Ambiguity and Logic

FREDERIC SCHICK

CONTENTS

vii

1

LIVING WITH AMBIGUITY

S AY we are fully informed. Say we know all we could pos-
sibly know. Still, there remains ambiguity. What we now
do is ambiguous, and what that will bring about is too, and so
is all that would have happened if we had done something else
instead. How we act in any setting depends on how we there
get around this, on how we disambiguate there. And our later
making sense of our actions calls for our knowing how we did it.

1

Let me begin with some stories that may help to bring that out.
The first will be about me, and it will do me little credit.

When this happened, I was thirty and on my first good job.
I then had two particular friends – call them Adam and Bob.
Adam was lively and good-looking. Women liked him and he
liked women. Bob too liked women, but they cared for him less,
and he ached for what Adam had. He would always ask about
Adam, hoping at least to feed fantasies, but I knew nothing he
wanted to hear, so I couldn't oblige him.

Then, one day, I did. To his "What's new with Adam?" I said
"He moved; he had to." Bob asked why. "Because it was three
o'clock in the morning and he had the music on loud, and
the landlord came up from downstairs" – I was making all this
up – ". . . and found him in bed with two women and evicted
him."

Bob was staggered and went home in a sweat. I thought
it funny and went to tell Adam, but it turned out he wasn't

in town, and he still was out of town when I saw Bob the next time. Right off, Bob asked about Adam, and I made up a new story for him about how Adam had thrown a big party to celebrate his new apartment and how the noise got the police to be called and how they responded to what the guests, both dressed and undressed, were doing. Bob could hardly breathe for excitement.

A few days later, Adam called. He began with "Are you out of your mind?", and with that the scales fell from my eyes. I had no answer to give him. I was now the one who was staggered; what I had done did now seem crazy and I could scarcely credit having done it. All of us taught at a straight-laced college, and Adam was up for promotion just then. The stories I was telling about him were enough to get him fired instead. I had known that all along. I knew I was putting his job at risk but I did it anyway, and thought it funny as I did it. I saw it as a joke, as pulling Bob's leg, as horsing around with friends.

I had seen what I did as a joke; now I saw it as a kind of betrayal. But that made no change in what I *believed* I had done, in what I *knew* about that. Neither did it change what I wanted. I wanted before to joke with my friends, and I still wanted that. I wanted now not to injure them, but I had wanted that before too. What changed was how I *saw* things – again, Adam's call refocused my mind: it gave me a new perspective. Had I seen things that way before, I would not have done what I did. Still, how could a change of perspective alone have unsettled what moved me to act? The usual theory of motivation – the usual theory of reasons for action – speaks of beliefs and desires only. How did the way I came to see things connect with the beliefs and desires I had?

Enough about Adam and Bob. Let me turn to an incident I have discussed at length before,[1] this one reported by George Orwell in an essay on the Spanish Civil War. Orwell tells of lying in wait in a field one day, hoping for a chance to shoot at some soldiers in the trenches ahead. For a long time, no one

appeared. Then some planes flew over, which took the Fascists by surprise, there was much shouting and blowing of whistles, and a man

> . . . jumped out of the trench and ran along the top of the parapet in full view. He was half-dressed and was holding up his trousers with both hands as he ran. I refrained from shooting at him. . . . I did not shoot partly because of that detail about the trousers. I had come here to shoot at "Fascists"; but a man holding up his trousers isn't a "Fascist," he is visibly a fellow-creature, similar to yourself, and you don't feel like shooting at him.[2]

Orwell wanted to "shoot at Fascists" and he believed he now could do it. On the belief/desire theory, he had a solid reason for shooting. What then was it about those pants that got him to put down his gun? Orwell answers that question: "a man holding up his trousers . . . is visibly a fellow-creature, similar to yourself." I take it the pants were down to his knees, and that Orwell is saying that someone half-naked and "visibly" human had to be seen as human. Before the man jumped out of the trench, Orwell had seen his firing at him as shooting at a fascist, which he wanted to do. The soldier's half-naked predicament was for him a wake-up call – like Adam's call to me. He then saw his firing at him as his shooting at a fellow-creature, and this he didn't "feel like" doing.

He had, of course, known all along that, under their pants, the fascists were human. He had never *faced* that fact, never fully confronted it, but how did his not having faced it weaken the force of his knowing it? And how could he want to kill a fascist and also not *feel like* doing it? How can a change in a person's seeings undercut what he wants to do?

