Most Ways I Could Move: Bennett's Act/Omission Distinction and the Behaviour Space

Fiona Woollard
University of Southampton
f.woollard@soton.ac.uk

The distinction between action and omission is of interest in both theoretical and practical philosophy. We use this distinction daily in our descriptions of behaviour and appeal to it in moral judgements. However, the very nature of the act/omission distinction is as yet unclear. Jonathan Bennett's account of the distinction in terms of positive and negative facts is one of the most promising attempts to give an analysis of the ontological distinction between action and omission. According to Bennett's account, an upshot is the result of an agent's action if and only if the relevant fact about her conduct is positive. A proposition about an agent's conduct is positive if and only if most possible movements of the agent would not have made that proposition true. However, Bennett's account will fail unless it is possible to make sense of claims about 'most possible movements of the agent'. We need a way of comparing the size of subsets of the behaviour space (the set of possible movements). I argue that Bennett's own method of comparison is unsatisfactory. I present an alternative method of comparing subsets of the behaviour space.

1. The act/omission distinction

Some things occur because I do something; other things occur because I do not do something. The words on this page appeared because I typed them. The dirty dishes remained in the sink because I did not wash them. This distinction is known as the act/omission distinction: something is the result of my action if (and only if) it occurred because I did something; it is the result of my omission if (and only if) it occurred because I did not do something. The act/omission distinction is one of the fundamental ways we divide up the world of human behaviour. It is of great interest in both theoretical and practical philosophy. It is a natural target of investigation for those working in the ontology of action, but is also relevant to causation, where we face the task of explaining whether and how omissions can

Mind, Vol. 120 . 477 . January 2011
doi:10.1093/mind/fzr012

be causes, and the law, which treats harmful action and harmful omission very differently.

My interest in the act/omission distinction began in considering its role in normative ethics. Commonsense morality seems to assign significance to the act/omission distinction, particularly in cases involving harm to others. It seems to matter morally whether harm comes to another because I did something or because I did not do something. It is obviously wrong to send poisoned food to people in underdeveloped countries. However, it seems permissible not to send them food even if they will die without it. Good reason is required for failing to throw a life preserver to a drowning man, but a much stronger reason is required to justify holding his head under water or moving the life preserver out of his reach. Harmful actions (sending the poisoned food, holding down the man's head, moving the life-preserver) seem harder to justify than harmful omissions (not sending healthy food, not throwing the life-preserver).[1]

It seems that if the act/omission distinction (or some distinction like it) is not of moral significance, then we need to make radical revisions to commonsense morality. Either harmful omissions are much harder to justify than we think they are — so we should be doing much more to help the needy — or harmful actions are much easier to justify than we think they are — so we may kill to protect our personal projects. Such adjustments will leave morality either far more demanding than commonsense suggests or far more permissive. So a morality which did not assign moral significance to the act/omission distinction would be highly counterintuitive. None the less, it is far from obvious that the act/omission distinction can bear this moral

[1] These examples are frequently cited in discussion of the moral significance of the act/omission distinction and the doing/allowing distinction. The poisoned food case is referred to by Philippa Foot (Foot 1967, p. 273) and Jonathan Glover (Glover 1977, p. 93) among others. Versions of the life-preserver case are mentioned by Warren Quinn (Quinn 1989, p. 368) and Kadri Vihvelin and Terrance Tomkow (Vihvelin and Tomkow 2005, p. 184). My focus in this paper is the act/omission distinction rather than the doing/allowing distinction. In the example, moving the life-preserver out of the drowning man's reach may count as merely allowing harm. This is easier to justify than holding the man's head under water until he drowns. For example, it may be permissible to retrieve the life-preserver to save your own life or to save the lives of five others, but impermissible to hold the man's head under water for this reason. None the less, moving the life-preserver still involves action. Moreover, this harmful action seems harder to justify than merely failing to throw him a life preserver. It would be impermissible to move the life-preserver out of the man's reach to save one other person's life or to avoid sacrificing your savings. This suggests (a) that the act/omission distinction and the doing/allowing distinction are not the same, for an agent may allow harm through an action; (b) that commonsense morality treats both distinctions as morally significant. I thank the anonymous referee whose comments made me clarify this.

weight. When the consequences are the same in each case, how can it matter whether these consequences occur because I did something or because I did not do something?

The best way to approach these issues is to investigate what the difference between action and omission *is*. What does it mean to say that a given upshot occurred because I did something rather than because I did not do something? Why do typing, posting food, and grabbing a life-preserver count as actions while leaving the dirty dishes, not sending food, and not throwing a life-preserver count as omissions? It is only when we understand the nature of the act/omission distinction that we will be able to start working out whether it matters morally whether an outcome occurs because of action or omission, whether we should draw a legal distinction between actions and omissions, and whether, and how, an omission can be a cause.

One of the most promising accounts of the act/omission distinction is Jonathan Bennett's account based on the distinction between positive and negative facts, presented first in his Tanner Lectures and then in his book *The Act Itself* (Bennett 1981, 1995).[2] Bennett illustrates his account using two examples, Push and Stayback, in which a vehicle is on a slope leading to a cliff edge. In Push, the vehicle is standing at the top of the slope; Agent pushes it, and it rolls over the cliff edge to its destruction. In Stayback, the vehicle is already rolling; Agent could, but does not, interpose a rock that would stop it, and the vehicle rolls over the cliff edge to its destruction (Bennett 1995, p. 67). The destruction of the vehicle is the result of Agent's action in Push, but the result of mere omission in Stayback.

According to Bennett's account, an upshot of an agent's behaviour is the result of an action if and only if the relevant fact about the agent's behaviour is positive; an upshot is the result of an omission if and only if the relevant fact about the agent's behaviour is negative. When an upshot is the result of an agent's action or omission, the upshot depends in some way on some fact about the agent's conduct. I will say that a fact upon which the upshot depends in this way is *relevant* to the upshot. Cases of pre-emption have taught us that the dependence will not be simple counterfactual dependence.[3] The

[2] Bennett himself argues that we should abandon talk of acts and omissions and refer instead to negative and positive facts about an agent's behaviour (Bennett 1995, pp. 29–88). I think it is more fruitful to retain the widely used terminology of acts and omissions and analyse these concepts in terms of positive and negative facts.

[3] The literature on causation is full of examples of preemption and overdetermination. For a classic discussion, see Lewis 1986.

vehicle's destruction is the result of Agent's push even if someone else would have pushed the vehicle down the slope if Agent had not done so. So there are interesting issues about how exactly to specify the relationship of relevance.[4] There are also interesting issues about how to pick out *the* relevant fact. For the purposes of this paper I will rely on our intuitive ability to pick out the relevant fact about an agent's conduct. In Push, the relevant fact about Agent's conduct is that he pushed the vehicle. We say that the vehicle was destroyed because Agent pushed it. We would not say that the vehicle was destroyed because Agent did not read a magazine or because he pushed with his right hand. In Stayback, the relevant fact about Agent's conduct is that he did not interpose the rock. We say that the vehicle was destroyed because Agent did not place the rock in its path. We would not say that the vehicle was destroyed because Agent performed a cartwheel.

A positive fact about an agent's behaviour tells us what he has done: it picks out some action, x, and tells us that he has performed x. In contrast, a negative fact about an agent's behaviour tells us what he has not done: it picks out some action, y, and tells us that it is not true that he has performed y. The fact that Agent pushed the vehicle in Push is a positive fact about his conduct. The fact that Agent did not interpose the rock in Stayback is a negative fact about his conduct: it is the negation of the positive proposition 'Agent interposed the rock'.

To avoid circularity, we need an account of the distinction between positive and negative facts about an agent's conduct which does not appeal to the notion of picking out an action. Bennett's account is based on the idea that the positive/negative distinction is a distinction in how informative the relevant fact is. Positive facts tell us something fairly definite, pinning us down to a small set of alternatives. In contrast, negative propositions do not tell us very much about the world, only ruling out the relatively small set of alternatives corresponding to the positive proposition that has been negated.

