

PROBLEMS

1. Explain how, given a prospect, P , valued by two citizens as $(0, 0)$ and one, Q , valued as $(1, 0)$, we can construct prospects valued as $(u, 0)$, for every u between 0 and 1.
2. Harsanyi's proof introduces $(0, 0)$ in equation (3) and $(0, 1)$ and $(1, 0)$ in equation (4). The former was grounded on the fact that $(0, 0)$ represented the worst prospect for each citizen, the latter on the anonymity condition. How can we ground (3) and (4) if we use the special prospects assumption to prove Harsanyi's theorem?
3. Prove that the distributable goods assumption implies the special prospects assumption.
4. Explain how the distributable goods assumption postulates infinitely many options.
5. Suppose that the social planner in the fertilization example could introduce new prospects by promising to give the citizens small amounts of money. Could he construct the special prospects by (1) selecting one of the original prospects, A , to be $(0, 0, \dots, 0)$ and (2) introducing a new prospect, B , that is just like A except that the first citizen receives enough money in B to prefer it to A , introducing a prospect, C , that is just like A except that now the second citizen receives enough money to ensure that she prefers C to A , and so on? What empirical assumptions would have to be true for this to work?

6-4c. Interpersonal Comparisons of Utility

Although utilitarianism has had and continues to have many adherents among economists and philosophers, there are several major problems with the view that must be considered by anyone seriously concerned with the problem of social choice. The one I will address here—the problem of interpersonal comparisons of utility—is closely tied to our previous discussion of utility. Although this difficulty arises most acutely with respect to utilitarianism, it affects several other approaches to social choice as well.

Mary and Sam are trying to choose a joint vacation. They have been considering going to the seashore, camping in the local mountains or visiting the museums in a nearby city. It turns out that these all cost about the same, so monetary considerations are not relevant. Mary prefers going to the seashore to going to the museums and that in turn to camping. Sam's preferences are the exact opposite. However, going to the seashore is the only alternative that Mary finds bearable, although she feels more negative about going to the mountains than to the museums. Each choice is fine with Sam, although he would much prefer going to the mountains. Mary and Sam know all this. Were they an ordinary couple with a modicum of mutual respect and goodwill, one of them would suggest going to the seashore, since the gain in happiness for Mary would more than offset Sam's slight disappointment. If Sam resisted this suggestion, Mary would rightly resent his selfishness.

But *this* Mary and Sam are not an ordinary couple for they have been to business school and have learned about utilities. They note that although it ap-

SOCIAL CHOICES

pears that Mary has stronger preferences concerning the three options under consideration than Sam does, it does not follow that more *total* utility will be achieved by opting for her first choice. So Mary and Sam decide to develop utility scales to measure their preferences more accurately. Instead of restricting themselves to just the three options under consideration they rank a very wide range of alternatives, so that they can see how their preferences for the three particular items fit into the general scheme of their likes and dislikes. For ease in computation they both use 0 to 100 scales. The amounts of utility scored by each option under this system are shown in table 6-18. Upon seeing this Sam

6-18	Mary	Sam	Total
Seashore	20	86	106
Museums	10	93	103
Mountains	9	100	109

insists on having his first choice—camping in the mountains. This turn of events causes Mary to become quite irked. Sam is a songwriter who spends much of his time in activities he finds enjoyable. She works in the complaints office of a large department store; her customers are aggressive and nasty. This is her annual vacation, and for her the difference between the seashore and the mountains crosses the threshold between the bearable and the intolerable. She feels that her “right” to an emotionally recuperative vacation will be violated by following this utilitarian scheme.

Mary believes she knows Sam well. Despite the fact that he spends most of his time in activities he enjoys, he never enjoys *anything* very much. Nor for that matter does he dislike anything much. His moods are so constant that Mary sometimes wonders if he is really human. Fortunately for her, he is open to reason, and she reasons with him as follows: “Look, Sam. We shouldn’t have used the same units on our utility scales. My preferences are so intense in comparison with yours that my scale should range between 0 and 1,000, if yours ranges between 0 and 100. If that change is made, the total utilities become seashore 286, museums 193, and mountains 190, and it is clear that we should go to the seashore.”

Sam responds in a quiet but unrelenting tone: “Your reasoning, Mary, is correct but it is based on a false premise. You think that my preferences are rather weak, but the fact is I feel things quite deeply. I have been brought up in a culture very different from yours and have been trained to avoid emotional outbursts. In my family it was considered unseemly to jump with joy or to scream with anger or to weep when sad. But I have strong feelings all the same. And even if I did not, I do not think that extra weight should be given in a utilitarian calculation to those who are capable of more intense preferences. After all, each person is due as much consideration as any other.”