The third story here is fictional, though it recalls an actual case and is formally like many others.[3] Jack and Jill are at a company banquet. Recent employees and the youngest ones there, they are seated in a corner of the room, where they notice, while the others are eating, that they hadn't been served.

The kitchen had run out of food. They were least likely to make a fuss and so had been picked for doing without.

Jill suggests that they leave and get burgers. Jack is firmly opposed. He says he has paid $50 for dinner and won't pay a nickel more. Jill says the $50 is spent and gone; the question is whether to leave and get a $5 burger or starve. Jack insists she has that wrong: the question is whether to pay $55 and get just a burger or starve. Fifty-five dollars is too much for dinner, never mind for a burger dinner – he prefers to starve. They argue this back and forth. In the end, they go out to eat.

How did Jill move Jack? (It was too early for hunger to have done it.) He had a belief/desire reason both for staying and for leaving. He wanted not to pay $55 for dinner and he knew he wouldn't pay that if he stayed. But he also wanted a burger and knew that leaving meant getting one. Jill didn't change his beliefs; she told him nothing he didn't know. Nor did she get him to agree with her judgment of what a burger was worth. He agreed with her all along that a burger was worth $5. What she did was to get him to stop seeing that burger as a $55 dinner – she got him to see it as a *$5* dinner. She changed his perspective on leaving to get it. But how did his new view of leaving unsettle the reason he had for staying?

2

I have presented three cases – my leg-pulling prank, Orwell's putting down his gun, Jack's leaving to get a burger. In none does the belief/desire theory fully explain what happened. On that two-factor theory, a person has a reason for taking action *a* where he wants to take an action of a certain sort *b* and he believes *a* is of sort *b*. I wanted to play a joke on Bob and believed I was doing that. But I also wanted to be a proper friend of Adam's and knew that this called for some self-restraint. I had as good a two-factor reason for holding back as for doing what I did. Why then did I not hold back? Likewise for Orwell

and for Jack. They too had belief/desire reasons both for doing what they did and for doing the opposite. Why then did they act in one way and not in the other instead?

The belief/desire theory of reasons must be refined in some way, and I think we can't just refine it by speaking of *intensities* of desire. We can't just say that Orwell wanted to avoid his killing "fellow creatures" more than he wanted to kill Fascists, that Jack wanted a burger more than he wanted not to pay $55 for dinner, etc. That would only give us more questions. Why should a man's holding up his pants have changed the relative strengths of Orwell's desires? Why should that phone call from Adam have changed the relative strengths of mine? Did our desires (their strengths) change at all? No doubt they *may* have changed, but why should one think that they did?

I will take a different line. In my report of each of these cases, I described the people in them as getting a new perspective on things. I spoke of how some event or discussion changed the way they *saw* things. What was it there that changed? I had seen my teasing Bob as a joke, but that wasn't like my seeing that his face became flushed. Orwell's seeing the fascist before him as a fellow human being – his seeing that to shoot would be to shoot a human being – wasn't like seeing he was running. The *seeing* in these cases was conceptual. We might describe it as a *conceiving* of the action (of the joshing or the shooting), and this conceptual seeing or understanding calls for being fitted-in somehow. I propose replacing the two-factor, belief/desire theory of reasons with this three-factor refinement of it, that a person has a reason for taking action *a* where he wants to take an action of a certain sort *b* (one that *leads to* or *brings about* or simply *is* a *b*) and he believes *a* is of sort *b* – *and* he sees or understands *a* as being of that sort (as leading to or bringing about . . . *b*). A reason for taking action *a* is a belief-and-desire-plus-seeing.

More should be said about *seeings* or *understandings*.[4] Let me just say here that we can, at any moment, see *a* in one way only, though that way may be compound; *a* might be seen as of sort

b-and-*b'*-and-*b"*. And that our seeing *a* in some way implies that we have the corresponding belief. If we see *a* as of sort *b*, we must *believe* it is of that sort, though not all we believe about it has to enter our view of it. Isn't this still then a two-factor theory, the factors being *seeing* and *desire*, the belief here implicit in the seeing? I will continue to speak of three factors because our beliefs, desires, and seeings are distinct mental states that we have and they all figure in our reasons. But, again, our beliefs and our seeings aren't logically independent. (If we see *a* as a *b*, we must *believe* it is a *b*, perhaps a *b*-and-*b'*.)

