Weighing Lives

John Broome

OXFORD
UNIVERSITY PRESS

7

Separability of times

In chapter 1, I posed the general problem for this book as setting a value on a distribution of wellbeing. Wellbeing is distributed across a grid, whose dimensions are people, time and states of nature, and all this wellbeing comes together to determine the overall value of the distribution. How? Now I am ready to start work on this problem.

I said I would particularly concentrate on two-dimensional distributions of people and times. Examples are shown in figure 11 on page 12 and figure 12 on page 24. How is wellbeing aggregated across a two-dimensional grid?

To approach this problem, we might naturally hope to aggregate wellbeing across the dimensions separately, taking them one at a time. We might hope to do this in either order: first across times and then across people, or first across people and then across times. In section 7.1, I shall describe these two routes to aggregation in a bit more detail. I call them, respectively, 'the people route' and 'the snapshot route'. The rest of this chapter explains why I shall not take the snapshot route. Section 7.2 presents a *prima facie* case against it. Section 7.3 describes and rejects a counter-case. Section 7.4 describes a different sort of counter-case, which section 7.5 rejects in turn. By the end of the chapter, I hope the snapshot route will seem unattractive. At any rate, I do not think we have good grounds for taking it. Chapter 8 starts work on the people route, which I prefer.

7.1 Routes to aggregation

The idea of the people route is this. A distribution gives each person some amount of temporal wellbeing at each time in her life. We first determine a value of the distribution for each person, by aggregating across all the times in the person's life. In one of my diagrams, we first scan across the horizontal lines. We arrive at a valuation for each person by looking along the line that represents the person in the

diagram. In some way we compound together all the temporal wellbeing we find spread along this line. We do this for every person – for every line in the diagram. That is the first step in the aggregation. The second step is to put together all these valuations – one for each person – that we have arrived at in the first step. This is a step of aggregation across people, and it brings us to an overall value for the distribution.

On the other hand, to take the snapshot route to aggregation, we first judge the value of the distribution at each time, by aggregating across all the people at that time. In a diagram, we scan a vertical rather than a horizontal line. We take everyone's temporal wellbeing at a single time, and put together all these amounts of wellbeing, to determine how good the distribution is at that time. This gives us what I shall call a 'snapshot' valuation of the distribution at that time. We make a snapshot valuation for every time. Then we aggregate all these snapshot valuations across time, to determine the overall value of the distribution.

Each of these approaches involves an assumption. Take the snapshot route. It implicitly assumes that a snapshot valuation makes sense. Scanning the distribution along the vertical line that represents a particular time, and looking only at all the people's wellbeing at that time, it assumes we can arrive at a valuation for that time. That is to say, it assumes we can value the distribution of wellbeing at that time independently of the rest of the distribution – independently of people's wellbeings at other times. It assumes aggregation is possible at any time independently of other times. Moreover, the temporal aggregates we arrive at – the snapshot valuations for each time – must be enough, taken together, to determine the overall value of the distribution.

This two-part assumption is technically known as 'separability of times'.[1] It can be expressed in terms of the value function that I set up in section 2.3. This function represents the goodness of the distribution in terms of the condition of each person at each time. It has this form in general:

(7.1.1) $v(g_1^1, g_2^1, \ldots g_T^1, g_1^2, g_2^2, \ldots g_T^2, g_1^3, g_2^3, \ldots g_T^3, \ldots)$.

If, and only if, times are separable in the distribution, this function can be expressed in the form:

(7.1.2) $\underline{v}(v_1(g_1^1, g_1^2, g_1^3\ldots), v_2(g_2^1, g_2^2, g_2^3\ldots), \ldots v_T(g_T^1, g_T^2, g_T^3\ldots))$,

where v_1, v_2, and so on are all real numbers, and the function $\underline{v}(\cdot)$ is increasing in each of its arguments. To put this more precisely: if and only if times are separable, there are functions $v_1(\cdot)$, $v_2(\cdot)$, and so on, such that the goodness of distributions can be represented by a value

function of the form (7.1.2), where $v(\cdot)$ is increasing in each of its arguments. (A function is said to be *increasing* in one of its arguments if and only if the value of the function increases whenever the value of the argument increases while the other arguments remain constant.)

