



Environment & Society



White Horse Press

Full citation:

Paden, Roger, "Urban Planning and Multiple Preference Schedules: On R.M. Hare's 'Contrasting Methods in Environmental Planning'." *Environmental Values* 8, no. 1, (1999): 55-73.
<http://www.environmentandsociety.org/node/5764>

Rights:

All rights reserved. © The White Horse Press 1999. Except for the quotation of short passages for the purpose of criticism or review, no part of this article may be reprinted or reproduced or utilised in any form or by any electronic, mechanical or other means, including photocopying or recording, or in any information storage or retrieval system, without permission from the publisher. For further information please see <http://www.whpress.co.uk/>

Urban Planning and Multiple Preference Schedules: On R.M. Hare's 'Contrasting Methods in Environmental Planning'

ROGER PADEN

*Department of Philosophy and Religious Studies
George Mason University
Fairfax, Virginia 22030-4444, USA*

ABSTRACT: This essay presents a critical analysis of Hare's article 'Contrasting Methods in Environmental Planning'. It argues that Hare has drawn an important distinction between two 'methods' used in both urban and environmental planning, and that Hare is correct in the conclusion of his argument that one of these methods, 'the trial-design method', is superior to the other, 'the means-end method'. However, this paper presents a new argument in support of that conclusion. This new argument is important for two reasons. First, it points to the existence of at least two different kinds of preference schedule. Second, it supports a type of decision making procedure to be used in 'multiple-client situations' different from the one envisioned by Hare. This procedure, oddly enough, resembles the procedures outlined by both Habermas and Rawls. However, it can be defended on recognisably utilitarian grounds.

KEYWORDS: Hare, Rawls, Habermas, urban planning, design, preference schedules, utilitarianism

I.

Sometimes we must live with the effects of our choices for many years. The kinds of choices made by urban planners as to how we should modify our physical environment, if indeed we should modify it at all, are among those with especially long-lasting consequences. These choices not only affect the quality of our experience – or, as many utilitarians would say, the degree to which our preferences are satisfied – for many years to come, but, because they also help shape the kind of people that we (and perhaps our children) will become, they play a significant role in determining the kinds of preferences that will be formed in the future. Choices of this kind, that is to say, not only have a direct effect on the quality of our experience, but they also have an indirect effect on the standards we use to make choices. Through the design of our environment, we also – for good or ill – indirectly design ourselves.

Environmental Values 8 (1999): 55–73
© 1999 The White Horse Press, Cambridge, UK.

Because these choices have such wide-ranging and long-lasting effects, it would seem especially important that they be made correctly. At the very least, it would seem that they should be made 'rationally', that is, in light of the available facts (including facts about our preferences) and with a detailed understanding of the impact that these decisions have on the future (in particular, on how they will affect our ability to satisfy future preferences). Given their complexity, it would seem best if we could approach these decisions systematically – ideally, through a method that would be designed to help us make fully rational decisions.

In 'Contrasting Methods of Environmental Planning', R. M. Hare attempts to distinguish and evaluate two such methods, which he calls, the 'means-end' method and the 'trial-design' method (Hare, 1989). These two methods, I believe, should be understood as 'ideal types' which define the range of methods open to the planner. Hare also distinguishes two professional situations within which the planner must work. In one of these situations, which I will refer to as 'single-client situations', the interests of only one individual will be affected. In the other, 'multiple-client situations', the interests of many people will be affected. Hare argues that, ideally, the trial-design method should be used in both situations, although in multiple-client situations, this method needs to be supplemented by a kind of cost-benefit analysis in order to ensure that the interests of all people are treated justly. I believe that Hare is correct both in the importance that he attaches to his methodological distinction and in his conclusion that the trial-design method is to be preferred in both situations. However, I believe that there are even better reasons than those discussed by Hare for preferring that method. In this essay, after discussing Hare's distinctions and arguments, I will present my own argument as to why the trial-design should be preferred in both situations. I will then use this argument as a basis to characterise some of the duties of the urban planner and to show how a planner might best discharge those duties.

II.

Given that Hare is a preference utilitarian (Hare, 1981), he takes the goal of planning to be the overall maximisation of preference-satisfaction. Given this goal, he argues that, to be counted as rational, any method used by environmental planners must distinguish values (preferences) from facts. This needs to be done because these two elements must enter into rational planning decisions in different ways. Both of the methods he considers make this distinction, but they differ as to how they allow values to enter into the decision making process. The

URBAN PLANNING...

two methods can be easily distinguished by reference to the design professions that make the greatest use of these methods.