Our beliefs and desires alone don't give us a reason for action. Only a belief and a desire plus a related seeing do that. Before the soldier appeared, Orwell had a reason to shoot. He thought that shooting would kill a fascist, he wanted to kill a fascist, and he saw his shooting as his killing a fascist. But facing the man holding up his pants, he no longer saw it so, and his reason came apart; he still wanted to kill a fascist, but he didn't "feel like" killing this one. Jack had seen paying $5 for a burger as the last payment on a huge bill for his dinner; when he ceased to see it that way, he lost his reason for not leaving. In the case of my joke with Bob, I came to stop seeing it as a joke. I then no longer had the reason that had led me to make it.

Putting all this differently, every action might be seen in any number of ways: in that sense, each is *ambiguous*.[5] Orwell might have seen the shooting as his shooting a fascist or as shooting a "fellow creature." He did in fact see it the latter way, and seeing it that way disambiguated it for him. He disambiguated *not* then shooting as *not* hurting a fellow man. Given what he believed and wanted (he wanted to avoid hurting fellow creatures), he had a reason to put down his gun. A person has a reason for action *a* where he believes *a* is of sort *b*, he wants to take an action of that sort, and he sees, he *disambiguates a* as *b*.

It has been said that any action can have different *meanings* for people, for different people at the same time and the same person at different times. Different perspectives yield different

meanings. To see an action in this way or that – to disambiguate what we might do – is to assign a certain *meaning* to it.[6] Thus whether or not we take action *a* depends on what it *means* to us, and the point might be put like this: one of the factors of our reason for *a* is the meaning *a* has for us. I will avoid the language of *meanings*, but it may give some readers a handle on my concept of reasons.

3

What I have said about reasons for actions carries over to reasons for choices. Our choices often have reasons too, and how we see those choices figures in our reasons for making them. Why did we *choose* to take action *a*? Perhaps we chose it because we believed that it was a *rational* action, we wanted to take a rational action, and we saw it as being rational. But our seeings sometimes also enter in a second way here, for they bear on the question of which of the actions we might take would be rational.

Suppose you have some benefit you must give either to Jack or to Jill. It must go to one or the other. It can't be divided between them. Say that the benefit is a job, that you don't care which of them gets it, and that you have three options: you might give it right off to Jack, you might give it right off to Jill, or you might toss a coin. Many people would reach for the coin. That is the fair way to do it, and most people, in such cases, prefer (and take steps) to be fair.

This scenario has troubled choice theorists for a number of years. The problem they find is this. You set the same utility on Jack's getting the job as on its going to Jill; let that utility be *x*. If you toss the coin, the job is as likely to go to one as to the other. The expected utility of tossing that coin is therefore $\frac{1}{2}x + \frac{1}{2}x$, and this too is *x*. So you should be indifferent between tossing the coin and either outright hire. If you *aren't* indifferent – if you insist on tossing that coin – you have turned

against rationality. Does this mean that being rational calls for indifference between fairness and unfairness?[7]

Here is a very different question that is formally similar. (The difference is only that there are two options in this and not three.) The question comes up near the end of Tolstoy's *War and Peace*. Two commanders of guerillas are discussing what they do with the prisoners they take. Denisov sends his prisoners to the regular army camp, many days' march away. Dolohov does not send his off. He says,

> "You send off a hundred prisoners and hardly more than a couple of dozen arrive. The rest either die of starvation or get killed. So isn't it just as well to make short work of them? ..."

> "That's not the point [says Denisov]. ... I don't care to have their lives on my conscience. You say they die on the road. All right. Only it's not my doing."[8]

Let the prospects be even more grim: *all* the prisoners die on the road. Looking just at these people's lives, the outcomes are then exactly the same whether the prisoners are sent off or shot. When we add the risk to the troops sent to guard the prisoners and we go by the outcomes, we have no choice but to shoot them. That would be Dolohov's view of it. Denisov wouldn't shoot even here, which squares with Dolohov's thinking him a squeamish, soft-headed fool.

Must we agree with Dolohov's judgment? Rational choosers go by the outcomes: they are consequentialists. Yes, but that doesn't settle it, for we must look at *how* they go by them, how their thinking connects with those outcomes, how certain values they set on the outcomes get wired-up in their thinking.

On what do our desires, preferences, and utilities – collectively, our values – focus? What are the *objects* of those values? Say that I want to be king of France. There is then something here I want (something my wanting takes as its *object*). But that something can't be the situation of my now being the king, for there is no such situation. What I want is that *I*

become king, and this describes or *reports* a situation, actual or just possible. A report of a situation sometimes is called a *proposition*, and we can say that the values we have take propositions as objects.

