

TIME IN HUMAN EXPERIENCE

By Jonathan Bennett

This paper consists of eight reflections on the place of *time* in the human condition. Without locking together into a single argument, they are, I hope, interlinked enough to make this more than a mere ramble. The paper began life as the Lewis Burke Frumkes Lecture at New York University in November 2001, which was to be for ‘an audience that will have intellectual interests, but might not consist wholly or even primarily of professional philosophers’.

At every waking moment, each of our senses delivers a little drop of news about the world - an atom of sensory intake - and in time the drops become a torrent. The great flood of news that we call our experience does not sweep us away, because the *thoughts* we bring to bear upon it give us a grounding, hold us steady. Those thoughts hang together; they are not just a second flood of atoms. A torrent of inputs from experience accompanied by a torrent of loose thoughts about them would be no use to us; it would still be a blooming buzzing confusion, with the thoughts merely adding to the turbulence, like being swept down the Mississippi past its junction with the Missouri.

Philosophy investigates our thoughts, trying to see how they fit together into a dependable structure. One way to do that is to imagine our experience’s being different from how it actually is in some basic ways, and to consider how that would affect our thinking. Some imagined change in what the world presents us with may require a change in our thoughts; and the interesting part comes when we alter a thread in the web of thought - relocating it in the web, or pulling it right out as no longer needed - and see what else moves with it.

Why should we do this? Adapting Alexander Pope, I answer that *one* proper study of mankind is man.

1. Famously, time has a *direction*: we all think so, though it has proved hard to say exactly what this means. Some physicists think that their discipline can explain what time’s direction consists in, and I’ll bet it can. But I am concerned with temporal direction as a feature of our lives as lived, a feature of our experience. There ought to be something we can say about this without the help of physics.

Hoping to get some grip on *what* role temporal direction has in our world and in our thoughts about it, let us imagine the world's temporal direction to be reversed, and see what happens. We can easily sketch the surface of a world resembling ours but with time's direction reversed; we recall seeing films run backwards, with dust and bricks collecting and assembling themselves into a chimney, or the like.

Chimneys are only a start of our thought experiment; let's introduce people into it. In the reversed world, a human being comes into existence by the gathering of bits of matter in the ground or in a crematorium, followed by further events that assemble the bits into a structure that works biologically. These are varied. For some people the structuring events occur quietly in a hospital bed: those folk come gently out of that good night. For others the events occur violently in a burning vehicle, and so on. But all humans go out of existence in the same way, by shrinking, entering a woman's womb, and through nine months being absorbed into her body. With care, we can tell this story - the part belonging to physics and biology - without running into incoherence.

When we try to bring human behaviour into the reversed-world story, we find that much of it makes no human sense. Think of a reversed game of golf, or dinner party, or sight-seeing tour: we can imagine the behaviour, but cannot give a coherent account of what the people *would be up to* in behaving like that. But sometimes we are not defeated in that way, and one large class of such cases can be identified: when actual behaviour is any kind of *making* or *destroying*, its temporal dual fits the other member of the pair: constructing becomes dismantling, healing becomes wounding, and so on.

Harming and helping are a pair like that. In harming someone you drive his condition from better to worse, while helping is the reverse of that; so if we take an actual episode in which harm is done, its temporal dual should make sense, understood in terms of help. War provides grist for this mill. A battle involves a long, steady move from better to worse. One side, or both, may have decent strategic objectives; but the battle itself is a steady drive of deliberately injurious human activity through which just about everything and everybody ends up worse off than they were at the start.

Run *that* backwards and you get a protracted move from worse to better - a saga of pure help. Kurt Vonnegut brilliantly exploited this fact in his novel

Slaughterhouse Five. Its central character, a man whose relationship to time is screwed up, sees a film showing the 1945 allied bombing raid on Dresden - a normal film, but he experiences it backwards. Here is Vonnegut's account of what he experiences:

American planes, full of holes and wounded men and corpses, took off backwards from an airfield in England. Over France, a few German fighter planes flew at them backwards, sucked bullets and shell fragments from some of the planes and crewmen. They did the same for the wrecked American bombers on the ground, and those planes flew up backwards to join the formation.

The formation flew backwards over a German city that was in flames. The bombers opened their bay doors, exerted a miraculous magnetism which shrunk the fires, gathered them into cylindrical steel containers, and lifted the containers into the bellies of the planes. The containers were stored neatly in racks. The Germans below had miraculous devices of their own, which were long steel tubes. They used them to suck more fragments from the crewmen and planes. But there were still a few wounded Americans, and some of the bombers were in bad repair. Over France, though, German fighters came up again, made everything and everybody as good as new.

