

Reluctance to Vaccinate: Omission Bias and Ambiguity

ILANA RITOV AND JONATHAN BARON

University of Pennsylvania, U.S.A.

ABSTRACT

Subjects are reluctant to vaccinate a (hypothetical) child when the vaccination itself can cause death, even when this is much less likely than death from the disease prevented. This effect is even greater when there is a 'risk group' for death (with its overall probability held constant), even though the test for membership in the risk group is unavailable. This effect cannot be explained in terms of a tendency to assume that the child is in the risk group. A risk group for death from the disease has no effect on reluctance to vaccinate. The reluctance is an example of omission bias (Spranca, Minsk & Baron, in press), an overgeneralization of a distinction between commissions and omissions to a case in which it is irrelevant. Likewise, it would ordinarily be prudent to find out whether a child is in a risk group before acting, but in this case it is impossible, so knowledge of the existence of the risk group is irrelevant. The risk-group effect is consistent with Frisch & Baron's (1988) interpretation of ambiguity.

KEY WORDS Omission bias Ambiguity Vaccination

The present study concerns the role of two biases in hypothetical decisions about vaccinations. One bias is the tendency to favor omissions over commissions, especially when either one might cause harm. We show that some people think that it is worse to vaccinate a child when the vaccination can cause harm than not to vaccinate, even though vaccination reduces the risk of harm overall. The other bias is the tendency to withhold action when missing information about probabilities is salient — such as whether the child is in a risk group susceptible to harm from the vaccine — even though the missing information cannot be obtained. We show that this bias is found even when the overall probability of each outcome is clearly constant across the conditions compared. We take both of these effects to be overgeneralizations of principles or heuristics that are generally useful to situations in which they are not useful.

Consider first what we shall call *omission bias*, the tendency to favor omissions (such as letting someone die) over otherwise equivalent commissions (such as killing someone actively). In most cases, we have good reasons for the distinction between omissions and commissions: omissions may result from ignorance, and commissions usually do not; commissions usually involve more malicious intentions than the corresponding omissions; and commissions usually involve more effort, itself a sign of stronger intentions. In addition, when people know that harmful omissions are socially acceptable, they look out for themselves; this self-help principle is, arguably, sometimes the most efficient way to prevent harm.

In some cases, however, these relevant differences between omissions and commissions seem to be absent. For example, choices about euthanasia usually involve similar intentions whether the euthanasia is active (e.g., from a lethal drug) or passive (e.g., orders *not* to resuscitate). In such cases, when knowledge and intentions are held constant, omissions and commissions are morally equivalent (see Spranca, Minsk & Baron, in press, for discussion). Yet many people continue to treat them differently — not everyone, to be sure, but enough people to influence policy decisions. We suggest that these people are often overgeneralizing the distinction to cases in which it does not apply.

The intuition that commissions are worse, valid as it may be in most cases, is no longer valid when knowledge and intention are known to be the same for both omission and commission or when a decision maker must choose between an omission and a commission, knowing the consequences of both (as in the studies reported here). If you have a choice of killing 5 or letting 10 people die, assuming (for present purposes) that all are drawn at random from the same population, you should kill the five. Each member of the population has twice the chance of death from your omission than from your commission, and each would therefore prefer you to act. If you choose not to act, you are hurting all by going against their preferences.

Any principle that tries to justify the omission here would have to have a strong justification, for that principle will have a price in lives. When attempts are made to formulate a principle that can justify a bias towards omissions, the very distinction between omission and commission becomes unclear, and the distinctions that can be maintained have no clear moral significance (Bennett, 1966, 1981 — for example, one distinction is that there are more ways of not doing something than doing it, yet the number of ways of doing something has no normative significance). Arguments in favor of the distinction (e.g., Kagan, 1988; Kamm, 1986; Steinbock, 1980) fall back on intuitive judgments about cases. The correctness of intuition, however, is exactly what is at issue: to appeal to intuition is to beg the question.

One might argue that intuition is relevant because it affects the regret that people feel about different outcomes. Active killing of one person might cause more regret than passive killing of two. Our answer to this is that the regret is felt by the decision maker, not those who die, so the use of this argument as a justification is a kind of selfishness. In addition, the difference in regret might not be inevitable: it might disappear with a change in the person's view of omissions and commissions, so a general change in this view might still be justified.