Propositions may be thought of as (in part) like abstractions, like numbers or relations: they have no spatial or temporal locus but can, where true, be locally instanced. They are then instanced by what they report, which makes them also unlike abstractions: there are many instances of the relation *larger than*, but a (simple, noncompound) proposition can't report more than one situation. If *p* and *q* report different situations, *p* and *q* are different propositions.[9] Where we both want that *Jack gets the job* but we are thinking of different Jacks, what we want is different too. Still (again), we cannot say that what we want are the situations themselves, or that what we want is built up somehow of situations (or of certain Jacks plus . . .). We can distinguish what we want from what we don't by what is "out there" somewhere, but we can't go on to infer that what we want is "not in our heads" – unless we mean only that propositions, as abstract objects, are nowhere at all.[10]

About the outcome of an action, this holds that the outcome isn't valued in its brute natural state. It is valued *as it appears in these or those propositions*. Putting the point another way: we value an outcome always *under these or those reports of it*. Or, in terms of *descriptions*: what we believe about an outcome allows for many descriptions of it, and we sometimes set different values on the same outcomes *under different descriptions*. Which of these values then enter our thinking – which of them get wired up? I think the answer has to be this, that the values that count for us are those that we set on the outcomes described as we see those outcomes.[11]

A rational chooser goes by the outcomes, by certain values he sets on them. But these values are the ones he sets on the outcomes *as he sees them*; other people, equally rational, who see them differently, may make different choices. Dolohov saw all

the possible outcomes in terms of how many prisoners would wind up dead; that argued for their being shot. Denisov too saw the outcome of sending them off to the camp as their dying on the road. But he saw the outcome of shooting them not as their dying but as his having killed them. He resisted the latter outcome, *as he saw that outcome*; he didn't want it "on his conscience." He preferred the other, *as he saw that*. Both Dolohov and Denisov chose rationally. They both went by the outcomes, by how they saw the outcomes. The difference in what they did had to do with *how* they saw the outcomes.

So too in our hiring story. Are you being soft-headed if you choose to toss that coin, if you *prefer* to toss it? In this case also, that depends. If you see the outcomes solely in terms of which person gets the job, tossing the coin is pointless for you. For since you don't care which of them gets it, you value all the outcomes of tossing and not-tossing, as you see them, the same. But if you see the possible outcomes of tossing the coin in a fuller way – as Jack's or Jill's getting that job because of how fairness worked out in their case – and if you value fairness, you prefer the so-seen outcomes of tossing the coin to those of not tossing (as you see *those*). If you are rational, you will toss that coin; you then *have* to toss it. No soft-headedness there.

<div align="center">4</div>

What about someone thinking of how certain other people are choosing? If he ignores how the others see things, he will often misjudge them. He will take his own perspective to be shared by these other people and will judge the rationality of their choices by whether he would have made them. That will then often lead him to think these others themselves not rational.

Here is what is called Zeckhauser's problem.[12] You are being held captive by a lunatic who threatens to force you to shoot yourself with a fully loaded six-shooter unless you pay him a ransom – if you pay, he will empty one chamber before

		1	2	3	4	5	6
I	N	*d*	*d*	*d*	*d*	*d*	*d*
	R	*d*	*d*	*d*	*d*	*d*	*l*
II	N	*l*	*l*	*l*	*l*	*l*	*d*
	R	*l*	*l*	*l*	*l*	*l*	*l*

Figure 1.1

you must spin the cylinder and shoot. That is his threat in the mornings. In the evenings, he relents; he empties five of the chambers and offers also to empty the sixth if you pay a ransom. There are two cases to think about here. You are willing to pay in both. Would you pay more in one than in the other? If so, in which case would you pay more?

If you think as most people do, you would pay more in the second than the first.[13] Paying more in the second case seems to make good sense, for in the first case your paying a ransom would only lower your risk of dying while in the second it would keep you alive. Still, on a common analysis, rationality argues against this. Whatever you are willing to pay in one case, you should be willing to pay in the other. For the expected-utility value of your paying-a-ransom option (of the proposition *I pay a ransom*) exceeds that of your not-ransoming option by the same in both cases.