When the bombers got back to their base, the steel cylinders were taken from the racks and shipped back to the United States of America, where factories were operating night and day, dismantling the cylinders, separating the dangerous contents into minerals. . . . The minerals were then shipped to specialists in remote areas. It was their business to put them into the ground, to hide them cleverly, so they would never hurt anybody ever again. (Kurt Vonnegut, Jr., *Slaughterhouse Five*, New York: Dell Publishing, 1969; from chapter four)

'...they would never hurt anybody ever again' - that moves me every time I read it; and yet it's a blunder, isn't it? In the context of *that* story the contents of the cylinders have hurt nobody; indeed, they have done nothing but good.

Still, the rest of the story makes a kind of sense, but only because *harm* and *help* are temporal duals of one another.

I have thought of Vonnegut's story as involving people who act for reasons: they act as they do *so as to* bring help, *aiming to* make things better. Let us take the motivational story further, and see where that leads us. Stretching history a bit, pretend that the actual raid on Dresden was launched as revenge for the less devastating raid, four years earlier, on defenceless Coventry.

On the face of it, the switch of temporal direction turns backward-looking to forward-looking revenge - harming someone because of harm he *will* do to you later. But in our reversed story, both the Coventry and the Dresden episodes are benign - pure cases of helping, healing, making better. So what I am supposing to have been the actual motivation - harming Dresden because of the past harm to Coventry - becomes helping Dresden because of the future help to Coventry.

That makes sense; but now things become harder. Anticipatory reward requires solid knowledge of the future. Quite generally, to make sense of a time-reversed version of a complete, thickly detailed story about human behaviour, we shall have to reverse the past/future asymmetry in our knowledge. Most motives - including revenge and reward - involve firm beliefs about the past, and in the reversed world these will be beliefs about the future. Let us look into what that switch involves.

2. We know a lot about small parts of the past, and much less about comparable parts of the future. Looking back, we have memory; looking forward, we have predictions in which we extrapolate from past to future. Predictions can be sure and true, and memories can be wrong; but still, broadly speaking, as time passes we expand what we know (the past) and contract what we do not know (the future).

In a time-reversed world, it would go the other way, wouldn't it? In that world, as time passes people gradually shift parts of the world from the known to the unknown, rather than from the unknown to the known as we do. Their knowledge is the temporal dual of ours: as time passes, we increase the stretch of time about which we can know fairly directly, and decrease the stretch for which we have to calculate predictions. For the reversed folk, on the other hand, the past is relatively unknown, and the more securely known is the future - so that as time passes their range of secure knowledge shrinks. Few of them know much about when or where

they became alive, for instance, but most are informed about when and where they will go out of existence. More mundanely, a person will be more likely to know what he will eat for lunch tomorrow than to know what he ate yesterday.

Consider now a reversed-world adult with his start in life behind him and his approach to annihilation ahead. Before him lies the period when he will be a toddler - losing his ability to walk and run competently, and losing the last of his grasp of language. He knows little about this period except by hearsay - the reports of people who will be adults at that time. But as it gets nearer he will know more, and know it more firmly; and when he actually *reaches* that toddler period, some of its episodes will loom up on him powerfully and in detail. As they approach, they will become more and more solidly real to him, until they swarm into his present and then immediately drop out of his consciousness.

This story seems to be growing ever more wild, but wild stories taken to extremes sometimes snap back into being tame after all. That may be the case here. Consider this theory:

Our most basic and immediate notion of how past differs from future comes from the difference between what we know fairly solidly and immediately and what we know only by inference. Given that some event is present to me at a time I call Dick, unknown to me at a nearby time Tom, and fairly vividly known to me at another nearby time Harry - that *is* the notion of the relevant times' occurring in the order: Tom, then Dick, then Harry. For us, the difference in our knowledge of them *defines* the difference between past and future.

That's the theory I invite you to consider. Obviously, for much of the past we have nothing but inference, just as we have for the future. But the suggested theory says that our *basic, central* notion of the difference between past and future comes from the difference between memory and inference.

This view about our sense of the slope of time, our sense of how past differs from future, first came to me from Augustine, Bishop of Hippo, who devotes the eleventh chapter of his *Confessions* to a stunningly acute reflection on the concept of time.

After ten chapters in which he unburdens his soul *to God*, Augustine starts chapter 11 by worrying about how a time-bound penitent can communicate with a

timeless God. From that he modulates gracefully into a more general inquiry into time and timelessness, and so we get this wonderful coda to the *Confessions*.