[4] When an agent is relevant to an upshot, this will often be because some fact about his behaviour is causally relevant to the upshot. However, relevance is not always causal relevance. There are non-causal consequences of behaviour. A non-causal consequence of an agent's behaviour is a state of affairs that occurs as a result of the agent's behaviour, but that is not causally connected to the agent's behaviour (see Bennett 1995, p. 127). Suppose an agent has promised not to sit down. If he sits down, his behaviour does not cause his promise to be broken; the connection between his sitting and the breaking of the promise is not a causal one. Rather, in sitting down he breaks his promise. The fact that he sat down, when conjoined with the fact that he has promised not to sit down, makes it the case that he has broken his promise. We can ask whether an agent broke a promise by action or omission.

To analyse what it is for a fact about an agent's behaviour to be positive (or negative), Bennett limits the set of alternatives under consideration. Rather than considering all states of affairs (which would be seriously problematic), Bennett considers all possible ways the agent could move. Instead of asking whether a fact is informative overall, Bennett asks whether a fact is informative about the agent's movements. If a fact about an agent's conduct is negative, this fact does not tell us much about his movements; whereas if a fact about an agent's conduct is positive, this fact tells us that he moved in one of a relatively small number of ways.[5] Thus:

> A proposition is a *negative proposition* about the conduct of an agent if and only if most possible movements of the agent's body are such that if he had moved that way, the proposition in question would have been true.
>
> A proposition is a *positive proposition* about the conduct of an agent if and only if most possible movements of the agent's body are such that if he had moved that way the proposition in question would not have been true. (Bennett 1995, pp. 91–2)

Bennett uses a square to represent the agent's behaviour space — all the ways the agent could have moved. If P is a negative proposition about an agent's behaviour, then P will correspond to a large subspace of the behaviour space.[6] This subspace will be much larger than the subspace corresponding to the proposition not-P. The converse will be true if the proposition about the agent is positive. This is illustrated in Figs 1 and 2.

In Push, the relevant fact about Agent's conduct is that he pushed the vehicle. Most of the things Agent could have done would not have involved pushing the vehicle. Intuitively, Push is represented by something like Fig. 3.

[5] Bennett's account analyses the act/omission distinction in terms of bodily movements. None the less, action is not the same thing as bodily movement. Whether a person is waving or simply stretching depends on his understanding of what he is doing as well as on his physical movements. A full theory of action should explain this. However, this does not affect the act/omission distinction. Whether I am fasting or just not eating depends on my understanding of my behaviour, but fasting is still an omission not an action. Fasting is not eating for certain reasons. The act/omission distinction is based on whether a certain fact picks out that the agent did something or that the agent did not do something. It is not affected by the fact that some ways of describing what we did or did not do stipulate that these acts or omissions must be performed with certain intentions.

[6] Each point in the square represents a proposition stating that the agent moved in some absolutely specific way. Subspaces of the behaviour space are areas of the square; each area corresponds to the disjunction of all the propositions represented by the points it contains. See Bennett 1995, p. 91.

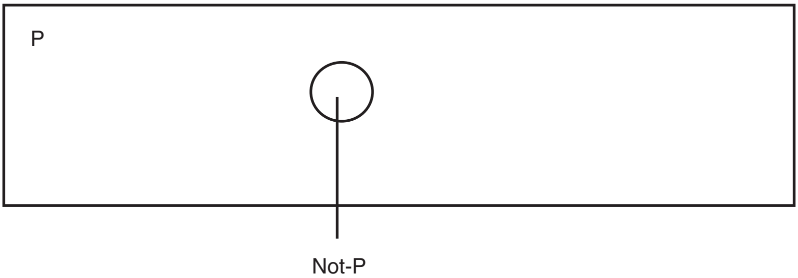

Fig. 1. P is a negative proposition.

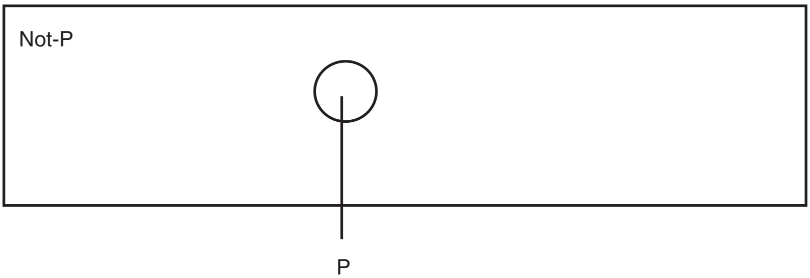

Fig. 2. P is a positive proposition.

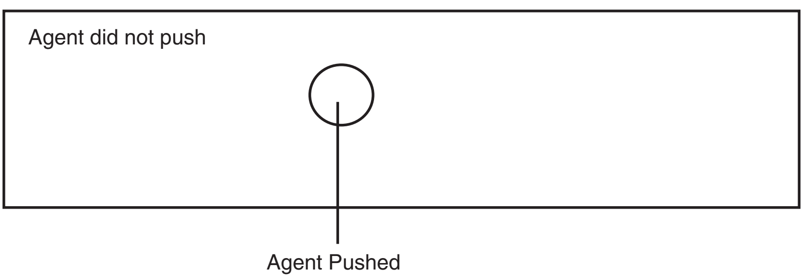

Fig. 3. Push.

The proposition 'Agent pushed the vehicle' corresponds to a very small subset of the behaviour space. It is a positive fact about Agent's behaviour. Thus, in Push, Agent is relevant to the vehicle's destruction through a positive fact about his behaviour. This fits the intuition that in Push Agent is relevant to the vehicle's destruction through action rather than omission.

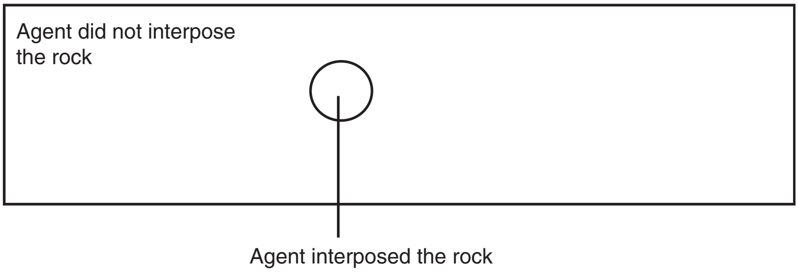

Fig. 4. Stayback.

In Stayback, the relevant fact about Agent's behaviour is that he did not interpose the rock. Most of the things Agent could have done would have involved failing to interpose the rock. The proposition 'Agent did not interpose the rock' will only be false in a very small number of cases — those cases where Agent does actually interpose the rock. Thus it seems intuitively as if the situation in Stayback will be represented by something like Fig. 4.

The proposition 'Agent did not interpose the rock' corresponds to a very large subset of the behaviour space. It is a negative fact about Agent's behaviour. Thus, in Stayback, Agent is relevant to the vehicle's destruction through a negative fact about his behaviour. This fits with the intuition that in Stayback Agent is relevant to the vehicle's destruction through an omission rather than an action.

Bennett's account is an attractive analysis of the act/omission distinction. It gives us the correct results in a wide range of cases. It seems to fulfil Bennett's desiderata of stating the distinction 'in terms which are clear, objective, and deeply grounded in the natures of things' (Bennett 1981, p. 48). Whether most ways the agent could have moved would have made a given proposition true does not seem to depend on how we describe things or think about things. Moreover, it is intuitively appealing to link the act/omission distinction to a positive/negative distinction analysed in terms of how informative the relevant fact is about the agent's movements. This fits with the idea that performing an action involves doing something fairly specific whereas omitting merely involves avoiding some fairly specific behaviour.

