THOUGHTS ON THE ESTABLISHMENT OF STANDARDS

Ralph L. Keeney

December 1974 WP-74-72

/

THOUGHTS ON THE ESTABLISHMENT OF STANDARDS

by

Ralph L. Keeney

The role of standards as I see it is to promote decision making that is in the interest of the public. For instance, air pollution standards might limit the sulfur content of fuels burned within a city or country. Without these standards, some individuals may burn cheaper higher sulfur fuels using the reasoning that "the little bit of sulfur dioxide contributed by me can't hurt that much". However, if everyone uses the high sulfur fuels, the general health of the public may deteriorate. Hence, standards are enacted to prohibit this situation from occurring.

Usually standards specify maximum or minimum limits in terms of either absolute amounts or flowrates. Examples of the former are limitations on heights of builidngs and ability to withstand an earthquake of specified magnitude. Under the latter are the air pollution standards of "parts per million" and radiation dosages due to nuclear facilities.

This short note attempts to support the contention that standards should be specified in light of (1) the public preferences and, (2) the alternatives available. The implications of any alternative must always have some degree of uncertainty, so in this sense, my contention is that standards <u>should</u> depend on preferences for and probabilities of the consequences of the available alternatives. This particular viewpoint will be explained in terms of a simple abstract example. The thoughts contained here are meant to illustrate some of the considerations and relationships that I feel are important in establishing standards. They are not meant as complete procedures for the process of setting standards.

1. Notation

Let X_1, X_2, ..., X_N be the attributes of interest to the public. For instance X_1 may be a health index, X_2 an economic index, etc. A specific level of X_i will be designated by x_i. Thus a consequence to the public may be written $\underline{x} \equiv (x_1, x_2, \ldots, x_N)$. The public's utility functions is $u(\underline{x})$, and for now, we will neglect the substantial problem of assessing this utility function. The point is that u is an appropriate index for maximizing an expected value in selecting among alternatives. We will designate alternatives by A_1, A_2, ..., A_j, ..., where in general, the number of alternatives may be infinite. The consequences of an alternative A_j can be described by the probability distribution $p_j(\underline{x})$.

Standards serve to limit the alternatives available, and in particular, they are established to "throw out" particularly bad alternatives. Let us suppose we must select a standard Q_k for society. This standard will make 'previously feasible' alternatives A_j illegal, and hence, eliminate them from further consideration. To be simple, suppose that the selection of Q_k restricts the feasible (i.e. the legal) alternatives to the set A_1, A_2, ..., A_k. Then the decision problem of choosing a standard is shown in Figure 1. Of course, in Figure 1 we have assumed complete compliance with the law which is a simplification we will accept since it does not alter our illustrations.

2. The Impact of Standards

To make the discussion more concrete, let A_j be the alternative where a nuclear power plant is designed to emit no more than j manrems per year. Then, Q_k can be the standard that a plant may emit a maximum of k manrems per year.

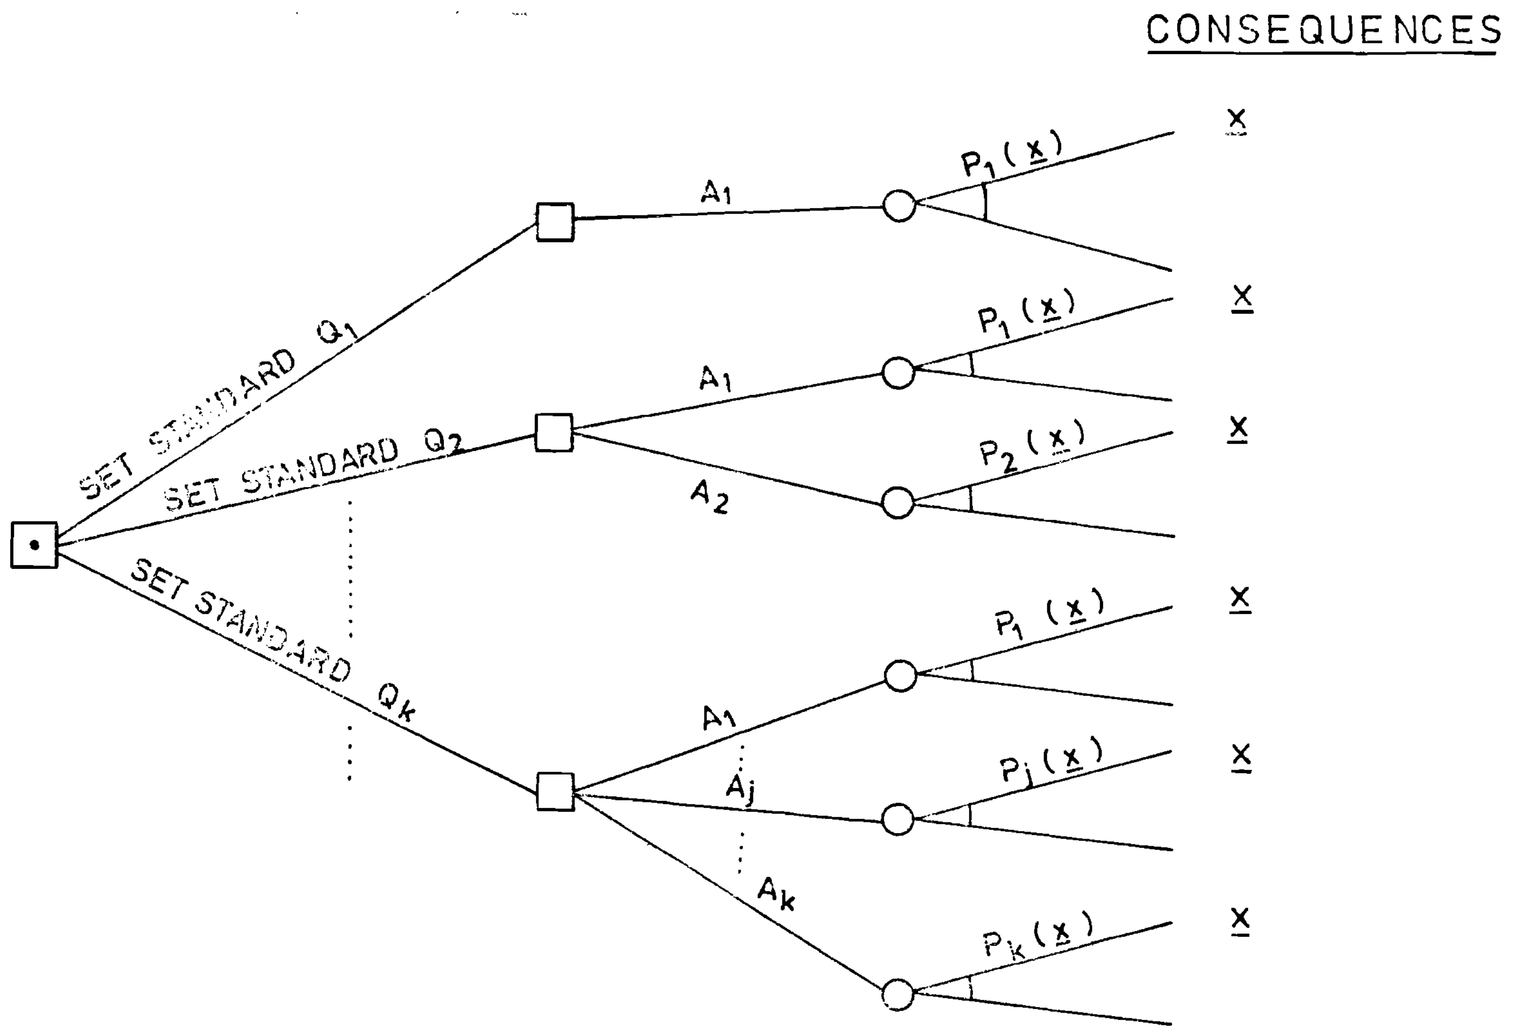

FIGURE 1. SELECTING A STANDARD Q_k

For this problem, suppose we have used society's utility function $u(\underline{x})$ and calculated the expected utility $E_u[A_j]$ for each alternative A_j. The results may be exhibited as in Figure 2a. In the absence of any standard then the optimal alternative for society is clearly A_{max}.