What did God do before he made heaven and earth? Some theologians, Augustine says, take refuge in a wise-crack: 'He was busy making hell, for people who ask awkward questions.' Wanting to improve on that, he eventually works his way to the answer that time did not pass before the creation: God made time itself along with the world; his creatures live in the time that he created, and God himself exists outside it.

Augustine was not the first to say this, but he was unique in how thoroughly he explored it. Plenty of theologians will say that God created time, and leave it at that; because *God* did it, we cannot expect to know how he did it or even what, exactly, it was that he did. Sir Thomas Browne wrote loftily: 'Time we can understand. It is but seven days older than ourselves, and its horoscope is the same as the world's.' I enjoy that, but it is conceptually complacent: Browne helped himself to the notion of time's being created along with the world, without asking what that could possibly mean.

Augustine stared that question in the face, starting with a famous remark that has, figuratively speaking, hung framed on the walls of analytic philosophers ever since: 'What is time? If no one asks me, I know; but if I try to explain it to a questioner, I do not know.'

Mostly he discusses the *reality* of time. At any moment, only the present is real; and that is a mere knife-edge, with no duration; so how - Augustine asks - do there get to be stretched-out periods of time? He answers that the stretched-outness of time consists in a set of facts about how realities are presented to our minds. Do not think that how we experience events somehow *keeps in step with* the temporal order in which they occur; rather, for them to occur temporally ordered *is* for them to be presented to our consciousness in one way rather than another. That answers the question 'What could it be for God to create time?', by saying 'He created time by creating conscious beings to whose minds realities were presented in a certain way'.

Although he does not explicitly put temporal *direction* front and centre, one can see what Augustine's view about that must be. From his standpoint, it is not a deep *fact* about us that we have one kind of access to the past and another to the

future; rather, the different kinds of access define the difference between past and future. He nearly says this outright in this charming passage:

Suppose that I am about to recite a psalm. Before I begin, my expectation is directed to the whole of it; but when I have begun, so much of it as I pluck off and drop away into the past becomes matter for my memory; and the whole energy of the action is divided between my memory of what I have said and my expectation of what I am still to say. But there is a present act of attention, through which what was future passes on its way to becoming past. The further I go in my recitation, the more my expectation is diminished and my memory lengthened, until the whole of my expectation is used up when the action is completed and has passed wholly into my memory. And what is true of the whole psalm is true for each part of it, and for each syllable; and likewise for any longer action . . . including the whole history of the human race.
(*Confessions* Bk 11, §28)

Augustine means, I think, to be illustrating how the difference between our two modes of access to non-present times *creates* or *defines* the difference between past and future. And that lets him say that God's being outside time consists in his knowing everything, immediately, all at once.

3. Thomas Hobbes wrote scornfully about this idea (chapter 46 of *Leviathan*). Theologians deny that eternity is 'an endless succession of time', he wrote, because if God were in time that would make trouble for 'many of their bold opinions concerning the incomprehensible nature of God'. They teach us, he reports with a curl to his lip, 'that eternity is the standing still of the present time, a *nunc-stans* [a standing now]; which neither they nor any else understand.' While I like Augustine's view about what it is to be *in* time, I am on Hobbes's side about the idea of being *outside* it. We can chat along fluently about the non-temporal existence of *numbers*, *propositions*, *relations*, and other abstract entities; but existence outside time for something that thinks and knows - that is a darker matter.

We seem to be quite unable to give ourselves any sense of what it *would be like* to have a timeless mode of being. I now explain why we cannot do this.

To get a sense of what it would be like for us if things were somehow different from how they are, we take fragments of our experience and assemble wholes out of them. We know what it would be like to have purely achromatic vision; for we have watched black-and-white movies, and have seen mountain landscapes whose whole palette is black, white and grey; and we can have the thought of a visual life that is, so far as colour is concerned, *all like that*. Or suppose we want to envisage experiencing an outer world which does not consist of hard physical objects but rather of smooth *waves* of reality of some kind. Never mind the physics. I am talking about the idea of the world's being *given* to us as wave-like, with the sort of immediacy with which it actually comes to us as full of knobbly things. We can get some sense of that, too, by focussing on the parts of our actual experience that pertain to fluids and jellies and clouds, and out of those materials trying to build a picture of a complete course of experience that presents us with an objective, outer world which is not organised in a thing-like manner.

I know of no other way for us to imagine alternative possibilities for ourselves. If that *is* our only way, then to get a sense of what it would be like to exist out of time we must focus on the parts of our lives that are not temporal, and out of those fragments assemble a picture of a way of being that is *all like that*.