Although various putative counterexamples have been put forward, I believe that a Bennett-style approach can deal with them.[7] I will

[7] Common counterexamples to Bennett's account fall into three main groups: (1) immobility counterexamples; (2) pre-emption counterexamples; (3) restricted-range-of-movement counterexamples. Immobility counterexamples trade on the intuition that staying still is an

not discuss these counterexamples here. In this paper, I focus on a more fundamental objection to Bennett's approach, an objection that if sound would prevent the whole analysis from getting off the ground. The objection is that we cannot compare the size of subsets of the behaviour space; it does not make sense to speak of 'most possible movements of an agent's body'. For Bennett, whether an agent is positively or only negatively relevant to a certain event depends on whether 'most possible movements of the agent's body' would have made the relevant proposition about the agent's conduct true. So if we cannot compare the size of subsets of the behaviour space, Bennett's account will not even make sense.

I argue that the method Bennett gives for comparing subsets of the behaviour space is not satisfactory. It does not enable us to compare the size of subsets of the behaviour space. Moreover, using this method could result in contradictory results such that it appears that a given subset is both smaller than and bigger than another subset. None the less, as I shall show, it is possible to compare the size of subsets of the behaviour space. It makes sense to speak of 'most possible movements of an agent's body'. This is an important result. Many people are (rightly!) suspicious of Bennett's method of comparing the size of subsets of the behaviour space and seem inclined to reject Bennett's overall approach for this reason. So showing that it does make sense to compare subsets of the behaviour space is an important step in giving Bennett's account the hearing it deserves.[8]

My way of comparing the size of subsets of the behaviour space uses quite complicated mathematics. It may seem absurd to propose that such complicated mathematical notions can be the basis for the act/omission distinction. After all, this distinction is used by almost everyone almost every day. It is crazy to suggest that we perform such

omission, thus if a significant upshot will occur if and only if the agent stays still, Bennett will classify the agent's relevance as positive when it is intuitively negative (see Dinello 1971, Locke 1982, Quinn 1989). Bennett has responded to this objection in *The Act Itself* (Bennett 1995, pp. 96–100, 112–14). I discuss pre-emption counterexamples in Woollard 2008, pp. 59–67. Restricted-range-of-movement counterexamples involve restrictions on the movements open to the agent. An almost-paralysed man can only make three movements: leaving his finger where it is, pressing it down, or pressing it further down. Either pressing-down motion will set off a button, calling the nurse. Bennett seems to classify calling the nurse as an omission. I suggest that we can avoid this problem by paying proper attention to what makes a movement count as 'possible' (see Woollard 2008, pp. 70–72).

[8] Judith Jarvis Thomson raises concerns about this in Thomson 1996, p. 550 and n. 7. I have also encountered this worry in conversation more times than I can count.

feats of mathematics before drawing this most commonplace of distinctions.

I do not propose that anyone actually performs these calculations before drawing distinctions between action and omission. My mathematical analysis is intended as a foundation that underlies a distinction that we are naturally equipped to detect (more or less) reliably. We can use mathematics and physics to calculate the point at which a bouncing ball could be caught. These calculations are fairly difficult. A dog certainly could not carry them out, yet a dog can catch a ball. The dog simply sees how the ball will bounce. I suggest that the mathematical measure underlies our instinctive judgements about 'most of the ways the agent could move' in the same way as the physical calculations underlie the dog's seeing where to catch the ball. We do not appeal to this mathematical measure when we draw the distinction, but it is important to know that we have this foundation to support our intuitive judgements.

Before beginning, I will note two things about my Bennett-style approach. This approach associates actions with positive facts about an agent's behaviour and omissions with negative facts about an agent's behaviour. At any given time, there are many different facts about an agent's behaviour. At this moment, I am working hard; I am not drinking wine; I am typing on the keyboard. Some of these facts are positive facts, some negative facts. It follows that I am currently performing many actions and many, many omissions.[9]

However, the account does not imply that actions must correspond to particular movements. Pushing the vehicle counts as an action even though there are quite a few different ways Agent could move his body in pushing the vehicle: he could push with his right hand or his left, from a standing start or with a run up, and so on. That Agent pushed the vehicle counts as a positive fact as long as most of the ways he could move would not make it true that he pushed the vehicle. However many different ways Agent could push the vehicle, there are far more ways he could move that would not involve pushing the vehicle. The same applies to actions that can

[9] Some of these actions and omissions are related in interesting ways. My working hard is a complex action constituted by my typing and other more basic actions. However, on the fact-based account such interconnections do not give rise to difficulties about how to individuate actions: we must simply ask whether the fact that I am working hard is a different fact from the fact that I am typing.

be constituted by what we would recognize as 'different movements'. A person might propose marriage by getting down on one knee, writing a message in the sky, or hiding a diamond ring in his lover's dessert. Given the conventions of our society, each of these very different movements counts as a way of proposing. Bennett's account still classifies proposing marriage as an action because there are far more ways a person can move his body without proposing marriage than ways of moving his body which do involve proposing marriage.[10]

2. Bennett's measure

Bennett proposes that we compare subsets of the behaviour space using the notion of 'specificity'. Two propositions about how the agent moves take up the same area of the behaviour space if and only if they are equally specific (Bennett 1995, p. 93). Bennett then suggests that we use such 'comparable pairs' to show that one subset is smaller than another (Bennett 1995, p. 95).

In Stayback, the vehicle will be destroyed unless Agent interposes a rock. Bennett asks us to consider the different ways Agent could interpose the rock. He claims: 'A few dozen pairwise contrary propositions would pretty well cover the possibilities, each identifying one fairly specific sort of movement which would get the rock into the vehicle's path' (Bennett 1995, p. 95). We can thus divide up the 'Interpose' subset of the behaviour space into a few dozen smaller regions, each representing a push or kick. Bennett suggests that we pair each of these regions up with an 'echo' in the 'Non-Interpose' subset: 'by which I mean a proposition which is very like it except its truth would not rescue the vehicle' (Bennett 1995, p. 95).[11] Bennett illustrates this as follows:

> If [Interpose] contains a proposition attributing to Agent a certain kind of movement with his left foot, an echo of it might attribute to him a similar movement of that foot but angled so that the foot misses the rock. (Bennett 1995, p. 95)

[10] I thank the anonymous referee who pressed me on this.

[11] Bennett uses the subsets 'Survive' and 'Destroy' instead of 'Interpose' and 'Non-Interpose'.

Pairing up each of the propositions with an echo, ensuring that the echoes are pairwise contraries so that their regions do not overlap, gives us something like the following picture:

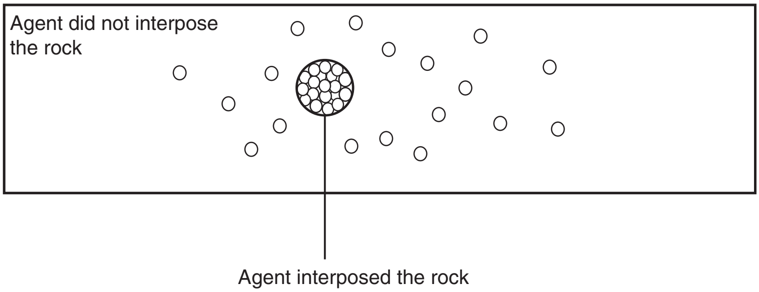

Fig. 5. Stayback.

Commenting on this picture, Bennett observes:

> My 'degree of specificity' criterion secures that the combined area [of the echoes] is the same as that of [Interpose]; and clearly they only take up a tiny proportion of [Non-Interpose]. (Bennett 1995, p. 96)

Initially, Bennett's argument seems convincing. However, his method of comparison depends upon three assumptions:

(1) We can make sense of the idea of two propositions being 'equally specific' in a way that means they will take up equal areas of the behaviour space

(2) The echoes that Bennett proposes will be equally specific in this way

(3) The Interpose subset can be covered by a finite number of fairly specific pairwise contrary propositions

I shall argue that we cannot accept these assumptions. To be able to pair off our propositions with equally specific echoes, we would have to use such detailed propositions that infinitely many would be required to cover the Interpose subset. On the other hand, to cover the subset with finitely many pairwise propositions we would need to use propositions of a level that would not allow us to match for specificity.

