

Ensuring Two Bird Deaths With One Throw

JOHN LESLIE

I. Symmetry

Crossing a plain, I come to what looks like a gigantic mirror. But pushing a hand against it, I feel flesh and not glass. The universe must be symmetric, the flesh that of my double—left-right reversed but otherwise a perfect replica.

The universe must also be fully deterministic, for how else could my double have moved exactly as I did? But never having seen freedom and determinism as incompatible, I find this untroubling. In fact I rejoice. My decision-making ability, I tell myself, is not limited by any admixture of quantum randomness such as would make me in part a puppet pulled by the strings of Chance.

Do I not actually have twice the power I earlier seemed to have? Besides governing my own hands and legs, do I not govern others too? Cannot I make my double kick, wave, clap, whenever I want? Cannot I throw two stones instead of one?—the stone I hurl to kill a bird, and the second stone which my double hurls simultaneously? If ever I controlled what my image did in a mirror, why deny that I control what my double does?

Well, grounds for denying it are easy to find. Yes, the universe-halves interact. I can smash against my double painfully, and while unable to enter his half of the universe (he keeps getting in the way!) I have no problem in seeing him since light waves pass through each other instead of colliding and bouncing back. Yet he and I are in crucial respects causally independent. Seeing a snake sneaking up on him, it is not of that particular snake that I feel afraid but of the other which I thereby know to be sneaking up on me. It is my fear which makes me run, my double's fear which makes him do likewise. I do not genuinely cause him to run by myself running.

Maybe my double moves as he does "because" I move as I do. It might even be said that the heaviness of the rock in my hand "causally explains" why he drops his rock when I drop mine. Rather similar senses of "because" and "causally explain" are recognized by Frank Jackson and Philip Pettit (1988, esp. pp. 392-3). The sameness of the force on two electrons "causally explains" why they accelerate at the same rate, they write. And why was the price of lamb identical at two auctions? "Because," they say (and it is surely a "because" of Ordinary English), "the relationship between supply and demand was the same." Yet the fact remains that the stone killing the bird in the other half of the universe is accelerated by my double's hand, not mine. I do not throw two stones.

Still, by choosing to run I can *ensure* that my double runs. While I cannot strictly speaking make him throw stones, I can *make it certain* that he will. I have only to decide firmly to throw stones myself.

Granted that the universe is symmetrical, the making certain would be no mere matter of my giving myself evidence that my double would throw stones, evidence which made *me* certain that he would. It is not just a psychological or epistemic certainty that *I* can make—a question of personal confidence or of the right to be confident. Instead this truly is a matter of being able to ensure, to guarantee, to make something *be* certain. Not, of course, of being able to ensure or guarantee causally, but rather of being able to ensure, guarantee, in the fairly straightforward (and definitely strong) sense that *if I were to throw stones then he would*. I could no doubt have psychological or evidential assurance that he would, yet there would also be the fact that, yes indeed, he *would*. He could not fail to, granted that he really was my double.

Ugly though this terminology is, let us speak of my “quasi-causing” whatever I ensure or make certain in the above-recognized senses. The “quasi-” signals that this is not true causation. Still, it is just as good as true causation for guaranteeing that things get done. If I want a bird killed in the other universe-half, all I need do is throw a stone appropriately in my half.

Some morals of all this have been hinted at by David Lewis in “Prisoners’ Dilemma Is a Newcomb Problem” (1979).¹ Lewis, however, does not himself bow before what I take to be the overwhelming force of some of his own points. I shall therefore develop somewhat similar points with only occasional reference to his paper.

2. Newcomb

Here is a version of Newcomb’s Problem. A first box, transparent, contains a thousand dollars; a second, opaque, perhaps contains a million. Shall I open just the opaque box, or both? It might seem obvious that I should open both. Grabbing the visible thousand could not cause any concealed million to vanish. Failure to grab it could not cause any million to jump into being. But I have been in this situation twenty thousand times before, each time deciding at the last moment, often as randomly as I could, whether to open both boxes or just the opaque one. On opening the opaque box I had always found one of two messages signed “Professor Laplace”. Message A: “You opened both boxes and, predicting you would, I’ve put nothing in this one.” Message B: “You opened this box only. So predicting, I placed a million dollars in it.” Now, in every single case the professor had predicted correctly.