Unable to resolve their dispute, Mary and Sam visit their business school professor, a man of great wisdom and wealth. He ponders their problem for a while and tells them that they need to talk with a philosopher. Mary and Sam leave him in utter despair and later agree to make their choice as most other couples would. They go to the seashore.

This parable illustrates most of the difficulties raised by interpersonal comparisons of utilities. These include epistemological, metaphysical, and ethical issues. For instance, are such comparisons based on an objective reality? Is Mary's preference for the seashore *really* stronger than Sam's for the mountains? Or is she just a more vocal person, as Sam suggests? If some people's preferences are in fact stronger than others', how could we *know* this? Does it make any more sense to compare Sam's preferences with Mary's than it does to compare a dog's preference for steak bones with a horse's preference for oats? Finally, even if we answer all these questions affirmatively, is it morally proper to respond to such information in making social choices?

Some utilitarians have argued that such problems will dissolve once psychology develops an adequate theory of the emotions. In particular, if as recent neurophysiology suggests, our emotional life can be explained in terms of brain processes, there is an underlying reality at which interpersonal comparisons hint and eventually it will be known. Furthermore, the argument continues, there can be no question that we do make such interpersonal comparisons in our daily lives. Anyone who bases a decision on the varying preference intensities of the people affected implicitly makes such comparisons by his or her very choice. Thus we should not raise artificial philosophical quibbles against a practice that is so well entrenched in human affairs.

But this does not end the matter. Let us grant that a neurophysiological account of the emotions is possible and that we do react to preference "intensities" in making certain choices. This still does not advance the cause of utilitarianism. For compare the neurophysiological account of perception with the hoped-for account of the emotions. Neurophysiology tells us quite a bit about perception: how the eye works, what parts of the brain are concerned with sight, what one can or cannot see under certain conditions, and so on. However, in discussing perception, this science speaks entirely in terms of neurological phenomena or observable human behavior. When neurophysiologists tell us that under certain conditions people suffer from double vision, their criterion for double vision is given in terms of patients' responses to questions and tests, and, possibly, measurements of the behavior of the patients' brains. But no one who thinks there is more to perceiving than responding to external stimulations in certain ways or having certain brain waves will take the neurophysiologists' account as the complete account of perceiving. Similarly, a neurophysiological account of the emotions will tell us about laughing, crying, fighting, or kissing, and, perhaps, brain patterns associated with them; but it will have nothing more to add for those who believe that joy, sadness, anger, and love are not totally reducible to our behavioral displays and bodily states. And that will leave us

SOCIAL CHOICES

right in the midst of Mary and Sam's debate over whether there is more to emotions than our displays of them.

The second part of the utilitarian argument focused on our choices made in response to preference "intensities." Consider the case of an employer who must choose between two equally qualified employees to promote. Let us assume that everything about their contributions to the firm, their length of service, personal financial needs, and so forth, is the same, and that the employer favors neither one on purely personal grounds. He summons both employees to his office for separate conversations. The first is an impassive type who allows that he would be pleased to be promoted. The second, on the other hand, effusively tells the employer how long he has hoped for the promotion, how he and his family have discussed it many times. The employer promotes the second employee and later explains to a friend that he did so because "it meant so much more" to the second.

Next consider the case of a politician who must decide whether to demolish a block of old houses to make room for a new library. The residents of the houses are old and feeble, and the sponsors of the library are young and quite vocal. The politician studies the matter and finds that the case for either side is equally compelling. Then delegates from both sides visit him. The old are so weak that it is all they can do to speak, but the young forcefully plead their case and indicate how grateful they will be if he decides in their favor. The politician finds it politically expedient to favor the young.

Those who believe in the interpersonal comparison of utilities will grant that the two cases have been correctly described: The employer weighed the utilities of his two employees; the politician simply responded to political pressure. But to those who are skeptical about interpersonal comparisons of utility, the difference between the two cases will appear chimerical. In both cases, they will argue, the decision maker is simply behaving in accordance with cultural conditioning to respond in certain ways to the actions of others; the second employee's effusiveness is just as much a form of pressure as the political activists'. Both forms of behavior are intended to manipulate the persons at whom they are directed. Thus the claim that we make interpersonal comparisons of utility in making our everyday social choices has no substance for we merely respond to pressures in making these choices.

Let us turn from this general philosophical debate, which is likely to remain unresolved for some time, to examine the more technical aspects of interpersonal comparisons of utility. It will be useful for us to distinguish interpersonal comparisons of utility *levels* from interpersonal comparisons of utility *increments*. Mary's resentment of the utilitarian choice that favored Sam's preferences was based on her current utility level being lower than Sam's. Despite the greater overall utility gains to be obtained from going to the mountains, she felt raising her utility level should be given priority. Social planners influenced by Rawls would also want to compare the utility levels of various citizens in order to identify the worst off. On the other hand, those concerned with how much additional utility one prospect affords in comparison with another are concerned

with utility *increments*. Sam's case for going to the mountains was based on comparing the incremental utility it gave him with that yielded by the other alternatives. So long as we rank alternatives by simply asking which one yields more utility, we are concerned with utility increments rather than utility levels.