Suppose a standard Q_L indicated in Figure 2b, was set such that the maximum legal emission was L manrems per year. Then since the feasible alternatives are A_1, A_2, ..., A_L, from Figure 2a, the best alternative is clearly A_L with expected utility $E_u[A_L]$. Note that this alternative is less desirable than the alternative A_{max}. It has a smaller expected utility. On the other hand, if standard Q_H of Figure 2b is in force which allows plants designed to emit less than H manrems per year, where H is larger than the emissions of alternative A_{max}, then A_{max} is still clearly the best alternative. A graph of the expected utility of the optimal alternative as a function of the standard--set in this case as a maximum emission level--is shown in Figure 2b.

So from Figures 2a and 2b, one fact is simply clear provided that one is interested in selecting the alternative that is best for society. If one has society's utility function as well as the implications of all the alternatives, and if expected utility is to be used in selecting an alternative, then standards can in no way improve upon the decision and can in some circumstances actually force a suboptimal alternative to be chosen.

The need for standards is created by the fact that the decision makers and influencial groups influencing a decision do not have interests that coincide precisely with society's interests. In the terminology used here, the utility functions of these individuals and groups may be different from society's utility function u.

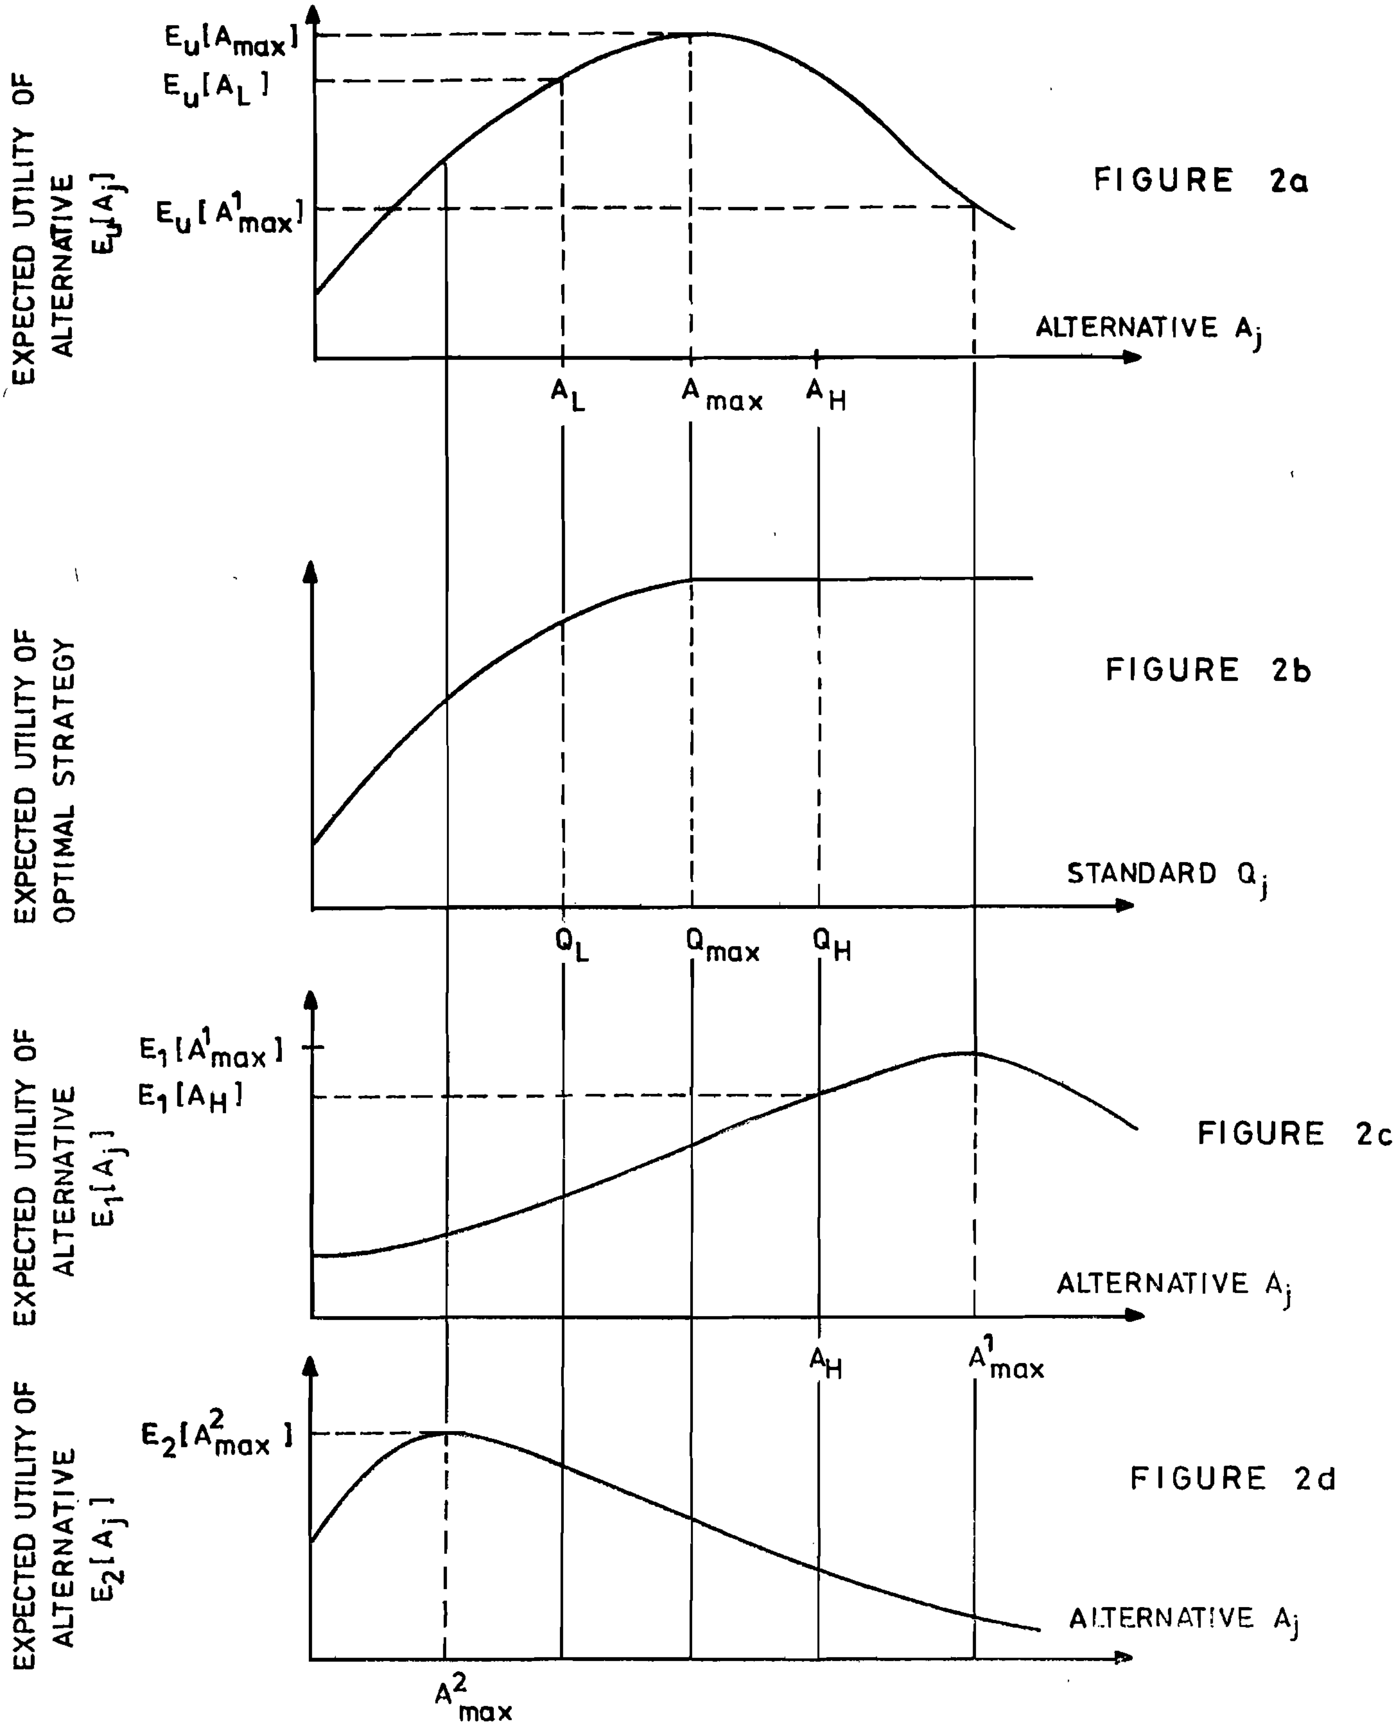

FIGURE 2. OPTIMAL ALTERNATIVES

Again to be over simplistic to make a point, suppose the body which constructs and operates the nuclear power plants, referred to here as the "operator", has a utility function $u_1(\underline{x})$. The "environmentalists" who wish to protect the environment have utility function $u_2(\underline{x})$. The expected utilities of the alternatives A_j using both u_1 and u_2 are plotted in Figures 2c and 2d respectively.

What happens if society does not implement standards and allows the group with utility function u_1 to make the decision? They should choose the alternative A^1_{max} indiciated in figure 2c since it has the highest expected utility for them. Note however from Figure 2a that A^1_{max} has an expected utility for society that is far below that of A_{max}, the best alternative for society.

If this same group is allowed to make the decision, but it must conform to a standard Q_H, then again using the group's utility function u_1, their optimal decision is found to be A_H. Hence the utility accruing to society by this decision is $E_u[A_H]$, which is much better than the $E_u[A^1_{max}]$ that would have resulted with no standard. Well the point is probably clear, society should in this case put into effect a standard Q_{max} as defined by Figure 2b.

Now let us consider what happens if group 2, the "environmentalists", have the power to select an alternative. From Figure 2d, it is clear that they should select alternative A^2_{max}. Note however that this will obviously lead to a utility $E_u[A^2_{max}]$ to society that is far below the maximum utility $E_u[A_{max}]$ for society. One might superficially argue "why can stronger standards on radiation levels lead to worse social alternatives, certainly less radiation is preferred to more radiation?" The reason, which we will try to illustrate clearly by a more detailed example in the next section, is that very low radiation levels will force poor performance on other objectives (e.g. costs will become very high).