His predictive successes could seem to give strong grounds for opening only the opaque box. Yet the grounds for opening both boxes could also seem as strong as ever.

It is widely held that there would be no problem if one could know for sure that

¹ On the excuse that you need at least two prisoners to get the dilemma, I shall follow Lewis in writing “Prisoners’ Dilemma” although “Prisoner’s” is what one normally sees.

the professor would never make wrong predictions, since then the only rational policy would be to open just the opaque box. Opening the transparent box would make you know you would not get the million; you then might as well leave the opaque box unopened. Opening just the opaque box would make you know you had got the million; without bothering to look inside the box, you could carry it off to the bank.

However, many people consider that, given the slightest chance of the professor's predicting wrongly, one should "two-box" in seemingly paradoxical disregard of his previous predictive successes. Others, though, think that even had his predictions been wrong, say, fifteen per cent of the time, there would be good grounds for "one-boxing"—in seemingly equally paradoxical disregard of how the million dollars are either there or not there *already*.

There would be no paradox if the act of one-boxing (or of two-boxing, as the case might be) were known to cause the professor *to have predicted* that very act, the causation here working backwards in time, and if it were also known that this backwards causation would be followed, as a matter of ordinary forwards causation, by the professor's inserting or not inserting (as appropriate) of a million dollars. Two-boxing would then genuinely cause the opaque box to be empty of dollars. Its emptiness would indeed have been settled BEFORE any decision to two-box, yet it would have been settled BY that decision. All this sounds queer, yet what does that go to show? Merely that backwards causation would be queer, which nobody denies. Let us now agree that it would be so very queer that the idea of it had best be abandoned. But how then can Prof. Laplace predict so well?

Presumably he uses not backwards causation, but backwards quasi-causation.

3. Backwards quasi-causation

Newcomb's Problem might appear to show Determinism's wrongness. For suppose the cosmos were fully deterministic. A predicting demon, if able to know the details of some early situation without disturbing it, could then act like Prof. Laplace, making it rational to pay a small fortune for the privilege of opening the opaque box only. But, you might be tempted to argue, it could not be rational to pay for that privilege because two-boxing could not cause any million dollars to vanish; hence cosmic determinism cannot be possible. However, such an argument would overlook two things. First, that *quasi-causation* really would be every bit as good as *causation* for guaranteeing, say, a bird's death or a box's emptiness. (Want a chicken's neck wrung in the other half of a fully symmetric universe? Just wring a chicken's neck in your own half.) And second, that *backwards quasi-causation* can avoid various problems which attend *backwards causation*: for instance the problem that the past might be considered "simply not there any more" in a way making it impossible to influence its details.

Quasi-causation is a matter of replicas. I quasi-cause *my double* to throw stones by throwing them myself (and, of course, my double also quasi-causes *me*

to throw them: unlike true causation, quasi-causation is reciprocal). Relationships of *backwards* quasi-causation are ones in which persons or situations stand to *earlier* replicas of them. Prof. Laplace's successes could be explained if the cosmos were fully deterministic and if he had an immensely accurate replica of me, a replica developing more speedily than I did.

Let us specify that he has a computer model of me rather than an atom-for-atom biological replica. Because the model's developments are speeded up, we might say that they quasi-cause my later developments. But it would be equally correct to state that *my developments* quasi-cause *the model's earlier developments*. I quasi-cause the choices which my computer replica has already made. (Yes, it was caused to make those choices by its previous states; but this does not deny that it was *quasi-caused* to make them by my states, which were in its future. Being caused by one affair is compatible with being quasi-caused by some other affair entirely or by an infinite number of such other affairs, i.e. by the states of an infinite number of replicas.) Were I to open the transparent box, this would quasi-cause my computer replica *to have opened* a computer replica of that box—since Prof. Laplace's model of me would be far from reliable unless accompanied by models of the things with which I interact, including the boxes. And the professor, who would have noted the opening of that box-replica, would then have caused the opaque box to contain no money. Conclusion: If wanting a million dollars I had better leave the transparent box untouched.