I begin by considering (1), the claim that we can make sense of two movements being equally specific. Bennett argues:

> This will not work with every determinable, e.g. with colours, because for them we have no agreed measure of specificity; but we have such measures

for space and time, and thus for movements and for specificity of propositions about movements. (Bennett 1995, p. 93)

We can determine when two propositions about an object's position in space are equally specific. Let P be the proposition that the object lies in a cube with sides of length 1 cm centred on coordinates (1, 3, 5). Any proposition which states that the object lies within a cube with sides of length 1 cm is exactly as specific as P about the object's position. Generally, two propositions about an object's position in space are equally specific if the areas within which they locate the object are the same size. Statements about the speed of an object at a given time and about the direction of movement can also be compared for specificity in fairly obvious ways.

Unfortunately, this method is not easily extended to propositions about bodily movements. Bodily movements involve different parts of the body moving in different ways. Our ordinary ways of describing bodily movements bring in statements about space and time in different ways. There is no canonical way in which ordinary language builds statements about movement out of statements about space and time. This makes it difficult to extend the specificity criterion.

We might say that two propositions about movement are equally specific if they bring in propositions about space and time in the same way and to the same degree of specificity, and say that propositions that bring in space and time in different ways are incomparable. Thus *He walks northwards* is as specific as *He walks southwest*. *He moves his left leg forward at 1 cm/s^{-1} to the nearest cm/s^{-1}* is exactly as specific as *He moves his left leg forward at 6 cm/s^{-1} to the nearest cm/s^{-1}*. *He walks northwards* and *He moves his left leg forward at 1 cm/s^{-1}* are incomparable.

This strategy faces several problems. It may be unclear if two propositions bring in space and time in the same way. Can we compare the specificity of *He moves his* ***left*** *leg forward at 1 cm/s^{-1} to the nearest cm/s^{-1}* with that of *He moves his* ***right*** *leg forward at 1 cms^{-1} to the nearest cm/s^{-1}*? What about comparisons between propositions about the speed with which he moves his right leg and propositions about the speed with which he moves his left arm?

More seriously, this strategy may give contradictory results. Two subsets of the behaviour space may come out as the same size using one set of propositions for comparison, but as different sizes using another set. According to this strategy, *He moves his head vertically* *upwards* is to be equated with *He moves his head vertically downwards*. However, suppose that we consider more detailed propositions about the movements the agent must make with his body. Then there may be many more ways he could move so that his head ends up lower than there are ways he could move so that his head ends up higher.

We might try to find some canonical form for propositions about movement. We could try to find out all the dimensions of movement — all the different ways in which movements can vary. If we have a proposition in canonical form, we can see how specific it is about each of these dimensions of movement. We then say that two propositions about movements are equally specific if they are equally specific about each dimension at each moment of time. We will only be able to say that two propositions are equally specific if they are in canonical form.

However, the suggestion that we can only match up propositions in canonical form casts Bennett's last two assumptions into doubt. First, it seems unlikely that we will be able to divide the Interpose subset of the behaviour space into a finite number of propositions in canonical form. Propositions in canonical form pick out a set of movements using the dimensions of movement. At each point in time, *t*, for each dimension, they give a set of values. A movement will be picked out by the proposition if and only if, at each point of time, in every dimension, its value along that dimension is one of the given set. It does not seem that we can divide the set of ways Agent could interpose the rock into a finite number of subsets of this kind.

Secondly, Bennett's 'echo' movements may not be exactly as specific as their originals. We can imagine movements that are 'very like' our original movements, but which do not result in the same outcome. We can imagine movements that would 'make it look as though Agent were trying but failing to interpose the rock' (Bennett 1995, p. 95). However, we cannot give the canonical form of these echoes and check they have the same degree of specificity as the originals.

We could be sure that the echo proposition is exactly as specific as its original if we considered only single points in the behaviour space. A single point in the behaviour space represents 'a proposition attributing to [Agent] some absolutely specific way of moving' (Bennett 1995, p. 91). We can match each absolutely specific possible movement in Interpose with another absolutely specific 'echo' movement, which is very like the original but differs just enough so that Agent does not interpose the rock.

However, this leads us to another problem. It seems likely that there are infinitely many, slightly different, ways Agent could move that would involve interposing the rock. If so, there are infinitely many members of Interpose. If there are infinitely many members of Interpose, the fact that we can match up each member of Interpose with a member of Non-Interpose does not tell us anything about the relative sizes of the two subsets. In an infinite space, matching up members does not imply that two subsets are the same size in the sense that we are interested in. Every member of the real number line, R, can be matched up with a member of $(0,1)$, the open interval from 0 to 1, without any overlap. Both R and $(0,1)$ have infinitely many members, so they both have infinite cardinality.

It seems likely that any subset of the behaviour space that corresponds to a proposition such as 'Agent interposes the rock' or 'Agent does not interpose the rock' will also have infinite cardinality. If this is so, when we talk about proportions of ways the agent can move, we do not want to compare cardinalities. We want some other way of measuring the size of subsets. We will not get a measure of this kind by simply pairing up members of subsets.

Of course, these criticisms only apply if the behaviour space is infinite. If the behaviour space were finite, Interpose would also be finite. We could pair up each member of Interpose with an echo in Non-Interpose using finitely many pairs. As there would be only finitely many pairs, we could deduce that Interpose is the same size as the subset of echoes in Non-Interpose. Because the echoes obviously would not take up all of Non-Interpose, we could deduce that Interpose is much smaller than Non-Interpose.

Bennett argues that the behaviour space must be finite (Bennett 1995, p. 93). He notes that how Agent moves at time T is determined by the intersection of: (a) the neuronal events that could have occurred in his brain; and (b) the facts about other relevant factors, such as wind, gravity, etc., and non-neuronal internal factors such as the temperature of his blood. The members of (b) are not under his voluntary control, so variations in the behaviour space come entirely from (a). Bennett argues that the number of distinct possibilities in (a) is finite 'because it is determined by how many neurons Agent has and how many relevantly different states each can be in, both numbers being finite' (Bennett 1995, p. 93). He concludes: 'So there are, after all, only finitely many points in the Agent's behaviour space and our metric can be got by counting them.'

Variations in brain states may be more complicated than Bennett assumes.[12] Whether there are finitely many possible brain states is an empirical question that will not be answered by philosophers, but it would be preferable if analysis of the act/omission distinction did not depend upon a controversial account of neurophysiology. Moreover, even if the number of possible brain states is finite, we should still expect the behaviour space to be infinite. Bennett suggests that the behaviour space should represent 'ways Agent could move at time T' (Bennett 1995, p. 91). However, the kind of movement we are interested in — movements such as pushing a vehicle or interposing a rock — take time. This realization puts the final nail in the coffin for Bennett's argument that the behaviour space is finite. Even if there were only a finite number of possible states in which Agent's neurons could be in at time T, we are not only concerned with what happens during that instant. We are concerned with Agent's behaviour over a period of time. How Agent moves during a period of time will depend upon the neuronal events over that whole period. His movement is a function of neuronal states over an interval of time: unless time is granular, there will be infinitely many such functions even if there are only finitely many neuronal states.

3. A new measure on the behaviour space

Bennett's measure on the behaviour space faces some serious problems. None the less, I believe that Bennett's overall approach is sound. I shall now suggest a new way of conceiving of the behaviour space and of comparing the size of its subsets. Using this new method of comparison, we can still make sense of Bennett's account of the act/omission distinction.

I begin by representing members of the behaviour space mathematically. The behaviour space is made up of all the ways the agent could move his body between t_1 and t_2. Humans move by lengthening and contracting their muscles. The changes in length of a muscle over a time period can be represented by a function. Such a function assigns to each moment in the time period a value representing how far the muscle is extended at that moment in time. We can put these functions together to form a more complex function that represents the overall movement by showing us how far each muscle is extended at each instant in time.

[12] For discussion, see Penrose 1989, p. 511.