Consider the more realistic situation where standards will in fact be established. Two of the powerful groups to be heard in the process of setting the standards may be groups 1 and 2. Group 1, based on its analysis summarized in Figure 2c, should fight for a standard greater than Q^1_{max} as this will then not hinder them if in fact they have the power to make future decisions. Group 2, the 'environmentalists' on the other hand would fight a standard Q^2_{max}, which still allows them to follow their optimum alternative A^2_{max}. However, if the 'environmentalists' had any idea of the "power plant builder's" preferences, they would realize that a standard A^2_{max} would also force the power plant builders to prefer option A^2_{max} to all their available options.

The issues are probably reasonably clear. What would be best for society would be to establish two standards, a minimum and a maximum, which are respectively, slightly less and slightly greater than the standard Q_{max}. This would limit anyone making decisions to select alternatives close to society's optimal alternative. But it doesn't seem particularly reasonable to put a minimum level on radiation levels, when as we have said, less radiation is always preferred to more. Other procedures to achieve the same effect are considered in the next section.

3. A Two-Dimensional Example

One of the crucial issues in setting standards is tradeoffs--both technological tradeoffs and preference tradeoffs. To illustrate this, suppose only two attributes, call them X and Y, are important to society and that society's utility function is $u(x,y)$. Although we will work through some of this example in the abstract, one specific context might be as follows. The government, acting in society's interest, is to establish standards concerning nuclear power plant siting. The 'operator', group 1, and 'environmentalists', group 2, are both involved in setting the standards. Once the standards are set, the operator is permitted to site plants where it wishes as long as they satisfy the standards. One can think of attribute X as manrems per year and attribute Y as cost per kilowatt hour of electricity to the consumer.

Using our abstract model, the feasible consequence space in the absence of standards is given in Figure 3a. A natural minimum for each of the attributes is zero. However the attribute levels have no obvious maximums. Three alternatives--or to be more specific, impact regions of alternatives--are shown in Figure 3a also. With two attributes, a probability distribution $p_j(x,y)$ is needed to specify an alternative A_j. The circles are meant to indicate the region of X, Y space where $p_j(x,y)$ is non-zero for any particular alternative. A technological tradeoff curve, indicated by the dashed line, is also shown in Figure 3a. This curve has two interpretations for our purposes. First, if one neglects uncertainty for a moment--or from another viewpoint, after the uncertainty has been resolved--the technological tradeoff curve says that if one has a consequence at one point on the curve, it is technologically feasible to move to any other point on the curve. However, our choice is between alternatives involving uncertainty, so the technological tradeoff curve can be thought of as indicating the locus of the expected values of x and y--assume probabilistic independence--for the range of possible alternatives. Clearly, some of

