

LIFE AND DEATH

KEN BINMORE*

Abstract: What is the value of a life? How should we regard death? This paper uses the methods of economics to defend some of the views of Epicurus against the utilitarian approach that welfare economics takes for granted.

Keywords: Epicurus, ataraxia, timing of decisions, utilitarianism, modelling death

1. EPICURUS

What is the value of a life? How should we regard death? Answers to such questions have presumably been sought since the first human being realized that death is not something that only happens to others.

Philosophical debate on the nature of life and death continues to this day without any clear consensus emerging. But when ordinary people speak of taking a philosophical attitude to life and death, they usually intend some variant of the views of Epicurus that provide the focus for this paper (Long and Sedley 1987; Gaskin 1995; Warren 2004). To paraphrase Philodemus: Those who have achieved true happiness have already experienced all the happiness that eternity has to offer; death need not be feared because we won't be there to experience it (Henry 2009).

What is true happiness? In spite of the bad press to which Epicurus has been subjected down the ages, he was far from being the kind of hedonist who advocates thoughtless indulgence in the pleasures of the table or the flesh. In contrast to even straight-laced utilitarians like Jeremy Bentham (1987) who argued that poetry has no more intrinsic worth than the tavern game of push-pin, Epicurus held that philosophical reflection can be used to alter our states of mind so as to achieve a greater level of happiness than mindlessly following our untutored urges.

* Philosophy Department, Bristol University, Woodland Road, Bristol BS8 1TB, UK. Email: k.binmore@ucl.ac.uk

There are utilitarians (notably John Stuart Mill) who agree on the value of reflection, but orthodox economists currently ignore such refinements. Sophisticates sometimes quote David Hume's (1978) dictum that 'reason is and ought only to be the slave of the passions' in defence of this position. Such an attitude is alien to Epicurus's idea that true happiness lies in the tranquillity that comes from freeing oneself from anxiety and pain, and that philosophical reflection can allow such freedom to be achieved in spite of all the unavoidable ill fortune that the world may bring down upon one's head.

A contemporary Hume would doubtless agree with a contemporary Epicurus that reason can properly be used to modify internal psychological states to improve one's happiness, but how are utilitarian sceptics to be persuaded to venture beyond their defensive laager of theorems and proofs? I agree with John Rawls that the answer is not to abandon the deductive precision that mathematicians prize. As Rawls (1972: 121) says of ethics: 'We should strive for a kind of moral geometry with all the rigor which this name connotes.' I follow Rawls by applying the same approach to life in general that he took when offering his version of egalitarianism as an alternative moral system to utilitarianism, but few equations appear in the text because they are available elsewhere (Binmore 2009). This paper only examines how standard axiom systems need to be modified to generate conclusions on how we should value life and death that are closer to Epicurus than to Bentham and his successors.

Section 2 offers two versions of the philosophy of Epicurus: a mundane version and a transcendental version. Section 3 reviews the Von Neumann and Morgenstern theory of rational decisions made in risky situations and discusses how this can be adapted when time enters the picture. Sections 4 and 5 discuss how various difficulties can be dealt with by insisting that utilities are assigned to states of mind (rather than physical objects). The theorem of Section 6 is the heart of the paper. It explains how adopting an Epicurean approach allows the Benthamite evaluation of a life as a discounted sum of utilities can be replaced by an evaluation that takes account only of the best and the worst episodes in a life. Section 7 explores some implications for Epicurean philosophy.

2. EPICUREAN ASPIRATIONS

In seeking to reconstruct the doctrine of the historical Epicurus from the fragments that remain of his writings and the commentaries of his friends and enemies (Gaskin 1995), it is easy to be distracted by his personal idiosyncrasies. For example, Epicurus apparently took little interest in sex, feeling that it is better avoided (although not necessarily harmful). One may also grumble that it is easier for a popular folk hero (as Epicurus was in his time) to discount the acclaim of the world than for those of us whose efforts are largely ignored. For such reasons, I think it best to avoid

textual disputes about what the historical Epicurus may or may not have believed, and to propose two modern versions of his doctrine, of which I find the first much easier to endorse.

2.1 Mundane version

There is no evidence that we have immortal souls or that there are supernatural entities who care how we conduct our lives. Our bodies are survival machines that evolved for no grander purpose than that of replicating our genes. Our minds are merely part of the control mechanism. The segment of our minds of which we are conscious evolved to evaluate the likely long-term costs and benefits of actions taken in complicated situations, especially those in which we interact with other people. As such, our rationality is indeed the slave of our passions, being guided by signals from a part of the brain to which we have little or no conscious access. Modern neuroscientists know a good deal about the biochemical mechanisms, but I shall follow the tradition of referring to the signals in terms of pleasure or pain.

In asking how best to live one's life, a naturalist does not pretend that we can transcend such facts of human nature. A modern Epicurus would accordingly urge us not to fight our biology, but to go along with it. We are programmed to seek pleasure and avoid pain, so use your reasoning power to this end. When you feel anxious, for example, it seems reasonable to conjecture that your body is telling you to pay attention to your feasible set of possible actions. So consider your feasible set very closely. Even when it turns out that no feasible action can help, the procedure can nevertheless be wonderfully effective in lifting anxiety.

Nobody is upset at the idea of borrowing such ideas from cognitive therapy, but sceptics become restive when it is proposed that one apply similar techniques to the evaluation of long-term costs and benefits. Economists, for example, often hold that preferences are beyond rational discussion, being somehow given as part of a person's identity. It is true that psychological studies suggest that some 60% of the variation in personality indices is attributable to genetic factors beyond our control, but the other 40% is almost entirely explained by our childhood peer group (Harris 1998). It therefore seems reasonable to presume that much of what we value or despise, the things we strive to achieve or avoid, are similarly determined by our culture. It may not be easy to transcend this cultural conditioning, but Epicurus is only one of a long line of sages who have urged that the attempt to make oneself anew is worth the effort.