Spranca *et al.* (in press) found that many subjects considered commissions that caused harm to be morally worse than omissions that caused harm, even with intention held constant. For example, active deception was considered morally worse than withholding the truth, even when the actor's intention to deceive was judged to be the same in the two cases. Spranca *et al.* also asked subjects to evaluate two options from a decision maker's point of view: a treatment that would cure a disease but cause death with a .15 probability, or no treatment, with the disease itself causing death with a .20 probability. In 13% of the cases, subjects chose no treatment because (they said) they did not want to be responsible for causing deaths through their decision. (When the probabilities were reversed, subjects preferred the treatment in only 2% of the cases, and, in general, when subjects rated the desirability of both options, the relative desirability of the lower death rate was higher when that was associated with inaction than when it was associated with action.) This result was equally strong whether the decision was made from the point of view of a physician, a patient (deciding for himself), or a public-health official deciding for many patients. In the present study, we extend this result, using different examples and a different method.

This omission bias is related to other phenomena. Kahneman & Miller (1986) point out that commissions lead to greater regret than omissions when a fortuitous bad outcome occurs. Demonstrations of the status-quo bias and related biases (Knetsch, Thaler & Kahneman, 1988; Samuelson

& Zeckhauser, 1988; Viscusi, Huber & Magat, 1987) usually confound the status quo with an omission. For example, when the willingness to pay to remove a risk is less than the willingness to accept payment to bear the risk, changing the status quo requires an action (accepting or paying) in both cases. Ritov & Baron (in press) have unconfounded the status-quo effect from omission bias in these contexts by asking subjects whether they would act in order to prevent a change from the status quo or whether they would feel worse when a bad outcome resulted from failure to take such action (vs. acting to maintain the status quo). In both of these situations, we found that the omission-commission distinction is the critical one, not the preservation of the status quo.

We do not mean to suggest that people are always biased toward omission. Under some conditions, for example, when the decision maker is in a position of responsibility, people show the opposite bias (Ritov, Hodes & Baron, 1989). Most subjects in the studies of Spranca *et al.*, and in the studies reported here, show no bias. A substantial minority, however, can influence public policy (e.g., on active vs. passive euthanasia) or can affect overall rates of cooperation, as in a vaccination program.

The bias toward omissions does not seem to have a single explanation (Spranca *et al.*, 1990). Many subjects justify the distinction by arguing that omissions are not causes (despite the fact that they affect the probability of outcomes relative to the alternative option). Some of these subjects do not hold themselves responsible for outcomes that would have occurred if they were absent or ignorant, despite the fact that they were not absent and not ignorant. The use of omissions as a reference point also seems to play a small role, so that harms caused by omissions are seen as foregone gains, which are less aversive than pure losses caused by commission (as is consistent with norm theory, Kahneman & Miller, 1986, and the loss aversion assumption of prospect theory, Kahneman & Tversky, 1984).

Consider next the effect of salient missing information. Frisch and Baron (1988) have argued that the effects of 'ambiguity' on decision making can be described in terms of the salience of missing information. For example, in a situation first described by Ellsberg (1961), people told that they will win a prize if a red ball is drawn will prefer to draw from an urn with 50 red balls and 50 blue balls, rather than an urn with an unknown proportion of red and blue balls. Here, the proportion of red balls is a salient piece of missing information in the second case. Subjects do not think about other missing information that would be just as useful, for example, information about the proportion of red balls in the region of the first urn from which the ball will be drawn. The perception of missing information can incline people toward inaction because they feel a desire to seek the information before doing anything else. When the information is not available, however, this desire must be left unsatisfied.

Frisch and Baron argue that the tendency to withhold action when information is missing can account for other effects of ambiguity, such as the effects of conflict among experts who estimate probabilities (Kunreuther & Hogarth, 1989). Brun & Teigen (1990) recently provided some evidence consistent with this view: subjects prefer guessing the outcome of an uncertain event before it has occurred to guessing after it occurred but before they know it.