First, we number each of the muscles that the agent can control directly. If there are N muscles under the agent's direct control, these muscles will be numbered from 1 to N. It does not matter what order we number the muscles in. What is important is that each muscle under the agent's direct control is assigned a number. I will label the right biceps 'muscle number 1'.

For each natural number n from 1 to N, we want $f_n(t)$ to be the function representing how much muscle number n has extended or contracted at each point in time during the movement. So $f_1(t)$ represents how far the right biceps is extended (compared with its position at the start of the movement) at each point in the movement. Suppose that the movement involves bending the forearm upwards towards the shoulder — what weightlifters call a biceps curl. To find $f_1(t)$, we first note the original length of the biceps. Suppose it starts at length eight units. The biceps contracts to bring the forearm upwards towards the shoulder. Five seconds into the movement, the biceps is at length four units. To find the value of $f_1(5)$, we find the difference between the original length of the muscle and its length at $t=5$. We take the length at $t=5$ and subtract the original length. So $f_1(5)=4-8=-4$. At $t=5$, the biceps has contracted so that it is four units shorter than it was at the beginning of the movement. We represent the movement of the biceps over the whole time period by noting down the length of the biceps at t and subtracting the original muscle length to find out how much further the muscle has extended. For any t, $f_1(t)=l_1(t)-v_1$, where $l_1(t)$ is the length of the biceps at time t and v_1 is the original length of the biceps.

To represent the movement of the rest of the body, we perform the same process with each of the other muscles. We note down v_n (the original length of muscle n) and $l_n(t)$ (the length of muscle n at t). To find out the additional extension of the muscle at t, we subtract the original length from the current length: $f_n(t)=l_n(t)-v_n$. We now have a function for each muscle representing the extension and contraction of that muscle during the time period.

To represent the entire movement, we put these functions together in an ordered list or N-tuple $(f_1(t), f_2(t), f_3(t), \ldots, f_n(t), \ldots, f_N(t))$. The first member of the ordered list, $f_1(t)$, represents the additional extension of muscle 1 at time t; the second member, $f_2(t)$, represents the additional extension of muscle 2 at t; the n^{th} member, $f_n(t)$, represents the additional extension of muscle n at t. We then represent the whole movement using the function, f, that matches each time, t, in the interval with the appropriate ordered list. More formally:

$$f : [t_1, t_2] \rightarrow \mathrm{R}^{\mathrm{N}}$$

$$f(t) = (f_1(t), f_2(t), f_3(t), \ldots, f_n(t), \ldots, f_{\mathrm{N}}(t))$$

Where for each $n : v_n =$ length of muscle n at t_1; $l_n(t) =$ length of muscle n at t; $f_n(t) = l_n(t) - v_n$; R^{N} is the set of N-tuples of real numbers. Each movement is represented by a unique function of this kind telling us exactly how far each of the agent's muscles would be extended at each moment.

We now want to compare the size of subsets of the behaviour space. The behaviour space is infinitely dense, so simply pairing up members of the two subsets will not work. Luckily, it is still sometimes possible to compare the size of subsets when we are dealing with infinitely dense spaces. Consider Fig. 6. Subset S1 is obviously bigger than subset S2. There are no more points in S1 than in S2; both subsets contain infinitely many points. However, S1 obviously takes up a bigger area than S2. We can demonstrate this using a method that is similar to Bennett's 'echo' idea. We start by thinking of a circle of a fixed radius r. We see how many such circles are needed to cover S2. Call this number $\mathrm{M}_r(\mathrm{S2})$. It is clear that (so long as r is small enough) we will always need more than $\mathrm{M}_r(\mathrm{S2})$ circles of radius r to cover S1. In Fig. 7, S2 is covered by 10 circles of radius 5 mm, it is clear that 10 circles of the same size will not come close to covering S1. Intuitively, whenever r is small enough, we will require more circles of size r to cover S1 if and only if subset S1 is bigger than subset S2.

Formally:

> For any real number r, let $\mathrm{M}_r(\mathrm{S}) =$ the minimum number of circles of radius r by which S can be covered (note that circles are permitted to overlap)
>
> $\mathrm{S1} > \mathrm{S2}$ if and only if there is some positive real number, α, such that for any r between 0 and α, $\mathrm{M}_r(\mathrm{S1}) > \mathrm{M}_r(\mathrm{S2})$

In other words, S1 is bigger than S2 iff the minimum number of circles of a given radius needed to cover S1 is always larger than the minimum number of circles of that radius needed to cover S2 (so long as the radius is small enough).

We can now extend this idea to compare the size of subsets in other spaces. We used circles with a fixed radius to compare the size of subsets of 2-dimensional space. We need to extend the idea of circles of fixed radius to other spaces.

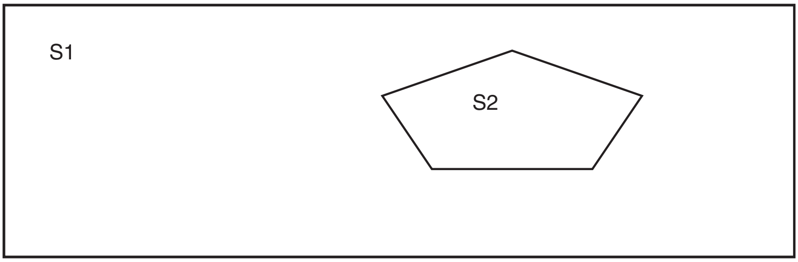

Fig. 6. S1 is greater than S2.

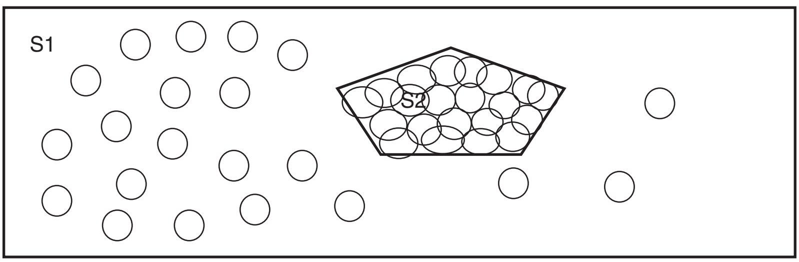

Fig. 7. S2 covered by circles of radius *r*.

First, we set up a way of making sense of the 'distance' between two points in our new set, extending this idea of distance beyond its natural domain of the real line. This is done by using a metric function. A metric function on a space is a function that for any two points in that space gives a real number that represents the 'distance' between them. Once such a function is defined, we define 'r-balls' which are the analogue of circles of radius r. Just as a circle of radius r is the set of points which are less than r away from the circle's central point, an r-ball around a given point x is the set of points which are less than r units 'distance' away from x.

Suppose that S1 and S2 are subsets of a bounded metric space, we can say that:

> For any real number, r, let $M_r(S)$ = the minimum number of r-balls by which S can be covered (note that balls are permitted to overlap)
>
> S1 > S2 if and only if there is some positive real number, α, such that for any r between 0 and α, $M_r(S1) > M_r(S2)$

Any 'strictly greater than' relation must have certain features: if S1 is greater than S2 then S2 cannot be greater than S1; no subset can be strictly greater than itself; if S1 is greater than S2 and S2 is greater than S3, then S1 must be greater than S3. The relation must be asymmetric, irreflexive, and transitive. It can be shown fairly easily that in any space in which $M_r(S)$ is well defined, the relation defined above meets these conditions.

We can now apply this to the behaviour space. First we need to define a metric function on the behaviour space, a function that, for any two points on the behaviour space, gives us a real number that represents the distance between them. This is quite easy to do. Suppose that we have two different possible movements of the agent's body (members of the behaviour space). We represent the first possible movement by a function, *f*, with $f(t) = (f_1(t), f_2(t), f_3(t), \ldots, f_n(t), \ldots, f_N(t))$, where $f_n(t)$ is how far muscle *n* would have extended or contracted at *t* if the agent had made this first possible movement. We represent the second possible movement by another function, *g*, with $g(t) = (g_1(t), g_2(t), g_3(t), \ldots, g_n(t), \ldots, g_N(t))$, where $g_n(t)$ is how far muscle *n* would have extended or contracted at *t* if the agent had made this second possible movement. It can be shown quite easily that the function $d(f, g) = \sqrt{(\sum_{n=1}^{N} \int_0^1 (f_n - g_n)^2 dt)}$ fulfils the conditions for metric functions. (For details see Appendix I.)

