

*Utility, Publicity and Manipulation* \*  
Adrian M. S. Piper

In our dealings with young children, we often get them to do or think things by arranging their environments in certain ways; by dissembling, simplifying, or ambiguating the facts in answer to their queries; by carefully selecting the states of affairs, behavior of others, and utterances to which they shall be privy. We rightly justify these practices by pointing out a child's malleability, and the necessity of paying close attention to formative influences during its years of growth. This filtering of influences is necessary, we point out, if children are ever to reach a degree of maturity and inner stability that will enable them to understand and cope adequately with the complexities, contradictions, and difficulties of the world from which we now seek to shield them. Thus a child's eventual state of competence, maturity, and autonomy adequately justifies our current practices of manipulation and selection of his environment: such practices are rightly held to be ultimately in the child's best interests as an adult. These truistic remarks motivate the following discussion.

I

In *The Methods of Ethics*<sup>1</sup>, Sidgwick considers some differences in the strategies of action and decision appropriate to the Utilitarian living in an ideal social community, and again in the actual one. In an ideal community of enlightened Utilitarians, he claims, no one would be justified in secretly acting in some way not sanctioned by the accepted moral rules. For even in cases where it would seem that one was justified on grounds of utility in excepting oneself from some such rule, that one *was* justified would simply mean that certain qualifications should be added to the rule to cover the exigencies of that type of situation, and thus that these qualifications would apply in all cases relevantly similar to one's own:

It is evident, that if these reasons are valid for any person, they are valid for all persons; in fact, that they establish the expediency of a new rule... more complicated than the old one; a rule which the Utilitarian, as such, should desire to be universally obeyed<sup>2</sup>... If therefore we were all enlightened Utilitarians, it would be impossible for anyone to justify himself in making false statements while admitting it to be inexpedient for persons similarly conditioned to make them; *as he would have no*

---

\* I am indebted to Professors John Rawls and David Auerbach for trenchant criticisms on an earlier version of this paper. They are not, of course, responsible for my blunders.

<sup>1</sup> Book IV, Chapter V, Section 3.

<sup>2</sup> *Ibid.*, p. 485

*grounds for believing that persons similarly conditioned would act differently from himself.*<sup>3</sup> (my emphasis)

This last clause is ambiguous but significant. One reading of it would say that a Utilitarian in some situation would expect other Utilitarians to act the same when 'similarly conditioned' because all Utilitarians would react the same way under some particular set of conditions, i.e. that all would reason similarly and thus act similarly. Here 'similarly conditioned' would have to mean similar in all respects relevant to the making of one particular decision, i.e. similar in personal makeup as well as in circumstances. This is not an unacceptable interpretation of the passage, but it is uncharitable in that it implicitly ascribes to Sidgwick the view that in a thoroughgoing Utilitarian society everyone is essentially alike, hence that a similarity in situation suffices to determine a similarity of response. While this may in fact be a valid implication of the Utilitarian doctrine in its ideal form<sup>4</sup>, it is not likely that Sidgwick would accede to it.

A weaker and more sympathetic reading would construe Sidgwick as meaning that if everyone commonly acknowledged holding to Utilitarian principles, then *my* reasons for acting in a certain way will be acknowledged as valid by everyone, even though no one else can, strictly speaking, be 'similarly conditioned' just as I am. Here my expectation that others would behave similarly if similarly conditioned is actually an expectation that, since we all share and mutually recognize the same moral principles, others would condone and support my action as being that which they would perform if they were, hypothetically speaking, in my shoes.

The first reading explains Sidgwick's claim in terms of an assumed uniformity of motives, beliefs, and responses among Utilitarians -- not a clearly desirable condition to impose on the ideal state. The second explains it in terms of an assumed publicity of mutual acknowledgement of Utilitarian principles as binding on all individuals in the community. The latter would seem more faithful to Sidgwick's intended meaning. Sidgwick can then be understood in this passage as asserting not simply the truism that an exception to a rule which ranges over some class of cases itself ranges over some class of cases. He is asserting that if everyone justified his actions on grounds of utility, these grounds would be acknowledged as valid and accessible to everyone in any situation; that any action consistently and adequately justified on these grounds could be expected by the agent to receive validation by others in the community. It is in this sense, then, that the

---

<sup>3</sup> *Ibid.*, 488

<sup>4</sup> In fact I suspect it is, though I will not try to argue this here.

rule in question would acquire a qualifying clause, and for this reason, seemingly, that Sidgwick sees the principles of Utilitarianism as public. They would be public in the sense that we could not know what someone had done without thereby knowing why.

This is a consequence of Sidgwick's conception of the ideal community as consisting of what are essentially Act-Utilitarians (thus I will use 'Utilitarianism' and 'Act-Utilitarianism' indifferently in discussing Sidgwick's Utilitarianism and its implications unless specifications are explicitly made to the contrary). Although moral rules are held in common, the decision to follow or not follow them is made on Act-Utilitarian grounds.<sup>5</sup> For even where an apparently Rule-Utilitarian stance is adopted (e.g. where Sidgwick appraises the utility value of commonsense moral rules), this is done on the grounds that the overall utility of following and promulgating the rule outweighs the personal disutility of doing so. But clearly this does not preclude the case -- without begging the question -- where, in the estimation of the Utilitarian, the overall utility of controverting the rule is in fact greater than that of following it. In this latter case, it would seem to clearly conflict with Utilitarian *first* principles to follow the rule.<sup>6</sup> So if everyone agrees to follow these rules, this is because it is commonly recognized as useful to do so. But because justification by utility is accessible to all members of this community, the very acknowledgement of an agent's situation as being of a certain kind will determine that an exception should be made, for so would any Utilitarian reason who has access to the facts -- including the agent himself.

Under actual circumstances, however, the case for Sidgwick is somewhat different:

the Utilitarian may have no doubt that in a community consisting generally of enlightened Utilitarians, these grounds for exceptional ethical treatment would be regarded as valid; still he may... doubt whether the more refined and complicated rule which recognizes such exceptions is adapted for the community in which he is actually living;

---

<sup>5</sup> *Op. cit.*, pp. 486-490, *passim*.