Epicurus would be asking us today whether we really will achieve satisfaction by making money or otherwise getting ahead in the world. Are pop stars or football heroes more content than other folk? How about gourmets and aesthetes? What of sexual athletes? Or kings and presidents? How about Bill Gates or Donald Trump? What can they do

with all their money but buy status? Are things any different for the winners of literary prizes or Nobel Laureates? And what if all one gets is the foreman's job at last? Or a spell occupying the office of the dean?

Epicurus thought that we can use the rational part of our minds to re-educate ourselves about the costs and benefits of seeking to come out on top in the rat race. His methodology amounts to asking oneself what pleasures or pains one would *actually* feel in various states of the world. In particular, if you failed to attain some previously prized objective, how much pain would you genuinely suffer?

Viewed from this perspective, the benefits of worldly success or intemperate consumption seem almost absurdly small compared with the felicity enjoyed by small children in a loving and safe household. So why not seek to create such a household for yourself? Within such a household, we would seek to satisfy no more than our basic biological requirements – including a family life, the companionship required by a social animal, and the intellectual stimulation needed by a thinking being. Which does not imply that worldly honours or a good night out need be rejected if offered without obligation; only that the costs normally incurred in striving after such benefits are excessive when properly evaluated.

2.2 Transcendental version

In this version, philosophical reflection is seen as generating more than an everyday state of contentment. It is thought capable of achieving an unimprovable state of true happiness or ataraxy, that frees adepts from concern with events in the physical world, so that the joys and sorrows that trouble ordinary folk cease to be of any relevance to their states of mind.

Rather like enlightenment in Buddhism, to achieve ataraxia is to change your outlook permanently. Never again need you fear that your tranquillity will be disturbed by intrusions from the outside world. Time ceases to be relevant. To experience ataraxy for an instant is the same as enjoying ataraxy for all eternity. Suffering can be borne without distress. Death ceases to have any significance.

Sceptics would like to see better evidence for such claims than is ever offered. For example, it is said that Epicurus died from kidney stones, but remained cheerful to the end. However, the account given by Diogenes Laertius of his death makes me suspect that Epicurus very sensibly committed suicide when he could no longer bear the pain.¹ But nothing prevents even a sceptic from envisaging ataraxy as an ideal state towards

¹ Epicurus is sometimes said to have been hostile to suicide, but what survives in his own writings are only some scathing remarks about those who advance more than one reason for taking their own lives.

which one may strive without any expectation of actually attaining the ideal.

3. MODELLING DECISIONS

Philo seeks to live the examined life. How should we model his decision problem?





3.1 Decisions under risk

Economists draw tables to model decision problems. Each row in such a table represents a strategy that Philo may choose. The collection of all the rows is called Philo's feasible set (of choices). The columns usually represent all the possible states of the world that may be relevant. The entries in the table represent all the possible outcomes that can result as a consequence of Philo's choice of a strategy and Nature's choice of a state of the world.

For example, if Philo were deciding on which horse to bet his life-savings in a race, the rows would represent all the horses in the race on which he might bet. The columns would represent all the possible winners of the race. The possible outcomes would be the amounts in dollars that Philo might win or lose depending on which horse he chooses to bet, and which horse turns out to win the race. If the probability that each horse will win is known, we can then apply Von Neumann and Morgenstern's (1944) famous result that Philo will behave irrationally unless he chooses the strategy that maximizes his expected (long-run average) VN&M utility.

How are VN&M utilities assigned to the outcomes that appear as entries in the decision table? The standard method is to assign utilities of zero and one to the worst and best outcomes respectively. Philo is then asked to consider all lotteries in which only the worst and best outcomes are allowed. To find the VN&M utility of some other outcome, Philo adjusts the probability of winning in such a lottery until he is indifferent between the new outcome and the adjusted lottery. The VN&M utility of the new outcome is then the probability of getting the best outcome in this adjusted lottery.

Assigning utilities of zero and one to the worst and best outcomes is only one of many ways of creating a VN&M utility scale. One can always create a new scale in the same way that one moves from degrees Celsius to degrees Fahrenheit by making a different choice of Philo's zero and unit. In the rest of this paper, we suggest rescaling Philo's utilities by assigning a utility of zero on a new scale to the state of being dead and a utility of one to the (unattainable?) ataraxic state. From an Epicurean standpoint, these states can be regarded as being the same for everybody, and so it will make

	young	old
work		
play		

	young	old
work	1/2	3/4
play	1	0

	young	old
work	1/8	7/16
play	3/4	-1/2

FIGURE 1. (Colour online) Decision tables.

sense to compare Philo's rescaled utilities with those of any other person that have been rescaled in the same way (Binmore 2007: 555).

The VN&M theory of rational choice under risk is not beyond criticism (Binmore 2009), but is taken for granted in this paper. It will also be assumed that objective probabilities can be assigned to all possible states of the world. Those who prefer rival theories of rational choice under uncertainty are free to replace VN&M's theory by their own in what follows.

3.2 Decisions over time

The brief account offered above of VN&M's theory proceeds as though everything happens at a single instant. What happens if the strategies available to Philo consist of future life-plans whose choice will partly determine what happens to him throughout the future?

A strategy now has to be understood as a plan of action that specifies what Philo would do under all possible contingencies that might arise now and in the future. Nature similarly chooses a strategy that the literature often calls a state of the world. Such a state of the world does not refer to how the world is at some particular time, but is a specification of everything external to Philo that might be relevant to him. Once Nature has chosen a state of the world and Philo has committed himself to a life-plan, all the future events of Philo's life are then fixed.