We then need to show that $M_r(S)$ is well defined on the behaviour space — in other words, that there is one and only one value of $M_r(S)$ for each subset, S, of the behaviour space. For a given *r*, $M_r(S)$ will be well defined if the behaviour space can be covered by a finite number of *r*-balls. If the behaviour space can be covered by a finite number of *r*-balls, then any subset S of the behaviour space can be covered by a finite number of *r*-balls. If S can be covered by a finite number of *r*-balls, then either there will be some number, N, of *r*-balls such that S can be covered by N *r*-balls but S cannot be covered by N − 1 *r*-balls (in which case $M_r(S) = N$) or S will not need any *r*-balls to cover it (in which case S is the empty set and $M_r(S) = 0$). So to show that $M_r(S)$ is well defined, all I need to do is to show that for any *r*, the behaviour space can be covered by a finite number of *r*-balls. This is fairly difficult, but possible, to prove. (For details see Appendix II.)

So I have given a way of representing the members of the behaviour space — the possible movements of the agent's body. Each possible movement is uniquely represented by a function from moments in the time period to an ordered list or N-tuple, $(f_1(t), f_2(t), f_3(t), \ldots,$

$f_n(t), \ldots, f_N(t)$). For each n, $f_n(t)$ is the difference between the original length of muscle n and its length at time t. Once the members of the behaviour space are represented in this way, we can make sense of the thought that one subset of the behaviour space is bigger than another. I used a metric function, $d(f, g)$, to represent the 'distance' between any two members of the behaviour space, f and g. This allowed me to define r-balls, which are like 'circles' of radius r. Subset S1 is bigger than S2 if and only if, for any r that is sufficiently small, more r-balls are needed to cover S1 than S2.

We now need to show that our method of comparing areas of the behaviour space gives results that match our intuitions. In the Stayback example given above, the Interpose subset was clearly much smaller than the Non-Interpose subset. Most of the ways Agent could have moved his body would not have involved Agent's interposing the rock. Does our method of comparing subsets give the right result in cases like Stayback?

4. Matching our intuitions

We can show that one subset of the behaviour space is bigger than another using a method that is similar to Bennett's 'echo' method. Bennett's idea was to show that one subset was larger than another by matching up movements in one subset with 'echoes' of equal size in the other. This method failed because there was no way of checking that the echo movement took up the same amount of space as the original. However, we know that an open r-ball in S1 will take up the same space as an open r-ball in S2. If we can match r-balls in one subset with balls of equal dimension in the other subset, where no two balls overlap unless their counterparts do, we can give sense to Bennett's idea of an 'echo'.

We match r-balls to r-balls using a rigid function. A function is rigid if it does not change the distance between two points. If F is a rigid function, then the distance between F(a) and F(b) will be the same as the distance between a and b. This means that a rigid function will map an r-ball to another r-ball. Suppose we can define a rigid function on S2. Then we will have matched S2 to a subset of S1, which we will call S2*. It can be proved quite easily that $M_r(S2) = M_r(S2^*)$. So S2 will be smaller than S1 if and only if S2* is smaller than S1. Thus, to show that S2 is smaller than S1, we need only to show that S2* is a proper subset of S1.

Any function that adds or subtracts the same value at time *t* from the same place in the *n*-tuple of every member of the behaviour space is a rigid function. So long as we always add the same value at the same place and time to every member of the behaviour space, the distances will not change.[13]

In Stayback, to interpose the rock Agent must kick the rock into the path of the vehicle.[14] To do so, he must extend his leg to the rock. We can define a rigid function that takes every possible way he could kick the rock to a possible movement that involves his extending his leg but not quite far enough to kick the rock. We do this by decreasing the value at the places in the N-tuple that represents the extension of some of Agent's leg muscles. This decrease in value ensures that Agent's leg muscles are never extended enough for him to kick the rock. If we perform the same decreases in value to every member of an *r*-ball in Interpose, the distances between functions remain constant. By selecting the decreases in value appropriately we can ensure that the images of non-overlapping *r*-balls will not overlap.

This gives a set of rigid functions from *r*-balls in Interpose to non-overlapping *r*-balls in Non-Interpose. For any reasonably small value of *r*, another *r*-ball will be needed to cover Non-Interpose — the ways Agent could extend his leg almost but not quite far enough to kick the rock take up nowhere near the whole of Non-Interpose. So our method implies that, in Stayback, most of the ways Agent could move would not have involved interposing the rock. This matches our intuitive judgement. Thus, my method allows us to compare the size of subsets of the behaviour space and gives results that match our intuitive judgements in cases such as Stayback.

5. Conclusion

Bennett's original method for comparing subsets of the behaviour space is unsatisfactory: in an infinite behaviour space, Bennett's method may lead to contradictory results; Bennett's argument that the behaviour space is finite fails to take into account the fact that actions take place over intervals of time rather than in an instant.

[13] For details see Appendix III.

[14] He might also push or pick up the rock. However, as he has only finitely many limbs and each of the ways he can interpose the rock involve touching the rock with a limb, we can repeat the process described below to find a finite number of rigid functions from Interpose to Non-Interpose.

None the less, an alternative method is available. It is possible to compare the size of different subsets of the behaviour space. We can make sense of the claim that 'most of the ways an agent can move' would make a given proposition true. This puts to rest a persistent worry about Bennett's account of the distinction between positive and negative facts: it does make sense to say that a fact about an agent's behaviour is positive if and only if most of the ways the agent could move would not make the associated proposition true. Bennett's account thus remains one of the most promising attempts to analyse the act/omission distinction.

Bennett uses his analysis of the act/omission distinction to try to refute the claim that the act/omission distinction is morally significant (Bennett 1995, pp. 139–42). He states: 'The positive/negative distinction that I have defined *obviously* has no basic moral significance' (p. 102). Unlike Bennett, I claim that the act/omission distinction is morally significant — although much work is required to bring this moral significance to light.[15] However, whether one wants to argue that the act/omission distinction is morally significant or that it is not, getting clear on the nature of the distinction is an important part of the argument. Once we realise that we can sensibly speak about 'most ways' an agent could have moved, Bennett's proposal is in a good position to play that part.[16]

References

Bennett, Jonathan 1981: 'Morality and Consequences'. *The Tanner Lectures on Human Values*, 2, pp. 45–116.

—— 1995: *The Act Itself*. Oxford: Oxford University Press.

Dinello, Daniel 1971: 'On Killing and Letting Die'. *Analysis*, 31, pp. 83–6.

Foot, Philippa 1967: 'The Problem of Abortion and the Doctrine of Double Effect'. In Steinbock and Norcross 1994, pp. 266–79. Originally published in *Oxford Review*, 5, pp. 5–15.

Glover, Jonathan 1977: *Causing Death and Saving Lives*. London: Pelican Books; rpnt 1990, London: Penguin Books.

[15] See Woollard 2008, Chs 7, 8, and 9.

[16] Part of this work was completed in the course of a PhD funded by an AHRC Doctoral Award. I thank the AHRC for their financial support. I would also like to thank Damien Sereni for invaluable help, both in reminding me of the basics of topology and for fruitful discussions about how best to apply this to the behaviour space. I also thank Brad Hooker, John Cottingham, and James Lenman for discussion and advice.