Engineers typically use the means-end method, working as follows: First, the engineer is assigned some goal, together with some design specifications: Build a bridge spanning this canyon, capable of carrying some predicted load, in a way that does not obstruct a valued 'scenic view', and at the lowest possible price. The engineer's job is to take these specifications (values or preferences) and combine them with her knowledge of the site and of materials and design elements (facts) in order to design a bridge which satisfies the specifications in the most efficient way possible. This method could be represented by the following diagram, which I have adapted from Hare's essay:

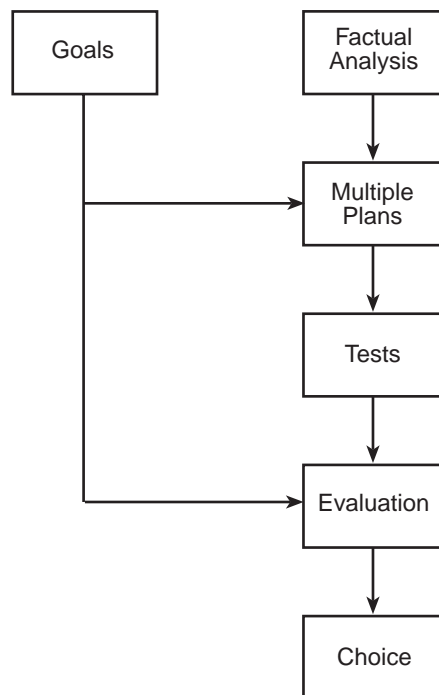


FIGURE 1. Means-end method

This diagram represents both the temporal and the logical dimensions of the decision making process. Time is represented vertically: The higher the box in the diagram, the earlier the stage in the decision making process. The horizontal separation of 'facts' and 'values', on the other hand, is intended to represent the logical distinction between facts and values. As is indicated by this diagram, in the means-ends method, values are determined separately from – and, as it were, simultaneously with – the facts, at the beginning of the project. Next, the engineer is to use her knowledge of what can be done, together with her knowledge of her client's values, to develop a number of alternative plans. These plans are then subjected to several tests. Finally, the engineer uses the results of these tests, together with her knowledge of her client's values, to evaluate the various plans. At this point she decides which plan to adopt. Of course, this diagram represents an idealisation of the decision making process, and, in the real world, various stages might be skipped or remain implicit.

Architects, according to Hare, typically use the trial-design method. According to this method, after some discussions with her client concerning his initial goals (values), an architect uses her knowledge of structures, materials, and zoning laws (facts), to rough out several preliminary sketches ('trial-designs') that will be presented (along with a rough cost-estimate for each plan) to the client who, in light of his preferences, will select one of the preliminary design for full development. Only then does the architect produce the final detailed plans from which the project will be built. This method could be represented by the following diagram, which I have also adapted from Hare's essay:




FIGURE 2. Trial-design method

I believe, however, that this diagram actually obscures the details of this decision making procedure. Therefore, I would like to substitute the following diagram for it:

URBAN PLANNING...

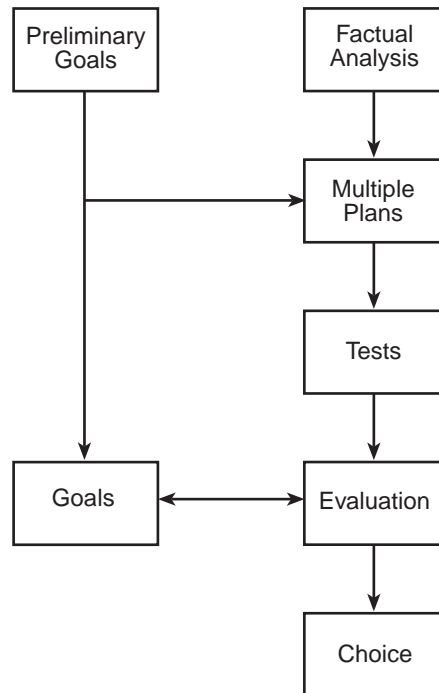


FIGURE 3. Modified trial-design method

The advantage of this last diagram is that it not only visually retains the distinction between facts and values found in the first diagram, but it also reveals more clearly the various steps that (ideally) make up the trial-design method.

Together, Figures 1 and 3 reveal that the difference between the two methods lies in how they incorporate values into the decision making process. As the two diagrams indicate, there are actually two points at which values enter into this process. Both the engineer, using the means-end method, and the architect, using the trial-design method, must make an original determination of the client's goals at the beginning of the process. (Clearly, however, the engineer must do this in much greater detail.) At first, both use these values in the same way – to select a limited set of alternative plans for initial development and testing. Following this initial evaluation, however, the two methods diverge. In the means-end method, the engineer, using her knowledge of the client's values, would evaluate the results of these tests in order to decide which plan will

maximally satisfy the client's preferences. She would then select and further develop that plan, from which the project would be built. In the trial-design method, on the other hand, the architect (or, in a competition, a number of architects) would present several alternatives together with the test results (if any) to the client, at which point the client would make his final decision. At this point, the client is free to apply whatever standard he wishes to this final decision. No doubt, the values he uses to make this decision might be identical to his original values (and he might even feel some social or economic pressure to use the values he originally outlined), but in this method he is free to change his mind even at the last minute about what he believes important and make his choice accordingly. In fact, it should be expected that the client will change his preferences in light of his encounter with the architect's various trial designs. (Thus, the two headed arrow connecting 'Goals' and 'Evaluations' in Figure 3.) Only after he has made his decision would the architect develop the detailed plans from which the project would be built.

