instances of the far more general class of problems of strategic interaction in the production of outcomes. In such interaction, it makes little or no sense to say that you or I choose an outcome. We have interests in the potential outcomes. But what we choose is merely a strategy for getting to these. Hence, the commonsense moral language that links responsibility and choice or action is partly out of place for choices of strategies. At the very least, it is complicated by the plural, complex causation of outcomes.

Collective action and organizational problems raise the same issues for individual moral responsibility. In either context, my particular action or choice may not be the cause of our outcome and therefore I am not responsible for that outcome in the simple commonsense meaning of the term. Indeed, it is a fallacy of composition to suppose the notion of responsibility generalizes from simple choice to interactive choice situations. For such choice problems, deontological

action theories are in need of further articulation. It seems likely that relevant articulation must rely heavily on consequentialist considerations. Hence, strategic interaction captures the problem of institutional design: Its moral assessment is inherently consequentialist.[12]

Practical resolutions. How do we resolve assignments of responsibility in practice? This is a descriptive question about what we do, not what we ought to do. But an understanding of actual practices may inform our normative arguments, even if only by suggesting a range of possibilities. It is to this descriptive issue, as it arises in the law, that H. L. A. Hart and Tony Honoré address their concern in Causation in the Law.[13] It is also to a descriptive question that contemporary economic analysis of the value of inputs in production is addressed.

In the law, a standard device is to focus on what was anomalous in the causal background of a result. Law commonly confounds actions and their effects as, for example, in the definition of murder as death brought about in certain ways with relevant intentions. Simple theories of direct causation may only poorly fit the actions that, taken together, bring about a particular death. I will take up more extensively the determinate nature of legal rules that may be arbitrary.

Business firms may commonly assess causal relations only at the design stage and then ignore the problem of individual-level causal responsibility, at least for most individuals, for what the firm creates or the profits it makes. For the individual-level they focus instead on market allocations of prices for various causal inputs. They do not pay workers according to what the workers contribute causally but according to the market wages for such work. Those wages can change dramatically for demographic reasons of supply of workers having nothing to do with economic valuations of the firms' products.

In many contexts, proxy measures of inputs are used. For example, A did not act in certain ways, and is therefore held responsible for not helping to prevent some harm. Focusing on

such proxy measures, particularly on kinds of actions, threatens to make our focus not causal but deontological, especially when we begin to focus on symbolic rather than efficacious action. But there may be good epistemological reasons for using proxy measures, as, for example, the test of alcohol concentration in the blood for a driver or an airline pilot. Stochastically, we may know there is a strong correlation between alcohol levels and incidence of accidents. Although we will not typically know that your high alcohol level means you will have an accident, we can still be sure that going after all drivers with high alcohol levels will reduce the rate of accidents.

In a variant of proxy measures of inputs, we may assign responsibility to someone — for example, in an organization — in order to create incentives for relevant action. We know that a lot of certain kinds of activity will produce the results we want and we wish to motivate such activity. We cannot assign shares of the responsibility for the final result to each activity or person, but we can enhance the prospect of achieving that result.

The striking fact is that all these resolutions of our problem are irrelevant for standard individual-level moral theory except in a contingent way. My following the incentives offered by my organization or reprimanding you for violating some proxy measure of output cannot be justified merely by the kinds of actions these are but must be justified with reference to consequences. To handle all of our individual actions within institutions in standard consequentialist moral theory we would be well served by a genuine capacity for the composition of causes that would allow us to assign shares of responsibility for various outcomes. Proxy, pragmatic, and incentive devices are beside the point for such direct moral responsibility. And we surely have no ready moral analog of the economist's sly substitution of supply-and-demand price for causal contribution.

The legal model of responsibility

Law is fundamentally unlike morality in the structure of its rules. In law we may conventionally select one of many possible rules for settling some class of issues. The rule need not be an a priori logical inference — it can be arbitrary. Often its main burden is to be determinate. All of this is possible in law because there are authoritative bodies to enforce application of legal rules, no matter that they may be arbitrary. Unless we turn moral restrictions into legal requirements, or we concede determination of morality to a religious authority, or we have a strong communal sense of unquestioned morals, there is no authoritative body to enforce an arbitrary moral rule. Note that legal, religious, and communitarian theories are typically command theories of the right, at least in part. The state, god, or the community commands rules, therefore it is right to follow those rules. If our morality is not under the duress of some authority's command, we will want our moral rules to be persuasively well-grounded and not arbitrary.