Lewis, after pointing the way to this conclusion by speaking of replicas superbly adequate for predictive purposes, still refuses to accept it. He writes that "some—I, for one—think it is rational to take the thousand no matter how reliable the predictive process may be. Our reason is that one thereby gets a thousand more than he would if he declined, since he would get his million or not regardless of whether he took his thousand." Yet this seems to me wrong when we have backwards quasi-causation followed by forwards causation. Taking the thousand would quasi-cause an earlier taking of a replica-thousand, which would in turn cause the million to be lost.

4. Prisoners

Next day, the professor has organised a Prisoners' Dilemma. I sit opposite a computer replicating me and my surroundings so well that its screen appears to be a mirror. Before me are buttons marked CONFESS and DON'T. If I push my CONFESS button and my computer replica pushes his, I shall get a thousand minutes of freedom from prison. But if I push DON'T while my replica does the same then I shall do far better, getting a million free minutes. Now, I am further told that if I push CONFESS while my replica pushes DON'T then I shall get the million minutes *plus* the thousand minutes, whereas if I push DON'T while my replica pushes CONFESS then I shall get no free minutes at all; but why bother my head with all this? By putting out my tongue and biting on it while waving my

arms anticlockwise, I have become convinced that the professor has done a fine job of modelling me: my replica will be virtually sure to do whatever I do. (I may actually have sped my hand towards the CONFESS button suddenly while feigning to be fully occupied with nose-blowing—the speeding forward of my replica’s hand at the same instant then stopping me in my tracks.) Surely, then, the sensible thing is to push DON’T.

It might be protested that, just conceivably, the computer will go wrong. Now, suppose I push CONFESS. If the replica pushes DON’T, I shall be glad of having acted as I did. I shall have earned one million one thousand minutes of freedom—a thousand minutes more than if I had pushed DON’T. If, on the other hand, the replica also pushes CONFESS, I shall have gained a thousand minutes of freedom whereas I should have gained none by pushing DON’T. So CONFESS *dominates* DON’T, no matter how slim the chance of the computer going wrong. CONFESS *earns an extra thousand free minutes regardless of what my replica does*. It is thus the best choice even if I see my replica’s hand speeding towards the replica CONFESS button in seeming imitation of how my own hand speeds. Since there must always be some chance that the replication will break down at the last moment, DON’T may be what the replica will actually push. Must I not argue like that, to be rational?

Not at all, I would say. It would be wrong to be motivated by a tiny chance that the replication would break down, disregarding the fact that if it did not—as is far more likely—then the only possibilities would be (a) that both I and my replica would push CONFESS (result: a mere thousand minutes of freedom), and (b) that both would push DON’T (result: a million minutes of it).

Still, what can have gone wrong with the Dominance Argument that pushing CONFESS would be best no matter what my replica did? Its error, I suggest, is in treating the choices of each agent as if they were disconnected from those of the other. Now, they are indeed disconnected causally, yet suppose for a moment that they were not. Suppose that pushing a particular button would *cause* my replica to act similarly, almost certainly. Could I still use the Dominance Argument? No. Pushing CONFESS would almost certainly earn me a mere thousand minutes of freedom whereas pushing DON’T would almost certainly earn a million minutes. The idea of “doing better no matter what one’s replica does” is inapplicable when the replica’s behaviour is almost sure to be controlled by one’s own. But recall now that causation and quasi-causation, control and quasi-control, can be equally good at guaranteeing things. Remember the wringing of the chicken’s neck. If you were actually faced by a replica doing exactly what you did over many hours, would you use the Dominance Argument? Waving your arms about in fantastic ways, you see your replica doing likewise as faithfully as any mirror image would. Do you now move your hand to the CONFESS button, helpfully labelled “Just a thousand free minutes if you and the replica both push this”, instead of to the other button labelled “A million free minutes if you both push it”? Do you say, “Since this is mere quasi-causation, not real causation, Logic tells me to push CONFESS”? Presumably not.

Notice, though, that if contemplating the affair abstractly—and particularly if constructing four little compartments showing the four sets of possibilities (that both push CONFESS; that both push DON'T; etcetera)—then it is easy to be impressed by the Dominance Argument even when *actual causation* is at work. Let us say that instead of a computer replica there is an actor trying to imitate you. (A genuine case of causation, this.) Is there not a tiny chance that he will fail, pushing CONFESS when you push DON'T, or *vice versa*? Construct four little compartments. If you push CONFESS, the compartments can seem to tell you, then you will be glad to have done so regardless of what the actor does. Follow the Dominance Argument without thinking any further. You will have plenty of time for thought while in prison.