All like *what?* *What* fragments? In this case the technique cannot get started, because *all* of our inner lives are temporally ordered, not just over-all but also down to the finest detail. We have no atemporal fragments out of which to build; no ground to stand on while we try to get a sense of a non-temporal way of being.

So temporality lies *deep* in our thought because it spreads *wide* in our experience. We cannot think our way down to a level where time does not apply, because no parts of our experience, however small or odd, lie outside time.

4. As Augustine and I have both emphasized, memory is quite unlike expectation. At least with regard to the events in one's own life, one normally knows *more* about the past than about the future; and *how* one knows it is different. But the *how* part of each story is complex.

On the face of it, for the future I have only prediction - working out what will happen on the basis what has happened up to now, taking the past and present as *signs* of what will happen in the future. Yet some of one's knowledge of the future

is not like that. I now *know* that in a few seconds' time I shall raise my left hand - there! - I knew I was going to do that. How did I know? It was not ordinary, predictive, inferential knowledge, based on reading the signs - like the knowledge I might have that I am about to throw up, or about to lose my temper. It is *not* the case that I detected in myself some sign that my hand was going to go up, and on that basis predicted that I would raise my hand - like looking at the clouds and predicting rain. The lives of athletes would be even harder than they are if *that* were our only way to know how we are going to move. This is a big topic, about which I have no more to say.

For my knowledge of the past I must rely upon *traces* of past events, as when charred trunks now tell us that there was a forest fire earlier. In our world it happens that traces of the past are easier to read than signs of what is to come. (In the temporally reversed world, signs will be easier to read than traces. I didn't mention this back there, because we had enough on our plates without it.) One special kind of trace is written records; an even more special kind is *memory*. Memories *are* indeed traces: how could I have a memory of entering this room if it weren't that the experience of entering laid down some trace to which I now have access? This view that memories are *caused* by the events that they are memories of was espoused by Leibniz, who apparently didn't convince anybody of it. It then dropped out of sight until Martin and Deutscher rediscovered it 35 years ago, and convinced us all.

My present interest, though, is not in the *likeness* between memories and traces of other kinds, but rather in the seeming *unlikeness*. On the face of it, memory-knowledge seems to be less like knowledge gained from reading traces than it is like sense-perception. My present knowledge that a blue heron flew past my study window three days ago seems to be significantly like my knowledge that there are now people in this room: in neither case do I figure it out from bits of evidence; in each, I simply *see* that it is so.

That's on the face of it. But further reflection shows that our memory-knowledge is strikingly *unlike* sense-perception. I was reminded of this by comments that Tobias Wolff made in the *NY Times* about Senator Bob Kerrey's youthful sortie in Viet Nam, in which civilians were killed. Although Kerrey's account clashed with that of another participant, Wolff thought that neither man

was lying: 'My instinct is', he wrote, 'that each of them believes the the story his memory has been telling him all these years.' He enlarged on that as follows:

We tend to think of memory as a camera, or a tape recorder, where the past can be filed intact and called up at will. But memory is none of these things. Memory is a storyteller, and like all storytellers it imposes form on the raw mass of experience. It creates shape and meaning by emphasizing some things and leaving others out. It finds connections between events, suggests cause and effect, makes each of us the central figure in an epic journey toward darkness or light.

This rings true, doesn't it? Never mind suppressed and recovered memories, and other such pathologies; even normal memory involves a constructing, a trying-to-make-sense.

We can imagine being better rememberers, ones whom Toby Wolff's description does not so closely fit; but one part of it cannot be imagined away like that, because we are stuck with it - it is *necessarily* true. Immanuel Kant discovered that the central, healthy, normal functioning of memory *must* involve a lot of reconstruction in the light of cause and effect.

He raised a terrific question: given that I remember seeing a flash of lightning and hearing a trumpet, how can I remember the order in which these occurred? Events don't come with their 'when' written on their faces, Kant said; we must afterwards *put* them into the order in which they actually occurred. Even if in many ways our memory *does* function like a camera or tape-recorder, it cannot record *when* something occurred: the tape-recorder might tell you how she sounded, the camera how she looked; but neither can tell you *when* you heard and saw her.

Sometimes we can get the order right in either of two ways that Kant does not mention. One: we might remember experiencing one of the events while remembering the other. We might, for example, remember hearing the trumpet while thinking 'It would have been nice if this trumpet had sounded a bit earlier, while that lightning was striking'. Two: we might remember a pair of events as parts of a single brief event - for example, remembering the experience of flash-then-sound - like remembering da-da-da-*daah* rather than *daah*-da-da-da. That gives us the ordering only of pairs of events that are close enough in time to be

contained in a single *now*. We can extend its range by linking such pairs into longer chains of events of which we have continuous memories; but most remembered pairs of events are not linked in that way. So these two ways of remembering the order of events cover only a little of the ground. The great remainder has to be dealt with in the manner that Kant describes.