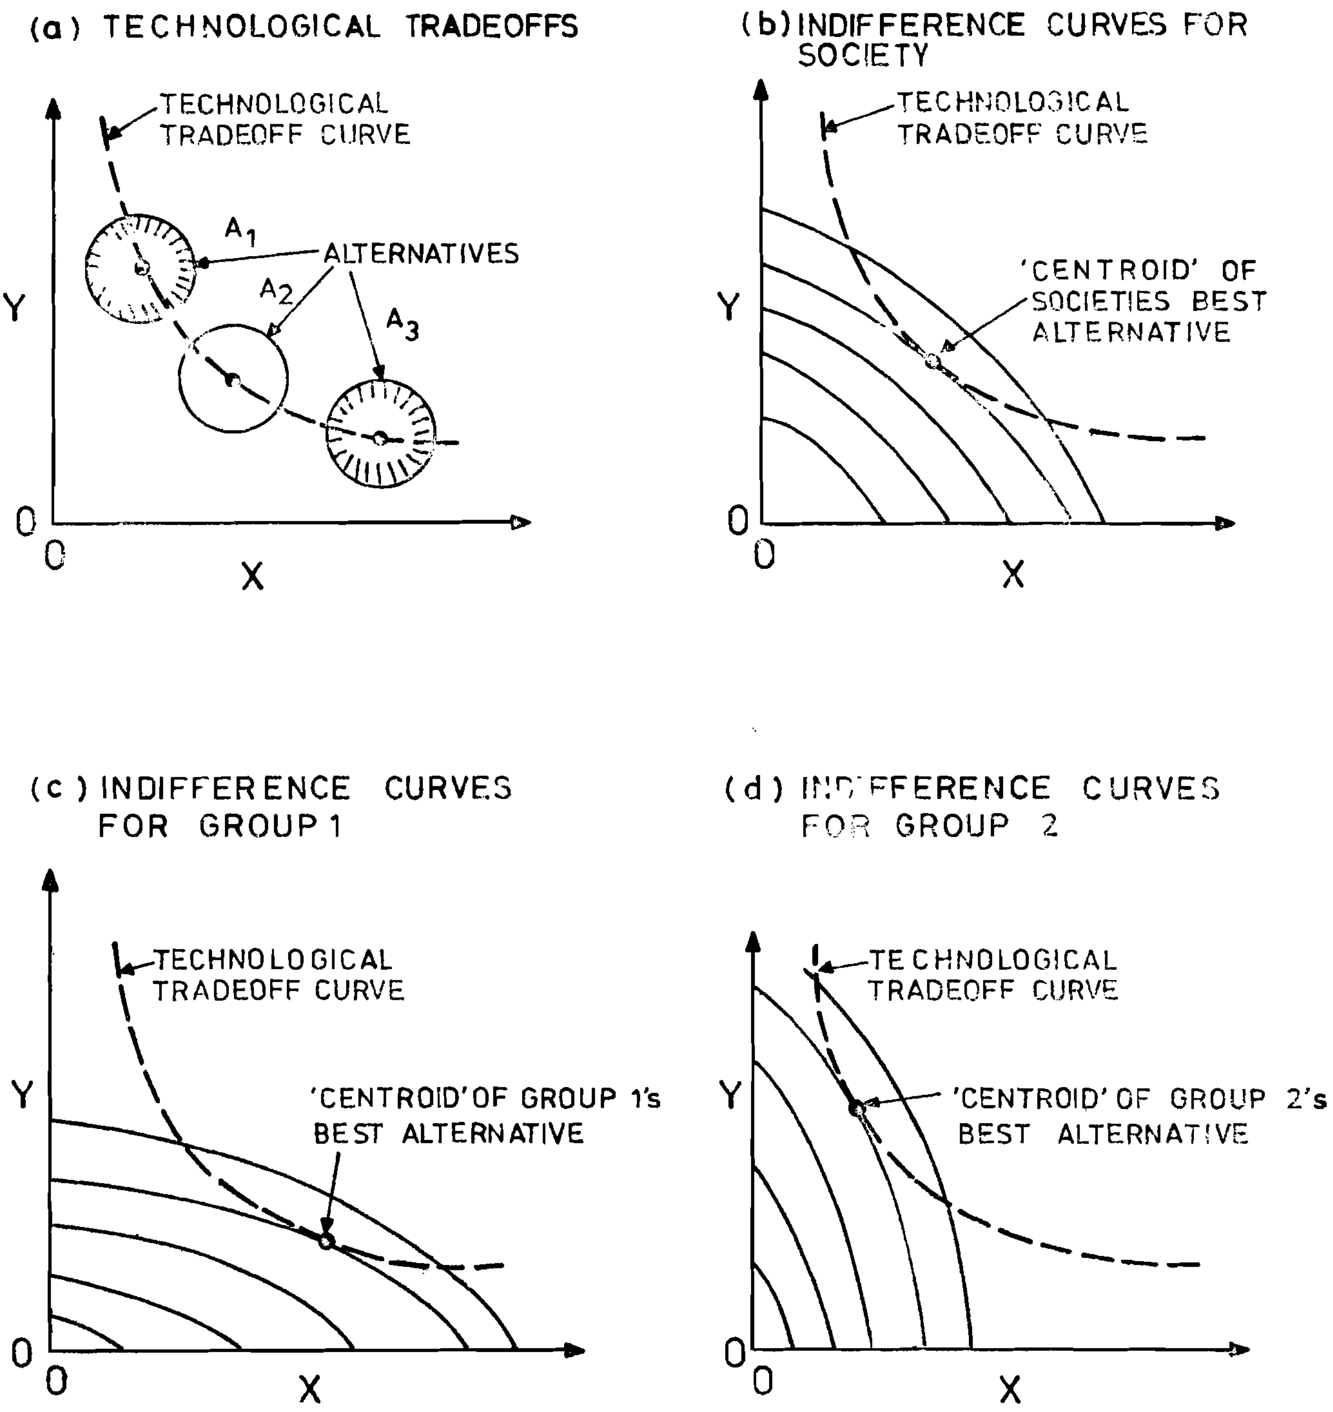

FIGURE 3. TECHNOLOGICAL AND PREFERENCE TRADEOFFS

this is a bit simplistic, but we wish to illustrate some ideas here with as few complications as possible.

Parts b, c, and d of Figure 3 indicate the preference tradeoffs of society, group 1, and group 2 respectively. The solid lines are indifference curves, and in all cases we will assume preferences are decreasing in both attributes X and Y. That is, smaller x is preferred to larger x and smaller y is preferred to larger y in all cases. Note that this would be the case if X designated manrems per year and Y cost per kilowatt hour.

Superimposed on Figures 3b, 3c and 3d is the technological tradeoff curve, and what I've chosen to call the 'centroid' of the best alternatives for society and the two groups. All of this is subject to no standards. Given the preference structures of society and of groups 1 and 2 as indicated in Figure 3, one might calculate the respective expected utilities for these three entities and find out they are just as illustrated in Figures 2a, 2c, and 2d respectively.

Essentially, Figures 2a, 2c, and 2d present plots of the expected utility to the entity involved as the alternative moves along the technological tradeoff curve. The 'essentially' is a qualifier because uncertainty is neglected in this interpretation, but the sense of the statement should be clear.

The aspect about the sets of indifference curves in Figure 3 which make them different is that society and groups 1 and 2 have different preference tradeoffs. To clearly illustrate this let us assume the utility functions of each of the three entities is of the same additive form. Thus society's utility function is

$$u(x,y) = k_X u_X(x) + k_Y u_Y(y) \quad , \tag{1}$$

where u_X and u_Y are utility functions over attributes X and Y respectively with their origins at zero, and k_X and k_Y are positive scaling constants. Similarly, the utility functions for groups 1 and 2 are of the form

$$u^i(x,y) = k_X^i u_X^i(x) + k_Y^i u_Y^i(y) \quad , \qquad i = 1,\ 2, \tag{2}$$

where i indicates the group. Furthermore, let us assume that the conditional utility functions over the single attributes are identical for all three entities as shown in Figure 4.