5. Real life

The computer replications above-described would be extraordinarily hard to produce in practice. Any efforts I made to choose between two-boxing and one-boxing in a random way could well be influenced by trivial stimuli from my surroundings. Prof. Laplace's computer would therefore have to model those surroundings in immense detail. Also, any attempt by him to discover every detail of my brain at any instant, so as to be able to predict what I would choose, would seemingly necessitate frying it with X-rays. Even if it did not, there would be the difficulties to which Heisenberg pointed. The more exact the professor's knowledge of the positions of my brain's particles, the less would he know about their momenta. Further, there would be difficulties of the kind investigated by Chaos Theory: the slightest difference in two situations often quickly leads them to develop very differently. Bouncings of billiard balls on a frictionless table would soon be affected markedly by a spectator's gravitational pull. And again, quite apart from whether Prof. Laplace had any ability to forecast that my cerebral activity would take place in this or that fashion, it might be that quantum physics was such that there just *was not yet* (or was not yet "in anything more than a trivial, semantic way"²) any fact of the matter with respect to exactly how it would take place.

What possible relevance to real life could we find here, then?

The answer is supplied by Lewis in the paper cited earlier. Any conclusions drawn from the case of superbly accurate replicas should be extended also to that of markedly imperfect replicas: crude computer models, maybe, which yield predictions successful at a level only a little better than chance. Now, says Lewis, "the most readily available sort of replica of me is simply another person, placed in a replica of my predicament". If I am in a Newcomb predicament then putting you in the same predicament and seeing how you decide might give useful guid-

² If a particular atom of radium decayed in 1940 then Ordinary English seemingly forces us to say that in 1939 it was a fact that it would decay next year. But this does not give us the non-trivial truth that whether it would decay next year was already settled.

ance on how I would probably decide, given that you and I are somewhat alike. And Prisoners' Dilemma predicaments—predicaments involving two or more people perhaps much alike and each placed in the same difficult situation—can be strikingly similar to Newcomb predicaments. Individuals in both kinds of predicament can feel powerful but conflicting tugs: (a) the tug of the Dominance Argument that by two-boxing or pushing CONFESS one gets, say, an extra thousand dollars or thousand minutes of freedom, regardless of what is done by a computer replica of oneself or by the other biologically ordinary prisoner who is a replica (of a sort) of oneself, and (b) the tug of the Expected Utility Argument which asks us to be guided instead by the fact that one-boxing or pushing DON'T would almost certainly maximise the gains, in dollars or minutes of freedom or whatever, of any large number of people confronting the same sort of choice. Now, as Lewis observes, Prisoners' Dilemmas "are deplorably common in real life".

In particular, real life very deplorably often includes situations in which you think of some other person and yourself as much alike and still wonder whether to trust him or her to co-operate in doing what would yield the best result for the pair of you.

This could make the Dominance Argument conclusion reached by Lewis—viz. that one should two-box "no matter how reliable the predictive process may be" and that, correspondingly, "ratting" (confessing) in a Prisoners' Dilemma situation is rational "no matter how much alike the two partners may be, and no matter how certain they may be that they will decide alike", because "one is better off if he rats than he would be if he did not, since he would be ratted on or not regardless of whether he ratted"—itself deplorable.

Many people, it is true, have thought that the Expected Utility Argument becomes powerful only when it is certain or virtually certain that a Newcomb Predictor will predict correctly or that one's fellow prisoner (or prisoner-analogue) will replicate one's choices. If they were right then the relevance to real life would be very slight. However, they would seem to be wrong. What if Prof. Laplace's computer, instead of being 99.999% reliable in its modelling of my choices, were reliable only 92% of the time because of random voltage irregularities? I could still have strong reasons to avoid two-boxing or pushing CONFESS. (Watching your replica on the computer screen, would you push CONFESS just on the 8% chance that a voltage irregularity would lead the replica to push DON'T ?) And there is no significant difference, so far as concerns what it is rational for me to choose, between Computer One which gets wrong results 8% of the time because its electricity supply is irregular and Computer Two which gets wrong results 8% of the time because its replica of me is crude. Using the language of quasi-causation, we can say that quasi-causation comes in varying degrees of perfection and that even a quasi-causation which was markedly imperfect, very largely unreliable, could be important.