As this comparison indicates, the two methods are remarkably similar. Nevertheless, these methods *are* different, and significantly so. Most obviously, the two methods differ on their answers to a number of 'political' or 'professional' questions (such as who – the designer or the client – is to make the final decision; or, more exactly, on when the client should cede final control of the project to the designer). As a result, these answers partially determine the professional identities and styles of practice of both architects and engineers. More importantly, however, these methods differ on the question of when the determining values or preferences should be fixed. In the means-end method, the values are determined *prior to* and *independently of* the plans, while in the trial-design method, values are finally fixed only *after* the preliminary plans have been sketched out. I would argue that the answers to the 'political' questions – and, therefore, the professional identity of architects and engineers – depend in large part on the answers each profession gives to this more fundamental value question. For example, it is because, in the means-end method, values are determined in advance, that the engineer can proceed on her own initiative from the development of alternative plans to the selection of the final plan. Indeed, the professional authority, status, and *modus operandi* of the engineer rests on the claim that she possesses a specific technical skill which her clients lack, that of efficiently matching means to ends. That this is a technical skill is demonstrated by the fact that, ideally, it could be performed in a mechanical, even deductive fashion. As Hare puts it, quoting one advocate of the means-end method,

the ultimate extension of this idea ... would be one in which a computer would be given a statement of goals ... together with [a statement of] the facts, and then it would ... produce the best plan (Hare, 1989: 222).

On the other hand, according to the trial-design method – the method associated with the practice of architecture – because the goals are fixed only after the plans

URBAN PLANNING...

have been developed – ‘The crucial evaluation comes at the end of the process, not at the beginning’ (Hare, 1989: 224) – there can be no mechanical movement from the development of various alternatives to the final plan. As a result, the architect, unlike the engineer, cannot claim the authority to make the final design decision and, therefore, her profession must be structured differently from the engineer’s.

III.

Hare believes that the trial-design method should be used by urban planners. He begins his argument in support of this position with a consideration of single-client situations. He develops his argument by means of an extended example in which a single individual (who is in this matter completely self-interested) hires an architect to redesign his kitchen. If a planner hired to do this job is to employ the means-end method, it would, of course, be necessary for her to determine – at the beginning and ‘once and for all’ (Hare, 1989: 221) – her client’s goals. However, Hare argues, to do this, she would have to do several things, all of which are problematic: she would have to make a list of her client’s goals, which, no doubt would include such things as ‘convenience, economy, beauty, hygiene, and so on’ (Hare, 1989: 218); she would have to determine how much of each goal the client desires (e.g., how convenient does he want the kitchen be); she would have to determine the relative value of each goal so that she can make the correct decisions when it comes time to make the inevitable compromises; and finally, she would have to come up with the design that would maximally satisfy her client’s preferences.

The fact that the client inevitably has a number of goals which cannot all be fully and simultaneously realised by a single design requires the planner to decide which of these preferences is to be satisfied or, more likely, the degree to which each will be satisfied. Therefore, if the planner wants to maximally satisfy her client’s preferences and if she wants to make her decision rationally, she must determine the relative importance of these values to her client. To do this, as Hare puts it,

the goals have ... to be stated in very simple [commensurable] terms, and such that the extent to which they are realized is not only quantifiable but quantifiable in a way that enables us to compare the realization of one goal with that of another on a common scale (which in practice has to be that of money) (Hare, 1989: 222).

Only if the planner does this can she reliably make the inevitable trade-offs in such a way as to maximise preference-satisfaction. Thus, the decision to adopt the means-end method commits the rational planner to making use of what is known as ‘cost-benefit analysis’. Unfortunately, Hare argues, there are many problems inherent in that form of analysis. For example, as Hare notes, many

values, such as beauty, resist quantification into monetary units and, as a result, in practice are often simply omitted from the decision process. Moreover, even with respect to those preferences that can be so quantified, because the planner can only estimate the object and strength of her client's preferences indirectly (by deductions based on choices her client made at an earlier time), she will often make errors in characterising those preferences (Hare, 1989: 224-225). As a result, the use of this method will often lead the planner to make a wrong decision and, therefore, to fail to maximally satisfy her client's preferences.

The superiority of the trial-design method in this situation lies, according to Hare, in the fact that it does not require the planner to undertake the difficult and problematic task of measuring a client's values in monetary terms. In fact, this method does not require the planner to represent these values with any precision at all. Instead, after gaining a general appreciation of a client's values, the planner proceeds to draft a number of plans from which the client will make his choice. At no time does the planner have to state the client's values in commensurable units. In fact, in this method, the client, himself, does not have to state his values in such units. Instead, he simply makes his decision. However, given that the individual is directly aware of his preferences and given that all people are rational preference-maximisers, it follows that, if the client is fully apprised of the effects of each plan, he will inevitably choose that plan which will maximally satisfy his preferences. Therefore, because the trial-design method does not make use of cost-benefit analysis, but instead allows the client to make the final decision directly, it is not subject to the problems inherent in the means-end method. Therefore, in single-client situations, the trial-design method should be used.