<sup>6</sup> For an in-depth and comprehensive examination of this and related issues, cf. Lyons, D., *Forms and Limits of Utilitarianism* (Oxford: Clarendon Press, 1965), esp. Chapter IV.C. Also see Hodgson, D.H., *Consequences of Utilitarianism* (Oxford: Clarendon Press, 1967), pp. 3-7, and Singer, P. 'Is Act-Utilitarianism Self-Defeating?' *Purit. Rev.*, V. 61, January, 1972, esp. p. 565

and whether the attempt to introduce it is not likely to do more harm by weakening current morality than good by improving its quality.<sup>7</sup>

While the justification for conforming or failing to conform to accepted moral rules is the same for the Utilitarian in the actual as in the ideal community there is an asymmetry with respect to the accessibility of his principles to others. In the actual, implicitly non-Utilitarian community, the Utilitarian must consider not only the effects of following or not following commonly-accepted moral precepts, but also the comparative utility of letting others know the grounds for his decision. For the Utilitarian does not, presumably, do the same things, for the same reasons, as others do in this situation. So whenever his considered actions diverge from those enjoined by the moral rules of the community, the Utilitarian must weigh the utility of this divergence as such, in addition to the utility of the act itself. As Sidgwick argues, the disruptive effects on others of this divergence may well lead the Utilitarian to conclude that the greatest utility would be served either by performing his action secretly, or performing it publicly and lying about his reasons for doing so. For in the latter case as well, publicizing the Utilitarian doctrine might undermine general conformity to useful moral precepts even more strongly than his seemingly immoral act, which is at least amenable to moral or legal sanction. Sidgwick seems to imply that, in a non-ideal, non-Utilitarian society, the principles of Utilitarianism should not be propagated at all in their most general form, for their effects on the general community may well be insidious to moral conduct if commonly acknowledged:

the opinion that secrecy may render an action right which would not otherwise be so should itself be kept comparatively secret; and similarly it seems expedient that the doctrine that esoteric morality is expedient should itself be kept esoteric... And thus a Utilitarian may reasonably desire, on Utilitarian principles that some of his conclusions should be rejected by mankind generally.<sup>8</sup>

Thus Sidgwick claims not only the validity of a covert application of Utilitarian principles to yield grounds for secret exemption of oneself from some moral precept, but also the validity of secretly adopting these principles themselves. Both are justified on Utilitarian grounds. So when Mill in *Utilitarianism*<sup>9</sup> dismisses the possibility of such exemption as an objection to Utilitarianism because no doctrine can be formulated which successfully rules it out in all cases, he seems to miss the real point of the objection, which is the

---

<sup>7</sup> *Methods*, p. 489

<sup>8</sup> *Ibid.*, p. 490.

<sup>9</sup> Chapter 2, Paragraph 25.

*ubiquity* of the Utilitarian doctrine: all moral conceptions must admit the possibility of exception in practice, but Utilitarianism is unique in rationalizing such exception in theory.

But given Sidgwick's classical formulation, we can see why he should take the strategy he does. For if the community knows that agent S is a Utilitarian, they know he justifies all his actions with reference to their utility. And then it is easy for them to infer that S will conform to or exempt himself from publicly-held moral rules as it maximizes utility to do either. But if the community knows when supporting these rules would not, in S's view, maximize utility, they can infer when he will secretly exempt himself from them. But if they know this, S has clearly failed to act secretly, hence failed to maximize utility, and failed to publicly uphold the moral principles of the community. So the Utilitarian must either forego forming the normal human relationships relative to which his actions and beliefs would necessarily be public (to degrees varying with the extent of personal involvement in the relationship), which is of questionable utility, or else the agent must adhere to his Utilitarian convictions covertly. The latter seems to be the more expedient strategy. The Utilitarian cannot, then, make public his convictions without undermining both commonly-accepted non-Utilitarian moral precepts and his own attempts to maximize utility in a non-ideal situation.

That this thoroughgoing policy of secrecy suggests a difficulty in theory about bridging the gap between the non-ideal and the ideal societies will surely be noted. Unless the Utilitarian is prepared to deny any utility at all to conforming to non-Utilitarian precepts, it is hard to see what his strategy might be for bringing a community from a non-ideal to an ideal state, since he cannot, without overall loss of utility, publicize his convictions in the non-ideal one.<sup>10</sup>

## II

Unfortunately, this problem extends to the so-called ideal society as well. We will now see that even if the Utilitarian could make everyone else a Utilitarian suddenly, without working through the near-insurmountable obstacles of transition just described, such a community still would not be completely viable for the same kinds of reasons. In Sidgwick's brief adumbration of the ideal community, recall that he says only that everyone is a Utilitarian, hence that everyone justifies his actions according to the same principles. But we saw that he does not explicitly say that everyone

---

<sup>10</sup> Sidgwick seems to be sympathetic to this conclusion. Cf., e.g., pp. 474-5, 480-2, 484-6, 489.

acknowledges these principles publicly. On the sympathetic reading, he seems to assume that this follows from the universal applicability of the principles themselves. Now let us look at two varieties of the ideal Utilitarian society, one where the publicity condition holds and one where it does not, in order to see whether either alternative will yield the model Sidgwick has in mind.

First let us try to characterize more fully a variant of the ideal Act-Utilitarian society in which the publicity condition does not hold, i.e. in which it is not common knowledge that everyone is an Act-Utilitarian. While everyone in fact adopts Act-Utilitarianism as his only rule of conduct (where by 'utility' let us understand, roughly, the maximization of happiness, without filling this in any further for the moment), each person does not explicitly recognize others as so doing. Hence although everyone attempts to promote the greatest social utility through his actions, no one views this rule of conduct as the commonly-held one. Thus each is motivated by benevolence towards the rest of society, but no one is conceived as explicitly sharing these benevolent purposes with anyone else. This is not to say that each conceives the others as selfish and only himself as benevolent. Rather, it is that benevolence is so much an all-pervasive but unarticulated motive of conduct that no one self-consciously conceives of himself or of others in this way. We should try to imagine a situation in which benevolence is so ingrained in behavior that there are no circumstances under which conscious articulation of it is required: it is, let us say, too much of a truism to be worthy of mention. We can think of benevolence in the ideal Act-Utilitarian community as analogous to the motive of self-support in our own. Though we have many reasons and motives for choosing a particular plan of life or vocation, that we should do something with our lives that will insure our own survival is unquestioned; so much so that it rarely figures in an explanation of why we chose as we did.

The extent of each person's benevolence is, let us suppose, constrained by his own adoption of the Utilitarian doctrine: that is, in estimating the sum of social utility to be achieved by any action, each person automatically gives equal weight to his own happiness as to those of others.<sup>11</sup>

---

<sup>11</sup> I add this proviso in order to avoid the problem inherent in the notion of perfect altruism, knowing full well the import of Williams' convincing argument to the effect that a rational Utilitarian must be willing to abdicate the maximisation of his own happiness - here, those convictions, feelings, plans, and projects with which he most deeply identifies - when doing so would increase the net balance of happiness for the community ('A Critique of Utilitarianism' in Smart and Williams, *Utilitarianism: For and Against* (Cambridge U.: 1973), Ch. 5). While Williams' case is well made, it is only

Insofar as this is an ideal society, we must also imagine it to be stable, well-ordered, and otherwise successfully operated in the absence of the publicity condition. We must, above all, assume it to be reasonably well-coordinated: in performing the act with the best consequences, each member takes into account the probable behavior of others, and the necessity of insuring against conflicting self-defeating acts. Thus we can assume, with Sidgwick, that in general the members of this society concur in following certain commonsense moral precepts and rules of thumb on Act-Utilitarian grounds. This is to stipulate that Act-Utilitarian deliberation will, in the absence of the publicity condition, generate these precepts as conventions. Now since everyone is an Act-Utilitarian, hence reasons similarly with regard to the consequences of actions, each person will have no trouble in predicting or assessing the outcome of the behavior of others when deciding what to do. For although they do not assume that each acts from Act-Utilitarian convictions, they do consider the effects of each others' behavior. To each member of this society, the others behave *as if* they were Act-Utilitarians in the minimal sense that their actions have, and are recognized by others to have, best consequences under the circumstances.