On our hike we stopped for a snack, and we paused to admire a peregrine falcon; the two episodes were well separated, we do not remember everything that happened between them, and during neither episode were we remembering the other. How, then, can our memory tell us which came first, the snack or the bird? It does this, Kant says, by making a judgment about the order in which they *had* to occur: we order our past in time with help from the notion of causal necessity. For example, we recall eating the snack with the sun in our eyes, and straining to see the bird in fading light. Or we recall admiring the bird on a rock near the summit, and we know that we couldn't have got that high without having a snack on the way.

Clocks often solve such problems, but they are *parts* of Kant's solution, not *rivals* to it. I recall that while I had my snack my watch said 3 o'clock, and when I watched the falcon it said 5 o'clock; but those two watch-readings are merely two more remembered events which I have to put in their proper order. Watches are designed to make this easy - we know that the short hand on a properly handled watch can't get to 5 without passing 3; so I easily order the watch-hand events and, through them, the associated episodes in the hike. All this can be re-applied to calendars.

This strikes me as a nice bit of philosophy, which fits our experience well. Try it for yourself some time: think of two episodes in your past life whose dates you don't recall, and try to decide which came first. If you succeed, I predict, it will be by realising that only one ordering makes causal sense.

5. I want us now to think about one kind of knowledge of the future - namely, predictive knowledge of how people will behave. Not the timeless knowledge that Augustine attributes to God, but rather a kind of foreknowledge about behaviour that might be available to mere humans.

How will this boulder roll if I kick it down the slope just *so*? In theory I could answer this by knowing enough about the boulder's weight and shape, the

contours and texture of the hillside, and so on. Well, one might similarly be able to know how a person will act a minute from now by knowing enough about the lay-out of his body, especially his brain, and of course the lay-out of the environment that will act upon him in the next minute. No humanly possible scheme of prediction could tell one for sure how a person will behave a week from now; for that depends upon too much else - how the wind will blow, what the neighbours will say, how the stock-market will move - and it would be crazy to think that *we* could ever securely predict all that.

What about the thought that it is *in principle* predictable? That comes from the thought of the world as being strictly determined, causally controlled down to the finest detail, so that its entire future is settled already, and has only to unroll inevitably, in accordance with rigid causal laws. Many people feel undermined by that thought, even if they are sure that nobody could actually do the predicting; and the books and articles arguing that they are right to feel like that are matched by others arguing that they are wrong to do so.

Well, we have it on good authority that our world is *not* deterministic: its basic physical laws do not determine outcomes but only make them probable, and so our futures are not entirely encapsulated in the present. That is my excuse for running away from this topic, though it does not kill the topic itself. Even if determinism is false, there are challenging problems about what it implies, what trouble we *would* be in *if* it were true.

Problems arise even from the more modest idea of limited short-term predictability, such as may be possible even if determinism is false. When I am wondering whether to turn right or left at the end of the lane, I find it unsettling to think that someone who knew enough, and could compute fast enough, might know already which way I am going to turn. Why does that thought upset me?

Well, it seems to me that when I *decide* what to do, I thereby close a question which was genuinely open until just then: how I am to turn remains *unsettled* until I make my decision, which *settles* it. But if my decision was - or even could have been - predicted a moment or two ahead of time, then the question was settled *back then*. Even if my wondering and then deciding which way to turn were required parts of the entire causal chain, I am still left with a sense that if they were securely predictable then they are not what I have thought them to be - the closers of something that was wholly open.

It bothers me that I cannot crystallize this line of thought into something sharply literal, not depending on the metaphorical terms 'open' and 'closed'. I cannot say any more about this, but I shall address a certain worry over predictability that may be what lies behind the open/closed one.

It is the thought that if I knew myself well enough, and could calculate fast enough, I could tackle my own practical problems in the role of a predictor, pushing aside the role of decision-maker. Instead of wondering whether to turn left or right, I could take a quick look at the dials and print-outs from my brain and *predict* that I was going to turn to the left, say. That is a threat, all right. My human dignity as I see it essentially involves my being a creature that wonders what to do, and then decides. I hate the idea that I have this role only because I am, relatively speaking, ignorant and slow.