Hare also argues for the use of the trial-design method in those multiple-client situations in which many people with conflicting values are involved – as they always are in large-scale planning projects. In effect, Hare argues that multi-client situations require a complex two-stage decision making process. In the first stage, the preferences of each affected individual must be determined. As Hare has already argued, this is best done by using the trial-design method. It would be best, therefore, for the planner to develop a number of trial-designs that could be presented to the public so that each individual could make his or her own decision. Once these individual decisions are made, it is then necessary to combine them together to make the final political decision. Of course, as a preference utilitarian, Hare holds that the best and most just decision will be the one that maximises overall preference satisfaction. The problem for Hare is to develop a method of making political decisions that will reliably lead to this result.

One political decision making procedure that seems likely to accomplish this goal is some kind of participatory democracy. Of course, one common understanding of democratic voting procedures is that they serve as a way of determining the will (the 'preferences') of the people. Thus, submitting planning

decisions to a public referendum (after everyone has had a chance to make his or her own decision between the several trial-designs developed by the planner), can be understood as a way of determining the public's overall preference. Unfortunately, as Hare points out, democratic procedures are notoriously insensitive to the strength of each individual's preference: While each vote can indicate an individual's overall preference, it cannot indicate the strength of that preference. As a result, a widely-but-in-each-case-weakly-favoured option would always receive more votes than a less-widely-but-in-each-case-very-strongly-favoured option. Unfortunately, in these cases the selection of the former option by democratic vote would not maximise preference-satisfaction. Hare also presents two other objections to democratic procedures. First, he argues that many people are often not very knowledgeable with respect to planning issues and that, therefore, their votes will not reflect their 'informed preferences' (however slight). Second, he argues that democratic decision making procedures cannot formally take into account the preferences of the young or of those members of future generations. For these reasons, democratic decisions cannot be expected to reliably maximise overall preference satisfaction.

As a result, Hare argues, the planner must act 'paternalistically', if she can 'get away with it' (Hare, 1989: 232) in order to ensure that preferences are maximally satisfied. To do this, it will be necessary first to determine the relative (interpersonal) strength of each person's overall preference. As this is what cost-benefit analysis attempts to do, ideally, it must play a role in planning in multiple-client situations. Therefore, Hare argues, cost-benefit analysis should be employed as a supplement to the trial-design method to help ensure that the correct design is chosen (Hare, 1989: 230-232). Hare does not outline any procedure that would incorporate both cost-benefit analysis and the trial-design method, but several procedures could be imagined. One such procedure would make use of 'focus groups'. In this procedure, a design team would first attempt to determine which constituencies would be affected by a proposed project. They would then assemble a set of focus groups to represent these constituencies, and, in discussions with these groups, determine the preliminary values that would guide the development of the trial-designs. After testing these designs, the planners would return to the groups, explain the various trial-designs in detail in order to fully educate the groups on the likely long-term effects of each design, and then ask each focus group to rate each design. These ratings, together perhaps with other information about each constituency's past choices, would be used to determine the strength of their preferences. Having thereby determined the strength of each group's preferences for each design, the design team could then determine which design would, if implemented, maximise preference-satisfaction, in effect, by multiplying the strength of each group's preferences by the size of the constituencies represented by that group. The design team would then have to convince the municipal decision making body that the design chosen by this

method is the best option. Alternatively, the planner can ensure the correct outcome by requiring the municipal decision making body to give the team the power to make the final decision as a condition of doing the work. Although Hare might argue about the details, it is clear that he believes that a procedure combining the trial-design method with cost-benefit analysis would be superior to either the means-end method or a purely democratic trial-design method.

IV.

In this section, I will offer another kind of argument for the use of the trial-design method. I will begin by developing an alternative account of why the trial-design should be used in the single-client situation, and, like Hare, I will develop my argument through the use of an extended example involving the remodelling of a kitchen. Suppose that I hire a means-end oriented architect to help redesign my badly rundown and outdated kitchen. In order to gather data on my preferences, she asks me to record, over the space of a month, my preferences concerning the kitchen. To do this, she suggests that I keep a pad and pencil in the kitchen and keep a running account of the preferences I form as I putter about. What would such a record look like? Clearly, it would be shaped largely by my day-to-day experiences in the kitchen: I note angrily, with sponge in hand, that it is hard to clean; on another day, I realise that the refrigerator is a bit far from the cupboard; one night, I am struck with how difficult it is to set the flame on the stove; later, I note that it lacks a phone or a microwave oven. At the end of a month, the architect could take my comments, compile them to indicate the relative strength of my various preferences, and, using her professional knowledge, design a kitchen so as to maximally satisfy this revealed preference schedule.