Look at Figure 1.1. I and II are the two cases. The columns stand for the chambers that might be either loaded or not. If you fire a loaded chamber, the outcome is death (*d* is *I die*); if

11

	red (1)	white (89)	blue (10)
A	$1,000,000	$1,000,000	$1,000,000
I			
B	$0	$1,000,000	$5,000,000

	red (1)	white (89)	blue (10)
C	$1,000,000	$0	$1,000,000
II			
D	$0	$0	$5,000,000

Figure 1.2

you fire an empty one, the outcome is life (*l* is *I live*). R is the option of ransoming, N is that of not-ransoming. In case I, the expected-utility value of N is *u(d)*; that of R is 5/6*u(d)* + 1/6*u(l)*. Subtracting the latter from the former yields 1/6*u(d)* − 1/6*u(l)*. In case II, the expected utility of N is 5/6*u(l)* + 1/6*u(d)*; that of R is *u(l)*. The difference again is 1/6*u(d)* − 1/6*u(l)*. So you should be willing to pay the same in both cases for the better option R. Since most people would pay more in the second, most people wouldn't be rational.

Before I put in a word for these people, let me describe a related problem, the so-called Allais' paradox.[14] This appears in Figure 1.2. Again there are two cases, in each of which you have two options. In case I, if you take option A, you will get a million dollars. If you take B, what you will get depends on the color of a ball that will be drawn from an urn. If that ball is red, you will get nothing; if it is white, you will get a million dollars; if it is blue, you will get five million. In case II, if you take option C, you will get a million dollars if the ball drawn is either red or blue; if it is white, you will get nothing. If you

12

take D, you will get five million dollars if the ball is blue; if it is not, you will get nothing. You know that there are 100 balls, one of them red, 89 white, and 10 blue.

Which options would you choose in these cases? It seems that most people choose A in case I and choose D in case II. But notice that, in both these cases, if a white ball is drawn, the outcomes of both options are the same. So, in both cases, which of the options has the greater expected utility depends just on the outcomes if that ball is either red or blue; and in these two situations, cases I and II are identical. This means that the expected utility of A can't be greater than that of B unless that of C is greater than that of D. Again: most people choose A and D. So if we think of rationality as the maximizing of expected utility, here too most people aren't rational.[15]

Is this judgment correct? It assumes that the people involved believe that the outcomes are those announced and that they value money (the more, the better), and this needn't be doubted. But it also assumes that they see the outcomes as these are described to them, in money terms only. That is the way those making the judgment see the outcomes here, but perhaps the A-and-D choosers see these outcomes differently.

How do the choosers of A and D account for their choosing as they do? They say that, in case II, both options are risky and that, all in all, they prefer D, but that, in I, though B is risky, A is a sure million dollars. If they chose B and the ball is red, they would be going home empty-handed when they could, whatever the color, have had a cool million if they chose A – they say that, in I, that settles it for them. This suggests that, looking at B, they don't see its zero-outcome in money terms only but also (in part) counterfactually, as that zero-money payoff minus the opportunity cost. They see it not as *I get nothing* but as *I get nothing when I was sure to get a million if I had chosen A*, and they set a lower utility on the latter proposition than on the former.

Their issues should then be reported as in Figure 1.3, which is like Figure 1.2 except for the *$0/w*-entry, *$0/w* being short for

		red (1)	white (89)	blue (10)
I	A	*$1,000,000*	*$1,000,000*	*$1,000,000*
	B	*$0/w*	*$1,000,000*	*$5,000,000*
II	C	*$1,000,000*	*$0*	*$1,000,000*
	D	*$0*	*$0*	*$5,000,000*

Figure 1.3

I get nothing when I was sure to get a million if I had chosen A. If they set a sufficiently low utility on *$0/w*, A has a greater expected utility for them than does B. If, in addition, the expected utility of D exceeds that of C, their choosing both A and D is rational. People who fault these choices, who hold that they *can't* be rational, see the outcomes as in Figure 1.2 and assume that the choosers do too: they are imposing their own way of seeing on the people they are faulting.

   Likewise in Zeckhauser's problem. Many people will pay much more to get the only loaded chamber emptied than to get their captor to empty just one of the loaded six. Can these people be rational? Yes, for they may see the dying-outcome in case II not as *I die* but as *I die when I was sure to live if I had paid the ransom*. If so, their issues should be reported as in Figure 1.4, in which *d/w* is *I die when I was sure to live*. . . .[16] Suppose, as is likely, they set a lower utility on *d/w* than on *d* simpliciter. The expected utility of R then exceeds that of N in case II by more than it does in case I, which calls for paying more in II than in I. The critics who fault the thinking of those who would indeed pay more are seeing the outcomes as in Figure 1.1 and failing

Living with Ambiguity

		1	2	3	4	5	6
I	N	*d*	*d*	*d*	*d*	*d*	*d*
	R	*d*	*d*	*d*	*d*	*d*	*l*
II	N	*l*	*l*	*l*	*l*	*l*	*d/w*
	R	*l*	*l*	*l*	*l*	*l*	*l*

Figure 1.4

to allow for their being seen as these people may in fact see them.