If Philo knew for sure what was going to happen in the external world at all future times (so that only one state of the world is possible for him), we could model his problem as a table in which the rows continue to be strategies (life-plans) but in which the columns now represent different time periods. Figure 1 shows a simple example. Philo must choose whether to be an ant or a grasshopper. He will enjoy himself more by playing when young but will prefer to have chosen to have worked when old.

The rows in Figure 1 represent possible strategies. The columns represent different time periods. The faces that appear as entries in the left table represent different states of mind. The grumpy face, for example, represents the state of mind of an old person who idled away his youth.

The middle table is the same except that the states of mind have been replaced by their VN&M utilities computed by the standard method. For example, Philo is assumed to be indifferent between the smiling state and a lottery in which he ends up in the laughing state with probability 3/4 and the grumpy state with probability 1/4. His utility for the smiling state is therefore 3/4. The right table shows a possible rescaling of these utilities in which the state of being dead is assigned a utility of zero and the ataraxic state is assigned a utility of one. (The formula used to transform the old utility x into the new utility y is $y = (5x - 2)/4$.) Note that Philo then prefers being dead to his grumpy state.

Philo will normally face a decision table with many more rows and columns than Figure 1. In the absence of uncertainties, each row will consist of a sequence of determinate states of mind. When uncertainties are present, each entry will be a *lottery* over such states of mind. If Philo is rational, such a decision problem is equivalent to another decision problem in which his choice of strategy results in a simple lottery whose prizes consist of sequences of *determinate* states of mind.²

To solve his problem in the presence of uncertainty using the VN&M theory, it is therefore enough for Philo to be able to assign V&NM utilities to all possible sequences of determinate states of mind. We have already assumed that he is able to assign VN&M utilities to each individual state of mind, and so Philo's problem reduces to what economists call an aggregation problem. How is he to assign a VN&M utility to a stream of VN&M utilities experienced at successive times? If we can succeed in coming up with such a number it will be called the value of the life summarized by the utility stream.

It is orthodox in economics to take the overall utility of a life to be a discounted sum of the utilities in the utility stream to which the life is reduced. (To discount is to treat future utilities as being less valuable than present utilities.) Future utilities are usually discounted at a fixed rate that corresponds to the interest you might earn per period if utility were money invested in a bank.³

3.3 Utilitarianism

We have looked at the case in which the columns in a decision table are interpreted as different states of the world, and also at the case in which the columns are interpreted as different time periods. One could also treat Philo as a social planner deciding which citizen gets how much of whatever is available. The columns in the appropriate decision table

² The same reasoning is used in game theory to show that all the chance moves in a game may be consolidated into a single chance move with which the game begins.

³ If the notional interest rate is r , then a utility at time t is discounted by being multiplied by $1/(1+r)^t$. A fixed interest rate is necessary if time inconsistencies are to be avoided.




FIGURE 2. (Colour online) A life.

would then be interpreted as the different citizens in Philo's society. Assumptions similar to those that lead Philo to value a strategy by working out a weighted sum of the entries in the corresponding row in the first two cases then lead to a similar conclusion in the third case. Since the third conclusion is what utilitarianism recommends, it seems reasonable to refer to all three conclusions as utilitarian.

In this paper, the plan is to stick with the utilitarian approach to uncertainty but to deny the utilitarian approach to time. I have written at length elsewhere defending a version of the Rawlsian alternative to utilitarian social aggregation, but social issues are mentioned only in passing in this paper (Binmore 2005).

3.4 Modelling death

Epicureans cannot regard death as a state of mind because they hold that corpses do not have minds. It is nevertheless convenient to include the state of being dead in the decision table that Philo uses when evaluating possible lives because we would otherwise have to adopt a special modelling device in order to include suicide among his options. Needless to say, the state of being dead needs to be sharply distinguished from the possibly painful and distressing act of dying.

Figure 2 shows an example of a row in a decision table in which Philo dies after five periods. An empty circle in the later columns represent the state of being dead. Its ontological status resembles that of the empty set in formal mathematics. It is unlike the other symbols in not standing for a state of mind but for the absence of a state of mind.

Note that the Epicurean symmetry argument denies that there is any difference in the state of being dead and the state of never having been born.⁴ So empty circles have also been used to indicate the state of being as yet unborn.

We again need to choose a suitable zero and unit on the VN&M utility scale used to value a life if the valuations of different people are to be compared. The obvious choice of a zero is the life that is never lived, which

⁴ As David Hume on his deathbed cheerfully explained to a tactless James Boswell when the latter quizzed him on how atheists felt about dying.

is represented by a doubly infinite sequence of empty circles. If we allow the transcendental claim that the attainment of ataraxia at any point in any life induces maximum bliss, we can also assign a utility of one to any such life.

It is significant that only a finite number of states of mind are allowed in any life (reflecting the assumption that Philo knows that death is inevitable). For example, the Stern (2007) report on global warming uses a model with an infinite time horizon. Some discounting is then essential to a utilitarian approach because the sum of utilities would otherwise often be infinite. But what discount rate should be used? The answer matters a lot because different choices in the Stern analysis lead to different strategies for dealing with global warming being regarded as optimal. It is therefore not surprising that Stern's choice of discount rate has resulted in an ongoing debate about what is the 'correct' social discount rate (which seems to me an unanswerable question).

4. STATES OF MIND

Warren's (2004) instructive book surveys a variety of philosophical objections to the Epicurean view that you cannot be helped or harmed by events that have no impact on your experience. Nagel (1979) even argues that you can be hurt by events that take place after your death. Insofar as such critics would be prepared to endorse any modelling at all, they would presumably favour replacing my states of mind by states drawn from some larger space. However, I believe that at least some of their objections can be accommodated without taking such a drastic step by interpreting a state of mind very much more widely than seems to be usual.