Consider an arbitrary rule for assigning responsibility for harms that we might adopt in the law. If a corporation does harm, we might hold the corporation responsible for compensating the harmed and then let the organizations work our how the burden is to be disaggregated down to the individual members of the corporation. This device might be fine as a matter of determinate law. But as a matter of morality, this device is afflicted with two problems. First, there is moral hazard — the Board and the chief executive officer, who work out the disaggregation, have a strong interest in the outcome.

Second, the device is afflicted with the problem of extracting resources from people who cannot be sensibly seen as morally responsible for the relevant harm. For example, a firm might do something whose harmful effect becomes legally clear only decades later, when virtually all individuals who might be considered causally responsible for the harm are gone. Johns-Manville

was eventually held accountable for the harms that were inflicted on those who worked with its asbestos, as in the U. S. naval yards of World War II. Its liabilities bankrupted the firm, whose stock holders lost their investment. Suppose a small company's or church's retirement fund invested substantially in the stock shortly before news broke of its impending liability and the fund lost much of its retirees' reserves. That the system should work this way might be good for general productive efficiency or determinateness; of this way might be better than any alternative we think of because the alternatives require knowledge that is unavailable or too costly. The result is then right, but for large systemic reasons, not for reasons of the causal responsibility of the retirees for the asbestosis and mesothelioma caused by Johns-Manville's products.

Also consider a relatively abstract rule: Corporate responsibility for corporate harm should be disaggregated into personal responsibilities whose sum is equal to the corporate responsibility. Unfortunately, causal relations do not sum in this way. If we could disaggregate causally, the collective responsibility might be greater than, less than, or equal to individual responsibilities. Arbitrarily stipulating that they are equal may be a good legal rule. For one thing, it avoids the horrendous analytical problem of determining the full causal story and of parceling causal responsibility to individuals. And this may be a good moral reason for adopting the rule in the law. But, again, this means we have a systemic moral justification for setting corporate and individual responsibilities equal, not a reason of direct causal responsibility.

Institutional purposes

How should we define the purpose or purposes of an institution? We could define the purpose narrowly or broadly. But the critical question is, as it is for organizational responsibility, one of composition. Who defines the purpose of institution X? Its current leaders? Its founders?

Its entire membership, somehow determined? Can we discover an organization's purpose from an external vantage point with or without testimony from its members? And if we have such testimony, does it trump other claims we might make?

These are inherently difficult conceptual questions to which there are no easy answers. They are conceptually difficult the way the question of partial individual responsibility for a group or institutional action or result is difficult. There may be no way to braid the various contributions to an institution's purpose into a single, coherent strand. It is probably in the law that the problem of multiple determination of purpose has been most extensively addressed. Since it would be odd to speak of the intentions of the collection of people who were the relevant legislators, judges, and citizens, we can speak of the "meaning" of a law. The meaning of a law is how the law comes to be interpreted and understood and how it is filled out with tacit assumptions. Clearly, the focus of meaning is on the result of a long process, not merely the origin in, say, the many minds of a legislature. Such origin might be the more typical focus of intention.

In assessing purposes of institutions we might do no better than to think of the organizational analog of legal meaning. We thereby escape the compositional problem of determining the purpose as an output of some complex process with many causal contributions. The result is a naturalistic, descriptive account of purpose rather than a developmental, intentional account. But if the account is descriptively accurate, it is hard to see what would be gained from a causal, intentional account. The latter might be of interest if our problem were to assess the rightness of actions of current organizational members by reference to what was intended for them. For example, some theorists of constitutional interpretation are concerned with what they call original intent. One might justify this concern by noting that constitutional constraints tend to lead to mutually beneficial stable expectations.[14] Such a justification would be far less compelling for many

other institutions and organizations.

For many institutions and organizations we may have no difficulty in determining central purposes. Among the purposes of universities are education and the support of research. The purposes of police forces include the maintenance of order, of business firms the making of profits, of certain professional organizations the service of particular clienteles. As we expand the list of purposes for any of these organizations, we are apt to find less agreement. Moreover, we may often be convinced that some ostensibly other purpose is derivable from the central purposes. We may also suppose that some other purposes are themselves imposed from some larger context. For example, business firms might be constrained to further the larger social purpose of eliminating racial or sexual discrimination even if doing so might interfere with profitability.