All this just parallels what could be said about genuine causation. Suppose, as we did earlier, that the computer is replaced by an actor who tries to imitate me

but occasionally fails. Here we have *imperfect causation*, mere tendency to imitate. Pushing a particular button would cause the actor to push his similarly labelled button often but not always. Plainly, I could have good grounds for pushing DON'T even if the actor made mistakes 8% of the time; or 35% of the time; or even (in view of the huge disproportion between a million free minutes and a mere thousand free minutes) 49% of the time.

Note that the words "Prisoners' Dilemma" are nowadays applied even to cases in which two parties are in close causal contact but each remains unsure of how the other will act. Two nations on the brink of nuclear war, each fearing that the other would make a first strike to destroy missiles in their silos, might be in a Prisoners' Dilemma. The dilemma could be there even if the heads of state were in touch by television. Thus people could be very unfortunately mistaken when they recommend two-boxing in any Newcomb predicament in which the Predictor's success is less than *virtually certain*, or doing the equivalent of pushing CONFESS in Prisoners' Dilemma cases where it is appreciably unsure that the other party will do the equivalent of pushing DON'T. Robert Nozick could be sadly in error when he writes that one should take what is in both boxes because "in Newcomb's example there is the *illusion* of influence" and not real influence. The Expected Utility Argument, as he goes on to say, would on the contrary favour one-boxing even if the predictions were very unreliable, "Yet, I presume, if the probability of the being's predicting correctly were only .6, each of us would choose to take what is in both boxes" (Nozick 1970, pp. 135-6 and 140). But it is definitely untrue of me that I would take what is in both. Agreed, Newcomb's example would no more exemplify real influence than a fake diamond would be a real diamond; but it would exemplify *real quasi-influence* and that can be good for something, just as a fake diamond can. Real quasi-influence, genuine quasi-causation, must be taken into account during decision-making even when it is markedly imperfect: only 60% reliable, for instance. True enough, quasi-causation can provide only the illusion of influence. But equally, I am arguing, we could have only the illusion of dominance when genuine quasi-causation was present.

There can on the other hand be cases of merely illusory quasi-causation, much as there can be illusory fake diamonds (mirror images, maybe, of actual fake diamonds, in a conjuring trick all done by mirrors). And of course illusory quasi-causation is something by which our decisions should not be guided. (a) One case of it could involve two agents acting in much the same way but not in reflection of any basic similarities in their natures. Perhaps the one swallows twenty aspirins because thinking it will cure him while the other swallows twenty because hoping it will kill him. (b) A more complex case would involve similarities—in genetic factors, say—which ensured or tended to ensure (perhaps at only the 60% level) that two agents behaved similarly. Here the behaviour of each genuinely would *quasi-cause* (perfectly or imperfectly) the behaviour of the other, since it would be similar factors, genetic factors intrinsic to each agent, which *caused* or tended to cause each to behave in the way in question. (Quasi-causation of each agent's

behaviour by that of the other is fully compatible with *causation* by similar factors inherent in each agent. Quasi-causal correlations between the behaviours of individuals are not *magical*.) However, some might try to see in this situation a further kind of ensuring, namely, the ensuring by the one agent's behaviour of the fact that the other agent had a particular genetic make-up. Now, that particular sort of ensuring would be illusory. Besides not being genuine causation, it would not even be genuine quasi-causation. It could not be argued, e.g. that if possessing genes such as make their possessors develop a dread disease were well correlated with behaving in a given way (smoking cigarettes, perhaps) then your choosing to behave in that way would tend to quasi-cause (or, as I have been expressing such things, would quasi-cause "imperfectly") your identical twin to develop that dread disease! Genuine quasi-causation is quasi-causation of behaviour which has a basis in similar causal factors and NOT of those similar causal factors themselves.

