

INSTITUTIONAL KNOWLEDGE AND THE ROLE OF INSTITUTIONS

Russell Hardin
New York University

ABSTRACT

As a result of its special epistemological capacities, an institution can make some kinds of decisions far better than typical individuals can. Hence, there are things that institutions can do that individuals acting spontaneously could not do. Ex ante, we must want institutions to help us accomplish some purposes. This fact entails strong conclusions about the use of interpersonal comparisons of welfare.

CONTENTS

Melioration and Mutual Advantage	2
Strategic Interaction and Institutions	4
Institutional Epistemology.....	6
Institutions As Meliorative	9
Apparent Mutual Advantage	13
Conclusion.....	15

INSTITUTIONAL KNOWLEDGE AND THE ROLE OF INSTITUTIONS

Russell Hardin

New York University and the Center for Advanced Study in the Behavioral Sciences

The so-called Enormous Theorem gives an exhaustive classification of all finite simple groups. The full proof is in about 500 separate publications by more than 100 mathematicians. It constitutes 15,000 pages in various publications plus some unpublished results.¹ Some of its pieces were directly worked out by particular mathematicians who could be said to have an overview of at least those pieces; others were accomplished by computer and some of these are plausibly so complex or massive that no particular mathematician can claim to have an overview of them. Naturally, no single mathematician is thought to have a confident grasp of the whole theorem — all must rely on the reputed expertise of the individual contributors to the theorem, including computer programmers and chip designers.

The knowledge that institutions and big organizations have is partially analogous to this proof. No one has a full overview of the knowledge of such an institution or organization, and yet it seems meaningful in some perhaps odd sense to say that the institution or organization has a vast body of knowledge at its useful disposal.² Because many utilitarian purposes cannot be achieved without the use of large institutions, it follows that, not only is the achievement of such purposes not within the purview of a single person, it is not even subject to the overview of a single person.

This conclusion raises both epistemological and normative issues. If we cannot even

¹ For a brief discussion of the theorem and of current efforts to simplify it to perhaps a more manageable 3000 pages, see Scientific American (April 1996), p. 29.

² One might make an analogous claim of the nature of individual knowledge. It is compartmentalized in the brain and cannot be overseen by any part of the brain.

have an overview of an institution's knowledge and how it puts it to use, it would be absurd to say we know that the institution maximizes the good that it could do with respect to its mandate. Indeed, it may be impossible even to know whether an institution performs as well as an alternative might. Hence, the criterion of maximization is not relevant to utilitarian judgment of typical institutions. Nevertheless, institutions must typically act in ways that result in tradeoffs between individuals. Creating an institution or even adopting a particular policy may benefit all ex ante, but it will commonly lead to net costs for some ex post. A welfarist argument for institutions seems compelling, but it has strong implications for the necessity of making such tradeoffs.

Melioration and Mutual Advantage

What is the relevant criterion for judging the welfare effects of an institution against some alternative? There is one that may work in certain cases and another that might be the best compromise criterion in most cases. The sometimes workable criterion is mutual advantage and the compromise criterion is melioration. Mutual advantage is all that is left of standard maximizing utility theory if we cannot make interpersonal comparisons of welfare.³ But it is arguably also a reasonable criterion for judging welfare effects that are not cardinally measurable in a population whose members are all similarly situated.

Melioration is, of course, the central impetus of pragmatism in moral judgment. John Dewey's argument for melioration is that maximization requires an account of what is best, and he supposed we cannot be certain that we know what is best, which must, as with all empirical inquiries, remain an open question.⁴ One could argue that these are also the relevant criteria for judging the utilitarian value of individual actions, but that is not my purpose here, which is to

³ See Russell Hardin, Morality within the Limits of Reason (Chicago: University of Chicago Press, 1988), chapter 3.

⁴ John Dewey, The Quest for Certainty (New York: Minton, Balch, 1929), chap. 9.

focus on institutions and their particular problems. Those problems are largely a function of the peculiar capacities and incapacities of institutions.