4.1 Stand-alone states of mind

In this paper, a state of mind should be understood as encompassing more than a person's mood at a particular moment. The concentration of such neurochemicals as serotonin and dopamine in Philo's brain at any given instant will always be important, but the kind of hedonism that pays attention only to such vectors of pain or pleasure is not adequate to justify the rationality assumptions of the next section. For my purposes, a state of mind needs to embody anything whatever going on *inside Philo's head* that may be relevant to his evaluation of that state.

Philo's memories of the past are one such factor. His assessment of future prospects is another. What his senses are telling him right now is a third factor. However, it is important that none of these inputs to his decision processes need be regarded as accurate descriptions of the world. Indeed, Berkeley would observe that such a specification of a state of mind is compatible with there being no external world at all. For example, Philo

should not be assumed to have *direct* knowledge of the current time. He may remember that yesterday was Friday, but none of us are gifted with perfect recall. He may have the image of a clock on his retina, but perhaps he is hallucinating.

Without such stand-alone requirements, even hedonists might reasonably doubt the assumption that Philo cares only about the VN&M utilities of the states of mind that make up a life. If states of mind were not stand-alone entities, Philo would not be able to assign them VN&M utilities without directly taking into account all kinds of external events in his past, present or future. With the formulation proposed here, he can proceed without taking account of more than his internal representation of such events.

It is particularly important to note that replacing one state of mind in a possible life by another state of mind with the same VN&M utility will leave Philo's overall valuation of that life unchanged.

4.2 All things considered

It is also necessary to bear in mind that Philo is assumed to be sold on the notion that the unexamined life is not worth living. The ideas he entertains in all states of mind to be assessed are therefore fully considered. So he will not be found crying over spilt milk, or lamenting costs that are now sunk. His beliefs will not be influenced by his preferences, nor his preferences by his beliefs. In particular, his estimates of future prospects will be based on nothing but the evidence that he finds lodged in his memory. In brief, Philo is unlike most of the rest of us in enjoying states of mind that are free from all irrationality.

4.3 Making wills?

Epicurus is sometimes criticized for making a will on the grounds that it was inconsistent of him to care what happens after his death (Warren 2004). It would certainly be inconsistent for Epicurus to care what happens *to himself* after his death, but he doubtless felt the same love or affection for his family and friends that most of us feel for ours. Just as rational consideration led him to make provision for his own future welfare (as when laying in some barley bread for tomorrow's breakfast), so it would lead him to make provision for the future welfare of his loved ones. The fact that he would be dead when his will was implemented seems irrelevant, because what mattered to him when writing the will was his anticipation *at that time* of the deprivation his loved ones might suffer if he were to make no provision for them right now.

4.4 Thou knowest not what a day may bring forth

A number of other objections to the Epicurean view can perhaps be answered by distinguishing carefully between intrinsic preferences and instrumental preferences (Binmore 2009: 5). For example, if Philo chooses a hammer rather than a screwdriver, it is not usually because he has some inbuilt bias in favour of hammers, but because a hammer is better than a screwdriver at driving home a nail. The preference Philo reveals for a hammer is therefore instrumental.

An example in philosophy concerns John Stuart Mill's claim – echoed in the economics literature by Amartya Sen (2000) – that people care (or should care) directly about liberty (so that liberty becomes an argument in a person's intrinsic utility function). But Epicurus would ask why Philo should care about being deprived of the liberty to make certain choices if he would not have made those choices anyway. Philosophers who ask similar questions are suggesting that rational people should value liberty for instrumental reasons. It may be that whoever or whatever is choosing for you is certain to choose what you would have chosen yourself from today's feasible set, but who knows for sure what tomorrow's feasible set will look like? It is therefore wise to reserve the right to make all choices for yourself (if doing so is not too costly).

It is to avoid such debates that I have taken care to factor out all questions of uncertainty from the decision problem that Philo faces when considering his preferences over alternative possible lives. An example may help to explain why I think this important. Perhaps Philo's current happiness depends on his unshakeable belief that the love of his life is faithful to him, but she actually takes pleasure in betraying him at every opportunity with men who make a mock of everything for which he cares. Is he not harmed by this fact, even though news of her bad behaviour never finds its way into his state of mind? In assessing his innocent state of mind, should we not take this externality into account? If so, should those of us who aspire to contribute to human knowledge not also take into account the possibility that an asteroid is on course for an eventual collision with our planet that will destroy the human race and so nullify all our efforts?

People seem curiously untroubled by the second (not so very unlikely) of these suggestions, but react quite sharply to the first. Like Warren (2004), I suspect that the reason is the immediacy of the possibility that Philo will actually find out that his lover is unfaithful. In other words, the facts of life in the real world outweigh the story teller's insistence (and the assumptions of my model) that there is no possibility whatever that Philo will discover that he has been betrayed.

There is also a secondary reason for not looking outside Philo's own experience when evaluating his life. A rational Philo will not suppose that

he alone of the human race is immune to betrayal. He will weigh all the evidence and assign a (possibly very small) probability to the event that his loved one is unfaithful. His future behaviour will then be conditioned on the expectations he derives from this and other probabilities. In particular, he will decide how to react if he were to learn that the low probability event that his lover is unfaithful were actually realized. He would then be upset, as people are upset when they bet on horses that fall at the first jump, but decisions have to be judged not by whether one could have chosen better had one known beforehand what was going to happen, but by the evidence that was available at the time the decision was made. Kipling's poem *If* is eloquent on this (and other) features of the Epicurean attitude to life.