Previous studies of the effect of ambiguity on decision making have often failed to inform subjects explicitly that the probability of the outcome was unaffected by the ambiguity manipulation. Frisch (1988) has found that subjects in experiments such as Ellsberg's often do not know that the expected probability is constant across the conditions being compared. In the present experiments, we test the effects of missing information directly by holding constant the probability of the outcome, a vaccine-related injury. We simply call attention to one factor that can influence the probability of such an injury, membership in a 'risk group' for the injury. 'Ambiguity' is therefore manipulated even though the probability of the outcome in question remains exactly known. We therefore test the role of ambiguity itself, unconfounded by subjects' beliefs about the effects of ambiguity on probability.

Both omission-commission and missing information are involved in public policy. A classic case in which the bias toward omissions affected policy was the argument that not seeding a hurricane could be justified, even though seeding would lead to less damage, because the damage would be felt by different people, to whom the decision makers would be 'responsible' (Howard, Matheson & North, 1972). Our legal system, as well, honors the distinction even when it seems irrelevant: we hold manufacturers strictly liable for damages that result from their decisions to make certain products, but we do not hold them liable at all for decisions not to produce the products (e.g., new vaccines). Similarly, we seem to put more effort into reducing risk when the risk is not well known but small (e.g., products of genetic engineering) than when the risk is well known and large (e.g., radon).

In all the experiments reported here, the basic task was the following: Subjects were presented with a hypothetical situation in which they had to make one of two decisions, either whether to vaccinate their child against an epidemic disease, or whether to support a law, requiring that all children be vaccinated. Naturally, the vaccine itself carries some risk.

The vaccination problems we present are modeled on the real case of DPT vaccine (diphtheria, pertussis, tetanus), which may cause permanent neurological injury in 1 dose out of about 300,000, far less than the damage formerly caused by pertussis (whooping cough) alone in infants. In 1987, the only manufacturer of DPT vaccine in the United States (Lederle) set aside 70% of the price of the vaccine as a reserve against tort claims (Inglehart, 1987). Likewise, the Sabin vaccine occasionally causes polio, although it is on the whole safer than the Salk vaccine, which sometimes fails to prevent polio. The producer of the Sabin vaccine has been held liable for such cases (Inglehart, 1987), although no suits have been brought against the producer of the Salk vaccine. More generally, manufacturers are liable for harmful effects of their actions but not for harmful effects of their inactions (Huber, 1988). Ambiguity about possible side effects further reduces the willingness of companies and their insurers to move forward with new products (Kunreuther & Hogarth, 1989).

EXPERIMENT 1

METHOD

Subjects were 53 undergraduates recruited with a sign placed on a prominent campus walkway and paid \$5 per hour.

Subjects were presented with a situation in which a disease kills 10 out of 10,000 children. A vaccine, which costs \$2 per child, can prevent the disease in everyone, but the vaccine itself has side effects that kill some children. The children that die from the side effects are not necessarily the same ones who would die from the disease. Subjects were given a table of different possible values of the risk of death from side effects, ranging from 0 to 9 out of 10,000, the 'net decrease in probability of death' provided by each level of risk, and the 'cost per life saved.' The 'net decrease' and cost per life, respectively, ranged from 10/10,000 and \$2,000, when the death rate from side effects was 0, to 1/10,000 and \$20,000, when the death rate from side effects was 9. Subjects were asked the maximum level of risk that should be tolerated by the government in order to institute a compulsory vaccination program.

In the 'same children' condition, subjects were told,

Suppose it were discovered that the children who are susceptible to death from flu are the same ones who are susceptible to death from the side effects of the vaccine. Thus, . . . the 'net decrease' would represent actual lives saved, children who would have died from flu if they had not been given the vaccine. There would be no children who would die from the vaccine who would not have died anyway.

Two other items were identical except that the cost was \$10 instead of \$2 per child, with corresponding increases in the cost per life. Additional items addressed other issues that are not relevant to this paper.

RESULTS

In the basic condition, in which the children who died from the vaccine were not necessarily the same as those who would die from the disease, 57% of the answers ranged between 1 and 8 per 10,000; 23% thought that no risk should be tolerated; and 9% gave answers of 9 (or 10, the maximum possible risk). (The remaining 6 gave uninterpretable answers or failed to answer this item.)