These utility functions are consistent with the indifference curves shown in Figure 3. The difference in these indifference curves is provided by the difference in the ratios k_X/k_Y, k_X^1/k_Y^1, and k_X^2/k_Y^2. Specifically, to be consistent with the indifference curves of Figure 3,

$$\frac{k_X^2}{k_Y^2} > \frac{k_X}{k_Y} > \frac{k_X^1}{k_Y^1} \quad . \tag{3}$$

Thus, since the conditional utility functions are identical, we can conclude that at any point (x,y), group 2 would be willing to allow Y to increase more than society would in order to get a specified reduction in X. Similarly, group 1 would be willing to allow X to increase more than society would in order to reduce Y by a fixed amount. Again note, all this might reasonably be consistent with X being radiation levels and Y being energy cost.

Setting Standards on X

A normal practice followed in setting standards is to set a standard for one attribute. Let us suppose a standard x_s is chosen which says "It is illegal to have X levels (i.e. radiation) greater than x_s". This

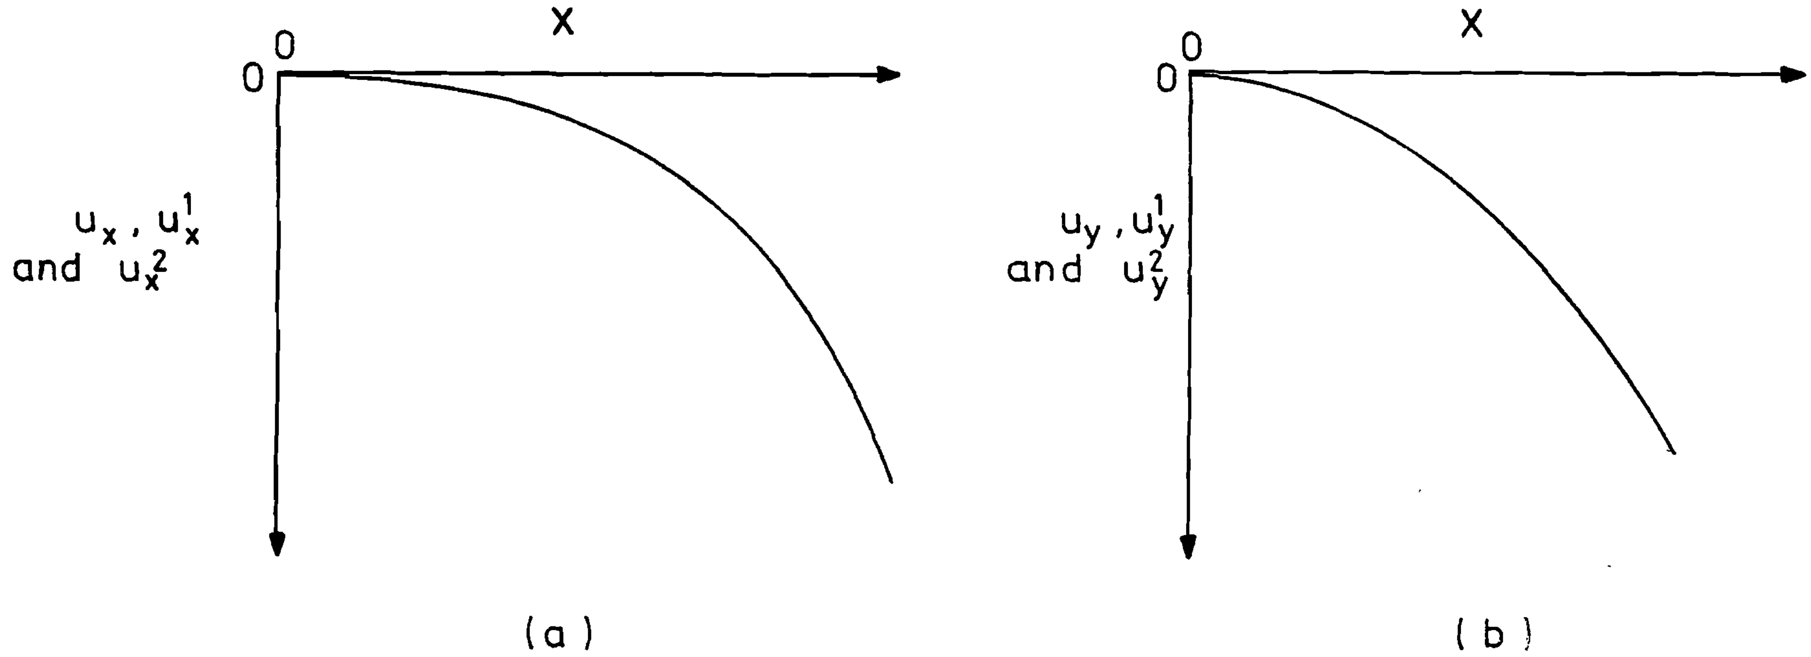

FIGURE 4. CONDITIONAL UTILITY FUNCTIONS FOR X AND Y

limits the alternatives as illustrated in Figure 5. Alternative A_1 is certainly legal since there is no possibility that the standard is exceeded. On the other hand, alternative A_4 will certainly lead to an illegal level of X so it can be discarded. Alternatives A_2 and A_3 each have a chance of resulting in a legal level of X. An 'operator' may consider A_2 to be a feasible option since it will very likely result in a 'legal' consequence, whereas A_3 would probably be rejected since an 'illegal' consequence would likely result. In any situation, where an illegal consequence did result, there would likely be an extra cost or some type of penalty involved in altering the consequence to meet the standard.

Given that the rules as we have established them--that is, a standard is set by government and then group 1 makes decisions--the objective of the government should try to set a standard to cause group 1 to select the alternative leading to society's best alternative indicated in Figure 3b. This is the same thing as setting the standard to lead to the alternative A_{max} in Figure 2a. The problem raised in the last section was how to protect oneself from too strong a standard on X.

It is clear from Figures 3d and 2d that group 2, if it had its own way, would select a standard leading to A^2_{max}. Such a standard would be too strong from society's viewpoint and lead to an alternative less desirable to society than A_{max}. But as we indicated, it seems unreasonable to establish a minimum standard on X, stating in fact that radiation must exceed a certain amount. Given the conditional utility function for X as shown in Figure 4a--one in which all parties agree--a minimum standard seems ridiculous.