5. FREE WILL AND DETERMINISM

What would Epicurus have made of free will, an idea that was apparently unknown to the ancients? I suspect that his down-to-earth naturalism would have led him to endorse Spinoza's (2009) view: 'Men believe themselves to be free, simply because they are conscious of their actions, and unconscious of the causes whereby those actions are determined.' As Dennett (1987) has explained, there is no contradiction in holding this view and simultaneously modelling individuals as intentionally making decisions in accordance with their preferences and beliefs. In fact, although most economists would be surprised at the thought, their model of rational choice would seem to take this position for granted.

If Philo is assumed to be rational, he will choose the outcome he prefers most from the set of outcomes that are feasible. A sufficiently well-informed onlooker will therefore be able to predict his choice in advance, and hence may reasonably regard Philo's choice as being deterministic. But this will not be how it seems to Philo before he has considered what choices are feasible and how much he likes the different possible outcomes, especially if so doing requires complicated calculations. While thinking these matters through, Philo will regard himself as free to make any feasible choice. He will later still hold that the choice his rationality eventually led him to make was entirely free, since it was constrained only by external feasibility issues. If challenged on the grounds that he was actually the helpless victim of his rationality, he might even paraphrase Kant (2007): 'Man needs a master to break his will so that he might be free.'

The issue of free will is introduced here only because there is an apparent tension between our conception of Philo freely choosing a life-plan and the deterministic model of a life we have imputed to him. For a given state of the world, our model of a life makes the succession of states of mind that constitute Philo's life completely determinate once he has

chosen a life-plan and Nature has chosen a state of the world. It is true that each state of mind will require that Philo take some action or think some thought that is a proximate cause of his moving from that state of mind to his next state of mind, but these actions or thoughts are determined by Philo's overall life-plan.

5.1 Timing?

One may reasonably ask *when* Philo's chooses his life-plan. So far, we have proceeded as though Philo has a second self who looks on from the outside while comparing different possible lives. But Philo has to choose while living his life. One then has to ask whether Philo is allowed to have second thoughts later. If so, why does he not act upon them?

Philo's life-plan can only be chosen after he has decided, perhaps through reading the works of the great philosophers, that he wants to live a rationally considered life. His attempts to examine himself will then lead him through a sequence of states of mind, at the end of which the current model assumes that he will have duplicated the reasoning outlined in this paper and chosen a life-plan. The moment of this choice marks the first period of his philosophical life. Events before this moment are not subsumed in the possible lives that Philo considers (and hence our earlier insistence that Philo's states of mind are to be regarded as free of irrationality). As Christian fundamentalists say, he is born again at this point.

However, we do not need to locate Philo's decision only at the outset of his philosophical life. He can be assumed to have the opportunity to renew or revise his calculations as his life proceeds. The next section will clarify why Philo would never thereby be led to change his life-plan (which is assumed to incorporate contingency plans for all that Nature might throw at him). However, he will always be free to do so (in the sense that nothing but his own reason stands in his way).

Finally, we need not assume that ataraxia is immediately attained after adopting a philosophical life. Just as someone who decides to devote his life to mathematics may still have to struggle long and hard to prove a good theorem, so Philo may have to struggle before something approaching ataraxia eventually comes in sight.

5.2 How it feels to have a life-plan

How does Philo's choice of life-plan influence his evaluation of a life governed by that choice?

Philo first lists all feasible life-plans, where feasibility is judged on nothing more exotic than his factual understanding of the physical world. The choice of such a life-plan may lead to many possible lives depending on the state of nature. It is important to recognize that each state of mind

in each of these lives will incorporate the chosen life-plan, because the Epicurean ideal of tranquillity depends partly on Philo feeling no anxiety about whether his current life-plan (or lack of life-plan) will need to be altered after reconsideration. In evaluating each life obtained in this way, Philo will therefore take account of the life-plan that generates the life in question. The result will be to reduce the life to a sequence of VN&M utilities. As described in the previous section, Philo will then choose a life-plan that maximizes his expected utility.

5.3 Complexity

The complexity of the problem Philo faces in choosing a life-plan is immense. A Philo living in the real world will therefore only be able to identify a number of general principles to guide his future actions. In assuming that Philo solves the problem completely, it will be appreciated that our model is therefore very much of an idealization.

6. RATIONAL CHOICE

The economic orthodoxy holds that rationality is concerned only with means and not with ends. This view originates with David Hume, who famously observed that it would not be *irrational* for him to prefer the destruction of the entire universe to scratching his finger.

Economic zealots like myself strip the concept of rationality down even further by refusing even to admit intrinsic preferences as an unexplained primitive. We argue instead that one should observe Philo's choice behaviour in some situations and then use his assumed rationality to predict his choice behaviour in others – an approach known as the theory of revealed preference (Binmore 2009). It then becomes transparent that the 'rationality axioms' to which economists of all stripes appeal only ensure that choices are made in a *consistent* manner.

The theory of revealed preference makes it irrelevant whether one speaks of consistent choice or rational preference. One can either regard consistent choice as the outcome of satisfying a rational preference or rational preference as a way of summarizing consistent choice. Utility functions are similarly just a way of summarizing rational preference (Binmore 2009). Different modes of expression are used in different contexts only to ease the exposition.

Far from celebrating the virtues of selfishness or hedonism as popular prejudice holds, orthodox economists therefore make no foundational assumptions at all about the psychological factors that lead people to make one choice rather than another. Such issues are regarded as empirical questions best left to psychologists or other social scientists. Rationality as understood in this paper therefore offers no substantive explanation of *why* Philo chooses one life-plan rather than another. It can only say that *if*

Philo would choose this or that life-plan from this or that set of feasible alternatives, *then* it would be inconsistent for him to choose various other life-plans when faced with various other feasible sets.