Melioration is not well defined in the work of pragmatists, but I will suppose it includes improvements that entail costs to some in order to benefit others in contexts in which we may suppose we can make at least rough interpersonal comparisons. For example, the change from central command economies in the former Soviet world to more nearly open market economies surely harms some while benefiting others. But one might suppose that continuing the Soviet-style central command economy in Russia indefinitely would imply continuing economic malaise that would reduce the life chances of most Russians of this generation and, considered *ex ante*, of all Russians in future generations. And one might conclude that the long-run losses from continuing the former system grossly outweigh the short-term benefits to some of continuing it.

An easier case is that of shifting from an agricultural to a broader economy. If eighty or ninety percent of the population are engaged in mostly subsistence agriculture, most of these people must live at very low levels of welfare. Although there would be losses to some of the present generation of farmers if a workforce is shifted from agriculture to more productive occupations, one might suppose the overall gains greatly outweigh the overall losses, even though the gainers and the losers would largely be different individuals. Making the shift would meliorate the conditions of life for people generally despite inflicting costs on some in the short run. That is to say, the overall long-run improvements would outweigh the interim losses.

Clearly, melioration defined in this way is distinct from mutual advantage. The two criteria might coincide in many instances, such as when we are all similarly situated and we all benefit from some change. And they might conflict in some situations, such as when we could make interpersonal comparisons to conclude that a possible mutual advantage move would be inferior to a meliorative move that would benefit some at cost to others. Mutual advantage is an

inherently incomplete criterion in the sense that there will be alternative states of affairs neither of which is better than the other by that criterion. Melioration will also be incomplete in contexts in which we cannot assess interpersonal comparisons of gains and losses accurately enough to claim that changing from one state of affairs to another is meliorating.

In general, we should want the same criterion to apply to all our choices, whether of personal actions, of institutions, or of policy choices by institutions. Choices of institutions, as in creating a constitution and a new form of government, are relatively unusual. Most of the time we merely rely on extant institutions and worry about their choices of actions or policies. Yet choices of institutions might commonly be mutually advantageous ex ante while choices of policies by institutions commonly will not be mutually advantageous even ex ante because losers and gainers will be identifiable ex ante.

Strategic Interaction and Institutions

An apparent disanalogy between institutional knowledge and the proof of the Enormous Theorem is that the latter can be broken down piecemeal and each piece can plausibly be worked on independently of the other pieces. An attempt to alter or redo a small piece of some institution, however, is often likely to ramify through the institution and affect other aspects of it. The parts of an institutional organization do not simply add their effects to those of the other parts. Rather, they interact. This means that we often cannot improve an institution except by trial and error, comparing each variant institution, as a whole, to the others. But this will commonly be impossible.

If the effects of the parts of an institution cannot be decoupled from the effects of other parts, institutional design is inherently strategic. In this respect, institutional structure is similar to human interaction more generally: One of my actions typically cannot be grounded in a simple

assessment of its direct consequences, but must rather be based on the expected result of what not only I but others do, including what others might be expected to do in reaction to my action. Hence, rule-utilitarianism is as incoherent for institutions as for individuals. As David Lyons has argued, rule-utilitarianism for individuals makes sense only if the actions mandated by the rules have clear consequential effects. But this means that the rules must be sufficiently articulate and varied as to allow them to produce the best consequences when applied to contexts that might be similar in broad respects but quite dissimilar in important ways. Furthermore, it means that the set of all rule-utilitarian rules must be ridiculously numerous, perhaps approaching the order of magnitude of the number of cases to which they are to be applied. Hence, it might commonly be more sensible to act from an assessment of consequences rather than from stylized rules.⁵

The fundamental principle of rule utilitarianism is that we somehow discover rules that generally entail particular consequences. The incoherence of even describing actions in terms of their consequences is instantly conspicuous in game theoretic representations of strategic interaction. In game theory, an action is a strategy choice, and such a choice leads to various possible consequences depending on what strategy choices others make.⁶ Both the best and the worst of an available range of results can follow from a single one of an individual's strategy choices. This is a problem that faces institutions both internally and externally as much as it does individuals externally. For example, a policy toward some group is likely to provoke external responses of changed behavior that then affect the value of the policy. More generally, policies are likely to have unintended consequences — or perhaps this is merely a less direct way of saying that policies, as is true of actions generally, commonly provoke responses.