However, it is important to emphasize that this state of things does not provide evidence to any member for thinking that everyone else *is* an Act-Utilitarian, for to act *as if* one was an Act-Utilitarian is often to adopt *prima facie* non-Utilitarian moral conventions when they have the best consequences, which in view of the benefits of coordination, will be a good part of the time. This means that one will be unable to distinguish Act-Utilitarians from, say, highly efficient Intuitionists, on the basis of behavior alone. To identify them *qua* Act-Utilitarians, we must know their intentions and their reasons for acting. But since, as we have said, this ideal society runs smoothly and acceptably in the absence of the publicity condition, members of this society will, by hypothesis, rarely be called upon to justify or explain their actions overtly; for the practices and conduct of each will mesh harmoniously with those of others. So the opportunity to discover the moral convictions on which they are based will be small indeed, if not non-existent. The situation bears comparison with a society in which traditional social roles and practices are, like the benevolent motive, so deeply embedded in the history of the society that talk of reasons and justification for them are otiose. Persons are conceived as inextricably dependent on these roles and practices in a way that practically vitiates the very possibility of calling them into question. In this sense, we may say that the ideal Act-Utilitarian society as a whole lacks *self-consciousness* -- not a stringent condition to impose, when a

---

by temporarily disregarding it that I can try to flesh out the notion of an ideal Act-Utilitarian society in a *prima facie* plausible way.

society is conceived as functioning smoothly: the roles and practices generated by rational Act-Utilitarian calculation are so embedded in the social structure that their justification is made unnecessary by the smooth and harmonious functioning of the social order itself.

Now if each Utilitarian had no reason to suppose that others shared his convictions, he would obviously have the same good reasons to assume the utility of covert actions in the ideal as in the actual case.<sup>12</sup> Further, he might again correctly assume the greater utility of his 'esoteric morality' than its public counterpart. These reasons would be grounded in the supposition that in this ideal society, a person's conduct would be fully informed by rational Utilitarian reasoning, but would differ essentially from our actual behavior only in its degree of efficiency and success in bringing about the best consequences. For Sidgwick, the reasons militating against making public the Utilitarian credo have nothing to do with peoples' actual relative perfections or irrationality, but with how any rational individual might reason in light of this doctrine, and the damaging consequences that would ensue if everyone, or most people, began to consistently reason this way. We will return to this problem shortly.

Even if our hypothetically-placed Utilitarian somehow found out that everyone else was also a Utilitarian, he might well judge even here that it would be better to maintain silence on this point, for fear of the disruptive effects of publicizing it. For note that to say that everyone is an Act-Utilitarian is not obviously to say that they act unanimously, but just to say that each tries<sup>13</sup> to bring about the best overall consequences through his action. Now if each has been following commonsense moral precepts in part on the supposition that they reflect the convictions of others and satisfy their valid expectations, we may well expect chaos to result when everyone's assumptions are thus falsified. This would seem to hold whether everyone is a Utilitarian or not.

---

<sup>12</sup> This seems to be J.J.C. Smart's conclusion as well. Cf. 'An Outline of Utilitarian Ethics', p. 50, in Smart and Williams.

<sup>13</sup> See Lyons' distinction between accepting and following the dictates of Act-Utilitarianism, pp. 151-2. My reasons for adopting the former, weaker version will shortly become evident.



## III

Now one might want to claim that as long as we are talking about an ideal society, there is no reason not to build the publicity criterion into it: this brings us to the second of the two cases, in which the publicity condition is assumed to hold. But there is good reason to think that there is in fact no consistent rendering of such a case. We have mentioned in passing the obstacles that must be overcome in getting from the actual to an ideal Utilitarian society; and the conclusions of III suggest that these pains of transition are considerably increased in severity when publicizing the Utilitarian doctrine is made part of the process. But it may be that even if we suppose this problem solved, the very concept of an ideal Act-Utilitarian society in which the publicity condition holds is impossible. This is one way of understanding Hodgson's interesting argument, which is the basis for his critique of Utilitarianism.<sup>14</sup> Essentially, he argues that truth-telling and promise-keeping would be impossible in a society where everyone was, and recognized each other as being, Act-Utilitarians. His argument is based on the assumption that one part of the utility of a great many types of social action involves the degree to which it satisfies the justified expectations of another. In an Act-Utilitarian society (to generalize from Hodgson's examples), no one could have valid expectations about another's action. An agent S would only do x if x had the greatest utility: and x would have the greatest utility only if it satisfied the recipient R's expectations. But R would expect x only if R believed that S would do x. Being equally rational, S would know this, hence would do x only if he believed that R expected x. But since S's doing x depends on knowing R's expectations, and R's expectations depends on knowing whether or not S will do x, R has no *prior* reason to expect S to do x. And since the utility of doing x depends on knowing R's expectations, S cannot determine whether or not doing x has greatest utility. So there is no *prima facie* reason to do x rather than  $\sim x$ . This dilemma holds for any act x that involves fulfillment or violation of someone else's expectations.

This brief adumbration of Hodgson's argument abstracts considerably from the actual cases he considers. But I think the fact that it can be generalized in this way shows why Singer's<sup>15</sup> attempted rebuttal of it does not work. Singer objects that Hodgson's argument is based 'not on the existence of a reason for lying or breaking a promise, but on the absence of a sufficient

---

<sup>14</sup> Hodgson, *op. cit.*, Ch. 2.

<sup>15</sup> Singer, *op.cit.*

reason for telling the truth or keeping a promise.<sup>16</sup> Most of the rest of his discussion is then devoted to examples within the context of an hypothetical Act-Utilitarian society that purport to demonstrate the existence of a sufficient reason of utility for telling the truth or keeping a promise. He then argues that in fact this is usually the case; and since, in an ideal situation (presumably of the general type discussed above, pp. 10-14), the motives we usually have for lying or breaking promises (self-interest, malevolence, pride, etc.) would be absent, S would usually tell the truth and R would usually expect him to; hence the general practice, with its concordant expectations, could after all be established.