So why does Philo choose one life-plan rather than another? The proximate cause is that his brain is in a certain state. We have various words that we use to describe our awareness of our states of mind, notably pleasure and pain, but the question is not properly answered by saying that Philo chooses the life-plan that seems to offer most pleasure or least pain, because one may then ask *why* Philo should feel pleasure in some contexts and pain in others. But in seeking an ultimate answer to such questions, where else is there to look but to Philo's genetic inheritance, the culture in which he was raised, and the accidents of his personal history (before his philosophical life began)? However, we are all sufficiently different that it seems pointless to pursue this line further. The best we can reasonably hope is to find a functional form for Philo's utility for possible lives that is common to all rational people but admits parameters that may differ between different individuals.

6.1 What does a rational valuation of a life look like?

We are assuming that when Philo chooses his life-plan, he has already adopted a rational outlook. This includes his subscribing to VN&M's theory of decision under risk, which he uses to reduce all possible lives to sequences of VN&M utilities. The question now is how rationality restricts the manner in which he evaluates such a sequence of VN&M utilities.

To answer this question using the theory of revealed preference, we need to forget for the moment the feasibility constraints and the uncertainty about the state of the world that will afflict his actual choice of life-plan. We ask instead a series of hypothetical questions about what choices Philo would make from all theoretically possible feasible sets of lives if he were certain that his choice would always be implemented. We then require that all such choices be consistent with each other. But what does consistency imply in such a choice context?

6.2 Milnor's assumptions

The consistency axioms to which I appeal were formulated by the mathematician John Milnor in 1954, but have been largely neglected in recent times. Milnor uses the matrix representation of a decision problem discussed in Section 3. Recall that the rows of the matrix represent all possible actions (strategies or life-plans) and the entries can be identified with the VN&M utilities of possible outcomes. The columns of Milnor's matrix represent states of the world, but we reinterpret them as subjective time periods.

For each such decision problem, Milnor assumes that Philo has a preference relation defined over the set of all actions. He then considers 10 different assumptions or axioms that might be imposed on this preference relation. We list below those of his assumptions that we need (preserving his own numeration and terminology).

1. *Ordering*: Philo has a full and transitive preference relation over set of all actions.⁵
2. *Symmetry*: Philo is indifferent to how the actions or time periods are labelled.
3. *Strong domination*: If each entry in the row representing to action strictly exceeds the corresponding entry in the row representing a second action, then Philo strictly prefers the first action to the second.
4. *Continuity*: Consider a sequence of decision problems with the same set of actions and time periods, in all of which one action is strictly preferred to a second action. If the sequence of matrices of outcomes converges, then its limiting value defines a new decision problem, in which the first action is weakly preferred to the second action.
6. *Row adjunction*: Philo does not change his preferences between the old actions if a new action becomes available.
8. *Column duplication*: Philo does not change his preferences if a new column is appended that is identical to one of the old columns.

6.3 Standard rationality assumptions

Assumptions 1, 3, 4 and 6 are standard. Assumption 6 is nowadays called the Independence of Irrelevant Alternatives (to which we have already implicitly appealed while discussing Mill on liberty).

6.4 Defence of symmetry

Assumption 2 seems harmless when applied to actions but has bite in the case of time periods since it rules out any discounting. The assumption relies on the requirement that states of mind have a stand-alone status. The actual calendar time at which Philo experiences a state of mind therefore makes no difference to him. In particular, we can swap two VN&M utilities in a sequence that makes up a possible life without changing Philo's valuation of the life.⁶

⁵ Philo's preferences are transitive if the fact that he prefers a to b and b to c implies that he prefers a to c .

⁶ This does not imply that Philo will be equally happy when young as when old if these terms are understood subjectively. Feeling the effects of rheumatism will be part of a state of mind typical of the elderly whatever time may be registered on an objective clock.

6.5 Defence of column duplication

We could get by without Symmetry, but Assumption 8 is the pivot on which my defense of the Epicurean approach turns. In justifying the assumption, it is necessary to distinguish subjective time from objective time. William James (1890: 609) observed that we do not register time internally as a point moving around an ideal clock. Our 'internal now' corresponds rather to a small but extended interval that James called the 'specious present' and this paper identifies with a state of mind. How long is Philo's specious present? Modern neuroscientists say that the answer to such questions of subjective timing depends on the frequency with which Philo has to process information (Eagleman 2008).

Imagine that Philo is contemplating how to evaluate what life-plan to choose if his choice were restricted to the rows of a particular decision table (with the state of the world given). In a column corresponding to a particular time period, he may observe that whatever life-plan he chooses, the next period will find him in a state of mind that is subjectively indistinguishable from his preceding state of mind. The objective passage of time involved in passing from one state of mind to the other will then be subjectively undetectable. Only an external observer will detect any difference between the two. So Philo might as well simplify his model by treating the two states of mind as one, because that is how they will seem to him from his internal perspective.⁷

Presumably we would treat sleep this way if we awoke as refreshed from one hour's sleep as from eight. Or to offer a more dramatic example, recall Omar Khayyam's yearning for 'infinite ecstasy, indefinitely prolonged'. If infinite ecstasy implies a mental state that cannot be improved, why should such a state need to be qualified by asking that it be indefinitely prolonged? If Omar's current ecstatic state leaves him concerned about when he might enter a less exalted state, the state cannot be unimprovable because it would be more enjoyable if he were not anxious about when it might end. Its length in objective time should therefore be irrelevant.

One might object that although Philo cannot distinguish between two subjectively identical states of mind while living his life, he will be able to distinguish between them when reviewing possible life-plans, but for him to make such a distinction would be to deny the Epicurean doctrine that only what people actually experience is relevant to their quality of life.

Theorem [Milnor 1954]. *The preceding six consistency assumptions imply that when Philo values a life (assigns it a VN&M utility), he takes into account*

⁷ Such reasoning would seem necessary in any case to explain the origin of the periods of time represented by each column in Philo's decision table. Why are they not shorter or longer?

nothing but the best and worst periods in the life (those assigned the largest and smallest VN&M utilities).