The far more complex and arguably more interesting problem for institutions, however,

⁵ David Lyons, Forms and Limits of Utilitarianism (Oxford: Oxford University Press, 1965).

⁶ Hardin, Morality within the Limits of Reason, pp. 68-70.

is the consequences of internal strategic interaction. Because the relations between parts of a complex institution or organization cannot be fully grasped by anyone, it is implausible to say that such an institution is designed in the straightforward sense in which one might say that a person designed a particular artifact. Rather, institutions are inherently, as Adam Ferguson noted, of human making but not of human design.⁷ To a large extent, institutions merely happen.

Yet, it is still true that we actively want to have certain kinds of institutions to make certain decisions and to implement them. Principal reasons for wanting them are that they have capacities individuals acting spontaneously lack. They can resolve problems that would pose insuperable obstacles of collective action for individuals. They can have the power to bring about extraordinary results and even to override what many individuals might attempt to do to block such results. For the present discussion, however, the most interesting capacity they have is epistemological. Their epistemology is, as are their organization and their power, enabling.

But their epistemology is also disabling in that it makes them incapable of addressing finer, variant or atypical effects on welfare. This is an important fact, but it will turn out to be, oddly, a positive fact in the account that follows. The virtual impossibility of having large-scale social institutions take interpersonal comparisons of welfare into account in many contexts will yield an argument in favor of taking such differences seriously.

Institutional Epistemology

The epistemological requirements of utilitarianism are different in kind from those of certain other moral theories. A rigorous Kantian theory that requires action according to specific rules or norms that declare some kinds of action wrong and other kinds right or even required

⁷ Adam Ferguson, An Essay on the History of Civil Society (New Brunswick, NJ: Transaction, 1980 [1767]), p. 122.

depends on an epistemology or a metaphysics that allows derivation from abstract principles of reason independently of specific knowledge of the practical world. Intuitionism and those parts of any moral theory that depend on direct, substantive intuitions that certain kinds of action are right or wrong, evidently require an unusual epistemology that no one has yet spelled out with care — although some philosophers have been morally certain that they had command of a relevant epistemology.⁸ Utilitarianism depends on an epistemology that is, in principle, similar to the epistemology for general practical knowledge,⁹ as does any other morality that is about consequences.¹⁰

For ordinary persons making utilitarian choices, however, philosophical epistemology, most of which has been argued for the context of the natural sciences and for the judgment of scientific knowledge, misdescribes judgments of what counts as knowledge. The street-level epistemology of ordinary persons can typically be described as an economic theory of knowledge. We come to know what we know because (1) it has value to us, (2) it was part of a larger plan of education for general or specific purposes, (3) it did have value to us in some past context, or (4) it came to us as a by-product of other activities. In the first case, we may seek the knowledge in order to improve the quality of a present choice. In the second case, we may have sought it in order to handle the kinds of decisions we now often face. In the latter three cases, the knowledge is happenstance knowledge in a moment in which we bring it to bear.¹¹ Whether we make good

⁸ H. A. Prichard supposed those who are properly brought up need only let their “moral capacities of thinking do their work” (Moral Obligation and Duty and Interest [London: Oxford University Press, 1968 (1912)], p. 17).

⁹ Russell Hardin, “Commonsense at the Foundations,” in Bart Schultz, ed., Essays on Henry Sidgwick (Cambridge, Cambridge University Press, 1992), pp. 143-160.

¹⁰ As Dewey argued in The Quest for Certainty, p. 274.

¹¹ Russell Hardin, “The Economics of Knowledge and Utilitarian Morality,” in Brad Hooker, ed., Rationality, Rules, and Utility: Essays on Richard Brandt’s Moral Philosophy (Boulder, Colo.: Westview Press), pp. 127-147.

choices, either morally or self-interestedly, turns heavily on whether the happenstance knowledge available to us while we are choosing is relevant and good.