But Singer's argument fails at two places. First, it is not at all clear that motives for lying and promise-breaking would be absent, even if we admit for the sake of argument that the ones he cites would be. The rational Utilitarian motives of benevolence and the desire to maximize utility are, as we have seen, more than sufficient to justify a lie for a Utilitarian under certain actual as well as ideal conditions; and in the absence of a rule prohibiting lying, we might expect such circumstances to multiply.<sup>17</sup> An attitude of paternalism might thus be expected to flourish, unless certain non-Utilitarian assumptions about the intrinsic worth and dignity of human beings, their freedom to undertake responsibility for their lives, etc. were to be built into the Act-Utilitarian ideal. We saw in III that in the absence of the publicity of the Utilitarian doctrine, paternalism and its consequences for deceptive conduct would certainly exist. Whether truth-telling and promise-keeping would predominate in the present instance needs to be demonstrated. I see no reason to *assume*, as Singer does, that they would. The positive consequences Singer claims for them follow only if they would be operative in these circumstances; and whether they would be or not is the very point at issue.

This brings out the second problem. Singer's supposition of the greater utility of telling the truth and keeping promises depends upon those practices being the ones that in actual fact have greater utility because of our reliance on them in getting about. The examples he discusses have in common the feature that, given our actual social habits and practices, truth-telling and promise-keeping generally have best consequences because we expect them, and can structure our plans around these well-founded expectations. But this clearly avoids the question of whether they could be expected to hold if we were, and acknowledged each other as, rational Act-Utilitarians.

---

<sup>16</sup> *Ibid.*, pp. 560-561

<sup>17</sup> Cf. Smart, pp. 50-51

Singer thus misses Hodgson's point. All our social conduct is in some way structured around these practices, so it is not difficult (as Singer demonstrates) to amass examples of this dependence. Nor is it difficult to imagine exceptional circumstances in which the utilities would seem to stack up on the other side. Indeed the general problem with any such examples of actions which are to be assessed in terms of their consequences is that the range of alternatives are too great for examples to prove anything. For if we assume the truism that any act can be regarded as having some good consequences, then we can say that for any act  $x$ , there is a weak disjoint range of causally possible consequences (both positive and negative) of  $x$ ,  $C_1 \dots C_n$ , and a disjoint range of possible situations  $S_1 \dots S_n$  in which  $x$  can occur. Also there is a mapping of  $C_i$  for any  $i$  onto some relevant subset  $S'$  of  $S_1 \dots S_n$  defining the class of situation in which  $C_i$  can occur, such that  $C_i$  is an actual consequence of  $x$  only if some  $S_i$  in  $S'$  obtains (this latter condition allows us to say both that the same consequences of  $x$  may obtain in different situations (i.e. if  $C_i$  is mapped onto  $S_i$  and also onto  $S_{j \neq i}$  and  $\sim(S_i \text{ iff } S_j)$ ), and that  $x$  may have many different but compatible consequences in the same situation (i.e. if  $C_i$  and  $C_{j \neq i}$  could both be mapped onto  $S_i$  and  $\sim(C_i \text{ iff } C_j)$ ), and that the utility value of any consequence of  $x$  depends on the circumstances in which it occurs or on the particular repercussions it has). But it follows from this that we may easily insure that an example can be constructed in support of doing  $x$  just by carefully selecting the situation in which  $x$  occurs with an eye to the positive consequences of  $x$  which result in that situation. Since there is always some class of situations in which  $x$  may be regarded as having positive consequences, some member of this class can always be selected in framing an example; similarly if one wishes to inveigh against doing  $x$  on the grounds of its negative consequences. So presenting the consequences of some positive-making situation support of  $x$  can never have any particular justificatory force, since this is, in effect, merely to load the dice in favor of these positive consequences. Thus the fact that Singer's examples (telling someone the right time, commending a film) seem to come down on the side of truth-telling and promise-keeping reflect little more than his own prejudices -- which, to be sure, we all share. But this is easy to do.

What is not so easy is to prevent these preferences from sullyng one's Utilitarian modes of reasoning; to try to assess the utility of these practices without a prior bias on the side of communication based on truth-telling or commitments based on promise-keeping. A refutation of Hodgson's argument would have to demonstrate that these practices could be established independently of such a bias, and consistently with the application of Utilitarian reasoning alone. For what Hodgson shows is that if we give no independent weight to prior expectations and habit patterns based on the prior existence of these practices, there is no way consistent with

Utilitarian reasoning of bringing them into the picture when we are trying to decide what to do.

I think we can see this a bit more clearly in the generalized outline presented above, where we have no information favoring the doing or forbearing of the act. If we treat the act as neutral with respect to our more ingrained actual social practices, it becomes clearer that Singer's 'sufficient reason' argument is inadequate. A sufficient reason to do any act *x* in an Act-Utilitarian society can only be that it maximized the general social utility -- and since there is no way of determining this in advance of the expectations aroused by doing it, there is no probability favoring its being done at all. Hodgson's analysis suggests that under such 'ideal' conditions, very little could *ever* get done.

A more sophisticated attempt to meet Hodgson's challenge has been made by Allan Gibbard.<sup>18</sup> He interprets Hodgson's argument as claiming that when good consequences depend on the coordination of actions, rational methods of promoting them are self-defeating unless kept secret.<sup>19</sup> He offers an example of two Act-Utilitarians who, having agreed to play tennis, deliberate about whether to keep their agreement or not, where each will come to the courts if and only if he thinks it is sufficiently likely that the other will. Gibbard argues that Hodgson's reasoning mistakenly infers from this type of instance that what an Act-Utilitarian should do *never* depends on what he has agreed to do, since he never has sufficient reason for believing that an agreement made with another will be kept. Thus the problem as Gibbard sees it is to show that making an agreement in an Act-Utilitarian society under certain circumstances could, after all, alter the expected consequences of the acts open to the two parties; and thus that the agreements that would be kept in such a society are most of those which an Act-Utilitarian would find it desirable to keep.<sup>20</sup>

Gibbard's argument to this effect is strategically similar to David Lewis' in 'Utilitarianism and Truthfulness'.<sup>21</sup> Both conceive the issue as a limited-alternative coordination problem (Lewis' example is of two rational Act-Utilitarians placed in different rooms who must choose whether to press the red button or the green; only if both push the same button will utility be maximized). Also, both argue that there may be a sufficient condition for

---

<sup>18</sup> 'Utilitarianism and Coordination', Ph.D. Dissertation, Harvard University, 1971.