Organizations often achieve their apparent purposes. Yet, it does not follow that we can assign individual causal responsibility for contributions to their collective products. What we might be able to do, instead, in some cases is determine what form an organization must take to accomplish its purpose and then hold officials responsible for fulfilling the specified duties of their offices. We would have only a relatively rough causal account of how the organization would achieve particular purposes. We would be unable in the cases of many office holders in the organization to show the causal relation between their particular actions and achieving the organization's purposes. The causal relations might often take the form of John Mackie's INUS conditions. One contributor to the grand result might be insufficient but necessary to a part of the organization that was unnecessary but sufficient to accomplish some task that helps achieve a purpose of the organization. Redundancy would be built into the system. We might not be able to factor any contributor out of the organization and find a resultant loss; yet we know that if we factor out enough contributors, the organization will fail.[15]

In the law, there is an additional consideration that the law be fairly applied. This cannot be achieved through ad hoc devices of particular legal office holders who override standard procedures. Could an organization which is not bound to the same degree by considerations of fairness more readily act from ad hoc assessments of the effectiveness of particular office holders? Plausibly. For example, we might think the correct general purpose for a company in the market is to secure profits by securing sales. If an office holder in a particular company happened to be an irritant to other office holders, a supervisor might reasonably conclude that higher productivity would follow from dismissing the irritating colleague. Unless there were a larger social purpose that governed the company's personnel relations and that trumped the supervisor's discretion, it could be moral within the institution to allow such discretion. There might be a company rule not to allow such discretion to supervisors — perhaps on the belief that it could not be reliably put to use for the benefit of the company's goals. Then it would be immoral within the institution for the supervisor to act with such discretion. The supervisor might successfully appeal to higher authorities in the organization for authorization to act with such discretion, and then it would be moral within the organization. But to do so merely on her own belief in the fit of her action with the fundamental profit-making purpose of her firm would be immoral within the firm.

The limits of moral design

Consider a near analog of institutional responsibility at the level of individual responsibility: stochastic problems of large-number or frequent interactions involving low levels of risks of large harms.[16] The difficulties of handling such problems in standard ethical discourse have led philosophers to treat them as matters of "moral luck."[17] For example, I may be one of thousands

of drunk drivers on the road tonight, but I am unlucky enough to lose control of my car just where there happens, oddly, to be a pedestrian, whom I kill. Many others merely weaved out of lane, jumped a curb, or spun out of control and then recovered to drive on without harm to anyone. If my pedestrian had not been there, I might have been harmless. It is in part, therefore, merely a matter of luck that I am held blameworthy while these other drivers are not.

An institution, such as the law, need not be bothered by the moral theorist's problem of moral luck. It can address such matters as drunk driving stochastically, counting the act of driving drunk as a crime. And it can choose to handle harms that occur according to an efficient rule, such as requiring drivers to have liability insurance to protect those with whom they have accidents, or according to a rule that sets up incentives to reduce the incidence of behavior that correlates with harm. Similarly, an organization that has rules for handling flammable or otherwise dangerous chemicals need not, when I violate those rules, defend the rules by claiming that I would have caused harm; it need only claim that, because following its rules has stochastic benefits, enforcing the rules has general benefits.

Such stochastic problems as drunk driving are morally serious not merely because of the actions involved but because of the plausible and likely consequences of the actions. However, any account of responsibility for such problems must focus not on actual results of action but on ex ante expectations of the results. For example, a drunk driver on a crowded superhighway in a city is culpable of a transgression, of an ex ante likelihood of harm to someone else. The same drunk driver on a secluded and abandoned military roadway in the desert or on a private racetrack may be much less culpable of such a transgression because there is so little ex ante chance of that driver's harming someone else. When we assign responsibility on ex ante suppositions, however, we may generally do so according to some proxy for the kinds of results we think like-

ly to follow. You may be a better and safer driver while legally drunk than I am while legally sober. Or you may be much more careful in handling a jug of explosive chemical than I am even though you violate our organization's rules and I do not. For various informational and epistemological reasons, however, institutions may not be able to take the possibility of such considerations into account in a very fine way and we may satisfice on relevant calculations by simply resorting to a legal definition of drunkenness and culpability for drunk driving or a clear rule on safe procedures.