Fortunately, this threat collapses under the pressure of reflection. Even if someone else could predict which way I shall decide to turn, the facts about me on which he based his prediction are not the ones I would confront if I tried to predict my own behaviour. Like him, I would attend to a plethora of events in my brain; but my data-set - unlike his - would include brain-events reflecting the fact that I was conducting this exercise; I could take those into account as well, but that taking-into-account would create yet further events, which might also be relevant to the outcome; so I should attend to them also, but *that* act of attention would you can see how the argument goes. It likens predicting oneself to chasing one's shadow; and it concludes that even if we are predictable, our role as deciders is secure. For a while it looked as though that role might be a mere product of our ignorance about ourselves, but now we see that it is more than that. Cure the ignorance and there is still no coherent story about how we could be forced or even invited to abdicate as deciders in favour of being predictors.

6. A special difficulty arises out of the possibility of someone else's predicting some of our behaviour. Suppose there exists a genius who predicts with great accuracy how you will behave in certain circumstances. He does not predict far in advance, and does not predict everything; but he securely foretells certain practical decisions that you will make in the near future, and acts on the basis of those predictions. The existence of such a person might induce in you anxiety,

depression, even panic; but I am concerned not with the *practical* or *emotional* difficulties he can cause you, but rather with a profound *conceptual* or *intellectual* nuisance that can arise out of his activities. It comes to us in the vehicle of the Newcomb problem.

I introduce it with a different supposition. Suppose that scientists have discovered a human gene which pretty reliably has two effects: lung cancer, and immoderate consumption of ginger. Neither the cancer nor the ginger-eating causes the other; they are joint effects of a single genetic cause, and that links them statistically so that each is a predictor for the other. Now, you learn those facts; you have been eating a lot of ginger, and you now know that your intake of ginger would count as a fairly good predictor for lung cancer. Should you give up ginger? Obviously not. If you don't have the gene, the link between ginger-eating and cancer has nothing to do with you. If you do have the gene, fighting what may be one of its effects won't lessen the other. We are comfortable with this conclusion, aren't we?

The Newcomb problem has the same structure, but a different content - one that creates acute *discomfort*. What is predicted in the Newcomb problem is not tumours or gluttony, but practical choices. Here goes.

You are confronted by two boxes, Small and Big, and invited to take either *Small* or or *both Small and Big*; it's your choice; you may *keep* whatever you *choose*. At the time of choice, you know the following things:

1. A genius has a so-far-perfect record in predicting what people will do given this choice; and he has predicted how you will choose.
2. This predictor has put money in one or both boxes; at the time of your choice their contents are settled.

Now for the two crucial things:

3. If he has predicted that you will take both boxes, he has put \$1000 in Small and nothing in Big.
4. If he has predicted that you will take Small only, he has put \$100,000 in Small and \$1,000,000 in Big.

Knowing that and no more, which choice should you make?

On one hand, your actual choice cannot affect what's in the boxes. There they sit, containing whatever the predictor has put into them. You may get more by

taking both, and can't get less; so you might as well take both. Imagine pointing to Big and saying 'That may have a million dollars in it, there for the taking; but I'm not going to take it'. Isn't that crazily feeble?

On the other hand, this predictor knows his stuff: people who took only box Small have always come away with an excellent lump of money (\$100,000), while those who took both boxes have come away with only a hundredth part of that.

There is an *enormous* literature on this - a seemingly endless debate between one-boxers and two-boxers. I have not read much of it, though I did recently see a piece entitled 'Why Aintcha Rich?' by David Lewis, in which he argued that the *rational* choice is the two-box one. Why are those who make that choice poorer than those who take only one box? Lewis answered that if the set-up described in the Newcomb problem really existed, this would be a world in which irrationality could sometimes *dependably* lead to better results than rationality.

Well, perhaps. But I think there may be a different moral to be drawn. In recent years a few philosophers have contended that some of the perennially disputed problems in philosophy have lasted so long because we are *not equipped to understand the right answers to them*; they mean that as a basic psychological or biological fact about us. An example that has been proposed is a hornet's nest of difficulties buzzing under the general label 'the mind-body problem'.

The general thesis that we can ask some questions whose proper answers lie beyond our reach strikes me as plausible, as does the applying of it to the mind-body problem; but I can't go into that now. I do suggest that the Newcomb problem points us to a further instance of the general thesis. I suspect that we are *not conceptually equipped to think straight* about our lives on the basis that foreknowledge of our conduct is operative in the world. I suspect this to be one of several respects in which the relative unknownness of the future is of the essence of the human condition.

7. Here we are with our pasts and futures. We can look both ways - reminiscently in one direction and expectantly in the other - and which way we look, and what we look for, and in what circumstances, tells a lot of the truth about us.