Now suppose that I hire a trial-design oriented architect. Such an architect would also begin by discussing my situation with me. However, that discussion might take a rather different course. Rather than focusing on the problems with my old kitchen, the discussion with this architect would focus on my evaluations of a broad range of alternative designs. It might proceed with the help of books or magazines illustrating alternative possibilities or the architect might simply try to draw out alternative ideas from me. Following this discussion in which the architect tries to learn more about me, as it were, than about my problems with the old kitchen, she would develop several more detailed, but still rough, alternative designs, explain them to me, and then give me a few weeks to pick a favoured design. This method, I believe, would cause me to approach the decision in a different manner. Rather than focusing on the problems of my old kitchen, this method would lead me to consider the opportunities afforded to me by some new possible kitchens. What, I would ask myself, would it be like to live in this kitchen rather than that one? This one looks efficient; I would not have to spend much time in it and could devote more of my time to my work. That one

is cheap; I could get it and go on a vacation with the money I would save. But, suppose, in the end I choose the more spacious, but more expensive and less efficient, kitchen that would allow me to cook for company; the one that would allow other people to join me in the preparation of large dinners or intimate brunches. It may be that I choose that one because the contemplation of alternatives required by this design process has made me aware of a (long-forgotten?) possibility – pushed aside by the daily struggle to become more productive – of the social and aesthetic pleasures of entertaining. In contemplating these designs, that is to say, I came to the realisation that I had become an uninteresting, fast-food-wolfing workaholic. Recognising that I am not the person that I want to be, I decide that a change in my life is in order, and resolve to use this decision, not just to redesign the kitchen, but as a first step in redesigning myself. I decide, that is to say, that I have become overly concerned with work and the values that it allows me to satisfy directly or indirectly (e.g., financial security, high professional reputation, achievement, ...), at the expense of another kind of life that allows for the satisfaction of a different set of values (e.g., stimulating social intercourse, sensual enjoyment, relaxation, ...); values which I had come to overlook in the rush of existence; values which I would prefer to see manifested in my life; values that I judge to be more worthy.

Although it can go very wrong – and despite the fact that the means-end method offers its own important insights – I believe that the trial-design method is to be preferred to the means-end method in single-client situations, not just because it is more likely to lead to maximal preference-satisfaction, but also because it encourages clients to make their decisions on the basis of the ‘right’ set of preferences. Assuming for the sake of argument that both worked perfectly after their own fashion, I would argue that the use of either method would lead me to maximise my preference-satisfaction. However, with the means-end method, I would have based my decision on – and maximally satisfied – what I will call (following Charles Taylor [1985]) my ‘brute’ preferences, while with the trial-design method, I would have made it on the basis of – and maximally satisfied – what I will call my ‘self-reflective’ preferences. Of course, I do not mean to imply by the use of these terms that the preferences which determined the first decision were not based on any form of reflection. Obviously, they were. However, unlike self-reflective preferences, they were not based on reflections about the *kind* of self or *kind* of person I want to be. Of course, what kind of person I am depends, in large part, on what kinds of values or preferences I have. Therefore, when I reflect on what kind of person I want to be, I must reflect on the kinds of preferences I want to have. Preferences that are endorsed by this kind of reflection are ‘self-reflective preferences’. Preferences that have not been so evaluated are ‘brute’.

Brute preferences differ from self-reflective preferences in many ways: They ‘feel’ different; They enter into the decision making process in different ways; And they play different roles in our lives. Brute preferences, for example, seem

to act upon us 'from the outside'. To say that a brute preference – it seems better to speak of a 'desire' in this context – is 'strong' is to say that I can detect its power through an introspective examination of its effect on me. The desire 'feels' different from other 'weaker' desires which it can override. It is more 'compelling'. At the limit, it 'sweeps' me in front of it; I feel that I lack the 'power' to resist it. This kind of strength is a (relational) quality of the desire itself, such that, through some sort of introspective examination of it (perhaps in comparison with other competing desires), I can come to an appreciation of its intensity. The kind of desires to which we most readily attribute this kind of strength are those desires that are associated with the satisfaction of somatically based drives, e.g., those for sexual gratification, food, nicotine, etc.

On the other hand, self-reflective preferences, unless they are uncommonly robust, seem to lack this 'felt' dimension almost entirely. They are, we might say, 'colourless'. It would be odd to speak of a self-reflective preference, for example to be more considerate of other people in the future, as 'compelling' me to act kindly. For that matter, it would seem odd to attribute any such perceivable energetic quality to it at all. This is not because these preferences are essentially 'weak' in relation to brute preferences (although brute preferences will often override self-reflective preferences). Rather, it is because self-reflective preferences are qualitatively different from brute desires. They are more like judgments – and, as such, seem to arise from our thought – than they are like drives that impose themselves on our will. Because they are akin to judgments, they lack the phenomenal qualities of brute desires. It would, therefore, be something very much like a category mistake to attribute an experiential dimension such as 'felt strength' to them. To say that a self-reflective preference is 'strong', is to place it within a system of similar preferences and judge it to be highly important. The strength of such a preference is revealed not by close introspection, but by careful reflection. A preference is strong, in this sense, if upon reflection it is judged to be important, significant, or particularly worthy. Such a reflection is not based on an examination of some qualitative attribute of the preference itself, rather such reflection locates the preference within a system of preferences that together constitute a vision of a worthy self. Thus, I might have a strong preference to be more kind if I accept an ideal of self within which benevolence plays a significant role. It is in terms of this ideal that my self-reflective preferences find their place; through it, that they become commensurable; and from it that they gain their strength.