The utilitarian SWF produced by Harsanyi's theorem—call it H —deals with incremental utilities and ignores utility levels. To see this, recall that H ranks a prospect x above another y just in case the sum of the citizens' utilities for x is greater than that for y . That numerical relationship will be preserved if we add the same number to x and y . Indeed, it will be preserved if we add different numbers to the individual citizens' utilities for x and y , so long as we add the same numbers to their utilities for both x and y . Adding these numbers is tantamount to changing the zero points or *origins* of the citizens' utility scales. These are the reference points by which utility levels can be measured. Consequently, adding different numbers to the scales of different citizens is to shift their origins independently and, hence, to shift independently the basis for measuring their utility levels. But shifts of origins, whether in concert or independently, do not register with H at all. Thus it neither responds to individual utility levels nor presupposes interpersonal comparisons of them.

On the other hand, H does respond to changes in the units used to measure individual utility increments. Mary hoped to use this to turn the tables on Sam. She tried to force him to use a 0 to 100 scale (with 100 units) while she used a 0 to 1,000 scale. In effect, she wanted to multiply her original scale by one number (10) while multiplying Sam's by another (1). Since this would place a different item at the 1 spot on her scale, it would change the item that marks the *unit* of her scale.

We fix the origin of a utility scale when we fix its 0 point; we fix its unit when we fix its 1 point. As we have just seen, we change the origin of a scale when we add the same number to every value on the scale, and we change its unit when we multiply each value by the same positive number. It follows that an arbitrary positive linear transformation may shift both the origin and unit of a scale. A social choice method that responded to any (nonidentical) positive linear transformation of a citizen's utility scale would thus respond to both utility levels and utility increments. The following SWF, G , is an illustrative example. G ranks x above y if and only if (1) the utility level of the best-off person in x is greater than that of the best-off person in y , or (2) (in case x and y tie according to (1)) the total utility afforded by x exceeds that afforded by y . In case x and y tie on both counts, G ranks them as socially indifferent.

All the social choice methods we have discussed so far respond only to changes in utility scales that alter the positions of the citizens with respect to one another—either by moving the origin of one with respect to that of another or by changing the ratio of the number of units of one to that of another. None of these social choice methods responds to uniform changes affecting every citizen's scale in the same way. Thus we need not worry about whether we have found the “real” origins or the “real” units for our citizens' utility scales, so long as we have properly calibrated their scales *with respect to one another*.

SOCIAL CHOICES

Moreover, when using a social choice method that responds only to utility levels, we need concern ourselves only with whether the origins of our citizens' scales are properly aligned, that is, with the *interpersonal comparison of utility origins*. On the other hand, when using a social choice device, such as Harsanyi's *H*, which responds to utility increments, we must concentrate instead on the proper alignment of the scale units, that is, on the *interpersonal comparison of utility units*. Summing up, some social choice methods respond only to changes in utility origins, these presuppose the interpersonal comparability of utility origins. Other social choice methods respond to changes in utility units and presuppose interpersonal comparisons of utility units. Finally, some respond to both sorts of changes and, accordingly, presuppose both sorts of comparisons.

Returning to Mary and Sam, we now know that the problem of comparing their utilities may be reduced in their case to the problem of comparing their units. Mary believes she has a greater preference interval separating the seashore and the mountains than Sam has. Thus she objects to their both being assigned 0 to 100 scales, which represents her interval as eleven units long when Sam's is represented as fourteen units long.

Mary bases her complaint on the following interpersonal comparison of preferences:

Mary's preference for the seashore over the mountains is stronger than Sam's for the mountains over the seashore.

This has the general form

A's preference for x over y is greater than B's for z over w .

Now if we can make a case for such comparisons, we can solve the problem of the interpersonal comparison of utility units. For we can stipulate that

A's preference for x over y is exactly as great as B's for z over w

is to hold just in case neither one's preference is greater than the other's. Then we can select the distance between two items on some citizen's scale as our standard unit. Having done this we can rescale his utilities so that the item he most prefers of the two selected marks his 1 point and the other his 0 point. Then we can recalibrate everyone else's scales by picking items that are separated by the same preference interval and use them to mark their 0's and 1's. Two items x and y are separated on A's scale by the same preference interval as z and w are on B's scale just in case A's preference for x over y is exactly as great as B's for z over w .