Locke, Don 1982: 'The Choice Between Lives'. *Philosophy*, 57, pp. 453–75.
McMahan, Jeff 1998: 'A Challenge to Common Sense Morality'. *Ethics*, 108, pp. 394–418.
Penrose, Roger 1989: *The Emperor's New Mind*. Oxford: Oxford University Press.
Quinn, Warren S. 1989: 'Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing'. *Philosophical Review*, 98, pp. 287–312.
Steinbock, Bonnie and Alastair Norcross (eds) 1994: *Killing and Letting Die,* second edition. New York: Fordham University Press.
Thomson, Judith Jarvis 1996: 'Review: The Act Itself'. *Noûs*, 30, pp. 545–57.
Vihvelin, Kadri and Terence Tomkow 2005: 'The Dif'. *The Journal of Philosophy*, 102, pp. 183–205.
Woollard, Fiona 2008: *Doing and Allowing; Acting and Omitting.* PhD Thesis, University of Reading.

Appendix I. Proof that $d(f, g)$ is a metric function

A metric function must meet three conditions: (1) The zero condition: for any two members of the behaviour space f and g, the distance between f and g is zero if and only if f is identical to g; $d(f, g) = 0 \leftrightarrow f = g$. (2) Symmetry: for any two members of the behaviour space f and g, the distance between f and g is the same as the distance between g and f; $d(f, g) = d(g, f)$. (3) The Triangle Inequality: for any three members of the behaviour space, f, g, and h, the distance between f and h is less than or equal to the sum of the distances from f to g and from g to h; $d(f, h) \leq d(f, g) + d(g, h)$.

1. Proof that d fulfils (1)

Suppose $d(f, g) = 0$

$$\text{Than } \sqrt{\left(\sum_{n=1}^{N} \int_{0}^{1} (f_n - g_n)^2 \, dt\right)} = 0. \text{ So} \left(\sum_{n=1}^{N} \int_{0}^{1} (f_n - g_n)^2 \, dt\right) = 0.$$

But for all f, g in the behaviour space, for any n, $f_n(t)$, $g_n(t)$ are real numbers. So

$$(f_n - g_n)^2 \geq 0$$

$$\text{So} \int_0^1 ((f_n - g_n)^2 dt) \geq 0$$

$$\text{So} \left(\sum_{n=1}^{N} \int_0^1 (f_n - g_n)^2 dt \right) = 0 \text{ implies that for each } n, (f_n - g_n)^2 = 0$$

So for each n and each t, $(f_n(t) - g_n(t))^2 = 0$. So for each n and each t, $f_n(t) - g_n(t) = 0$. So for each n and each t, $f_n(t) = g_n(t)$. So $f = g$. So if $d(f, g) = 0$ then $f = g$.

Suppose $f = g$.

$$\text{Then } d(f,g) = \sqrt{\left(\sum_{n=1}^{N} \int_0^1 (f_n - g_n)^2 dt \right)} \text{ and } f_n = g_n$$

$$\text{So } d(f,g) = \sqrt{\left(\sum_{n=1}^{N} \int_0^1 (f_n - f_n)^2 dt \right)} = \sqrt{\left(\sum_{n=1}^{N} \int_0^1 0 \, dt \right)} = 0$$

So d fulfils (1), the zero condition: the distance between f and g is zero if and only if f is identical to g; $d(f, g) = 0 \leftrightarrow f = g$.

2. Proof that *d* fulfils (2)

$$d(f,g) = \sqrt{\left(\sum_{n=1}^{N} \int_0^1 (f_n - g_n)^2 dt \right)} = \sqrt{\left(\sum_{n=1}^{N} \int_0^1 (g_n - f_n)^2 dt \right)}$$

$$= d(g,f)$$

So d fulfils (2), the symmetry condition: the distance between f and g is the same as the distance between g and f; $d(f, g) = d(g, f)$.

3. Proof that *d* fulfils (3)

$$d(f,h) = \sqrt{\left(\sum_{n=1}^{N}\int_{0}^{1}(f_n - h_n)^2\,dt\right)}$$
$$= \sqrt{\left(\sum_{n=1}^{N}\int_{0}^{1}((f_n - g_n) + (g_n - h_n))^2\,dt\right)}$$
$$\leq \sqrt{\left(\sum_{n=1}^{N}\int_{0}^{1}((f_n - g_n)^2 + (g_n - h_n)^2)\,dt\right)}$$
$$= \sqrt{\left(\sum_{n=1}^{N}\int_{0}^{1}(f_n - g_n)^2\,dt + \sum_{n=1}^{N}\int_{0}^{1}(g_n - h_n)^2\,dt\right)}$$

Given any non-negative real numbers x and y, x and y have a non-negative root.

So we have $2\sqrt{x}\sqrt{y} \geq 0$.

So $x + 2\sqrt{x}\sqrt{y} + y \geq x + y$. So $\sqrt{(x + 2\sqrt{x}\sqrt{y} + y)} \geq \sqrt{(x + y)}$.

So $\sqrt{(\sqrt{x} + \sqrt{y})^2} \geq \sqrt{(x + y)}$. So $\sqrt{x} + \sqrt{y} \geq \sqrt{(x + y)}$.

$$\sum_{n=1}^{N}\int_{0}^{1}(f_n - g_n)^2\,dt \quad \text{and}$$
$$\sum_{n=1}^{N}\int_{0}^{1}(g_n - h_n)^2\,dt \text{ are non-negative real numbers.}$$

$$\text{So}\sqrt{\left(\sum_{n=1}^{N}\int_{0}^{1}(f_n - g_n)^2\,dt + \sum_{n=1}^{N}\int_{0}^{1}(g_n - h_n)^2\,dt\right)}$$
$$\leq \sqrt{\left(\sum_{n=1}^{N}\int_{0}^{1}(f_n - g_n)^2\,dt\right)} + \sqrt{\left(\sum_{n=1}^{N}\int_{0}^{1}(g_n - h_n)^2\,dt\right)}$$
$$= d(f,g) + d(g,h)$$

So $d(f,h) \leq d(f,g) + d(g,h)$. So d fulfils (3), the Triangle Inequality: the distance between f and h is less than or equal to the sum of the distances from f to g and from g to h.

Appendix II. Proof that the behaviour space can be covered by a finite number of *r*-balls

First, note two facts about human physiology. Human muscles take time to extend and contract. Moreover, there must be limits on how fast a given agent can extend or contract his muscles. Thus for any given muscle, n, and any given initial muscle length, l_n, there is a minimum length of time it takes to extend or contract this muscle from l_n to l_n^*. Additionally, for any given interval of time, $[t, t+\delta]$, there is a maximum distance from l_n beyond which the muscle cannot extend/contract in $[t, t+\delta]$ if it is at l_n at t. So there exists $S = \sup \{|f_n(s) - l_n| s \in [t, t+\delta], f_n(t) = l_n\}$. As the interval of time gets shorter, the maximum length of contraction gets shorter. In other words, we expect S to tend to zero as δ tends to zero. More than this, we expect $S(l_n)$ to tend uniformly to zero. So:

(P1) For all $r^* > 0$, there exists some $\delta_n \in \{1/k_n\colon k_n \in \mathbf{N}\}$ such that, for all $l_n \in R$, for all $f \in \mathbf{B}$ (for all functions in the behaviour space) and for all $t \in [0,1]$:

$$f_n(t) = l_n \rightarrow (\text{for all } s \in [t, t+\delta_n], |f_n(s) - l_n| < r^*)$$

In other words, given some real number r^*, we can choose a small number, δ_n, such that if muscle n is at l_n at some point in time, it will always be less than r^* away from l_n for an interval of δ_n around that time.