It is possible, therefore, to talk about two distinct kinds of preference schedules. It is possible to construct a 'brute preference schedule' upon which brute preferences could be ranked according to their phenomenal strength. Of course, the strengths of these preferences change quite rapidly, declining precipitously when satisfied only to re-emerge later according to some underlying biorhythm. However, it would be possible to rank them in a prudent fashion, in terms of their long-term power. On the other hand, self-reflective preferences,

because they are of a different type would not appear on this preference ranking at all or, if they did, they would typically appear to be extremely frail. On the other hand, brute preferences, even prudent ones, may be judged unimportant, 'weak', on a self-reflective preference schedule. One may, in fact, have strong self-reflective preferences that a brute preference not be satisfied and vice-versa. These two preference schedules, that is to say, not only rank different kinds of preferences, but they may often recommend conflicting actions.

An almost universally accepted intuition holds that it is better to lead a life based on self-reflection. Taylor has suggested that this ability to lead a life based on self-reflection is what distinguishes human from animal life (Taylor, 1985). It has also been suggested that it is just such self-reflection that marks the difference between wisdom and clever foolishness. John Rawls has gone so far as to incorporate this intuition into his definition of happiness, according to which a 'person is happy ... during those periods when he is successfully carrying through a rational plan [acting, that is, on the basis of a self-reflective preference schedule] and he is with reason confident that his efforts will come to fruition' (Rawls, 1971: 550). Although it is possible to raise some objections to Rawls's notion of a 'life-plan', it would seem true that one does lead a happier life to the degree that one makes successful decisions based on one's self-reflective preference schedule. Surely, such a life would have a direction and significance lacking in a life of 'satiation', that is a life dedicated to the maximal satisfaction of brute desires. However, although we may believe that we would be happier if we made our choices on the basis of our self-reflective preference schedules, we often find ourselves choosing to satisfy our stronger brute preferences. It is for this reason that we often try to make important decisions, not in the heat of the moment when stronger brute preferences may dominate, but upon cool reflection when our self-reflective preferences can assert themselves. In order to improve our chances of acting on the basis of our self-reflective preferences, we often design our lives to lessen the chance of a direct confrontation between our brute and self-reflective preferences. We adopt strategies rationally designed to prevent us from becoming aware of our brute preferences. In effect, we try to avoid having to fully represent those preferences to ourselves. Such strategies may, for example, lead us to associate only with non-smokers, or to avoid passing bakeries, or to put off the purchase of a TV. By adopting these strategies, we can avoid the influence of our brute preferences in the short-run. Perhaps, in the long-run, these strategies might even cause us to lose the desire to smoke or watch television. On a larger scale, if we make design decisions correctly, by basing them on our self-reflective preferences, we might be able to train many of our brute preferences to be more in line with our self-reflective preferences. In effect, we can use these decisions to redesign ourselves so that we develop more harmonious, more worthy selves. This, of course, requires that we take a paternalistic attitude towards our brute preferences, but, if we are successful, the result will be a 'happier life', as Rawls has defined it.

I would argue that it is because the trial-design method puts clients in situations in which they are more likely to make their decisions on the basis of self-reflective preferences, rather than on their brute preferences, that this method is to be preferred. By removing the decision from the daily pressures of brute preferences and allowing clients to contemplate their lives as a whole, the trial-design method allows clients to reflect on the course of their lives, thereby helping them make their choice on the basis of their self-reflective preferences. Although choice on the basis of self-reflective preferences cannot be guaranteed by this method, it can be understood as a kind of rational strategy designed to make such a choice more likely. To the degree that it succeeds, it will help clients live happy lives.

On the other hand, it is likely that the means-end method will produce brute satiation rather than happiness. This method is designed to maximise the satisfaction of those preferences which we have prior to the beginning of the design process, which I will call our 'revealed preferences'. However, we tend to reveal our preferences in a promiscuous manner: often, we act on the basis of our brute preferences; occasionally, we act on the basis of our self-reflective preferences. No doubt, the precise mix that is revealed differs from person to person, however, the means-end method cannot distinguish between these two kinds of preferences. Moreover, as our brute preferences tend to emerge in the rush of life – revealing themselves, while concealing our self-reflective preferences, as they did in the first kitchen example above – the means-end method will tend to be biased in favour of the 'wrong' kind of preferences. Thus, even if the measurement problems which Hare argued beset this method could be overcome so that it would reliably lead us to make choices that maximise our preference-satisfaction, it would tend to maximise the satisfaction of our brute preferences instead of our self-reflective preferences. As a result, while we would be 'sated', we would not be happy. The superiority of the trial-design method in single-client situations, therefore, lies not in the fact that it allows us to base our decisions *directly* on our preferences, but rather its superiority lies in the fact that it allows us to base our decisions on the *proper* preferences. Because of this, it leads us to make wise choices that, with luck, will allow us to live happy lives. Of course, I would not argue that this method is foolproof. It can lead to disastrous results. Therefore, it is important that the architect take great care to investigate the client's initial preferences, present a set of trial-designs that reflect these preferences, and avoid designs that are impractical. This requires that the architect be both disciplined and creative, willing to explore new alternatives and develop creative designs, while at the same time willing to work within the limits imposed by the client's self-reflective preferences and by current construction technologies. Although, even with all this, mistakes can be made, this method would seem to be most likely to produce the best possible results.