I have argued that the concept of reasons must find a role for how things are seen, that we must move from a two-factor theory to a three-factor theory of reasons. The cases here (and in the section just above) suggest a like refinement of the theory of rationality. A person choosing rationally chooses an option that maximizes his expected utility, the probability-weighted average of the utilities of its possible outcomes. But outcomes can be differently described (think of the *d* and *d/w* descriptions), and we set utilities on them only as described in this way or that. And we sometimes set different utilities on the same outcome differently described (again, think of *d* and *d/w*). Which of these utilities then enter our thinking? Those that focus on the outcome described *as we see that outcome*. So, yes, a person choosing rationally chooses an option that maximizes his expected utility, but that is now the probability-weighted average of the utilities of the outcomes *as he sees them*.[17]

15

Some writers don't agree that a given outcome might be differently described. They hold instead that different descriptions mark out different outcomes. John Broome speaks of *individuation*: every description (and thus every seeing) individuates a different outcome. That idea comes to this, that outcomes are different if and only if the propositions that describe them are different.[18]

For Broome, *d* and *d/w* are different outcomes, not the same outcome under different descriptions. Still, in some sense, they must be the same, the same under different descriptions. If just one of the chambers is loaded and you now fire that chamber, they would report the same *event*, the same *causal effect* of your firing; both *d* and *d/w* would report your death, and you only could have died once. (The "*/w*" clause doesn't bring in any added effect of your firing.) I hold that events are independent of how it is we describe them, that what marks them off from each other are their causal connections.[19] But no need to insist on causal individuation here. The point is that a rational person attends to the outcomes as he sees them, however they are marked out.

Pulling it all together now in terms of ambiguity, here are four points I have made. Our actions are always ambiguous, and so (*pace* Broome) are the possible outcomes of any action we might take. Whether we have a *reason* for *a* depends on how we disambiguate *a*, on how we *see* that action. Whether our choosing it would be *rational* depends on how we disambiguate its outcomes (and the outcomes of our other options). And theories of reasons and of rationality that ignore ambiguity often misjudge people.

5

Most theories today ignore it.[20] Still, some writers study some issues along the lines I have sketched. They show how, in this or that case, certain so-called deviant choices might be explained,

and held to be rational, if we took note of how the choosers interpret (how they *see*) the outcomes in that case.[21] There is also a common objection to this occasional way of thinking – to the approach that I am saying ought to be generalized.

The objection is that such thinking leaves the theory of rationality vacuous. Here is Amos Tversky on the subject: "In the absence of any constraints [on how the outcomes may be described], the consequences [outcomes] can always be interpreted so as to satisfy the axioms."[22] Here is Mark Machina: "... [I]nvoking a right [to reinterpret the outcomes] is tantamount to defending the expected utility model by rendering it irrefutable."[23] Whatever action a person chooses, he can always describe the outcomes so as to make his choice rational.

Allowing for different "interpretations" lets us endorse different choices. But the same holds for utilities and probabilities: different utility and probability settings also let us endorse different choices. Would it be rational for a Zeckhauser captive to refuse to pay any ransom whatever? Yes, if he sets the same utility on living as on dying. Would he be rational if he refused to pay a nickel in the first case (where paying would leave five chambers loaded) but would pay a million in the second (where then none would be loaded)? Yes, if he took the probability to be 1 that, if any chambers stayed loaded, he would trigger one of those. Could a person rationally choose B and C in Allais' problem? Yes, if he didn't care about money.

This doesn't render the theory vacuous; it doesn't say that whatever one chooses can be held to be rational. It says only that any choice would be rational for a chooser in some frame of mind. *Your* frame of mind very likely is different from those imagined above. You set more utility on living than on dying, your probabilities aren't paranoid, and you care about money. A rational person in *your* frame of mind must pay some ransom even in the first Zeckhauser case (either six or five chambers loaded). So if *you* won't pay anything, you are not being rational.