Perhaps an institutional authority could properly ask of the institution's office holders, "What if everyone did that?" In private life, I might be able to answer this rule-utilitarian query that my action did not disrupt the more general achievement while it brought benefits of its own. I might plausibly be able to show that perfect compliance with some ostensible rule would be suboptimal, that occasional violation of the rule would produce a better outcome.[18] One might suppose that if I have to answer an institutional query why I violated one of its rules, the very fact that my violation has been noticed in the institution suggests, prima facie, that it may have effects on the behavior of others in the institution through the example it presents. The institution's mobilization behind a particular rule of behavior, thereby standardizing it, might actually make it more likely to be optimal to adhere perfectly to the rule. But even if such an argument tended to be valid, so that there were a seeming difference between institutional and individual strictures, this difference would only entail a greater likelihood that the individual violation of institutional procedures or rules is genuinely counter to the purposes of the institution. On some occasions, the violation might still be justified even by the purposes of the institution.

Suppose my office requires certain actions from me to enable my institution to achieve its purpose. The requirement is causally deduced. Hence, there may be disagreement over how well

the requirement fits the institution's purpose. Suppose I share my institution's purpose and genuinely wish to further it. But I conclude that what I am supposed to do, under the institution's rules, hinders rather than furthers that purpose. Have I done wrong if I therefore violate the institution's rules to help the institution? Within the institution I may have done wrong. If, however, my action is externally judged, just as the purpose of the institution itself must be externally judged, it may seem right. From that perspective I can conclude I have done right. Yet, my institution may rightly judge that I have done wrong. The difference turns merely on the differential assessments of causal facts. There may be understandable differences in epistemological capacities for organizations and individuals that lead them to reach different conclusions.

For an organization, unfortunately, we can commonly expect incoherency in the internal requirements on various office holders because there will be no fully worked out causal account of how the organization's structure and procedures achieve its purposes. And we can probably expect many of the requirements to be little more than conventional in the sense that no one could give a causal account of how the requirements contribute to the organization's purposes. Hence, the claim that the official requirements of an institution constitute the internal morality of the organization might be on shaky ground, at least for some requirements, although perhaps not for others.

But note that the assertion that a requirement is merely conventionally, and not causally, justified can be misleading. As in the argument above for institutions in general and for similar epistemological reasons, we can conclude that, in my life as a citizen under the law, it can be right for me to violate the law and nevertheless right for officials of the law to punish me for my violation.[19] There is a perhaps urgent difference between the law and other kinds of institutions in this respect. Among the purposes of the law may be to achieve fairness and to reach finality.

Hence, individual divergences by officers of the law may actually be causally harmful to its purposes. It might matter less which of various resolutions is reached than that the same resolution is reached for all like cases. Hence, although no one resolution may be clearly best, consistent application of any one of several might be better than inconsistent application of each of these in some fraction of cases. Insofar as similar concerns enter in the design of other institutions, the fact that a particular requirement is merely conventionally defined need not block its moral force within the institution. But there must be many instances in varied organizations in which there is not even this reason of precedent and fairness to moralize conventionally adopted rules for behavior.

Individual vs. institutional morality

A common response to institutions and their moral problems is to suppose that morality is only for individuals and that institutions are a problem, not a solution. In some moral traditions, this seems to be a relatively natural inference, although it is unlikely to be a hard logical inference. In other traditions, the move from individual-level to institutional-level theories seems to be incumbent on the theories. For example, utilitarianism was institutional and personal from the beginning. Kantian theory has clear connections to institutional concerns already in Kant's own jurisprudence. And on a plausible reading, Aristotelian virtue theory requires functionalist derivation of personal virtues from the requisites for institutional roles. But much of what may be called commonsense ethics is virtually blind to institutions.

The novelist and philosopher Iris Murdoch judges people by their character, to which their actions give clues. She writes, "I daresay human wickedness is sometimes the product of a sort of conscious leeringly evil intent.... But more usually it is the product of a semi-deliberate

inattention, a sort of swooning relationship to time."[20] Moreover, "The wicked regard time as

discontinuous, the wicked dull their sense of natural causality. The good feel being as a total

dense mesh of tiny interconnections."[21] Apparently, Murdoch supposes much of the world shares

her views, at least instinctively, and she is may be right. For example, Bradley Pearson, the pro-

tagonist of Black Prince, concludes in the end that he is "condemned for being a certain awful

kind of person" rather than for a particular action.[22]