V.

There are two ways that the distinction between self-reflective and brute preferences can be used to construct arguments concerning the proper procedure to be used in multiple-client situations. As these arguments point to significantly different decision making procedures, it is important to keep them separate. The first closely parallels the argument developed above. According to this new argument, if the means-end method is applied to multiple-client situations, it will lead to decisions that, at best, will fairly maximise a community's revealed preferences. However, because this is not equivalent to procedures that will fairly maximise a community's self-reflective preferences, this result will not necessarily lead to a general increase in happiness. For example, Mark Sagoff has argued that the dependence on cost-benefit analysis in the making of planning decisions, characteristic of the means-end method, tends to promote the 'Americanisation of society'. It tends, that is to say, to promote the development of a kind of architecture and land-use patterns that promote the efficient gratification of our more 'base desires' at the expense of our 'higher values', and, as Sagoff points out, life within these structures is typically unsatisfying (Sagoff, 1988: 99-123). If my analysis is correct, this tendency is a result of the fact that the means-end method does not encourage self-reflection and, therefore, gives unwarranted weight to our brute preferences. If this is the case, then the means-end method should not be used in multiple-client situations. Instead, the trial-design method should be used in these situations because it encourages clients to base their decisions on self-reflective preferences. However, as Hare has pointed out, in order to ensure that this method actually maximises a community's self-reflective preferences, it is necessary, in multiple-client situations, to measure the strengths of the preferences held by different people, and to do this, it is necessary to make some limited use of cost-benefit analysis. It is clear that, if it were accepted, this argument would lead to the use of a 'paternalistic procedure' like the one discussed above.

Although this first argument makes use of the distinction between brute and self-reflective preferences, with its emphasis on maximisation and its use of cost-benefit analysis, it is still a recognisably utilitarian approach – albeit of a type closer to Mill's theory than to Hare's. The second argument diverges more radically from Hare's approach. This argument begins with the distinction between brute and self-reflective preferences, but it draws a further distinction between two different kinds of self-reflective preferences. This distinction is based on a distinction between 'private' reflection and 'public' reflection. According to this distinction, a 'self-reflective personal preference' is a preference concerning the kind of person I want to be, while a 'self-reflective political preference' is a kind of preference concerning the kind of society we want to form. Thus, a self-reflective personal preference differs from its political

counterpart in two ways. First, while the former is concerned with an ideal person, the second is concerned with an ideal society. Second, while the former can be held by a single individual, the latter is shared. Furthermore, a self-reflective political preference will be said to be 'well-formed' if it results from a process of uncoerced deliberation, such as describe by Habermas in his discussion of the 'ideal speech situation'. Given this distinction, it is possible to speak of two different self-reflective preference schedules, a political preference schedule and a personal preference schedule.

I would argue that, generally, political decisions based on well-formed self-reflective political preferences are superior to decisions based on other kinds of preferences. Because these decisions would be based on self-reflective preferences rather than brute preferences, they would tend to promote happiness rather than satiation. Because they would be based on political preferences rather than personal preferences, they would tend to promote a sense of social unity and a feeling that the society is a joint project. Finally, to the degree that they are well-formed, they would appear (at least) to be just. Therefore, a society in which design decisions are based on well-formed self-reflective political preferences would, all things being equal, be better than a society in which design decisions are based on other types of preferences.

Unfortunately, ideal speech situations are, by nature, ideal; and, therefore, it would be impossible to actually make design decisions based on truly well-formed self-reflective political preferences. However, it is possible to design a decision making procedure that would produce decisions that would be based on something very much like these preferences. One such procedure would make use of what might be called a 'deliberative group'. Unlike the focus groups described above, this (one) group would be composed of representatives of the various different constituencies that would be affected by the proposed project. Like the focus groups, the members of a deliberative group would be assigned the task of developing a set of preferences or values, acceptable to all the members of the group, that would define an ideal city, with an emphasis on those preferences concerning that aspect of such a city that would be affected by the kind of project envisioned. In addition to producing a list of these preferences, however, the group would be asked to develop a set of arguments they believe could be used to justify the selected preferences to each other and to the members of the various constituencies they represent. It would be stressed to the group that, because they are not yet in the process of selecting a design, there is no requirement that the preferences they select form a consistent whole or that all the underlying arguments begin with a consistent set of assumptions. Moreover, it would be stressed that the group is to try, as far as possible, to select general preferences or values that would apply to any project of this sort: they should not think exclusively about the specific project at hand, but instead, they should decide as if the preferences they select will influence an indefinite number of such projects. The design team would then take these preferences and the

underlying arguments and use them to guide the development of a number of trial-designs, tailored to the specific project at hand. These trial-designs would then be tested in a number of ways to discover how they would function in the real world. They would then be resubmitted to the deliberative group, together with a set of descriptions developed by the planning team on the basis of these tests of what it would be like to live in a community shaped by each design. These descriptions should be as realistic and vivid as time constraints permit and might include computer simulations, although they would always focus on the nature of the lives made possible by the various designs. The group would then be asked not simply to choose which design they prefer, but to develop arguments as to why that design is to be preferred. The arguments would not only seek to justify one design by arguing that it, more than any other, satisfies the preferences originally chosen by the group, but they might also seek to justify a design by arguing that those original preferences should be changed in light of some factor that became apparent through the consideration of how the various trial-designs fit the situation at hand. The group's final decision (or decisions, if they cannot achieve unanimity), together with the arguments that they developed and the trial-designs they rejected, would then be submitted to the political body empowered to make this decision. Presumably it would make its decision on democratic grounds.