In (7.1.2) we can interpret $v_1(\cdot)$, $v_2(\cdot)$, and so on as themselves value functions, representing snapshot valuations at each time. But (7.1.2) does not determine these functions very tightly. Any increasing transforms of them could be fitted into the separable formula (7.1.2) just as well. On page 27, I explained that this is the mark of an ordinal value function. Each of these functions represents the snapshot value of the distribution at a time, but only in an ordinal way.

The snapshot route to aggregation requires the assumption that times are separable. The people route requires a parallel assumption that people are separable. In formal terms, to assume people are separable is to assume that the value function can be expressed in this form:

$$\bar{v}(v^1(g_1^1, g_2^1, \ldots g_T^1), v^2(g_1^2, g_2^2, \ldots g_T^2), v^3(g_1^3, g_2^3, \ldots g_T^3), \ldots)$$

where the values of $v^1(\cdot)$, $v^2(\cdot)$, and so on are all real numbers, and the function $\bar{v}(\cdot)$ is increasing in each of its arguments. This is the formal condition, but its interpretation is not straightforward. I shall leave it to chapter 8, where it will also turn out that a slightly weaker assumption will suffice. I only mention it here for the sake of comparison. The rest of this chapter concentrates on the assumption that times are separable.

7.2 The objection to separability of times

Snapshot valuations are implicit in a great deal of our thinking about the progress of wellbeing. We talk about how well off people are in a country, and this means how well off they are in that country at the present time. We talk about improvement in wellbeing over time, and how inequality in wellbeing changes. All of this implicitly involves setting a value on the distribution of wellbeing at a time, and watching how that value develops over time. So even if we are not intending to take the second step of aggregating the snapshot values across time, snapshot values commonly have a place in our thinking.

However, I shall argue that separability of times is a dubious assumption, so snapshot valuations and the snapshot route to aggregation are both dubious. At least, I shall argue that separability of times commits us to much more than most of us would be willing to accept. My principal aim in this chapter is to make this point.

Table 1

		Times 1	Times 2
People	p	1	1
	q	Ω	Ω

Distribution *A*

		Times 1	Times 2
People	p	Ω	1
	q	1	Ω

Distribution *C*

		Times 1	Times 2
People	p	1	Ω
	q	Ω	1

Distribution *B*

		Times 1	Times 2
People	p	Ω	Ω
	q	1	1

Distribution *D*

I shall start with the example shown in table 1. (I often find it more convenient to present examples in tables rather than in diagrams like those in chapter 1.) For simplicity, the example contains only two people and two times. The four quadrants of the table show four different possible distributions. The cells show the two people's condition at each time. That is to say, each cell shows a person's wellbeing at a time, unless it contains an 'Ω' to indicate that the person is not alive at that time. Whenever a person is alive, in this example I have always made her wellbeing just one unit. This is for simplicity again.

In *A*, person p lives through both times and q never lives at all. In *B*, p lives for one time and then dies, and q lives instead for the second time. We might say that in *B*, q 'replaces' p. Similarly, in *D*, q lives through both times and p never lives at all. In *C*, p replaces q.

To construct this example, I have assumed that a particular person could come into existence at one time, or alternatively at a different time. This is possible. With modern reproductive technology, indeed, it is even practicable. At several places in this book I shall take its possibility for granted.

What can we say about the relative goodness of *A*, *B*, *C* and *D* in table 1? Compare *A* and *B* first. In the first period, these two distributions are identical: in both of them, p is alive and q is not. If times are separable,

we can judge between A and B on the basis of the two periods separately, and in the first period there is nothing to choose between them. Which is better therefore depends on the second period only. For the moment, let us not try and settle which is better in the second period.

Instead, let us now start to compare C and D. This pair of distributions is also identical in the first period. So if times are separable, which of the two is better depends on the second period only. Now, the difference between C and D in the second period is exactly the same as the difference between A and B in the second period. In the second period, A and C both have p alive and q not alive, and B and D both have q alive and p not alive.

It follows that A is better than B if and only if C is better than D, if times are separable. However, this is an intuitively implausible conclusion. A and D each have a particular advantage over B and C. They each contain one long life rather than two short lives. In A, p continues to live through both times, whereas in B she dies and is replaced by q. We normally think it better that a person continues to live rather than that she is replaced. So intuitively A may well be better than B and also D better than C. This is inconsistent with separability of times.