One might suppose Murdoch merely asserts a definition when she associates wickedness

with regarding time as discontinuous (which implies having a weak sense of personal identifica-

tion) and having a dull sense of natural causality. But that would not be an interesting move. The

extraordinary feature of Murdoch's view is that it makes morality a causal correlate of epistemol-

ogy and sense of personal identification. Many of us dull "our sense of natural causality" because

we do not have or believe in this view of the causal structure of the world, in which one's actions

have explicit, determinate effects. Maurice Merleau-Ponty supposed this view was an English

perversion of good sense.[23] The view of personal identity as "a total dense mesh of tiny intercon-

nections" is also somewhat insular and archaic, although it may be widely held. In speaking of

"being," Murdoch presumably means to imply not only identity but also identification with, or

commitment to, one's self.[24] Both Murdoch's epistemology and her views of personal identifica-

tion are reasonably up for debate independently of moral theory. It is plausible that one would

conclude that, epistemologically and as a matter of personal identification, one should be as Mur-

doch describes the wicked person. How odd — we are plausibly wicked out of good sense.

We require large organizations for the achievement of many of our purposes, including

moral purposes, including moral purposes that Murdoch would likely share. But such organiza-

tions inherently share the burden of the epistemology and identification of Murdoch's wicked.

We are therefore damned to wickedness if we strive to achieve certain moral purposes. Murdoch's vision is consistent with Hannah Arendt's view of the amorality of institutions and the banality of evil. Large institutions break the sense of personal causal connection and they commonly produce a semi-deliberate inattention. These tendencies are, indeed, part of their great value. Institutions can bring things about despite the weakness of personal causal efficacy. And they achieve efficiency by dealing with problems in routinized aggregates rather than through deliberate, individualized attention.

As any reader of Murdoch knows, many of her central characters, if judged by what they do, are rotten — often merely because they are weak and self-obsessed and, in their self-obsession, they have what Murdoch calls "a swooning relationship" to time and much else. One of the greatest values of institutions is to block many kinds of rotten action. They may or may not succeed in elevating character, but they can often do well with the material they have. Given that institutions are necessary for our lives, we should wish to understand morality by design and its limits.

Concluding remarks

In assessing the morality of an institution, our chief difficulty is compositional. We can, in our ordinary vocabulary, understand what we mean when we say that an individual is responsible for a certain action or that the individual acts morally or immorally. To say of a group or institution that it acts morally or immorally or is responsible is not completely contrary to sense, but it would be odd to think of an organization as an intentional being in the sense in which a person is an intentional being. The organization may be composed of intentional beings, but it is not one itself. Not least of the oddities of treating an organization as an intentional being would

be in assigning moral responsibility and even punishment to the organization without having
these be reductively applied to the individuals in the organization.

Despite the difficulty of assigning moral responsibility to an institution or organization
without assigning it to individuals in the organization, organizations have capacities that tran-
scend those of individuals. For example, organizations can have superior capacities to collect rel-
evant information and theory for handling their tasks and radically superior capacities to mobi-
lize resources to get things done. Organizational policies and actions may also often be much
more coherent and consistent over time than individual "policies" and actions. If responsibility is
typically associated with capacity, it seemingly makes sense to assign responsibility for out-
comes to organizations that bring them about. But moral theories are not sensibly brought to bear
on organizations and other institutions without reassessment of notions of responsibility.

In general, the internal morality of acting within an institution is not simply composed
from the ordinary morality of individuals. It is artificially constructed to a large extent. The
moral theory of institutions and of the behavior of institutional office holders must be derived
from the purposes of the institutions. Complex, plural causation and the stochastic nature of
many outcomes block the simple association of individual-level moral responsibility with causal
responsibility. Hence, an uninstitutionalized act-utilitarian account or a deontological account
based on individual-level action theory is largely irrelevant to institutions. We might conclude
that institutions are therefore outside the realm of moral discourse. But that would be perverse.
Without institutions we can achieve far too few of the moral purposes we have.

Notes

* Original version presented at the conference, "Institutional Design," Research School of the Social Sciences, Australian National University, Canberra, 16-17 July 1992. I thank the organizers of the conference for stimulating this paper, John Braithwaite for extensive written commentary, Michael Stocker for extensive oral commentary, and participants in the Russell Sage Foundation seminar, New York, 21 October 1992, a colloquium at the Department of Politics, New York University, February 1993, and the Legal Theory workshop at Columbia University Law School. And I thank the Russell Sage and Mellon Foundations for general support.