Would problems in this area have been settled long ago if the language of quasi-causation had been available? I doubt it. What is needed here is not a set of words but a willingness to tell oneself vivid stories about idealized situations so as to get one's intuitions running properly. (The story of the symmetric universe; the story about watching a computer-based replica of oneself; and so on.) But once having seen the point of these stories, it can I think help just a little to have the word "quasi-causation", ugly though it is. It can help the point not to be forgotten. It can help above all during dealings with *imperfect* quasi-causation—and almost all actual cases of quasi-causation fall into that category.

One of them is supplied by the Voting Paradox.

5. Voting

The Voting Paradox highlights a difficulty involved in many *co-operative actions*. Any individual vote will almost surely make no difference to a U.S. presidential election. Why vote, then? True, there is a slim chance that such an election will one day be decided by a single vote. But often one political party is so strong that the chance that a single vote will affect the outcome can reasonably be disregarded.³

Let us say that my party, the Democrats, appears sure to win. For me to vote could seem pointless or a morally disreputable waste of time. Yet it could also seem wrong to call it irrational for me to vote. For if voting were irrational then my fellow Democrats, if (as I like to think) all of them rational, would all of them choose not to vote, with results disastrous to the party.

How can the idea of quasi-causation help me here? Well, suppose my fellow Democrats were all exactly like me in all matters which influence voting or not voting, and suppose also that actual behaviour were determined fully by the com-

³ For this point, and a general discussion of the paradox, see Mackie 1985.

bination of those matters. If I now voted then I would be guaranteeing (not guaranteeing causally, but guaranteeing none the less: this is what perfect quasi-causation is all about) that all my fellow Democrats voted too, whereas if I failed to vote then not one of them would vote. My failure to vote would *cause* my party to lose only one vote yet it would *ensure* the loss of tens of millions. Now, while this is a highly idealized situation, we must not treat such situations as irrelevant to real life. We need to think hard about clear-cut cases before tackling the immensely messy cases which real life throws at us. I must tell myself again and again that IF all in my party were EXACTLY like me with respect to all factors influencing voting or not voting, and IF the universe were fully deterministic, then my bothering to vote would ensure (quasi-causally) that ALL the others voted too, whereas my not bothering would ensure that ALL the others abstained from voting. Having hammered this into my head I can launch out onto the real-life waters of *imperfect* quasi-causation, the waters of a world which may well not be deterministic and in which Democrats will often be unlike one another in respects influencing election-day behaviour.

No doubt many reasons could be found for voting in an election even if doing so in no way tended to ensure that others in one's party voted in that same election. It could, e.g., be a good thing to make your party's majority 15,390 rather than 15,389 because people tend to like voting for the party they think will win: a marginally larger majority in the present election can therefore translate into marginally more support in the next. Still, imperfect quasi-causation could supply by far the weightiest reason for voting. "If I bother to vote," I might well say to myself, "then I shall be quasi-causing ten, a hundred or a thousand people to do likewise, almost certainly." Not *causing*, since quasi-causation is not causation. And, although quasi-causation can count as *ensuring* (for remember that distant chicken's neck), perhaps not ensuring in the case of any particular person—because very imperfect quasi-causation, the ensuring not of voting but of a slight tendency to vote, may be all that one is likely to get. Yet even a very slight tendency to behave in a given fashion, if present in several million people, can make it almost certain that a hundred or more will in fact behave in that fashion.

7. Responsibility

The disconcertingly alert reader may have detected that my going to the polling booths could quasi-cause others who are like me in many respects, but who are alas Republicans, to go to those booths as well. No doubt my fellow Democrats will tend to be *more like* me than the Republicans are, so that my voting will at least be likely to quasi-cause *the proportion* of votes to swing in my party's favour; but one has to concede that quasi-causation is in a way less important than one might at first think. Who would have suspected that my struggling out of bed to vote Democrat could ensure (quasi-causally) not only that thousands of further bed-lovers struggled out and voted Democrat but also that other thousands strug-

gled out and voted Republican? Yet that is how things are, alas. In one way, of course, this means that quasi-causation is *more* important, not less. My action can ensure that more people get out of bed. But the extra importance is of a disappointing kind.