But an individual can have some kinds and bits of knowledge that an organization cannot sensibly be designed to have.¹² For example, I can know you and your tastes and the effects on your welfare of various policies. An organization for implementing national policy might not conceivably come to have such knowledge of anyone. And if it did have it of some individuals, it could not have it of all — and it might therefore, for procedural reasons that can have a utilitarian justification, be excluded from using the information even on behalf of the individuals for whom it has the knowledge.¹³ On the other hand, a policy organization can have command over massive data on the general population that few if any individuals could ever be expected to have, and it might therefore be able to make better judgments of the comparative overall impact of variant policies than individuals could.

An institution can also have command over distinctively institutional devices for reaching decisions. It can break tasks into parts that are addressed by different specialists and can pool specialized knowledge in hierarchical and group decision making. An institution is typically focused on a fairly limited range of problems to resolve. It can therefore invest much more in resolving that range of problems than any individual could justify investing. And an institution may have a regularized system for collecting the information it needs for its decisions, a system that might be able to collect information that no individual could hope to collect.

As a result of its special epistemological capacities, an institution can make some kinds of decisions far better than typical individuals can. Hence, there are things that an institution can do that no individual would be capable of doing, because the institution can know what it takes to

¹² Hardin, “Economics of Knowledge,” p. 131.

¹³ Institutions may commonly be expected to work better if they follow standard procedures to assure that they not be careless.

do those things. This means, of course, that the institution has moral capacities that individuals cannot have. Hence, the institution ought to do some things that individuals cannot do and that individuals therefore cannot be said to have a moral obligation to do. There are some things we can say institutions ought to do when we cannot sensibly say individuals ought to do those things.

Institutions As Meliorative

Commonly, not only do we need institutions to accomplish particular purposes, but we could claim that having an institution of a particular kind rather than not having one would benefit everyone. Hence, *ex ante*, we could give a mutual advantage argument for the existence of some institution. For example, we might readily say of the general framework of government that it serves mutual advantage in comparison to having no government. This is Hobbes's justification of government.¹⁴ However, we cannot generally claim of the details of government and of policy that these serve mutual advantage in comparison to alternative details or policies. Even for such details, however, an institution has available some kinds of knowledge that an individual could not be expected to command.

There is therefore, for epistemological reasons, a difference between what institutions and individuals can do, because they can only expect to do what they have the knowledge to do. Again, this epistemological difference entails a normative difference in what can be done. When we choose to use an institution to help accomplish some purpose, we inherently make a meliorative, not a mutual advantage, move for the reason that we then choose on average or aggregative principles that may run net losses for some together with net benefits for others. In essence, therefore, recourse to government or other complex institutions to accomplish our ends

¹⁴ Thomas Hobbes, Leviathan (Harmondsworth, Middlesex: Penguin, 1968 [1651]), p. 238 [94]; Russell Hardin, "Hobbesian Political Order," Political Theory 19 (May 1991): 156-180, and Liberalism, Constitutionalism, and Democracy (1995 mss.), chaps. 1 and 3.

is in general a meliorative but not mutual-advantage move. This claim is generally true despite the caveat, discussed in the next section below, that a choice of institution or policy might be grounded in such limited knowledge that it could appear to be mutually advantageous.

Consider an example of this general conclusion. It is a commonplace claim that the most important values in the common law are to achieve definitive resolutions of cases at hand, and to establish rules to guide future actions. We need resolutions of conflicts in order to let us go on with our lives. And we need rules that allow us to act in the confident expectation that our actions will stand against legal attack so that we may sensibly and confidently invest in our projects and our lives. What the judge actually does in the common law when a case arises in a new context is to establish a rule to guide future actors while treating the present litigants as though the rule had been in place when they acted. This accomplishes both the values of the common law. In criminal law, holding someone accountable *ex post delicto* to a new rule is objectionable. In civil law it is almost necessary. In the criminal law we do define new crimes and prescribe new punishments in response to new offenses, but we can do this without imposing punishments on those who gave us reason to revise the law. The miscreant merely goes without the new punishment. In the civil law, one or the other party to a complaint will be left in the lurch until a new rule is adopted and applied to their case.