<sup>19</sup> *Ibid.*, p. 156

<sup>20</sup> *Ibid.*, p. 159

<sup>21</sup> *Austr. J. Phil.*, V. 50, No. 1: May 1972.

solving the problem in an assumption independent of but compatible with the case as stated. For Lewis, it is consistent with the problem to stipulate that the parties will be truthful whenever it is best to instill in the other true beliefs about which one has knowledge.<sup>22</sup> Gibbard's independent premise is of the parties' common knowledge of whether their society has kept such agreements in the past, i.e. whether they have a history of conventions.<sup>23</sup> I want to look more closely at the latter answer, since Lewis' answer presupposes it as well.

Gibbard's account relies heavily on the analysis of the origin of a convention as the solution to a coordination problem supplied by Lewis in his book, *Convention: A Philosophical Study*.<sup>24</sup> For Lewis, expectations concerning the behavior of other parties in a coordination problem rely largely on precedent, i.e. reasoning inductively from relevantly salient solutions to a similar or analogous past coordination problem, to a most efficacious solution to the present one. The precedent in question may have been established deliberately or by chance, by agreement or tacitly.<sup>25</sup> The more points of similarity between the present coordination problem and its precedent, the more each party is justified in expecting the other party to concur in solving it in a similar or analogous way. Thus a solution established by precedent gives each party reason to assume a certain pattern of predictable behavior in the others, and calculate his own conduct accordingly.

But I do not believe Lewis' analysis answers the problem raised by Hodgson, hence that Gibbard's reliance upon it is a mistake. Both Gibbard and Lewis (in 'Utilitarianism. . .') treat the issue as a simple question of coordination between two alternatives, where determining the best overall consequences depends on the ability of each party to predict accurately what the other will do, hence on each having justified expectations about the other's behavior. The independent suppositions which Gibbard and Lewis introduce function to underwrite such expectations. But this is not quite the dilemma Hodgson raises. For Hodgson, the question is not whether I should do x or y, but what, among a nearly *unlimited* number of possibilities, I should do at all in the presence of another rational Act-Utilitarian. Thus Hodgson's formulation *precludes* the type of independent hypothesis adduced by Gibbard and Lewis that would permit their game-theoretical solutions. Let me articulate this claim in greater detail.

---

<sup>22</sup> *Ibid.*, p. 18

<sup>23</sup> *Op. cit.*, p. 164

<sup>24</sup> (Harvard U.: 1969), especially Chs. 1 and 2.

<sup>25</sup> *Ibid.*, pp. 33-36

If the alternatives for two rational Act-Utilitarians were between, say, going to the courts and staying at home, this would mean that each party, being fully rational, might simply construct the same coordination matrix, accurately working out the probabilities and the desirable risk for each alternative. Knowing the full rationality of both parties, each would rightly expect this process of reasoning to be fully replicated by the other, hence each would expect the other to arrive at the same solution which he himself did. A condition of this would of course be that both know and assign the same weights and probabilities to each alternative. Thus the expectations of each party would, in this case, have to be known to the other: these expectations would derive from their rationality and their accurate weightings and probability assignments to the positive consequences of each alternative. But because both Gibbard and Lewis give the parties access to the same precedent-setting information, each *can* rely on the other to assess the alternative in just this way.

This reveals more clearly the role of the consistent but independent assumptions utilized in Gibbard's and in Lewis' solutions. They amount to stipulating in advance the weight which each party can be expected to assign to the alternative of keeping the agreement to play tennis (or telling the truth about which button one has pressed). This shows how similar Gibbard's and Lewis' version of the problem and the solution are. For both, the problem arises when two parties must choose between two alternatives whose respective weights for each party is to the other, and where this preference depends, for each, on knowing the preference of the other. But clearly, any assumption which provides reason for weighting one alternative more highly than the other will solve the difficulty, as long as this assumption provides additional common knowledge to both parties. But now we are back to Singer's strategy: the assumptions in question build into the situation the expectation of each party that the other will weigh the consequences of keeping the agreement similarly, and most strongly. And we have already discussed some of the problems inherent in this approach.

The problem as Hodgson poses it is different. His argument will, I believe, be best served by rather extensive direct quotation:

an act which has been promised could have greater... utility than it would have had if it had not been promised, only if the promisee expects the act promised more than he would have done if it had not been promised. But... the promisee would have good reason for such greater expectation only if (in the promisor's belief) the promised act would have such greater utility...

The promisor would know that the greater expectation would be a condition precedent for the greater utility; and so would not believe that

the act would have greater utility unless he believed that the promisor had greater expectation... the promisee would know this, and so would not have greater expectation unless he believed that the promisor believed that he had greater expectation. And this, of course, the promisor would know.<sup>26</sup>

Here - as in his discussion of truth-telling - Hodgson does not, it seems to me, describe a limited-alternative coordination problem. Hodgson postulates as implicit in the Utilitarian doctrine that

A: If an act *x* (say truth-telling or promise-keeping) has greater utility, then *x* is expected;  
and that

B: If *x* is expected then *x* has greater utility.

We can see why Gibbard's supposition of existing Act-Utilitarian conventions, and Lewis' supposition of truthfulness will not do, insofar as they are attempting to replicate a line of reasoning open to a rational Act-Utilitarian. First, Hodgson's argument is that a rational Act-Utilitarian cannot calculate the utility of an act without first knowing whether it is expected (B); but must know its utility as a prior condition of knowing whether it is expected (A). So it won't do to simply suppose that it *is* expected, then calculate its utility, as Gibbard and Lewis seem to want to do. This is to put the cart before the horse; or, as Hodgson would say, to engage in mere bootstrap-pulling.

Second, that Hodgson takes truth-telling and promise-keeping themselves as examples for discussion, rather than instances of these, as Gibbard and Lewis do, is significant (if not essential); for this means that the question of whether acts *these* acts are possible must be settled before the question of whether any acts that presuppose them as conventions are possible. Gibbard's and Lewis' treatments make it possible to beg the very question Hodgson raises: Gibbard answers the question whether A should keep his tennis date by positing a history of such agreements, while Lewis answers the question of whether A should listen to B's advice to press the red button rather than the green by assuming truthfulness to hold whenever it has best consequences. But these answers presuppose that agreements *can* be made, and that truthfulness can, under certain circumstances *that might exist in an Act-Utilitarian society*, have best consequences. But to assume these conditions begs the question of whether, in an Act-Utilitarian society, such conventions could ever arise, and this is what Hodgson denies.