It is surely true that our intuition normally rates continued life better than replacement. For instance, we think it better to save the life of a baby if we can, rather than let her die, even if, were she to die, her parents would replace her with another baby. There may be limits to this intuition. I am not sure we would think it better to prolong a 100-year-old person's life for another 100 years, rather than have a new person live for 100 years. But we take prolonging life to be better than saving life in some cases, and that is enough to contradict separability of times.

To put it differently, we intuitively value longevity, at least within limits. Let me specify what I mean by valuing longevity. Suppose some total amount of time is lived by people, at some level of temporal wellbeing. For longevity to be valuable means that, given this fixed total of time and level of temporal wellbeing, it is better for the time to be divided up amongst fewer lives rather than amongst more lives.

In each quadrant of table 1, the total amount of life lived is the same, and the level of temporal wellbeing is constant. In A and D the time is all included in one life, whereas in B and C it is divided between two. So if longevity is valuable by my definition, A and D are better.

We normally value longevity. Yet the length of a person's life is not something that shows up when we look at times separately. It only appears when we look at several times together. So it is a value that

cannot show up in a snapshot valuation. This is a strong reason to doubt separability of times.

7.3 Separatism

The objection I raised against separability is based only on the intuitive thought that longevity is valuable. One possible response to this objection is simply to deny it. This is a hard-headed, intuitively implausible response, but it is possible. I shall call it the 'separatist' view. Separatism implies that times are separable.

When we look at the wellbeing of people at a single time – surveying a vertical line in one of my diagrams – we see a lot of little pieces of wellbeing, each belonging to a different person. In the snapshot, we do not see how the pieces of wellbeing are connected with other pieces in a horizontal direction. We do not see how pieces are packaged together to make up particular lives. That is why the length of people's lives does not appear in a snapshot valuation. Yet the horizontal packaging of wellbeing seems intuitively to make a real difference to value; it is better for wellbeing to be packaged into fewer rather than more lives. The separatist denies this.

Separatism implies that horizontal connections really make no difference. All that matters is the little pieces of wellbeing that appear at each time, not their packaging in lives. Consequently, longevity is not valuable – at least, not in the sense I defined. If a person's life is prolonged, and other things remain the same, the effect is to add pieces of wellbeing to the distribution at the end of her life. A separatist values those added pieces. So in this sense she values prolonging life. But she does not value the fact that the added pieces are joined on to the end of a life that already exists. To a separatist, they might just as well be put into a new life; replacement would be just as good as prolonging an existing life. If a person dies and a new person appears in her place to enjoy the same wellbeing, that is just as good as the first person's continuing to live.

The separatist opinion that horizontal connections make no difference is an extreme version of the view that personal identity does not matter, which was propounded by Derek Parfit in part III of *Reasons and Persons*. Parfit arrives at his own less extreme view on metaphysical grounds. I shall outline a metaphysical argument for separatism in section 15.2; this argument is simply an extension of Parfit's. On page 257, I shall explain how separatism can also follow from a particular

theory about the nature of wellbeing. So separatism can be given a philosophical basis, despite its implausibility.

I think separatism is a defensible view. However, it is innately implausible, and I do not think the defences are so convincing that I would rely on them. I shall not assume separatism in this book. I shall return to consider it in section 18.2.

I can mention now the simplest example of a separatist theory of value. It is a view I shall call *complete utilitarianism*. A complete utilitarian thinks the value of a distribution is simply the total of all the temporal wellbeing enjoyed at any time by anyone. To arrive at an overall value, we simply add wellbeing across the whole distribution: across people and across times. One distribution is better than another if and only if it has a greater total of wellbeing. A complete utilitarian does not care in any way about how wellbeing is distributed. For one thing, she does not care how it is packaged into individual lives. All that matters is wellbeing; who gets wellbeing is irrelevant.

7.4 Dispersing the value of longevity

In section 7.2 I suggested that times are not separable, and I am now considering responses to that suggestion. One counter-argument is the tough separatist one set out in section 7.3. But there is also a more pragmatic way to try and preserve separability in the face of the objection presented in section 7.2.

I said the value of longevity does not show up when we look at times separately, but only when we look at several times together. Consequently, it does not show up in snapshot valuations. But we can artificially make it show up. We can simply take a person's longevity – the length of her life – to be part of her temporal wellbeing at every time in her life. I shall mention a possible justification of this move on page 112–13. First I am going to compare it with other, parallel moves in the theory of value.