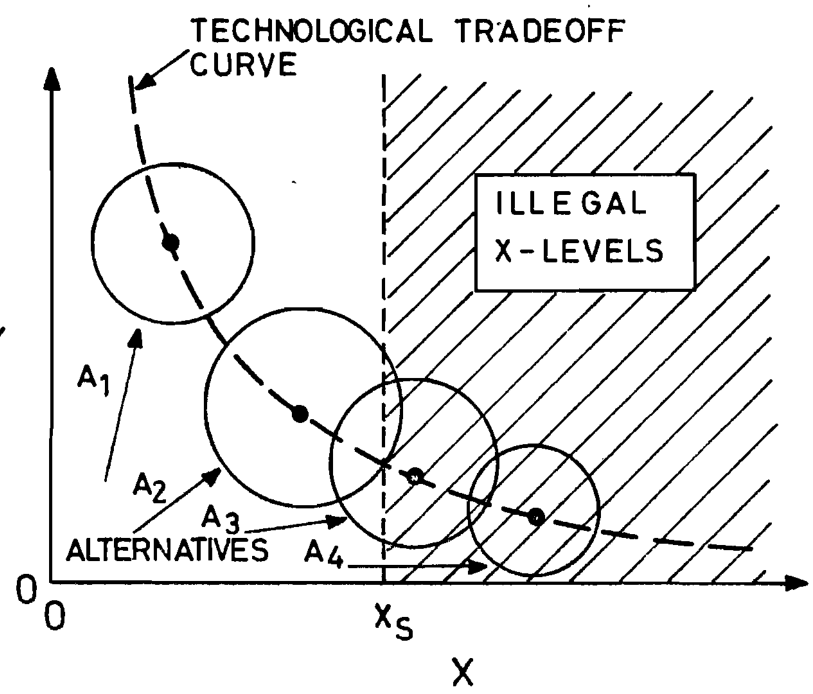

FIGURE 5.THE EFFECT OF A STANDARD ON ATTRIBUTE X

We will suggest two ways to address this issue, joint standards on X and Y and standards on society's preferences. Both of these approaches get at the issue of society's preference tradeoffs between X and Y.

4. Setting Joint Standards on X and Y

Recall that the greater Y becomes, the more undesirable it is to society, as illustrated in Figure 4b. Given this, it may be reasonable to set a maximum standard y_s which says "It is illegal to have Y levels (costs) greater than y_s".

Figure 6 illustrates four possibilities for the setting of joint standards on X and Y. The situation in Figure 6a is about right for society in the sense that the alternatives which remain feasible are those with an expected utility very near to the expected utility of society's optimal alternative. Figure 6b is the case where x_s is too low (i.e. too strict) and y_s is too high for society. It is the type of standards group 2 would obviously like, since it leaves as the only feasible alternatives, those alternatives near optimal for group 2. The standards in Figure 6c are just the reverse, x_s is too high and y_s too low. These are the standards that group 1 would like, because they promote alternatives close to optimal for group 1. Basically, group 1 is willing to accept high levels of X in order to keep Y low, whereas group 2 wants to keep the levels of X down at the expense of Y. Society is between these two groups.

Figure 6d represents the interesting situation where both standards are too stringent and no feasible alternatives exist. Such a situation can result--and has resulted--from trying to establish standards independent of the alternatives available. Looking at Figure 4, clearly we want both