Suppose we are given a positive real number, r. We need to show that the behaviour space can be covered by a finite number of r-balls. Let N be the number of muscles under the agent's control. For each $n \in \{1, 2, 3, \ldots, N\}$, we will use the following method. Choose some $S_n \in \mathbf{N}$ such that $S_n \geq (\sqrt{(m_n^2 2N)})/r$ where m_n is the maximum extension of muscle n. Let $r^* = r/\sqrt{2N}$. By (P1) we can then choose $\delta_n \in \{1/k_n\colon k_n \in \mathbf{N}\}$ such that, for all l_n, for all $f \in \mathbf{B}$ (for all functions in the behaviour space), and for all $t \in [0,1]$: $f_n(t) = l_n \rightarrow$ (for all $s \in [t, t+\delta_n]$, $|f_n(s) - l_n| < r^*$). Let $\mathbf{SL}_n = \{0, m_n/S_n, 2m_n/S_n, 3m_n/S_n, \ldots, m_n\}$. Let the set of selected time intervals, $\mathbf{T} = \{[0,\delta_n), [\delta_n, 2\delta_n), [2\delta_n, 3\delta_n), \ldots, [(k_n - 1)\delta_n, 1)\}$. We then have $S_n \times k_n$ values in the set: $\mathbf{G}\boldsymbol{n} = \{f_n(t) = c_1 X_{[0,\delta_n)} + c_2 X_{[\delta_n, 2\delta_n)} + c_3 X_{[2\delta_n, 3\delta_n)} + \ldots + c_{k_n - 1} X_{[(k_n - 1)\delta_n, 1)}\colon c_i \in \mathbf{SL}_n\}$. $X_{[a,b)} = 1$ if $a \leq x < b$ and 0 otherwise.

I will now show that given any function, f, in the behaviour space, there is some g_{nf} in $\mathbf{G}\boldsymbol{n}$ such that $d_n(f, g_{nf}) < r^2/N$ (where $d_n(f, g) = {}_0\int^1 (f_n - g_n)^2 dt)$). Given $f \in \mathbf{B}$, for each $i \in \{0, \ldots, k_n - 1\}$, let $f_{ni} = f_n(i\,\delta_n)$. Choose $c_{fni} \in \mathbf{SL}_n$ such that $|f_{ni} - c_{fni}| = \min \{|f_{ni} - c_i|\colon c_i \in \mathbf{SL}_n\}$.

Note that $|f_{ni} - c_{fni}| < m_n/S_n$. Then set

$$g_{nf} = c_{fn0}X_{[0,\delta_n)} + c_{fn1}X_{[\delta_n,2\delta_n)} + c_{fn3}X_{[2\delta_n,3\delta_n)} + \ldots + c_{fn(k_n-1)}X_{[(k_n-1)\delta_n,1)}$$

$$= \sum_{i=0}^{k_n-1} c_{fni}X_{[i\delta_n,(i+1)\delta_n)}$$

Set $h_{nf} = \sum_{i=0}^{k_n-1} f_{ni}X_{[i\delta_n,(i+1)\delta_n)}$

As d is a metric function, d_n is also a metric function. Thus by the Triangle Inequality, $d_n(f, g_{nf}) \leq d_n(f, h_{nf}) + d_n(h_{nf}, g_{nf})$.

$$d_n(f, h_{nf}) = \int_0^1 (f_n - h_{nf})^2 dt = \int_0^1 (f_n - \sum_{i=0}^{k_n-1} f_{ni}X_{[i\delta_n,(i+1)\delta_n)})^2 dt$$

$$= \sum_{i=0}^{k_n-1} \int_{i\delta_n}^{(i+1)\delta_n} (f_n(t) - f_{ni})^2 dt$$

But $f \in \mathbf{B}$, and $f_{ni} = f_n(i\,\delta_n)$. So by choice of δ_n, for all $s \in [i\,\delta_n, i\,\delta_n + \delta_n]$, $|f_n(s) - f_{ni}| < r^*$.

$$\text{Thus, } d_n(f, h_{nf}) < \sum_{i=0}^{k_n-1} \int_{i\delta_n}^{(i+1)\delta_n} r^{*2} dt = \sum_{i=0}^{k_n-1} r^{*2}\delta_n = k_n\delta_n\, r^{*2}.$$

But $\delta_n = 1/k_n$, so $k_n\,\delta_n = 1$. So $d_n(f, h_{nf}) < r^{*2} = (r/\sqrt{2N})^2 = r^2/2N$. So $d_n(f, h_{nf}) < r^2/2N$.

$$d_n(h_{nf}, g_{nf}) = \int_0^1 (h_{nf} - g_{nf})^2 dt$$

$$= \sum_{i=0}^{k_n-1} \int_{i\delta_n}^{(i+1)\delta_n} (h_{nf} - g_{nf})^2 dt$$

$$= \sum_{i=0}^{k_n-1} \int_{i\delta_n}^{(i+1)\delta_n} (f_{ni} - c_{ni})^2 dt$$

$$< \sum_{i=0}^{k_n-1} \left(\int_{i\delta_n}^{(i+1)\delta_n} (m_n/S_n)^2 dt \text{ (because } |f_{ni} - c_{fni}| < m_n/S_n \right)$$

$$= \sum_{i=0}^{k_n-1} (m_n/S_n)^2 \delta_n = k_n \delta_n (m_n/S_n)^2$$
$$\leq r^2/2N \text{ (because } S_n \geq (\surd(m_n^2 2N))/r) > 0$$

So $d_n(h_{nf}, g_{nf}) < r^2/2N$. So $d_n(f, g_{nf}) \leq d_n(f, h_{nf}) + d_n(h_{nf}, g_{nf}) < r^2/2N + r^2/2N = r^2/N$.

So given any function $f \in \mathbf{B}$, for each $n \in \{1, 2, 3, \ldots, N\}$, we can find g_{nf} in $\mathbf{G}n$ such that $d_n(f, g_{nf}) < r^2/N$ (where $d_n(f, g) = {}_0\int^1 (f_n - g_n)^2 dt$)).

Let $\mathbf{G} = \{g\colon g_n \in \mathbf{G}n$ for all $n \in \{1, 2, 3, \ldots, N\}\}$. The size of $\mathbf{G}$, $|\mathbf{G}| = \Pi^N{}_{n=1} (S_n K_n)$ — a finite number. For any $f \in \mathbf{B}$, let g_f be the function such that $(g_f)(t) = (g_{1f}(t), g_{2f}(t), \ldots, g_{Nf}(t))$ for all $t \in [0,1]$. For any $f \in \mathbf{B}$, g_f is a member of $\mathbf{G}$. $d(f, g_f) = \surd(\sum_{n=1}^N \int_0^1 (f_n - g_n)^2 dt) = \surd(\sum_{n=1}^N \int_0^1 (f_n - g_{nf})^2 dt) = \surd(\sum_{n=1}^N d_n(f, g_{nf})) < \sqrt{\sum_{n=1}^N r^2/N} = r$. So for all $f \in \mathbf{B}$, there is some g_f in $\mathbf{G}$ such that $d(f, g_f) < r$. So the set of r-balls: $\{Br(g)\colon g \in \mathbf{G}\}$ covers $\mathbf{B}$. $|\mathbf{G}|$ is finite, so $|\{Br(g)\colon g \in \mathbf{G}\}|$ is finite. So $\mathbf{B}$ can be covered by a finite set of r-balls.

So for any $r > 0$ and any subset S of the behaviour space, S can be covered by a finite number of balls and $M_r(S)$ is well defined.

Appendix III. Defining a rigid function on the behaviour space

Let F be a function on members of the behaviour space. Define F by $F(f) = f + h$ where h is a function from $[t_1, t_2] \rightarrow R^N$ such that if f is a member of the behaviour space, $f + h$ is a member of the behaviour space. Then given members of the behaviour space, f and g:

$$d(F(f), F(g)) = \surd\left(\sum_{n=1}^{N} \int_0^1 \left(F(f)_n - F(g)_n\right)^2 dt\right)$$
$$= \surd\left(\sum_{n=1}^{N} \int_0^1 \left((f+h)_n - (g+h)_n\right)^2 dt\right)$$
$$= \surd\left(\sum_{n=1}^{N} \int_0^1 \left((f_n + h_n) - (g_n + h_n)\right)^2 dt\right)$$
$$= \surd\left(\sum_{n=1}^{N} \int_0^1 \left(f_n - g_n\right)^2 dt\right) = d(f, g)$$

So F is a rigid function on the behaviour space.