Clearly, one of the litigants is likely to be a net loser from the decision in a case that establishes a new rule. Yet, we may generally conclude that it is better to establish that rule than not to do so and that, whatever rule we establish, we must resolve the dispute at hand. If the parties to the litigation in the present case are similarly situated, so that both might be expected to gain more in the long run from one rule than from an alternative rule, we might easily suppose the

resolution of the case to be mutually advantageous.¹⁵ But commonly we would have to conclude that the actual decision makes one of the litigants a net loser in comparison to how that litigant would have fared under an alternative rule.

The usual defense of the common law therefore is meliorative rather than mutual-advantage. It is grounded, perhaps only implicitly, in the view that some parties' losses are outweighed by other parties' gains. Analogously with Hobbes's defense of having government, one might reasonably claim that having definitive resolutions of cases makes all better off than if there were no such resolutions. But, again, one cannot claim that the details of a particular resolution make everyone better off.

A general problem of institutional decision making is that institutions must commonly make stochastic decisions about what is most likely to be true of typical individuals. Sometimes the decision can be classed as one of mutual advantage, sometimes not. For example, we might choose to vaccinate all children against some dreadful disease, such as smallpox before it was eradicated or polio still today. Yet we may know that the vaccination itself will likely give the disease to some of those vaccinated even while it protects the vast majority against contracting the disease.¹⁶ In this case, we might readily conclude that the policy of vaccination is mutually advantageous *ex ante*, because all children are similarly situated and face, so far as we know, similar risks of contracting the disease from the vaccination or of being subjected to an epidemic if

¹⁵ In the commercial law, this assumption might often be valid. In torts law, it is likely less commonly valid because some people are likely never to be the inflictors of a certain kind of harm but only the sufferers of it. For example, those who seldom or never drive but who are often at risk from cars as pedestrians are not in the same situation with respect to certain liability rules for vehicular traffic as those who often drive. See further, Russell Hardin, "Magic on the Frontier: The Norm of Efficiency in the Law," University of Pennsylvania Law Review (1996 forthcoming), pp. 1987-2020, at pp. 2010-11.

¹⁶ Russell Hardin, "Ethics and Stochastic Processes," Social Philosophy and Policy 7 (Autumn 1989): 69-80.

there is no vaccination.

But we might actually be able to reduce the total number of infections by vaccinating only some large fraction of all children. The live Salk polio vaccine yields some herd immunity, that is, protection of those not vaccinated. The incidentally protected herd might be at less risk of contracting the disease from vaccination than are those children whose vaccination protects the herd. And in any case, a particular disease may not be capable of spreading if only some fraction of the population is vaccinated. Now we could say that, again *ex ante*, it would be mutually advantageous to all children to vaccinate only, say, seventy percent of them. If we could get the percentage exactly right, it would then be a matter of indifference to each child whether she was or was not vaccinated, at least insofar as risk of disease is concerned. If, however, we are unsure of the right percentage and we wish to play safe and vaccinate more rather than fewer, those who are vaccinated might be put at greater risk than those who are not. The actual choosing of which children to vaccinate is then not a matter of mutual advantage even though selecting the percentage to vaccinate could be. (Of course, those children who do eventually contract polio do not *ex post* join in the mutual advantage of the vaccination program.)

Choosing which children to vaccinate has more the character of the choices made in a typical cost-benefit analysis, in which the costs to all are weighed against the benefits to all, although some may be net beneficiaries while others are net losers. It is hard even to imagine a workable alternative to cost-benefit analysis in many contexts, such as the siting of roads, although there can be compensations of the net losers to, in legalese, make them whole. But for obvious reasons of the problem of strategic misrepresentation of preferences in order to gain advantage, the assessment of costs and benefits must typically be somewhat stylized rather than personalized. Hence, standard compensations may not satisfy all the losers. As in Hobbes's argument for government, having a system of imminent domain and stylized assessments of costs

for public projects might be mutually advantageous in principle or ex ante. But its actual working in any particular instance is unlikely to be.