One kind of looking forward expresses discontent with the *present*, and an evasion of it. Chekov knew about this. In one of his great, sad comedies, three

sisters who feel thwarted in their lives caress the thought of the good times they will enter into when, as they hope, they leave their provincial town and settle in Moscow. In their case, the focus on a possible future is evasive and destructive; we distinguish it from the kind of forward-looking that goes with *activity*. Obviously, anyone actively engaged in some project - the pursuit of some goal - must look to the future in planning it, executing it, evaluating it as it unfolds. That is starkly unlike the passive hankering of the three sisters; but some philosophers have implied that the two are alike in that each expresses discontent with present. In this they have been making a curious mistake about purposeful activity. It seems blatant when clearly pointed out, but so do most philosophical errors; the cross we philosophers carry is that it is hard to show us to be wrong without making us look idiotic. The mistake I am considering, however, is subtle enough for some great philosophers to have fallen into it.

John Locke made the discontent mistake (as I shall call it) when he wrote that all purposeful activity arises from *uneasiness* - a vague technical term of his, referring to any kind of awareness that something is amiss and needs to be fixed. He had earlier bought into the popular doctrine that people are always moved to act by their judgments about what would be *best*; but he had come to realize the error in this.

He brings out its falsity, and his remedy, in his account of a compulsive drinker:

Let a drunkard see that his health decays, his estate wastes; discredit and diseases and the want of all things - even of his beloved drink - awaits him in the course he follows; yet the returns of uneasiness to miss his companions, the habitual thirst after his cups at the usual time, drives him to the tavern, though he has in his view the loss of health and wealth and perhaps of the joys of the after-life: The least of these is no inconsiderable good, but such as he admits is far greater than the tickling of his palate with a glass of wine or the idle chat of a soaking club. His trouble does not come from not viewing the greater good; for he sees and acknowledges it and, in the intervals between his drinking hours, he will resolve to pursue the greater good; but when the uneasiness of missing his accustomed delight returns, the greater

acknowledged good loses its hold, and the present uneasiness determines the will to the accustomed action . . . And thus he is from time to time in the state of that unhappy complainer, *I see and approve the better, I follow the worse*. (*Essay Concerning Human Understanding*, II.xxi.35)

This contains some truth, doesn't it? For most of us, alas, the mere judgment that *it would be best* if we acted in a certain way does not *ensure* our acting thus. Locke can reasonably suppose that for a value judgment to get leverage on our conduct there needs to be discomfort, a burr under the saddle, an uneasiness.

He goes further, however, generalizing from this to an across-the-board theory of human motivation: *all* voluntary behaviour, he holds, all conduct in which one tries to *change* something, proceeds from uneasiness, discontent with the status quo, a sense of something's needing to be remedied.

This can sound right: why would you cross the road if you were not in some way - ever so slightly, perhaps - discontented with the side you are on? I used to agree with Locke about this, and was led from it to a jeer at the Christian idea of heaven. Given the blissful contentment of the blessed host (I argued), none will suffer from uneasiness, nothing will need to be remedied, so no-one in heaven will be motivated to act; but our highest happiness lies in activity, and we cannot imagine a happy state of being in which we *do nothing*.

Spinoza made a different theological use of the discontent mistake. He was an atheist; but he used the name *Deus* - 'God' - for the entire world. This is pantheism: God as everything. Rather than flinging this dogmatically into the teeth of Christianity, Spinoza *argued* for it. If you start from Christian doctrine and think clearly and ruthlessly enough, he argued, you will find yourself driven his way. For a start, you will discover that some parts of Christianity conflict with others, so that some of your dogmas have to be dropped; and one that should be sacrificed, he argued, was the doctrine that God made man in his own image. That, Spinoza thought (reasonably enough), makes God out to be 'a man, or like a man', which he vigorously denied. He regarded as an infantile fantasy the notion of an omnipotent and eternal *person*, and his God has nothing of the personal about him/it. *Someone who truly loves God*, Spinoza memorably said, *will not try to get God to love him back*.

To talk the Christian into agreeing about this, Spinoza must get him to agree that God does not do things in the furtherance of plans or purposes. He argues for this by assuming that acting with a purpose involves lacking something, wanting to improve one's situation; whereas the Christian salutes God as being perfect and complete, and thus not in need of any fix. So the Christian should admit that God does not act purposively, Spinoza concludes. That takes a giant stride towards God's not being like a man, and a small step towards God's being the whole world.