However, now everything turns on our ability to compare one person's preference for one item over another with another person's preference for one item over another. In trying to come to grips with this, it is useful to turn to the problem of comparing one and the same person's preference for an item x over y with his preference for z over w . It turns out that Von Neumann-Morgenstern utility theory already provides a basis for such comparisons. For a person marks a greater preference interval between x and y than he does between z and w just in case $u(x) - u(y) > u(z) - u(w)$. But that in turn holds if and only if

$u(x) + u(w) > u(z) + u(y)$. That holds—as you can show in exercise 4—if and only if the person prefers the lottery L to L' , where

$$L = L(1/2, L(u(x), B, W), L(u(w), B, W))$$

$$L' = L(1/2, L(u(z), B, W), L(u(y), B, W)).$$

We now have placed the problem of the *intrapersonal* comparison of preference strengths on as firm a basis as that of utility theory. Unfortunately, there seems to be little hope for doing this for *interpersonal* comparisons. We cannot extend our lottery trick to the interpersonal case, because it presupposes that one and the same agent is choosing between lotteries. What is more, we can draw the interpersonal comparisons however we wish without affecting the intrapersonal ones. That is because no matter which positive linear transformation we use to draw interpersonal comparisons and convert one person's u -scale into his new u' -scale, $u(x) - u(y) > u(z) - u(w)$ holds if and only if $u'(x) - u'(y) > u'(z) - u'(w)$ does.

Thus we must take "A's preference for x over y is greater than B's for z over w " as a primitive notion that cannot be explicated in terms of our single-person utility theory. As we have indicated, we can use it to make enough interpersonal comparisons of preference to allow us to compare the utility units of different agents. Having compared their units, we can derive new interpersonal comparisons of preferences and check these for consistency with our initial judgments. To see how this works, suppose Mary and Sam have ranked a, b, c, d, e , and f and all lotteries constructed from them on 0 to 100 scales.

	0	10	11	15	20	30	35	88	100
	•	•	•	•	•	•	•	•	•
Mary	a	b	c		d		e		f
Sam	f	d		b				c	e

Now suppose we take Sam's preference for b over d as our standard unit. Also suppose we decide that Mary's preference for the lottery $L(1/2, c, b)$ over b is exactly as great as Sam's for b over d . That means that, according to the scales given above, 5 units of Sam's is actually the same as half a unit of Mary's. Suppose that to reflect this we transform Mary's scale by multiplying it by 10. Having done this, the distance on her scale between e and f is now 650 standard units whereas the distance between b and e on Sam's scale remains 85 standard units. An immediate consequence of this is that Mary's preference for f over e is greater than Sam's for e over b . If this consequence agrees with our independent comparisons of their preferences, our rescaling is confirmed. On the other hand, if it fails to agree with our comparisons, the comparison by which we standardized the two scales is inconsistent with our new one. In this case we must try to bring our comparisons in line with each other. If doing so happens to be very difficult or impossible, we should conclude that we do not know enough about Sam and Mary to compare their preference intensities and give up our attempt.

Although our procedure prevents us from using totally blind and irresponsible interpersonal comparisons of utility units, it makes sense only if it makes

SOCIAL CHOICES

sense to compare the preference intensities of different persons in the first place. If you are skeptical about that, nothing we have done will be of much assurance. One might try to respond to you by proposing links between interpersonal comparisons of preference intensities and overt behavior just as we forged a link between intrapersonal comparisons and choices between lotteries. However, as our parable illustrated, behavior is not always a reliable indicator of preference intensities. Thus the problem of the interpersonal comparison of utilities remains unresolved, and that is how we will leave it here.

PROBLEMS

1. Prove that if x yields more total utility than y , it continues to do so if we add different numbers to the individual citizens' utilities for x and y , so long as we add the same numbers to both x and y .
2. Show that the SWF G responds to both changes of origins and changes of units.
3. Explain why when we are concerned only with comparing the utility units of different persons, we can select any two items on any citizen's scale to serve as our standard unit and any other two separated by the same preference interval on each other citizen's scale to serve as his unit.
4. Calculate the expected utilities for L and L' and show that L is preferred to L' if and only if $u(x) - u(y) > u(z) - u(w)$.
5. Take Mary's preference for c over b as the standard unit, assume that her preference for c over b is exactly as great as Sam's for e over c , and transform his utility scale to bring it in line with Mary's.

6-5. References

An excellent treatment of most of the matters touched on in this chapter is to be found in *Sen*. *Arrow* is also valuable and contains the first proof of Arrow's theorem. *Luce and Raiffa* also cover most of the material of this chapter. *Smart and Williams* is a fine introduction to utilitarianism and its non-decision theoretic problems. For Harsanyi's theorem see *Fishburn*, *Harsanyi*, and *Resnik*. *Harsanyi*, *Sen*, and *Weirich* are good sources for material on the interpersonal comparison of utility.