As a matter of terminology, I shall keep the term 'separatism' for the hard-headed view that gives no value to longevity. So I do not count this new defence of separability as a version of separatism. It is an example of a strategy within the theory of value that I call the 'dispersion' of value. I shall take a moderately doubtful stance towards the application of dispersion to longevity. Perhaps dispersion can save separability of times, but it will face difficulties in doing so. In *Weighing Goods*, I endorsed dispersion in other contexts, and used it to defend separability

in those contexts. But I shall certainly not endorse dispersion for longevity. Consequently, I am not willing to assume separability of times. In this section, I shall explain the difficulties that dispersion encounters.

For the sake of comparison, let me start by mentioning a different application of this strategy. It defends separability of people, rather than separability of times. Separability of people can seem to be threatened by the value of equality. If equality is valuable, one might think its value cannot belong to any individual person, because it depends on a relation between people: it depends on the difference between some people's wellbeing and others'. It seems we first need to determine how well off each person is, before we can see how equal the distribution is. Consequently, the value of equality must be a purely social value, separate from each individual's wellbeing. So one might argue, and this argument threatens separability of people. It is developed in section 8.2 of this book.

Section 8.2 also answers this argument, using a counter-argument from *Weighing Goods*.[2] If equality is indeed a good thing, it is so because it is good for individual people. Conversely, if inequality is a bad thing, it is so because it is bad for individual people. The badness of inequality is not a separate negative value beyond people's own wellbeing. It is part of people's wellbeing itself. This argument is an example of dispersion. At first, equality seems to be a social value that cannot be captured within the wellbeing of individuals. But I took this apparently social value and dispersed it among the individuals. I argued it actually belongs to the individuals separately, despite first appearances.

It may be helpful if I compare dispersion with the assumption of distribution described in section 3.3. The assumption of distribution rules out goods that do not show up in the distribution of wellbeing. It rules out some types of good but not what I called on page 44 'pattern goods'. These are goods whose presence shows up in the pattern of the distribution, though not at any particular location in the distribution. Even if equality of wellbeing is not an individual good, it is certainly a pattern good. So it is consistent with the assumption of distribution.

However, the effect of separability is to rule out pattern goods; it requires all goods to appear at particular locations in the distribution. For example, if people are separable, we can evaluate a distribution by looking at each person separately; we do not have to look for patterns across people. The strategy of dispersion aims to support separability, so it needs to go beyond the assumption of distribution. It takes a particular good that at first seems a pattern good, and locates it at a

particular place, or several particular places, in the distribution. For instance, it locates the good of equality, which at first seems a pattern good, at particular people.

An essential part of this strategy of dispersion is to justify it.[3] I ought not really to call it a 'strategy' at all. That term suggests we theorists of value are free to disperse value as we like. But truly the strategy is to make the substantive claim that value actually is dispersed. Take any value that apparently does not belong to people as individuals, so it seems to contradict separability of people. Let it be the survival of a culture, say; suppose this is indeed valuable. A theorist could always disperse this value in a formal way by arbitrarily dividing it into parts and allocating each part to a person. She could then say that each person's allocated part was a component of that person's wellbeing. But this would not be a correct move unless the theorist could demonstrate that the survival of the culture was indeed part of the person's wellbeing.

The justification needs to be made for whatever specific value is in question. In the case of the survival of a culture, it might or might not be convincing. In the case of equality, I hope I made it convincing in *Weighing Goods*[4] by producing an account of the value of equality in terms of fairness. This account made it definitely an individual value.

Unless an instance of dispersion is genuine in this way, it will be no help with aggregation. If we know the aggregate value of a distribution, we can always disperse the value arbitrarily to people. But to do that, we would have to know the aggregate value first. Since we are trying to find aggregate value, this arbitrary manoeuvre would be pointless. On the other hand, if the value is genuinely a part of each individual's wellbeing, we should be able to assess the amount of it individually for each person, and then include it in an aggregation across people.

Is it justified to disperse the value of longevity to individual times? If we knew the aggregate value of a distribution, including the value of longevity, it could be arbitrarily divided up and dispersed among the times separately. So without doubt the aggregate can be arbitrarily dispersed in such a way that the dispersed parts could be re-aggregated to reach the correct result. But as I say, arbitrary dispersion is not enough. The value of longevity must genuinely be dispersed.