But Locke and Spinoza and my younger self were all wrong about this. When we are moved to act by a preference for a certain possible future over something else, it need not be a preference for it over *the present*. We may have no quarrel with how things are right now - no uneasiness, no burr under the saddle - but prefer a certain possible future to *another possible future*. Walking across the downs I suddenly find myself at the edge of a cliff, and I stop. I have no complaint against how I am situated; indeed, I *love* my present condition and place. I stop because I prefer a future in which I am alive and well to *a future* in which I am dead at the foot of the cliff. Another example: when I embarked on writing this paper, it was not to scratch an itch, to remove an uneasiness. It was because I preferred a future in which the paper had been written to one in which it had not. At times in the summer of 2001 the thought of the paper caused me discomfort because work remained to be done on it. But then the project caused the uneasiness, whereas according to Locke's theory the uneasiness always causes the project.

Indeed, what drives purposive activity is often not any kind of preference for some outcome but a sheer desire to *engage in that activity*. That is really why I write - not so as to have written a paper, but for the pleasure of the thinking and writing.

8. People look to the past in a variety of manners and moods: with forlorn longing, like the servant Firs in Chekov's *The Cherry Orchard*; with remorse, like someone realizing how much worse a parent he was to his children than they are to theirs; with calm satisfaction, like my father who on his deathbed said to me 'I think I have acquitted myself creditably'; with entranced pleasure at having all that material to reflect back upon, like Proust; and so on.

One way of looking towards the past, though I do not understand it well and have nothing enlightening to say about it, is my final topic. I bring it in because it preys on my mind. I have long been aware of it as active in some people's thoughts and feelings, and recently - no doubt because of my age - it has started to have a role in mine. It is a fascination with pastness as such - experiencing the past as a kind of enchanter *just because it is past*. Or because it *is* past and *was* present.

I have this feeling, in an acute form, when reading the diary of Samuel Pepys. Not continuously, but when he writes things like this:

Before going to bed, I stood writing of this day its passages - while a drum came by, beating of a strange manner of beat, now and then a single stroke; which my wife and I wondered at, what the meaning of it should be.

Stretching back three and a half centuries, and yet intensely *present*, this passage has on my mind a peculiar effect - the effect that I am struggling to convey to you. In another entry, Pepys writes:

I sat up till the bell-man came by with his bell, just under my window as I was writing of this very line, and cried "Past one of the clock, and a cold, frosty, windy morning."

That gives me the shivers, and not because it is about cold and frost.

As those examples indicate, this special feeling need not involve one's own personal past. Michael Frayn was onto that when he wrote:

Our nostalgia is often strongest and sweetest for things which never were part of our life at all. No one aches so much for the lost glories of the pre-1914 world as those who were just too young to remember it; or for the lost sweetness of summer weekends at the great country houses as those too poor ever to have been invited. (*Constructions*, London: Wildwood House, 1974, section 296)

That, I take it, concerns nostalgia - if that's the word for it - directed towards the past *as such*.

A photograph on a wall in my home shows a five-year-old girl washing medicine bottles in a tub, nearly a hundred years ago. As I sit in that room listening to music, I often gaze at that picture and am caught by the sheer idea of its representing something that was real and is now gone, something that is *of us*

and yet infinitely removed. The girl is my mother, and that feeds into my feelings about the picture; but they are also nourished by the other thing, the haunting idea of pastness as such, the idea of a person who did exist, and left traces which we can inspect now - including a trace that you are listening to now - and yet is as unavailable to us, as cut off from us, as Ahab or Macbeth.

Tennyson wrote eloquently about this. At one point in his long poem *The Princess*, a maid-in-waiting sings a song that begins 'Tears, idle tears, I know not what they mean', and expresses melancholy at the thought of the past *qua* past. I have no doubt that in this song Tennyson speaks for himself. He also speaks for me, both about the power of the tears and in the confession 'I know not what they mean'.

I like his calling 'the days that are no more' *fresh*, though I couldn't tell you why he does so. I do understand his calling them *sad*:

Fresh as the first beam glittering on a sail,
That brings our friends up from the underworld,
Sad as the last which reddens over one
That sinks with all we love below the verge;
So sad, so fresh, the days that are no more.

In the next stanza he brings in *sad* again and also *strange*. In a stunning comparison, he likens the sad strangeness of the *past* to the sad strangeness of the *present* to someone who has almost no *future*, the strangeness of a dying man's experience of the start of what may be his last day:

Ah, sad and strange as in dark summer dawns
The earliest pipe of half-awaken'd birds
To dying ears, when unto dying eyes
The casement slowly grows a glimmering square;
So sad, so strange, the days that are no more.