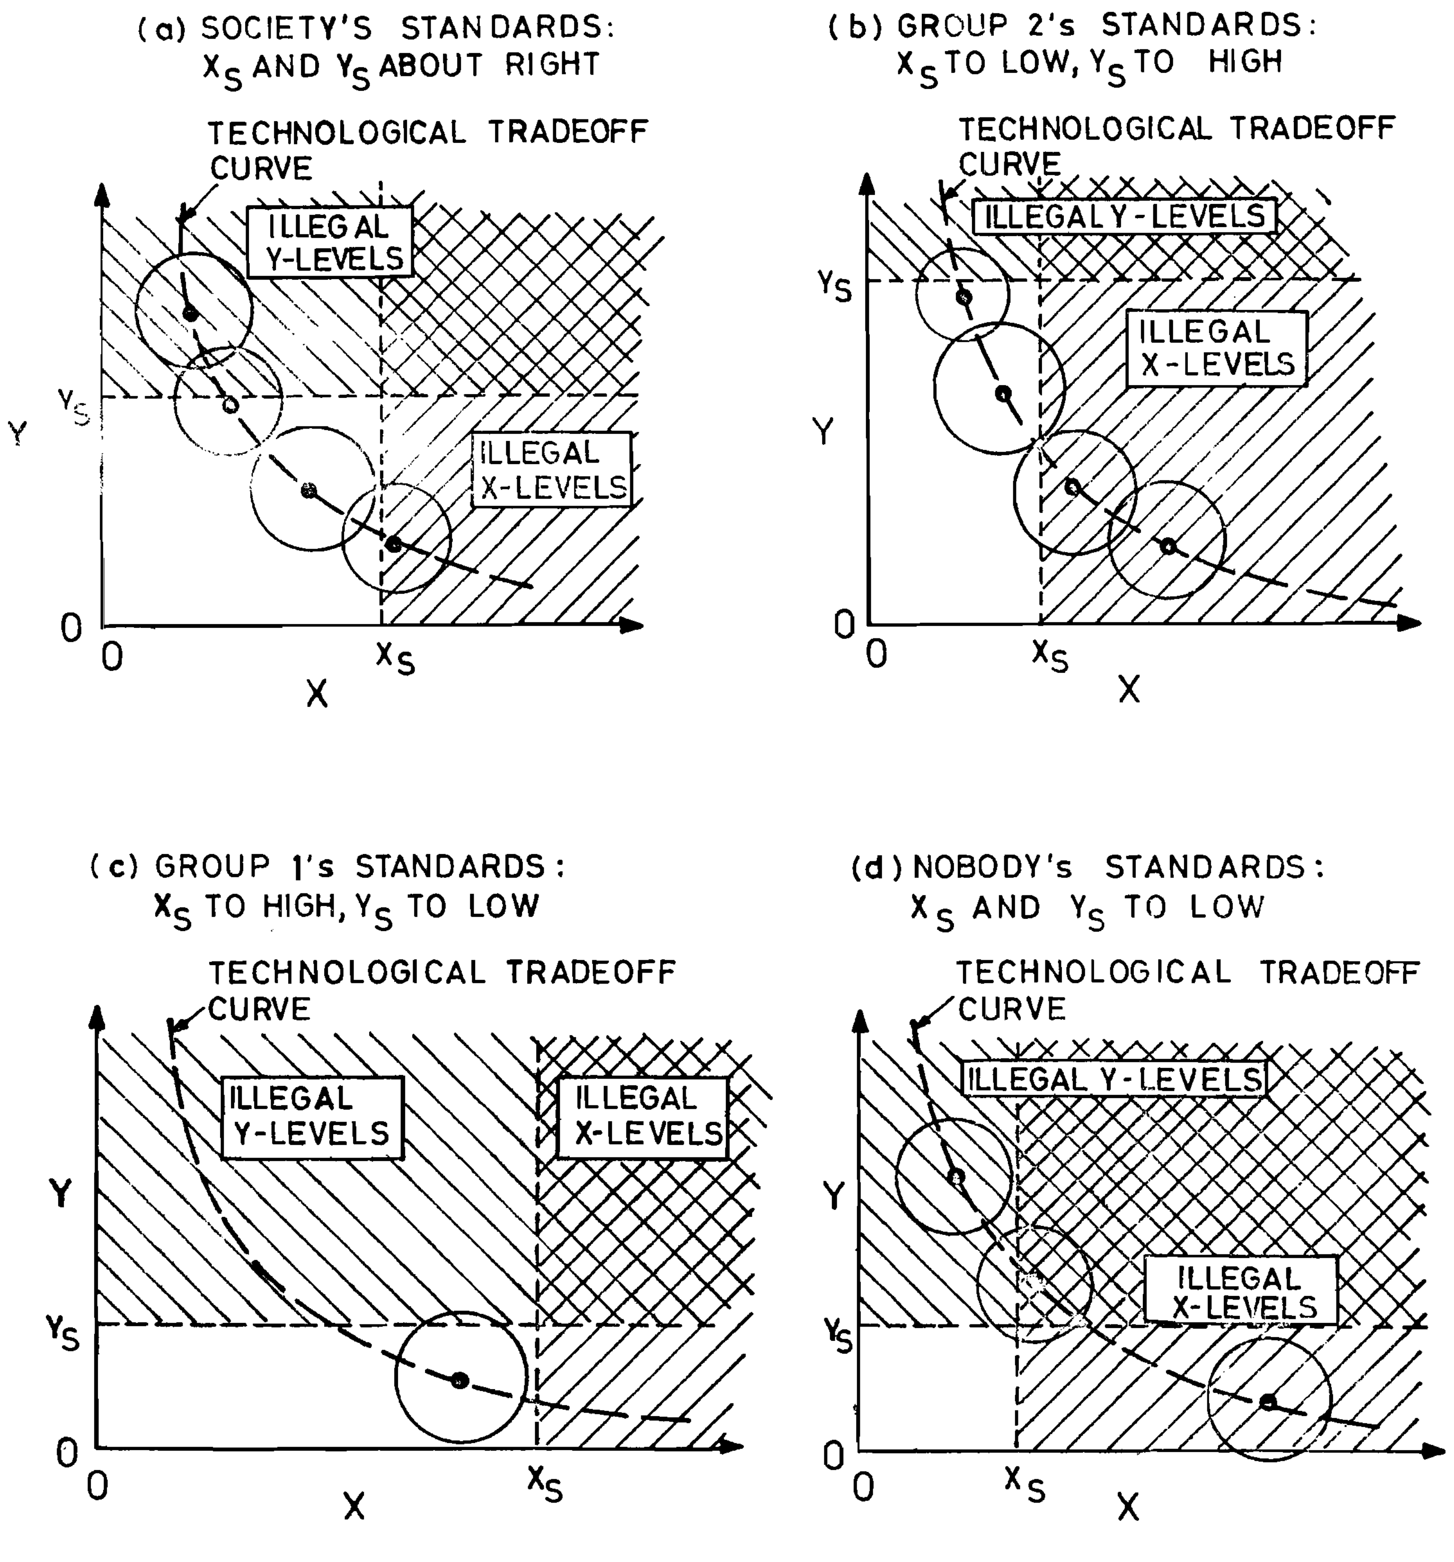

FIGURE 6. JOINT STANDARDS ON X AND Y

x and y to be as small as possible. However, the point is, that at some point, further reduction of one can only be achieved at an increase in the other. One must keep the technological tradeoffs, which are dictated by the available alternatives, in mind when setting standards. One might think of the technological tradeoffs as pushing the standards out (i.e. dictating high maximum standards) and of the preferences and preference tradeoffs as pushing the standards in (i.e. dictating lower maximum standards). The 'trick' is to balance these, which is clearly no mean task.

5. Standards on Society's Preference Tradeoffs

As an alternative to setting joint standards on X and Y, it may be more reasonable to dictate, the legal preference tradeoff between X and Y. Actually, what the government should do for society is just to legalize its preference structure u(x,y), and then no standards would be needed. We would have the situation as illustrated in Figure 3b, and society's best alternative should be chosen.

However, having said this, let us be a bit more realistic and assume that the government doesn't known society's utility function u(x,y), but that it has some idea of society's preference tradeoffs indicated in Figure 3b. If in addition, government has a good understanding of the technical tradeoffs near society's optimum alternative, then government need not use the complete u(x,y) as a standard. Refer to Figure 7a where we have duplicated Figure 3b and to Figure 7b which shows the same technological tradeoff curve but with a set of linear indifference curves. The point is that both sets of indifference curves lead to the same decision, the optimum for society.