How will Philo take account of the best and worst periods in a potential life? With an extra linearity assumption, Milnor shows that Philo will value a life using the Hurwicz criterion, which simply consists of taking a weighted average of its largest and smallest constituent VN&M utilities.⁸

If so, then it is easy to describe how Philo will choose a life-plan while ignorant of the state of the world. He will simply maximize the Hurwicz criterion using the *expected* values of the largest and smallest VN&M utilities in a life. Of course, as Philo learns more about the state of the world as his life proceeds, he will update the probabilities he uses to assess future contingencies and so these expectations will change with time.

6.6 Utilitarian alternative

To obtain the utilitarian alternative to the preceding theorem in Milnor's framework, one drops Continuity and replaces Column Duplication by what Milnor calls Column Linearity, which is a variant of the independence axiom of Bayesian decision theory (Binmore 2009: 158). Column Linearity says that Philo regards a decision table obtained by adding a constant to all the entries in one column as equivalent to the original decision table for the purposes of choosing an action. Milnor used Column Linearity to characterize Laplace's principle of insufficient reason (that two events should be assigned equal probabilities if we are entirely ignorant about which is more likely), but we need to consider the case in which the columns of a decision table represent time periods rather than uncertain states of the world. Milnor's argument then shows that Philo will value a life simply by summing all its constituent VN&M utilities. (To allow for discounting, it is necessary to abandon the symmetry axiom.)

6.7 Parfit and Rawls

When the columns of a decision table correspond to the citizens of a society, Parfit (1986) describes the result of applying the preceding utilitarian result to the size of a society as 'the repugnant conclusion' because it implies that one can always compensate for a very low quality of life for each individual citizen by making the number of citizens sufficiently large. It seems to me similarly repugnant for utilitarians to argue that Philo should be willing to tolerate a very low quality of life if this is what is necessary to live for a very long time.

⁸ The value of a life whose largest and smallest VN&M utilities are respectively C and c is then given by $aC + bc$, where a and b are non-negative constants that sum to one.

6.8 Time consistency

Under this heading, it only needs to be said that Philo will never have an incentive to change his original life-plan. When choosing his life-plan in the first place, he will look ahead and consider what new choice of life-plan for the remainder of his life he would wish to make under all possible natural contingencies at all future time-periods. He will then incorporate these future life-plans into his optimal life-plan for the whole of his philosophical life. It is of some importance that he will not thereby be trapped into less favourable best and worst periods than would be available if he were to choose an unalterable life-plan for his whole philosophical life once-and-for-all.

7. EPICUREAN INTERPRETATION

Section 2 offers mundane and transcendental versions of Epicureanism. The theorem of the preceding section is relevant to the mundane version in three ways.

7.1 MUNDANE VERSION

1. If the theorem applies to Philo, then his longevity is of no direct significance to him. He is indifferent between two lives that have the same best and worst periods, although one life may be long and the other short. Folk wisdom often commends a similar but less precise principle that what matters is not how long you live but how you live what life you have. As Victor Hugo put it: 'It is nothing to die; it is frightful not to live.'

2. Philo may value some states of mind as being worse than the state of being dead. Suicide will therefore sometimes be an optimal action. The extent to which suicide will seem optimal to Philo will depend on his estimate of what the future holds. If he thinks it likely that the best period of his life is still to come, then he may well be willing to endure much suffering in the present in order to enjoy this future benefit.

But what if Philo thinks it likely that the best period in his whole life is in the past? As Epicurus observed when dying, the memory of this past success will improve his valuations of later states of mind. So his future states of mind will be better than they would otherwise have been, but that is the only way his past success will influence his decision on whether to continue to live. What matters to him now are not the best and worst periods in his whole life, but the best and worst periods in his future life. When an ageing Epicurean is asked why he does not commit suicide, he

may therefore reasonably make the common-sense reply that he prefers living what future life is left to him than ending his life right now.⁹

3. To the extent that Philo succeeds in tying down a plan to follow for the rest of his life that does not need to be altered (very much) as a consequence of unforeseen circumstances, he eliminates the source of anxiety that derives from worrying about how to deal with future eventualities (notably the prospect of death). Philo has already answered the question of how best to live his life, so why keep worrying about it? Authors who speak of ataraxia usually seem to have something less mundane in mind, but freeing yourself from worries about whether you are wasting your life would seem worthwhile whatever elevated conception of ataraxia may be proposed.

7.2 Transcendental version

The transcendental version of Epicureanism generates a more elegant model of Philo's decision problem, because we have seen that a sufficiently widely interpreted conception of ataraxia provides a second anchoring point that can join the state of never having lived in providing a basis for making interpersonal comparisons between the valuations of lives made by different people.

The idea that ataraxia is subjectively timeless is also important to Epicureans (in accordance with the Column Duplication assumption). My own deeply satisfying experiences when immersed in thinking about mathematical or philosophical problems has certainly been timeless in this sense, and I presume that the same is true of Buddhists who report that certain meditation techniques can produce the state traditionally described as enlightenment.

But how does one distinguish between such apparently transcendental experiences and the (subjectively timeless) ecstatic states that are said to be achievable through the use of certain drugs? Epicurus argues that one needs to take account of the later consequences of such intemperance – a point captured in the model of this paper by the requirement that the highs of a life need to be balanced against the lows. Ataraxia is held not to induce such lows, but rather to improve all later states of mind. Some Epicureans even write as if the state of ataraxy is permanent once achieved (as Buddhists hold of enlightenment). After achieving the high of ataraxy, the only low that then need be considered is the non-state of being dead.