1. The political institutions in John Rawls's theory are clearly intended to achieve consequences: "to maximize human achievement in art, science, and culture" (John Rawls, A Theory of Justice [Cambridge, Mass.: 1971], p. 325). Also see Thomas M. Scanlon, "Liberty, Contract, and Contribution," in Gerald Dworkin, Gordon Bermant, and Peter G. Brown, eds., Markets and Morals (Washington: Hemisphere; and New York: Wiley, 1977), pp. 43-67, esp. pp. 61-65.

2. Rawls, "Two Concepts of Rules," Philosophical Review 64 (1955): 3-32; Russell Hardin, Morality within the Limits of Reason (Chicago: University of Chicago Press, 1988), pp. 100-105.

3. Mancur Olson, Jr., The Logic of Collective Action (Cambridge, Mass.: Harvard University Press, 1965); Russell Hardin, Collective Action (Baltimore, Md.: Johns Hopkins University Press for Resources for the Future, 1982).

4. Russell Hardin, "Ethics and Stochastic Processes," Social Philosophy and Policy 7 (Autumn 1989): 69-80.

5. Lon L. Fuller, The Morality of Law (New Haven, Conn.: Yale University Press, 1969

revised edit. [1964]), pp. 170-177.

6.   Russell Hardin, "The Artificial Duties of Contemporary Professionals," Social Service Review 64 (December 1991): 528-541; and "My University's Yacht: Morality and the Rule of Law," in Randy E. Barnett and Ian Shapiro, eds., NOMOS 26, The Rule of Law, New York: New York University Press, forthcoming.

7.   John Stuart Mill, A System of Logic, Ratiocinative and Inductive (any standard edition), book 3, chap. 6.

8.   Mill makes this point with reference to differences between chemistry and physics, A System of Logic, book 3, chap. 6, §1.

9.   For example, see Arthur Bentley, The Process of Government (Evanston, Ill.: Principia Press, 1949).

10.   The District of Columbia voted 78 percent for McGovern. No state other than Massachusetts topped 47 percent.

11.   The French engineer, Charles Joseph Minard, portrays this debacle in a masterpiece of graphic representation. See Edward R. Tufte, The Visual Display of Quantitative Information (Cheshire, Conn.: Graphics Press, 1983), pp. 40-41, 176.

12.   See further, Hardin, Morality within the Limits of Reason, pp. 68-70.

13.   H. L. A. Hart and Tony Honoré, Causation in the Law (Oxford: Oxford University Press, 1985 second edit. [1959]), especially chapters 3 and 4.

14.   Original-intent theorists also have deontological and purely intuitionist claims for the rightness of having today's citizens be governed according to the intentions of men from centuries earlier. I think none of these claims is compelling, only a consequentialist claim that addresses effects on present and future people can be compelling. One who finds a nonconsequentialist

claim compelling might have difficulties with the question, "Why obey the law?" If that, too, is a matter of intuition, how do we decide empirically that this is the law?

15.   Iris Murdoch provides a reverse analog in Black Prince (Chatto and Windus, 1973; Penguin, 1975❖), p. 326. Bradley Pearson is secretly at a vacation house where he has eloped with his beloved. He receives a telegram from Marloe asking him to telephone to learn of his sister's suicide. Such an event might seem to call for breaking off his elopement, although it was not yet consummated and might be fragile. Marloe should not have known how to find him, however, and had figured it out from a letter from a real estate agent in Pearson's desk. "And if that terrible telephone call," Pearson reasons as he determines to stay the night, "had been so little determined, so casually caused, it made it seem that much less real, that much easier to obliterate from history."

16.   For more extensive discussion, see Hardin, "Ethics and Stochastic Processes."

17.   Bernard Williams, "Moral Luck," pp. 20-39 in Williams, Moral Luck (Cambridge: Cambridge University Press, 1981); Thomas Nagel, "Moral Luck," pp. 24-38 in Nagel, Mortal Questions (Cambridge: Cambridge University Press, 1979).

18.   Hardin, Morality within the Limits of Reason, pp. 65-68.

19.   Hardin, "My University's Yacht."

20.   Murdoch, Black Prince, p. 189.

21.   ibid., p. 125.

22.   ibid., p. 387.

23.   Maurice Merleau-Ponty❖

24.   Identity over a lifetime is a perverse notion no matter how widely held it may be. See further, Hardin, Morality within the Limits of Reason, pp. 191-198.