That this issue cannot be solved by constructing a coordination matrix with limited alternatives can now perhaps be seen somewhat more clearly. If

---

<sup>26</sup> Hodgson, *op. cit.*, p. 41.

my account of Hodgson's argument is right, then in any confrontation between two Act-Utilitarians, neither can have *any* valid expectations about what the other will do. So any attempt to formulate the problem as a choice between two alternatives and assign weights and probabilities accordingly is bound to fail. For since both making promises and telling the truth are equally otiose, no act that depends upon the utterances of an Act-Utilitarian can have any credence, for no expectations will arise from them. And since an Act-Utilitarian society is one in which only those acts are performed which have the best consequences under the circumstances, no act conforming to the rule of Act-Utilitarianism, insofar as it depends on the expectations of others, can be performed. So no such act can be expected to be performed. This means that the possible alternatives of action open to two Act-Utilitarians in any situation are fairly unlimited. Within the constraints imposed by the particular circumstances and by human capacity, neither has *any* basis for predicting,

on probabilistic grounds, what the other will do. To the extent that doing *x* depends on its utility, *x*'s utility on whether *x* is expected, and whether *x* is expected on *x*'s utility, there can be no sufficient reason for expecting any *x* to be done. So the choice is not between, say, going to the courts and staying home -- for which a coordination could indeed be established. The choice is between going to the courts, staying home, walking the dog, breaking a window, doing a headstand, and the myriad other possibilities that exist between two Act-Utilitarians who have 'agreed' to play tennis.

Now Gibbard attempts to combat this dilemma by arguing in support of a 'teaching effect', i.e. that

Each would keep the agreement in order to teach others to expect him and other Act-Utilitarians to keep such agreements. In a society where everyone had just openly converted to Act-Utilitarianism, there would be occasions on which the act with the best consequences would be to teach others what to expect.<sup>27</sup>

He follows with specifications of the kind of agreement that would rationally be chosen for its teaching effect, e.g. that the agreed-upon act should have greatest utility only *if* it has been agreed upon; that it would be rational to perform it only if it would have different consequences when agreed upon than when performed independently of this agreement, etc. But in addition to skirting the question of how such an agreement could be made, much less kept, he seems to misunderstand Hodgson's argument. Hodgson denies that such a teaching-effect would be possible, i.e. that there would be

---

<sup>27</sup> Gibbard, *op. cit.*, p. 167



any way to *arouse* the requisite expectations among its members, on the grounds that since the good consequences of *any* act, including the act of keeping an agreement, would have to obtain independently of whether the act was expected the mere performance of the act - here, keeping the agreement - would give no reason to expect the act in the future.<sup>28</sup> Here I believe Hodgson makes a tactical error: one might argue in reply that if any act could be performed with best consequences, it would, seemingly, set a precedent for future such acts when coordination problems arose; thus we would be well on our way to a convention in the manner convincingly outlined by Lewis in *Convention*. I think that what Hodgson should have said is simply that the conditions under which an act must be performed in an Act-Utilitarian society as we have described it preclude the performance of any act which depends for its utility on its being expected, thus *a fortiori* precludes its having a precedent-setting teaching effect. This seems to me to be a valid implication of his argument. But there is, perhaps, a more circuitous route to the same conclusion.

Even if a promise were kept, say, by chance, it is nevertheless doubtful whether it would have the teaching-effect Gibbard needs. Let us distinguish between an act *x qua* act, and *qua* fulfillment of a promise; call the latter *x(P)*. Now Hodgson argues that a single performance of *x(P)* could arouse expectations of further such acts only if *x(P)* was taken to indicate that such acts can have best consequences. But *this* occurrence of *x(P)* would have best consequences only if

1. *x(P)* satisfied the promisee's expectations
2. *x(P)* aroused expectations of further such acts

We have already seen that the first condition must fail, since the promisee has no *prior* reason to expect this first promise to be kept. So *x(P)* automatically fails 2., for it cannot have best consequences under the circumstances. *x(P)* would have best consequences only if it aroused expectations, and it would arouse expectations only if it had best consequences. But here we have an independent *Utilitarian* reason for denying that it has best consequences, namely that *qua* promise, it fails to satisfy anyone's expectations. And since *x(P)* satisfies no one's expectations, it can, *qua* promise, arouse no one's expectations. This is perhaps a more perspicuous rendering along the lines of Hodgson's *ad infinitum* argument that shows why Gibbard's answer to it -- that 'sometimes keeping a promise will have best consequences because promise-keeping is expected'<sup>29</sup> won't

---

<sup>28</sup> Hodgson, *op. cit.*, p. 47

<sup>29</sup> Gibbard, *op. cit.*, p. 170.

work. The conditions under which promise keeping is expected have yet to be established. This variation on Hodgson's argument is based on its structural similarity to the previous argument against the possibility of performing any act the best consequences of which depend on satisfying someone's expectations. Substituting 'arousing expectations' for 'satisfying expectations' yields substantially similar reasoning in both cases, and reveals some of the paradoxes implicit in trying to do *only* those acts with the best consequences.

Gibbard seems to think that a fortuitous instance of promise-keeping will arouse expectations because it might, *qua* act  $x$ , have best consequences. But by hypothesis, these best consequences of  $x$  must obtain *independently* of the act's being a promised one; so while expectations might be aroused of  $x$ 's future performance when it again has best results, there is as yet no reason to expect any act to be performed as the fulfillment of a promise, i.e. as  $x(P)$ . The best consequences of  $x(P)$  must *include* arousing (and satisfying) expectations. One cannot give the best consequences of  $x$  as a *reason* for arousing expectations about  $x(P)$ . It is not as if one does  $x(P)$  as an unintended side-effect of doing  $x$ , sees that  $x(P)$  has best consequences, hence expects future promises to be kept. This would happen only if the best consequences of  $x(P)$  were identical to the best consequences of  $x$ . But they are not.  $x(P)$ , but not  $x$ , has best consequences only if it satisfies the expectations of the promisee, which is as we saw, impossible. So mere performance of  $x$  will not arouse expectations of future promise-keeping  $x(P)$ .

Now Gibbard might object that even if a coordination solution is not expected, it might nevertheless occur and be acknowledged as a solution; this might then provide sufficient reason for expecting it as a solution to future problems. This is basically Lewis argument for the origin of a convention, briefly adumbrated above; and the same objection to it is relevant. If the parties could originally *conceive* of the issue as a limited-alternative coordination problem, perhaps this line of argument would work. But we have seen that they cannot. Because a basis for expectations of other's behavior is lacking, no act possible under the circumstances has greater initial subjective probability than any other. Hence any act that may be performed cannot be regarded as a solution to the problem of whether to do  $x$  or  $y$ , since the question of what one should do cannot be made determinate in this way. So even if an act  $x$  was the solution to such a problem, the parties would not regard it as such because they would not conceive the situation in these terms. Because there could be no answer to the question of what act to perform, any act would be equally acceptable.

This is to argue, *contra* Sidgwick's claim, that the ideal Utilitarian community offers no more of a resolution of the publicity problem than does the actual one. So the Utilitarian cannot argue that his 'esoteric morality' is a

temporary practical measure and not intrinsic to the theory in its ideal realization, as Sidgwick wants to do. For in view of the implications of Hodgson's analysis, we seem forced to conclude that secrecy is a necessary ingredient in a viable Utilitarian doctrine.