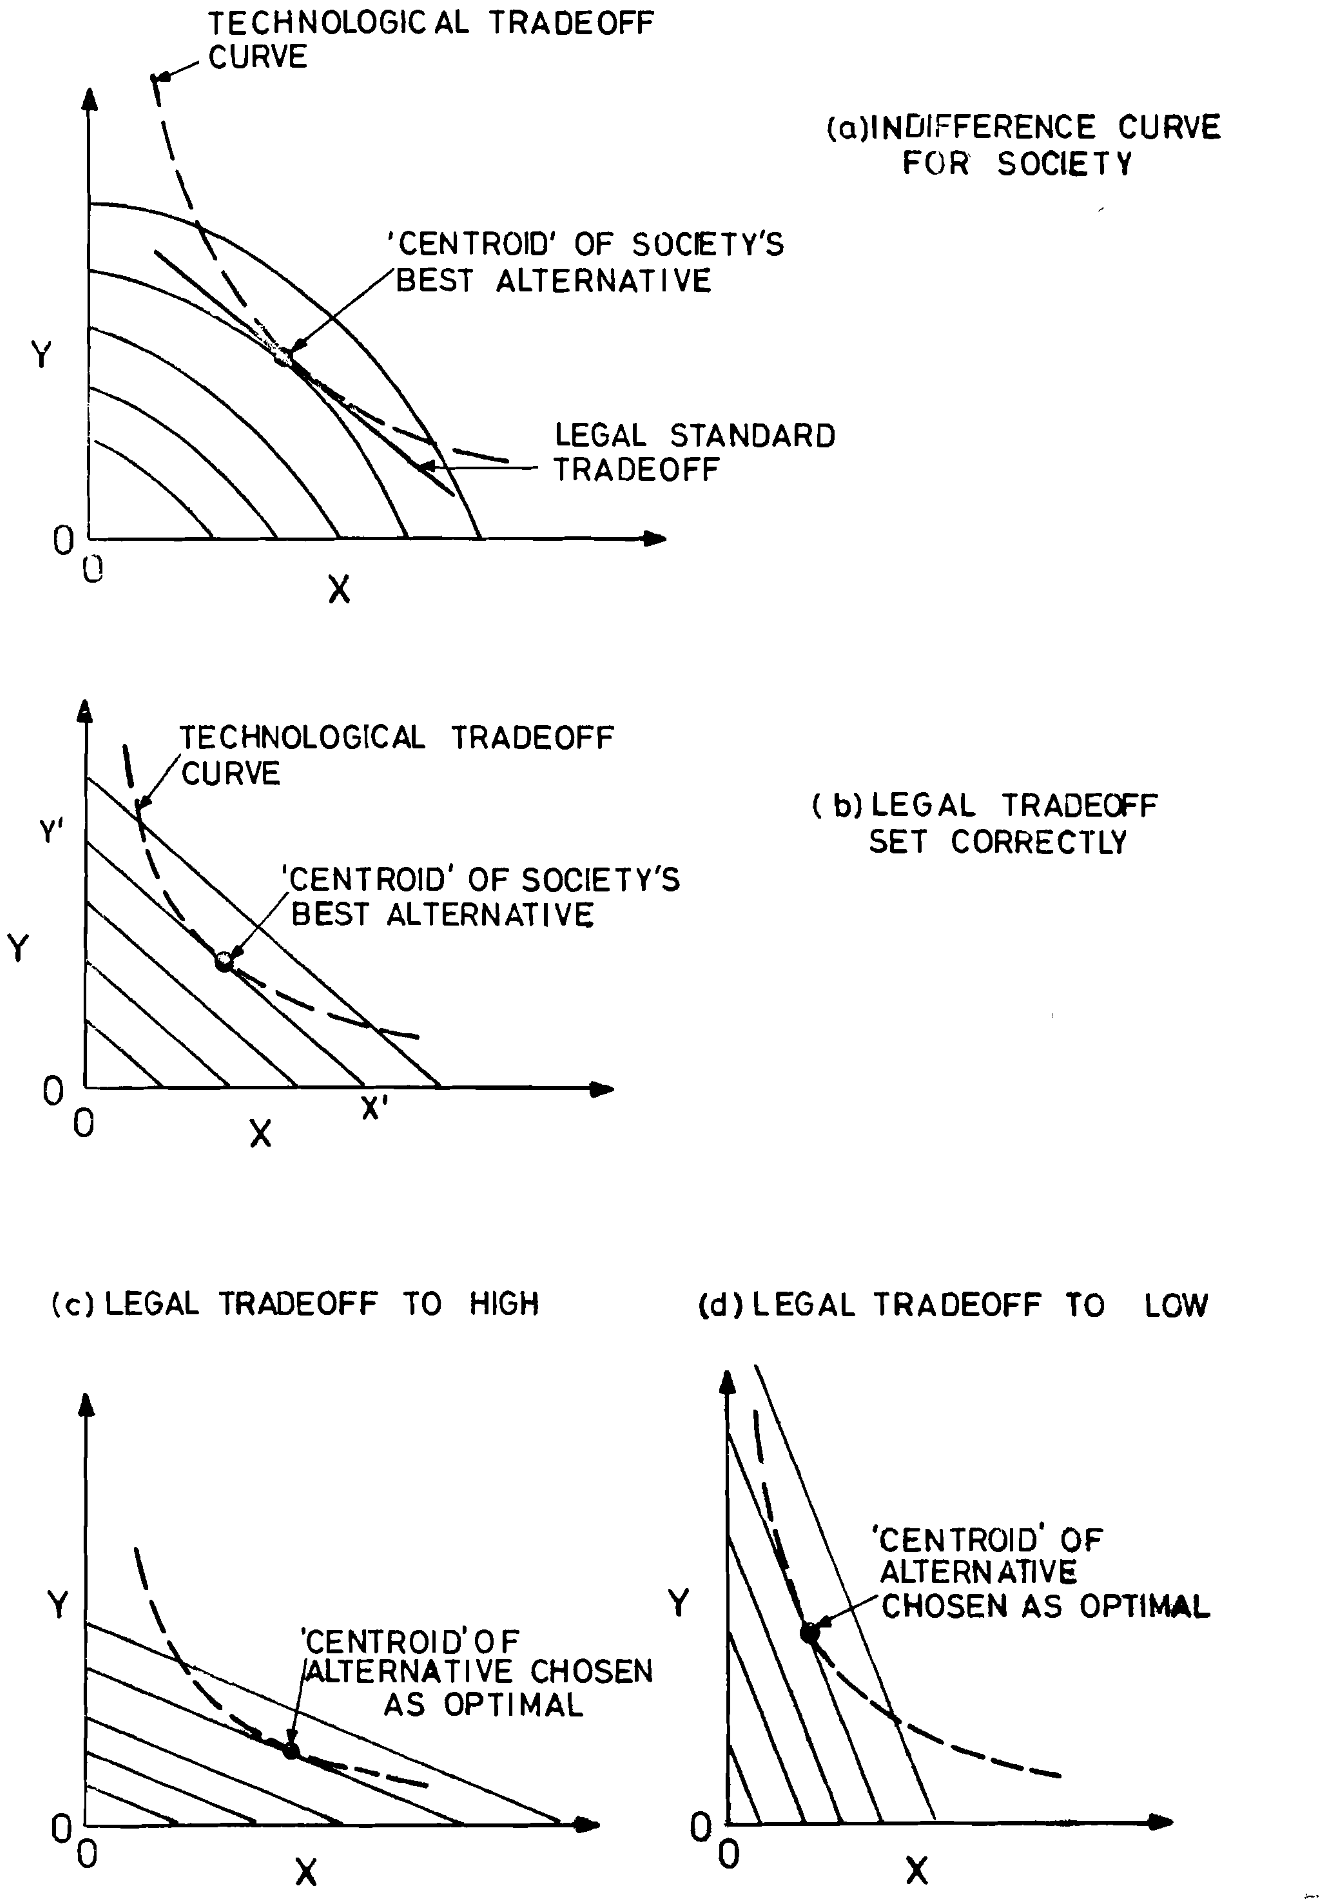

FIGURE 7 A STANDARD FOR PREFERENCES OF SOCIETY

To specify linear indifference curves only requires one constant, the rate of substitution between X and Y. If one defines an x' and y' as indicated in Figure 7b, then the rate of substitution between X and Y is y'/x' which will be defined as λ. This means that one is just willing to let X increase by λ units in order to reduce Y by one unit. Since the specification of λ defines the entire indifference structure indicated in Figure 7b, government could set the standard that "the legal rate of substitution between X and Y is λ". The standard that is best for society is that implied by the line tangent to both the technological tradeoff curve and society's indifference curve at the optimal alternative point in Figure 7a.

Figures 7c and 7d respectively indicate the situation where the legal tradeoff between X and Y is set too high and set too low. Group 1 would support the standard illustrated by Figure 7c and Group 2 would prefer a standard like that in Figure 7d.

6. Conclusions

In a society in which all public decisions are made for "the good of society", no standards are necessary, if in fact, society's preferences (i.e. utility function) are used in making decisions. The establishment and adherence to standards in this case can never lead to improved decisions for society. The argument is simple, with no standards, society's optimal decision is taken, and since standards only reduce the number of alternatives, no 'better' decision can be found.

However, usually the people responsible for making decisions affecting the public are not using society's utility function, but rather they have their own set of preferences. Thus standards are set to promote this group to in fact choose society's optimal decision. The standards attempt to eliminate all those alternatives which the decision making group prefers to society's optimum, so that in fact, the best remaining feasible (legal) alternative for the group is society's optimum.

In a two-attribute context, we examined the fairly typical process of establishing a standard on only one attribute. This is often done by considering only the impacts of various levels of that one attribute. For example, in a nuclear power context, radiation standards specifying maximum legal amounts might be set by considering only health impacts due to various radiation levels, and neglecting other important factors such as the cost of power and its impact on the quality of life, dependence on foreign power sources, etc. Oversimplifying, primary effects are considered, but secondary, tertiary, etc. effects are often neglected.

Our position is that standards should be set by considering

(1) the alternatives available (i.e. technological tradeoffs), and

(2) society's preferences structure (i.e. preference tradeoffs).

The alternatives, specified by probability distributions over the possible consequences are meant to capture all the effects. The preferences are needed to decide which set of possible effects, of those available, are preferred.

The implications of two types of standards were investigated. First, joint standards on the two attributes were considered. Here, to some degree, the technological tradeoffs dictate the 'absolute' level of the standards, and the society's preference tradeoffs are used to establish the 'relative' levels these standards should have in order to promote the choice of society's optimum.

The other type of standard concerned specifying society's 'official' tradeoff between the attributes. The official tradeoff indicates how much of one attribute one will give up to obtain a unit of the other attribute. This standard is set, as illustrated in Figure 7a, by jointly considering the technological tradeoffs and society's preference tradeoffs.