Notice, on the other hand, something else which very definitely adds to quasi-causation's significance. Quasi-causal responsibility *strongly resists being diluted by being shared*. Look once more at the idealized case in which all Democrats are exactly like me in all relevant respects. If I vote, all will vote; if I do not, none will. I thus carry quasi-causal responsibility for the entire Democratic vote or absence thereof. True, the same thing applies to all the other Democrats. Each one of them is similarly responsible for every Democratic vote or non-vote, mine included. (Quasi-causation is reciprocal, remember: in the symmetric universe I quasi-cause my double to kill a bird *and* he quasi-causes me to kill one. Were it not for such reciprocity, one would have to say that in an election in which each Democratic voter had quasi-caused many other Democrats to vote, the Democratic party must have gained infinitely many votes.) But the fact that the same thing applies to every Democrat is clearly no excuse for staying at home. I will have guaranteed disaster for the party if I do so.

Similarly, even in the messy, real-world case in which many Democrats are very unlike me, I can be quasi-causally responsible for hundreds of Democratic votes.

Questions of responsibility appear in a new light when these points are grasped. I have not managed to find any immensely strong reason for treating *being responsible quasi-causally for how hundreds have behaved* as entirely different, morally speaking, from *being responsible causally for how hundreds have behaved*. Never forget that quasi-causation can be every bit as good as causation for ensuring various outcomes.

Another way of expressing the point might be this: that the genuine distinction between real control and quasi-control cannot make *being responsible for something quasi-causally* into *being only quasi-responsible for it*. Quasi-causation gives startling force to Donne's "No man is an *Island*". This holds even outside cases of co-operative action towards a single result such as a Democratic victory. Murders are often very private affairs—but when, in the symmetric universe, I and my double each commit just the one very private murder, we each also thereby make another murder certain. Each of us ensures two deaths and not just one when he hurls his rock, fires his gun, or whatever, although he *causes* one death only. And even in his own universe-half each of us may be in a situation rather like that of many million other people—people (perhaps centuries earlier or later) facing temptations to murder—who are sufficiently like himself to give him some quasi-causal grip on them.

This entire way of thinking can appear crazy. Yet five minutes a day spent reflecting on quasi-causation—begin with murders committed in a fully symmetric universe and then move slowly on towards the case of the computer replica which replicates your button-pushing only poorly—might make you see it as

thoroughly commonsensical. And perhaps the germs of it are already present in thoughts influential in getting people into polling booths, thoughts on the lines of “What if everybody in my party stayed in bed?”.

8. *Marbles, wolves and men*

Just how widely should the concept of quasi-causation be applied? Must it be restricted to cases where highly intelligent beings are making choices? Or, in the symmetric universe, does the snake slithering up to my double quasi-cause the slitherings of the snake approaching me?

This strikes me as fairly arbitrary. Feel at liberty to talk of quasi-causation even in the case of two rather similar marbles dropped from a rather similar height onto rather similar surfaces, the rebounding of the one “quasi-causing” the other to rebound with rather similar velocity. Such talk would have little interest, though. It would bring to our attention nothing we had not grasped already. In fact—as Jackson and Pettit point out—there is already a well-established use of such words as “because” which covers such cases. Two squash balls rebound with similar speeds “because” they are alike made of rubber, are equally well warmed up, and so forth. Or—I suggest that this too has established itself as Ordinary English—the one squash ball will rebound to roughly eleven inches “because” the other has just done so.

Things may however get more interesting in the case of snakes, and still more so in that of a school of fish or a pack of wolves. To what extent do snakes *co-operate*? Fish can certainly appear to, a lot of the time, and wolves definitely do. Perhaps there are strong evolutionary pressures favouring quasi-causation, and hence also whatever genetic similarities underlie it. It could be evolutionarily advantageous for the fish to disregard minor stimuli, all then fleeing together when a stimulus passed a particular threshold, or for all the wolves to become eager to risk an attack on a bison when their individual bellies had reached a particular degree of emptiness.

A common objection is that all this is “science fiction”. There just are no things worth describing as computer-based replicas of human beings, neither are fish, or wolves, or humans, replicas of one another to any relevant extent! Examine two human brains in cross-section. Even in what is visible to the naked eye, those brains are recognisably different. But my reply is that quasi-causation of an important kind in no way depends on there being anything like identity of structure. To say that the propelling of an arrow by one bow quasi-causes another to propel a similar arrow with very similar velocity (or, if you prefer, that the one bow shoots an arrow at such and such a speed “because” the other does so) is not to declare that on slicing up those bows you would see no differences in the whorls and flowing lines of the wood. Again, humans themselves can be, in their behavioural tendencies, very good replicas of one another. Products of British public schools, or of military academies the world over; inhabitants of prisons or

of convents; members of right-wing political parties; do not many of these often say or do things quite ludicrously similar? We can even have considerable insight into why they act similarly, and not just when they are acting rationally. But this is not to suggest that slicing up their brains would at once show why.