(Note that the problems discussed immediately above need not be unique to institutional settings. An individual might also face choices that benefit some at the cost of others and might make calculations similar to those of an institution. But this is the daily fare of institutional decision making and of public policy, while it may a relatively small part of the lives of many of us as individuals.)

Apparent Mutual Advantage

There are contexts in which the argument that we should not make the interpersonal comparisons required for melioration seems implausible. For example, an intervention to stop a disastrous conflict might seem almost surely to entail net losses to some in order to bring general benefits to all. Almost everyone might agree that the end of conflict even at such individually specific costs would be a better state of affairs than the continuation of the conflict. Some might, of course, argue from a libertarian or other theory that no one should intervene even if such intervention would produce a better state of affairs. More generally, however, it seems plausible that all major social institutions and policies are likely to entail melioration that involves interpersonal comparisons or weighings that is not mutually advantageous despite the appearance of mutual advantage.

At the level of knowledge that many institutions can actually accommodate or claim to have, they typically act to the mutual advantage of those whose interests they address. They act from knowledge of typical or even statistical or average cases. More particular knowledge might show that there is wide variation around the average case and that a policy directed at the average case imposes net losses on some.

Many institutions similarly act from knowledge of typical cases but still make comparisons of welfare across different classes of such cases. These institutions may then act to meliorate rather than to serve mutual advantage. For example, hospitals are far more attentive to some problems than to others on the simple, and generally agreed, ground that those problems more seriously affect health and, therefore, welfare. Your rampant infection merits more attention than my sprained wrist. Similarly, an institutional body devising a tax scheme might make the scheme progressively more burdensome for those with higher incomes on the widely agreed view that the welfare effects of progressivity are generally good. One with special familiarity with a few particular persons might argue that the welfare effects of these policies on those persons are perverse in that they cost some more than they benefit others. But the institutions setting and implementing the policies might not be able to command the special knowledge to reach this conclusion.

In general, we face two quite separate problems of choice about collective matters. First, we create institutions and organizations to handle certain issues. Ex ante the creation of a particular institution may be genuinely mutually advantageous in our expectations (although the choice between alternative institutions may not be). Second, we — or our institutions — adopt specific policies and implement them. The choice and implementation of many actual policies will not be mutually advantageous. Our two choices, of institutions and of policies, are epistemologically analogous to a limited extent. Even though an institution might be expected by each person to be beneficial to herself, it may nevertheless be true that it would also be expected by everyone that it would bring net losses to some. My expectations of net benefit in such a case would be statistical. I might expect a range of possible effects on me, from net losses to net benefits, but with the expected value overall of net benefits.

Adopting an actual policy will be similar in the sense that it will be based on statistical

expectations of the costs and benefits of the policy. Those expectations might, again, come from a range of possible outcomes for particular individuals, with some of them losing overall while others benefit. The difference between the two — adopting a policy and creating a policy institution — would often be that actual individuals might be surer of where they are likely to come out from the policy. For example, a new policy of progressive taxation would affect many people in ways they could confidently predict.

Conclusion

If the argument here is correct, institutions, unlike individuals, are generally forced to make meliorative moves that are not mutually advantageous. You and I can agree to act in ways that serve our mutual advantage without making any claims of interpersonal tradeoffs. That is, indeed, the nature of ordinary exchange. I trade my x for your y and we are both better off. Perhaps one could reasonably say that you are made much better off while I am made only slightly better off. And perhaps one could even argue that we should not have exchanged but that you should merely have given me y and that your gift would have made us jointly better off than our trade would have. But the trade is still merely mutually advantageous and therefore utilitarian with respect to the status quo before the trade without any judgment of how much better off it makes the two of us together. Variance among the individuals that institutions address virtually guarantees that their actions will not have this character in actual fact even though an institution might be incapable of judging that its actions do impose net costs on any particular persons.