I doubt it is. I suspect that the value of longevity truly belongs to a life as a whole, and not to any of the separate times within the life. But I shall not insist on that, and I recognize that dispersion can be defended. Here is one possible defence.[5] In our lives we undertake various projects that take time to complete. Leaving a project uncompleted is a bad thing; it is better for us to complete each project, rather than not.

Moreover, the goodness of completing a project is located at all the time we are engaged on the project. If a project is completed, that gives a special value to the time we spend on it. It increases our wellbeing during that time. Conversely, if a project is not completed, all that time is devalued. Death generally cuts off some of our projects. The shorter our life, the fewer completions we shall achieve. This tends to devalue the time we are alive; it reduces our wellbeing while we live. In a long life we shall complete more, and this tends to add value to the time we are alive. Here is a possible explanation of why longevity is valuable, and it locates its value at the times we are alive. It gives longevity a dispersed value.

7.5 Can dispersion be successful?

I shall not try to assess the justification for dispersion. Instead, I shall show that, in any case, it is difficult to make dispersion do the job we require from it. It may not succeed in supporting separability of times against the challenge of longevity, which is its purpose.

The idea of dispersion is to divide up the value of a person's longevity and allocate the parts to the various times in the person's life. Let us try out this idea in the example of table 2. The cells of the table show the people's conditions at different times. Anyone who is alive at any time has the same level of wellbeing, 1. I mean this to be the level of her

Table 2

		Times			
		1	2	3	4
People	1	1	Ω	Ω	Ω
	2	Ω	1	Ω	Ω
	3	Ω	Ω	1	Ω
	4	Ω	Ω	Ω	1
	5	1	1	1	1

Distribution *A*

		Times			
		1	2	3	4
People	1	1	1	Ω	Ω
	2	1	1	Ω	Ω
	3	Ω	Ω	1	1
	4	Ω	Ω	1	1
	5	Ω	Ω	Ω	Ω

Distribution *B*

wellbeing without taking account of the value of longevity. I shall add this supposed value later.

In both *A* and *B*, if we take a snapshot of the distribution at any time, we see two people alive, each having wellbeing 1. In this example, I shall assume impartiality between the people. By this I mean that the bare identity of the people cannot make a difference to value. A snapshot taken at each time will therefore have to account *A* and *B* as equally good at the time.

If times are separable it would follow that, as a whole, each of these two distributions is equally good, provided we do not accord any dispersed value to longevity. A separatist will think they are indeed equally good, but that is because she gives no special value to longevity. If we do value longevity, we shall value *B* above *A*. In either distribution, the total time lived is eight periods. In *A* this total is divided amongst five people. In *B* it is divided amongst four people. If longevity is valuable, according to the definition I gave on page 108, *B* is better than *A*.

This conclusion has nothing to do with the inequality in *A*. In this example, let us ignore any value equality may have.

Can we reach the same conclusion in a separable fashion, by means of dispersion? This might depend on how, specifically, we disperse the value of longevity to individual times, and there are many ways of doing that. Whichever way we choose will need to be justified, but I shall now set aside questions of justification. I shall just try out some obvious ways of doing the dispersion.

Table 3

	Times			
	1	2	3	4
1	1+ a	Ω	Ω	Ω
2	Ω	1+a	Ω	Ω
People 3	Ω	Ω	1+a	Ω
4	Ω	Ω	Ω	1+a
5	1+4a	1+4a	1+4a	1+4a

Distribution *A*

	Times			
	1	2	3	4
1	1+2a	1+2a	Ω	Ω
2	1+2a	1+2a	Ω	Ω
People 3	Ω	Ω	1+2a	1+2a
4	Ω	Ω	1+2a	1+2a
5	Ω	Ω	Ω	Ω

Distribution *B*

One way is to suppose that the length of a person's life adds to her wellbeing at all times in her life, in proportion to its length. Table 3 shows the people's wellbeing adjusted according to this principle. (The number a is the constant of proportionality.) With the dispersion done, we can now try evaluating the distribution at each time separately. At each time, people's total wellbeing in *A* is greater than in *B*. So is their average wellbeing. If we assume separability of times, these facts strongly suggest *A* is better than *B* at every time. Consequently, they strongly suggest *A* is the better distribution. But if longevity is valuable, we have previously concluded the opposite: that *B* is the better distribution.