Quasi-causation need involve no more than this: that people, animals or things behave in ways sufficiently alike, on the basis not of magic but of similarities in their natures. What constitutes *sufficiency* here? Well, considerable behavioural likenesses can actually be found between a man and a frog. Both will withdraw a limb when a pin pricks it, for instance: they are to that extent each other's replicas. Whether such likenesses are perfect enough to make talk about quasi-causation profitable in, say, Evolutionary Theory or Decision Theory or Ethics, will depend on the details of the surrounding situations. There will of course be many situations (suggested-revolution-against-tyranny situations, for example) in which it would be giving very poor and immoral advice if one encouraged people to assume that many others would do whatever they did. Yet imagine a Prisoners' Dilemma in which the payoff for confessing, granted that one's fellow prisoner did not, was that just one minute of freedom was added to the million minutes which would have been earned by not confessing. Even a very tiny amount of quasi-causal correlation would be important here.⁴ It would be science fiction to declare that no such correlation could exist.

Similar remarks apply to the case of voting in a presidential election. Very, very imperfect quasi-causation could be enough to make one vote responsible for twenty others. In some cases a swing of twenty votes would be trivial; in others, not. Only when it was not, could there be much point in applying the language of quasi-causation to it.

It would no doubt be worthwhile to discover just how great are the quasi-causal correlations between various voters, or various fish, or between a fish and a voter, but efforts to discover such things would not be parts of a radically new research programme. Quasi-causal correlations are simply *non-magical behavioural similarities* under another name, and Psychology and Comparative Psychology are recognized disciplines.

Let it please not be complained that when several fish all react to a single stimulus by fleeing, or when colonels, momentarily out of causal contact with one another, all decide to refuse their obedience to a specially disgusting command, then these are cases of one and the same stimulus causing the flight of each individual fish, and of a training in military morals (or perhaps the survival of basic human decency despite such training) causing refusals to obey, *and not* of quasi-causation! For, as was insisted earlier, there need be no conflict between something's being responsible for some behaviour *causally* and something else's being

⁴ If the value of any free minutes is proportional to their number then the requisite correlation exists as soon as the reliability of predicting the behaviour of the one prisoner from seeing that of the other exceeds 0.5000005, where 0.5 would correspond to purely random correlation, i.e. a correlation of zero. (The formula for working this out is given by Lewis.) The figure 0.5000005 corresponds to a correlation of 0.000001. Perfect correlation would therefore be one million times better than this.

responsible for it *quasi-causally*. It is because the colonels all arrived at the military academy, years ago, with much the same indestructible minimum of human decency, or because their fine training at that academy made them the fine persons that they are, that every one of them decides to reject the command. Only because of their similarities in these matters of causation could each have any confidence that he was quasi-causing the others to act as he himself did.

Again, let it not be argued that decisions, to the extent that they reflect similarities in the natures which have been given to us by inheritance, training, and so forth, somehow cannot be decisions at all, or at any rate cannot be free decisions for which we should be held responsible.

Department of Philosophy
University of Guelph
Ontario N1G 2W1
Canada

JOHN LESLIE

REFERENCES

- Jackson, F. and Pettit, P. 1988: "Functionalism and Broad Content". *Mind*, 97, pp. 381-400.
- Lewis, D. 1979: "Prisoners' Dilemma Is a Newcomb Problem". *Philosophy and Public Affairs*, 8, pp. 235-240.
- Mackie, J. L. 1985: "Norms and Dilemmas". In *Persons and Values: Volume II* (Selected Papers of J.L.Mackie), Joan and Penelope Mackie, eds., Oxford: Clarendon Press, pp. 234-241.
- Nozick, R. 1970: "Newcomb's Problem and Two Principles of Choice", in *Essays in Honor of Carl G.Hempel*, N. Rescher, ed., Dordrecht: D.Reidel, pp. 114-146.