The role of courts of equity at their height in England was to handle the variant cases that were decided properly according to the law but that imposed supposedly unreasonable penalties on parties who were in some important respect unlike the standard parties before the law. One can imagine a similar device to accommodate claims that specialized knowledge outside

the normal purview of an institution would justify special treatment of someone. But one cannot coherently expect that an institution be able to treat every case as special in this way and still do a good job of accomplishing its purpose of serving our interests. Moreover, if an institution adopts a policy of allowing special appeals to “equity,” it must likely do so only for extreme cases, not for relatively close calls. Their capacity for streamlining and standardizing decision making is a very important part of what makes institutions useful to us. We cannot simultaneously wreck their standard procedures and expect them to continue to be as useful. Again, ex ante it is plausibly mutually advantageous that we create such institutions even though, in specific instances of their acting, they may inflict net costs on some and we know ex ante that it is extremely likely that they will do so.

One might wish to say that melioration approximates to mutual advantage in such contexts. But this does not relieve the fact that it is not actually equivalent to mutual advantage and that it requires de facto tradeoffs across individuals. Hence, a rigorous Paretian must oppose public policy altogether. The only genuinely rigorous Paretians in print are very hard-nosed libertarians, some of whom think we would manage to be as prosperous without institutions as we are with them. They therefore disagree fundamentally with Hobbes, whom they should see as their most threatening intellectual opponent. I think Hobbes has far the better case to make, even though he may exaggerate the extent to which efforts at reform threaten massive destruction from civil war and anarchy.¹⁷

But the general run of economists who claim to have inherited Pareto’s value theory and who sometimes assert the impossibility of making interpersonal comparisons of welfare are not Paretian in actual practice. They regularly prescribe both institutions and policies and they are the authors of cost-benefit analysis. The ardor with which economists have defended the

¹⁷ Hobbes, Leviathan, p. 238.

impossibility of interpersonal comparisons is more than matched by the energy with which they have urged policies requiring tradeoffs that could only be grounded on the assumption of such comparisons. Their inconsistency is no proof of the correctness of making such tradeoffs. But if Hobbes is even roughly right, the impossibility of escaping the stasis of ex ante failure to adopt institutions for collective purposes is a powerful argument for the correctness of such tradeoffs and of the de facto interpersonal comparisons they imply.

This conclusion suggests a slightly tricked up justification of such comparisons. Ex ante, we all would want a better regulated world in which life would be much better for most of us and, in our ex ante expectations, even for each of us. Once we have the institutions for such regulation, we cannot practically escape making tradeoffs between individuals — perhaps often unidentified and even unidentifiable individuals, but tradeoffs nevertheless. We cannot have the institutions that are ex ante justifiable unless we also take the eventual tradeoffs. By ex ante mutual advantage argument, we ought to have those institutions. But if we cannot have them without their eventual tradeoffs, then we ought also to have the tradeoffs. One of two contrary conclusions follows. The first is that the rectitude of interpersonal trade offs follows from an argument grounded only in mutual advantage. The second is that we should not have institutions at all because we cannot make interpersonal comparisons.

The second of these conclusions could be taken as an implication of the doctrine that ought implies can. If we metaphysically cannot make interpersonal comparisons, then we ought not. Hence, we also ought not justify a system that depends on making such comparisons and we should junk institutions. The first and contrary conclusion could be taken as an injunction to condone tradeoffs or, even more strongly, to take interpersonal comparisons seriously. I concur with the stronger version of this conclusion. I think we should (and do) take interpersonal comparisons seriously. I do not have a metaphysical argument to establish how such comparisons

can be given meaning in principle or an epistemological argument to show how we can know another's welfare to compare it to our own.

But the weaker version of the conclusion seems incontrovertible: We ought to condone interpersonal tradeoffs. The alternative — to suppose we ought not have institutions to enhance welfare — is preposterous. Faced with the comparative difficulties of supposing we ought not have such institutions and of supposing we ought not condone the requisite interpersonal tradeoffs, one who is motivated by the pragmatics of living well — for oneself and for others — must find the latter difficulty less implausible than the former.