To be sure, the conclusion that *A* is better than *B* is not inescapable. At any time, we need not aggregate across people simply by taking their total wellbeing, or their average wellbeing. But I cannot see what other procedure could be justified. True, *A* is an unequal distribution, whereas *B* is equal, and the value of equality might justify a different way of aggregating. But I have set this consideration aside because the value of longevity makes *B* better than *A* independently of the value of equality. So this way of doing the dispersion does not achieve what it aimed to achieve.

A second way is to suppose that a person's wellbeing at a time is increased by the amount of life she has in front of her. Table 4 shows this possibility. Again, with the dispersion done, let us try evaluating the distribution at each time separately. Distribution *A* turns out to be

Table 4

	Times			
	1	2	3	4
1	1	Ω	Ω	Ω
2	Ω	1	Ω	Ω
People 3	Ω	Ω	1	Ω
4	Ω	Ω	Ω	1
5	1+3a	1+2a	1+a	1

Distribution *A*

	Times			
	1	2	3	4
1	1+a	1	Ω	Ω
2	1+a	1	Ω	Ω
People 3	Ω	Ω	1+a	1
4	Ω	Ω	1+a	1
5	Ω	Ω	Ω	Ω

Distribution *B*

ahead of *B* in total and average wellbeing in two out of the four times, and behind *B* in only one. Once again, the conclusion that *A* is better is strongly suggested, if times are separable. Again, I can see no credible way of escaping this conclusion. So this way to disperse the value of longevity seems no more successful than the other in preserving separability.

I recognize that this case against dispersion is not conclusive. My arguments are not watertight, and I have not considered all the ways the dispersion might be done. Still, my example does show a fundamental difficulty in making dispersion successful. The trouble is that long-lived people live a long time. They live at many times, therefore. This means their wellbeing gets counted more often in a separable evaluation than does the wellbeing of shorter-lived people. If we disperse the value of longevity into people's wellbeing at times, the effect is that too much weight will be given to the lifetimes of the longer-lived people. For this reason, I doubt that dispersion can work.

The upshot of all this argument is that I doubt times are separable. Perhaps they are, and perhaps some argument can be found to show they are. But in this book, I shall not assume they are. I shall therefore not rely on snapshot valuations, and I shall avoid the snapshot route to aggregation.

Notes

1 Separability is explained in my *Weighing Goods*, chapter 4.
2 Section 9.4.
3 See the discussion of dispersion in *Weighing Goods*, pp. 191–2.
4 Section 9.4.
5 Rüdiger Bittner and Rolf Stoeker made this argument to me.

8
Separability of lives

In section 7.1, I described two alternative routes to aggregating wellbeing across a two-dimensional distribution of wellbeing. I then argued we should not follow the snapshot route. It relies on the assumption that times are separable, which I find dubious.

The alternative I favour is the people route. It also relies on a separability assumption. This chapter starts in section 8.1 by distinguishing two related separability assumptions: separability of people and separability of lives. Separability of lives is the weaker of the two, but it will turn out in section 8.4 to be sufficient for justifying the people route to aggregation. I shall postpone to chapter 13 a discussion of whether people are separable. But I shall argue that lives are separable in section 8.2 of this chapter.

The main basis of my argument is something I call the principle of personal good. Section 8.3 explains that this principle is inconsistent with discounting future wellbeing. If we accept discounting, that will force a modification on the principle. However, it will do no damage to the argument that lives are separable.

Section 8.5 introduces a technical assumption we shall need in due course.

8.1 Separability of people and separability of lives

To assume people are separable is to assume that the value function

$$v(g_1^1, g_2^1, \ldots g_T^1, g_1^2, g_2^2, \ldots g_T^2, g_1^3, g_2^3, \ldots g_T^3, \ldots)$$

can be expressed in this specific form:

(8.1.1) $\bar{v}(v^1(g_1^1, g_2^1, \ldots g_T^1), v^2(g_1^2, g_2^2, \ldots g_T^2), v^3(g_1^3, g_2^3, \ldots g_T^3), \ldots)$,

where the values of $v^1(\cdot)$, $v^2(\cdot)$, and so on are all real numbers, and the function $\bar{v}(\cdot)$ is increasing in each of its arguments. To put the assumption differently: there are functions $v^1(\cdot)$, $v^2(\cdot)$, and so on such that the