#### IV

Suppose we examine an intermediate case, midway between the ideal case and the actual situation. Assume that I am a consistent and fully rational Act-Utilitarian in a largely non-Utilitarian community, and that I reveal my convictions only to my closest and most trusted friends, who are not Act-Utilitarians. How can I expect this fact about me to influence our relationship, if at all? For one thing, this openness on my part will, I hope, be beneficial to our relationship, since it will presumably structure and confirm my friend's expectations of me as a Utilitarian, hence increase the security and affection of our friendship. But will it? If my friend knows I decide what to do on grounds of utility, he will justifiably infer that my openness with him is similarly a matter of policy; that I would not do it if it did not maximize social utility. But my concern with what maximizes social utility clearly transcends the particulars of our relationship; it is this larger goal which I consistently keep in mind, and in terms of which the quality of our relationship finds warrant. And if he knows this, he knows my honesty is not merely for the sake of our friendship -- not, that is, merely for the sake of my respect and affection for him, but for something in comparison with which the independent value of our friendship pales in significance. His knowing I think this way hardly seems beneficial to our friendship.

But again: I might sincerely advise my friend to do the most socially beneficial, utility- maximizing actions, and try to show him my own consistency in this regard, and the good effects that can be brought about; perhaps this will demonstrate my essential good will. But of course this will result only if I can somehow convince my friend that what I tell him to do and demonstrate by example should be done is in fact what I want him to do. For while he will not doubt that I see *some* course of action for getting him as being best, he has no assurance that I think that the best way of getting him to do it is by advising him to do it. I might think it best to advise him to do *x* in order to bring it about that he does *y* (where *y* is either incompatible with a side-effect of *x*, or part of doing *x*). This possibility will be particularly troublesome for him when it is clear that he disagrees with my advice.<sup>30</sup> In fact, my friend's suspicion on this point may be justifiably extended to all

---

<sup>30</sup> See Sidgwick's claim on p. 490 (quoted above, p. 5).

facets of our interaction: can he ever be sure that my responses to him are not intended to get him to do or think those things I think it best for him to do or think? It seems there is no way of insuring that even the most minimal conditions of moral dialogue are met. As Strawson points out,<sup>31</sup> I may seem to engage him in moral discourse without really doing so.

As far as I can see, this failure of moral dialogue arises independently of how 'utility' is defined. For part of the Utilitarian doctrine is the view that the maximization of social utility is given priority over principles grounded in an ideal of the person or a conception of right action, human dignity, or respect, which would regulate the quality of one's interactions with others more directly and determinately. If it was found that such principles produced greater utility only when applied selectively to an elite minority, a consistent Utilitarian could not hesitate to abandon them as a general practice. Indeed, if they were found not to produce greater utility *simpliciter*, the Utilitarian would of necessity abandon them altogether. Thus social utility, however defined, must be regarded as something over and above the intrinsic value of a friendship; such interactions are merely one means among others to the maximization of utility. As Sidgwick says, 'we perceive Friendship to be an important means to the Utilitarian end.'<sup>32</sup> The only available alternative to this conclusion seems to be to define 'utility' in a way that makes the above-mentioned deontological principles constitutive of the greatest happiness, rather than a means to it. But then it seems clear that we are no longer discussing the doctrine of Utilitarianism, as it is commonly understood.

In 'Freedom, Blame and Moral Community', Lawrence Stern<sup>33</sup> argues that Strawson fails to distinguish between calculation and manipulation in his concept of the objective attitude. He defines calculation as 'subjecting whatever feelings one has to the constraints of policy, to getting the result one is aiming at', while manipulation is 'subverting or bypassing another person's rational or moral capacities for the sake of some result.'<sup>34</sup> In therapy, for example, he points out that although calculation must enter into the attitude of the therapist toward the patient insofar as the patient's wellbeing is the

---

<sup>31</sup> Strawson, P.F., 'Freedom and Resentment' in *Freedom and Resentment and Other Essays* (Methuen: 1974). I think Strawson's distinction between the objective and the involved attitudes brings out nicely the difference between the perspective a consistent Utilitarian must assume and that of most other people.

<sup>32</sup> *Op. cit.*, p. 437

<sup>33</sup> *J. Phil.*, Vol. LXXI, No. 3, February, 14, 1974. PP. 72-84

<sup>34</sup> *Ibid.*, p. 74.

result being aimed at, manipulation need not be, since the therapist can make full use of the patient's rational or moral capacities in furthering this goal. Although the distinction seems to me to be well-taken, it is important to see how closely intertwined these two must be in the attitude of a consistent Utilitarian toward anyone else. Here, calculation *implies* manipulation. For in order to promote the result the Utilitarian is aiming at, *viz.* maximizing social utility, it will be necessary to bypass the other person's rational and moral capacities just in case publicly acknowledged agreement on the goal to be achieved is lacking -- which, as we have seen, must be true for the consistent Utilitarian in all cases. For example, the Utilitarian may enter into a friendship for reasons of utility, as Sidgwick suggests; but if the other person enters into it solely because he likes and respects the Utilitarian personally, and the Utilitarian knows this, it is unlikely that the latter will succeed in bringing about a commitment to the relationship from the former except by manipulation, by setting the other person to commit himself, without openly presenting his Utilitarian calculations of how to best maximize utility as a reason for doing so. For it would be unusual, to say the least, for two non-Utilitarians to commit themselves to being friends for reasons over and above the friendship itself. On the other hand, cases in which calculation would not necessarily imply manipulation are just those cases, e.g. business relationships, in which people are consciously committed to cooperation in some enterprise (aside from the pleasures of interaction for its own sake) the goals of which are mutually acknowledged. But since mutual acknowledgement and cooperation in the goals of Utilitarianism seem to lead to insuperable difficulties, the implication holds in this case.