However, Epicurus himself does not appear to have regarded himself as enjoying ataraxy on his deathbed. The remarks attributed to him seem more in line with someone looking back without regret over a life

⁹ Thales apparently replied that he was indifferent between living and dying (and so the fact that he continued to inhabit this vale of tears was no proof of his irrationality).

well-lived to a time at which he succeeded in approaching the ideal of ataraxia – much as mathematicians whose talent may have faded nevertheless continue to congratulate themselves on having proved a good theorem or two in their youth. Having come close to ataraxy, one need not continually strive thereafter for further reward but can confront one's end in the knowledge that death has no sting.

7.3 Ataraxy as an *absolute unit*

Von Neumann and Morgenster compared the utility scale they invented with temperature. Employing the Hurwicz criterion can then be seen as requiring the denial of the existence of an absolute zero. I have argued elsewhere against always insisting on such a denial, in which case arguments can be given that result in the weighted arithmetic mean of the Hurwicz criterion being replaced by a weighted geometric mean (Binmore 2009: 166).¹⁰ However, it seems inappropriate in the current context to argue for introducing an absolute zero on Philo's utility scale.¹¹

Instead I suggest that the (unachievable?) ideal of ataraxia can be used to determine an absolute unit – provided one is willing to adopt the transcendental view that achieving ataraxy at some time in a life makes everything else irrelevant. One can then represent the VN&M utility v of a life by the formula

$$v = 1 - (1 - c)^a(1 - C)^b,$$

where C is the utility of Philo's best period, c is the utility of his worst period, and the non-negative constants a and b (idiosyncratic to Philo) sum to one. Notice that the utility of a life in which ataraxy has been achieved then has a utility of one, and the utility of a life of someone who was never born (and so is always in the state of being dead) is zero.

8. CONCLUSION

Rawls (1972) famously argued in favour of clarifying our intuitions using the method of reflective equilibrium. When economists speak of modelling, they commonly mean much the same thing. This paper applies the method of reflective equilibrium to the problem of how to value a life. In particular, it examines how the assumptions necessary to justify the standard utilitarian valuation (in terms of a discounted sum of utilities)

¹⁰ The geometric mean of C and c corresponding to the weighted arithmetic mean $aC + bc$ is $C^a c^b$, where the non-negative constants a and b sum to one.

¹¹ The only obvious candidate is a life that is never lived, but we need to admit states of mind that Philo regards as worse than always having being dead (or never having been born).

need to be modified to generate what I call an Epicurean valuation (in terms of only the best and worst periods in a life). One might say that the essential factor is to replace an objective attitude to time by a subjective attitude, so that living long becomes less important than living well.

ACKNOWLEDGEMENTS

I am grateful to the European Research Council Seventh Framework Programme (FP7/2007–2013), ERC Grant agreement no. 295449 for financial support. Although we don't agree about everything, I am also grateful to Alex Voorhoeve for several helpful discussions.

REFERENCES

- Bentham, J. 1987 [1789]. An introduction to the principles of morals and legislation. In *Utilitarianism and Other Essays*, ed. A. Ryan, 65–112. Harmondsworth, UK: Penguin.
- Binmore, K. 2005. *Natural Justice*. New York, NY: Oxford University Press.
- Binmore, K. 2007. *Playing for Real*. New York, NY: Oxford University Press.
- Binmore, K. 2009. *Rational Decisions*. Princeton, NJ: Princeton University Press.
- Dennett, D. 1987. *The Intentional Stance*. Cambridge, MA: MIT Press.
- Eagleman, D. 2008. Human time perception and its illusions. *Current Opinion in Neurobiology* 18: 131–136.
- Gaskin, J. 1995. *The Epicurean Philosophers*. London: Everyman Library.
- Harris, J. 1998. *The Nurture Assumption: Why Children Turn Out the Way they Do*. New York, NY: Free Press.
- Henry, J. 2009. *Philodemus on Death*. Atlanta, GA: Society of Biblical Literature.
- Hume, D. 1978 [1739]. *A Treatise of Human Nature*. Oxford, Clarendon Press.
- James, W. 1890. *Principles of Psychology*. New York, NY: Henry Holt.
- Kant, I. 2007 [1788]. *Critique of Judgement*. Oxford: Oxford University Press.
- Long, A. and D. Sedley. 1987. *The Hellenistic Philosophers: Volume 1*. Cambridge: Cambridge University Press.
- Milnor, J. 1954. Games against nature. In *Decision Processes*, ed. R. Thrall, C. Coombs and R. Davies, 49–59. New York, NY: Wiley.
- Nagel, T. 1979. *Mortal Questions*. Cambridge: Cambridge University Press.
- Parfit, D. 1986. Overpopulation and the quality of life. In *Applied Ethics*, ed. P. Singer. New York, NY: Oxford University Press.
- Rawls, J. 1972. *A Theory of Justice*. Oxford: Oxford University Press.
- Sen, A. 2000. *Development as Freedom*. New York, NY: Knopf.
- Spinoza, B. de. 2009 [1677]. *Ethics*. New York, NY: Classic Books International.
- Stern, N. 2007. *The Economics of Climate Change: The Stern Review*. Cambridge: Cambridge University Press.
- Von Neumann, J. and O. Morgenstern. 1944. *The Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- Warren, J. 2004. *Facing Death: Epicurus and his Critics*. Oxford: Clarendon Press.

BIOGRAPHICAL INFORMATION

Kenneth G. Binmore is an Emeritus Professor of Economics at University College London. He currently holds Visiting Professorships of Economics at the Universities of Bristol and Warwick. He is a Visiting Professor of

Philosophy at the London School of Economics and a Research Fellow in Philosophy at the University of Bristol. He is the author of a number of books on mathematics, economics and philosophy, including those referenced above. His recent research includes theoretical and experimental work in decision theory and applications of evolutionary game theory in biology.