This procedure is designed, not to ensure that the outcome maximally satisfies all the community's preferences, but to ensure as far as possible that the decision is made on the proper grounds. The function of the deliberative group is not necessarily to make the final decision. Instead, it functions primarily to help select, highlight, and develop a set of political preferences to be used in making the final decision, and to provide a model of deliberation for those who will make the final decision. The hope is that if the group represents all aspects of the community and if they deliberate well, their findings will be difficult to set aside. Moreover, because they have developed and refined a set of arguments that appeal to the representative members of the group, and through them to the community at large, their deliberations will define the limits of the argument. Although the political decision making body would be free to ignore the group's work and recommendations, depending on how good a job the group did, this would be very difficult to do. Moreover, any debate about the group's recommendations would have to take those recommendations and arguments as a starting point, and might serve, in many cases, to further the deliberative process.

This procedure differs significantly from the paternalistic procedure described above. In particular, in this procedure only one advisory group is formed, and it is composed of representatives of all the different constituencies involved. Second, the function of this group is not to make a decision or even to express a preference. Instead, the group is to engage in a deliberative process and report on the substance of those deliberations. In effect, it is to act as a deliberative model, one that represents its community at its best. In addition to these practical

differences, this procedure is based on different theoretical grounds. The underlying theory, although still recognisably consequentialist and centred on preferences, is not utilitarian: this procedure does not seek to maximally satisfy all preferences, but instead seeks to ensure that decisions are made on the bases of the proper set of preferences. As a result, I believe, there are several reasons to prefer this approach. First, since it does not seek to maximise preference-satisfaction, it does not face the measurement problems inherent in utilitarian approaches to multiple-client situations. Second, this procedure is not as paternalistic as the other: although it seeks to limit the choices open to the controlling political body and to shape the debates of that body, it leaves the final decision in that body's hands to be made democratically. (Not only is this less paternalistic, but, in today's world, it is probably the only real alternative open to a planning team.) Finally, this procedure is free of one objectionable aspect of utilitarianism, namely its identification of justice with the maximisation of preference-satisfaction. Because, it does not define justice in this way, it would not require that the less-fortunate sacrifice to make the more-fortunate better-off, even if that increased utility. On the contrary, as the deliberative group resembles to some degree the parties in Rawls' original position, it might be expected that the design they choose would protect the interests of the least-fortunate. For all these reasons, I would argue, this procedure should be used in multiple-client situations.

VI.

It is interesting to note that in developing my arguments, I have gradually shifted from talking about the methods a 'planner' should use, to talking about the methods a 'design team' should use. This reflects a changed understanding of the role of the planner. If my arguments are sound, a good planner is not simply someone who possesses great technical skills, nor simply someone who can create designs that both satisfy the client's expressed desires and are intrinsically beautiful. In addition, a good planner must be able to identify the constituencies that will be effected by a project; form and facilitate the work of deliberative groups; develop trial-designs that are consistent with general values, but appropriate to specific contexts; imagine possible futures in detail; present, not just designs, but arguments to political bodies; help guide the resulting discussion; and, finally, develop the selected trial-design. The planner, therefore, must, exemplify the virtues of many professions, and it is unlikely that one person can do all this. If my argument is sound, however, a good planner must also exemplify the virtues of a good practical philosopher, someone who is able to uncover the structure of values which inform a community's practical reasoning. As such, a planner must not only be able to analyse arguments in order to uncover their basic assumptions, but the planner must be willing to let the members of a

URBAN PLANNING...

community develop their own values. Although the planner must possess a great deal of technical knowledge, it is important, at least at the beginning stages, that she act as no more than a facilitator and resource to the deliberative group. The planner does her greatest service to the community, if she allows it to deliberatively develop its own well-formed self-reflective political preference. If the planner gets involved in the design process early, fairly guides the deliberative group through the process described above, and provides all the technical and artistic help necessary, she has, I believe, the best chance of helping the community build a truly livable urban environment.

REFERENCES

- Hare, R. M. 1989. 'Contrasting Methods in Environmental Planning', in R. M. Hare (ed.) *Essays on Political Morality*, pp. 217-235. Oxford: Clarendon Press.
- Hare, R. M. 1981. *Moral Thinking: Its Levels, Method, and Point*. Oxford: Clarendon Press.
- Rawls, John 1971. *A Theory of Justice*. Cambridge: Harvard University Press.
- Sagoff, Mark 1988. *The Economy of the Earth: Philosophy, Law, and the Environment*. Cambridge: Cambridge University Press.
- Taylor, Charles, 1985. 'What is Human Agency', in Charles Taylor (ed.) *Human Agency and Language*, pp. 15-44. Cambridge: Cambridge University Press.