I believe that the possibility -- indeed, the necessity -- of a consistent policy of manipulation of others, and the calculation of their responses as variables in the service of a larger goal, reveals a serious problem with the very concept of a consistent Utilitarian doctrine. The problem is that it is not only what we normally regard as conflicts in moral behavior that is subject to Utilitarian evaluation. All modes of human interaction are susceptible to the question, 'Does this interaction/ relationship/ response/ association etc. maximize or further general social utility?' The ubiquity of the question is a function of the ubiquity of the goal. The first principle of Utilitarianism can be seen as a special case of the non-moral rationality principle of Efficient Means, in which the goal to be most expediently achieved is specified as that of maximizing social utility. Now normally the Principle of Efficient Means is called into use under circumstances that themselves determine whether or not the question, 'Does this act conduce to G?' is relevant; for most goals are such that not all actions, and not all circumstances will obviously bear on their realization. For example, if I wish to learn horseback riding, my taste for foreign films will not be a relevant consideration in trying to do so in any

obvious way. The nature of the goal G itself places certain practical constraints on that class of actions which are to be assessed for their expediency in bringing it about; hence G will not form some part of the purpose of every action an agent considers. Compare this to the Principle of Utility. Where G is 'the general social utility', this can be further fleshed out in any number of ways. What is important to note is that *any* more specific formula substituted in for it, e.g. 'Everyone's wellbeing', 'The general level of pleasure', 'The satisfaction of everyone's wants', etc. must be sufficiently general so as not to rule out the happiness, pleasure, satisfaction of wants (however these are suitably defined) in advance for some particular person, for this would be to decrease the total sum of utility. But since there is no prior way of determining what makes every person happy, satisfied, or what constitutes each person's pleasure, there is no prior way of selecting out any purposeful act occurring in a social context at all; for in either the fulfillment of its purposes or the consequences of this fulfillment, a want may or may not be satisfied, a pleasure may or may not be obtained, an expectation may or may not be violated, etc. The goal of maximizing social utility is so all-encompassing that any act performed in an interpersonal context may have consequences relevant to, or constitutive of, its realization.

Note that the contrast with the Principle of Efficient Means is merely a matter of degree. My taste in foreign films may, in theory, be relevant to my attempts to learn horseback riding. We could easily hypothesize a causal sequence, albeit an improbable one, in which this was the case. But here the goal is restricted enough so that we are justified on grounds of probability, at least, in assuming it not to be true. But the fact that it might be true here, and that it must be true relative to the Principle of Utility reflects a general feature of any purely teleological criterion for deciding what to do. G itself determines which actions are to be performed, hence provides the criterion for assessment or the worth of action. But there are no prior *methodological* conditions of relevance implicit in this criterion for which classes of actions are to be evaluated, for any such conditions could only be justified by assumptions about which acts are irrelevant to the realization of G -- and this is the very point to be determined by assessing a particular act. Usually there are conditions of relevance externally imposed, such as probabilistic criteria in the case of the Principle of Efficient Means. But this option is not open to the consistent Utilitarian. For any purposeful social act will have consequences that must be evaluated for their social utility.

This means that a concern with social utility must form some part of the *motivation* of a consistent Utilitarian in any interaction with others he engages in, indeed in any plan of action he undertakes: this is the sense in which Utilitarianism is the *only* rule of conduct for a rational Act-Utilitarian. It may

be that some such activities are then found or judged to be irrelevant to the furthering of social utility. But this can only be the consequence of an evaluation to which every action is initially susceptible. This reveals the extent to which calculation -- hence manipulation -- must inform the Utilitarian's every decision, action, and, deliberate response. So if people know that someone is a consistent Act-Utilitarian, they are bound to feel somewhat manipulated; somewhat suspicious of his manifestations of feeling, involvement with, or professed regard for them; and somewhat resentful of his attitude towards them. Clearly, it is more expedient for the Act Utilitarian that he keep his convictions a secret, both in this world and the ideal one.

## V

Now if the Utilitarian adopts the principle of thoroughgoing secrecy regarding his Utilitarian convictions, there is nothing to prevent his getting done a great deal of good. In fact, it is entirely possible that he will do better than most of us, for his actions will be more fully informed by rational Utilitarian deliberation. But he will stand in a unique and not wholly desirable relationship to everyone else in the world, whether or not they in fact share his convictions. He will, as it were, have to keep his own counsel on every occasion. He will be unable to reciprocate in certain attachments and dependencies on others, insofar as these involve commitment and trust in the form of expression or discussion of his deepest feelings and moral convictions; he will be unable to communicate them, and unable to find confirmation for or evaluation of them in the opinions of others. It is questionable how worthwhile a Utilitarian might then find this doctrine. For not only would it seem to necessitate a degree of alienation from others the psychological cost of which cannot be compensated. It also needs a rather strong, and probably incorrect, assumption about human psychology in order to be stable in the hierarchy of values of the agent. It needs, that is, to assume that a person's convictions can thrive on purely internal support; that a lack of confirmation and esteem by others of these convictions will not erode or weaken their importance and value in the Utilitarian's own mind. This is not to claim that our deepest convictions require public consensus in order to reassure us of their validity; it is just to question the sense in which moral principles can be believed to be the correct ones if they are in principle acknowledgeable by no one but the agent.

This has certain consequences, implicit in the above discussion, for how the consistent Utilitarian must regard other people. He must, without confiding them, both do what he sees as best promoting general utility, and also do what is necessary to get others to do the same. The telling asymmetry

of justification we mentioned near the beginning of this discussion thus reappears in a stronger form: the Utilitarian acts from well-reasoned motives that accord with his deepest convictions, while he requires and expects no such deliberation on the part of others. It is sufficient for his purposes that they perform the (from his perspective) requisite actions and have the requisite thought and responses. But however complex or reflective these may be, they will have no independent validity for the Utilitarian. He accords them weight only insofar as they coincide with his plan. That is, he views the opinions, feeling, and deliberations of other people as *instrumental* to his moral goals.

At the beginning of this paper I made some observations on how children are often raised, and on the rationale for doing so. It may not have been evident what this had to do with utility and publicity; now I want to try and bring out the connection. Just like the Utilitarian, parents have reasons of utility for not publicizing some of their intentions and beliefs to their children: they will be disruptive, misunderstood, have an untoward effect on psychological development, and so on. Like the Utilitarian, parents try indirectly to get their children to do things that the parents believe to be ultimately in the children's best interests. Like the Utilitarian, parents cannot require their children to make a considered judgment or mature confirmation of the validity of these beliefs. For this reason, parents -- like the Utilitarian -- can have a satisfying and affectionate relationship with their children, but do not expect to form the same complex relationship of affection, trust, dependence, and respect that is possible with a friend or an equal. Like the Utilitarian, the morally best act for a parent is often the one with the most favored consequences for others, i.e. the children: parents often feel that their beliefs and efforts will be sufficiently validated if only their children grow up to be happy, mature, and productive adults who have a minimal gratitude for their parents' efforts.

But at this point the analogy with the Utilitarian importantly fails. For we have seen that in fact there is *no* future state of things with reference to which the Utilitarian might justify his policy of secrecy and manipulation, and in light of which this policy might eventually be dispensed with and commonly validated, in retrospect, as a means to the worthwhile goal of moral maturity. That is, there is no point at which the attitude of the Utilitarian to the rest of the community can develop past the analogous attitude of the parent towards the child; no point at which the Utilitarian might eventually bear to others a relationship of mutual acknowledgement and respect as mature, autonomous, moral adults. The consistent Utilitarian, then, largely regards himself as if he were the only adult in a community of children.