In the 'same children' condition, 47% of the subjects (in contrast to 9% in the basic condition) said that the vaccine should be given at the maximum risk (9 per 10,000). In all, 68% tolerated higher risk in this same-children conditions than in the condition with the same cost, and only 4% tolerated lower risk in the same-children condition than the control ($p < .001$). Several subjects pointed out that the difference between these conditions was in whether the vaccine killed children who would not die from the disease in any event. Subjects who tolerated very little or no risk in the basic condition, or who said (in answer to another question) that they thought that the government had no right to compel anyone to have the vaccine, often commented that giving the vaccine on a large scale would involve causing the deaths of some children, which was wrong even if it meant that a greater number of children would be saved. For example, 'You can't force parents to give their kid a drug or vaccine that could cause the kid to die!' Subjects apparently are not inclined to consider deaths caused by failure to vaccinate (that would not occur in any case) as results of a decision, although they do consider deaths caused by vaccination as results of the decision, if they would not occur anyway.

Price had no effect on the results despite a fivefold change in cost per life. Four out of 53 subjects were less willing to vaccinate when the price was high, three were more willing, and the rest were equally willing. The bias toward omission therefore cannot be explained in terms of vaccination being more costly.

EXPERIMENT 2

The present experiment adds ambiguous situations in which the final outcome of the vaccination is dependent upon an unknown intermediate state. We expect that the subjective feeling of missing information will be more salient in such a situation, leading, in turn, to a stronger omission bias. As before, subjects were presented with a hypothetical situation, in which they had to decide whether to vaccinate their child against an epidemic flu. Several conditions were described, differing in the presence (or absence) of 'risk groups' in the population, with regard either to the flu or to the vaccine. In all cases, however, the information whether the child belongs to any of the risk groups was not available to the decision maker. We predicted that the presence of risk groups would make subjects less inclined to vaccinate, in spite of our emphasizing the fact that the overall probability of death is identical in all conditions.

We also compared the effect of ambiguity (missing information) with regard to risk of death from the vaccine to ambiguity with regard to death from the flu. To that end, we concluded two conditions: one condition with information only about a risk group for the vaccine and another condition with information only about a risk group for the flu.

Finally, we compared decisions for a hypothetical child of one's own with decision concerning a hypothetical law requiring vaccination for all children.

METHOD

Twenty-eight students were solicited as in Experiment 1.

Subjects went through all cases twice: once to make a personal decision whether they would vaccinate their child, and once to indicate their support of a law requiring vaccination. Half the subjects did the policy decision before the personal one.

The instructions to the questionnaire read:

In the state you live in, there had been several epidemics of a certain kind of flu, which can be fatal to children under 3. The probability of each child getting the flu is 1 in 10, but only 1 in 100 children who get the flu will die from it. This means that 10 out of 10,000 children will die.

A vaccine for this kind of flu has been developed and tested. The vaccine eliminates the probability of getting the flu. The vaccine, however, might cause side effects that are also sometimes fatal.

In the personal decision, subjects were instructed: 'Imagine that you are married and have one child, a one-year old. You wonder whether you should vaccinate your child. Your child will have a 10 in 10,000 chance of dying from the flu without the vaccination.' In the policy decision, they were instructed: 'Suppose that the state government is thinking of passing a law to require vaccination for all children, unless a physician thinks it is dangerous to the child's health. (Such laws exist in Pennsylvania and other states.) The question now is whether you would support such a law. If the law is not passed, the vaccine will not be offered, and 10 out of 10,000 children will die from the flu.'

For each case, subjects were asked to indicate the maximum overall death rate for vaccinated children for which they would be willing to vaccinate their child (or to support the law), using the following scale:

Would you vaccinate your child [support a law requiring vaccination] if the overall death rate for vaccinated children were (check those in which you would vaccinate [support the law]):

- 0 in 10,000
- 1 in 10,000
- 2 in 10,000
- 3 in 10,000
- 4 in 10,000
- 5 in 10,000
- 6 in 10,000
- 7 in 10,000
- 8 in 10,000
- 9 in 10,000
- 10 in 10,000

Eight different versions of the vaccination problem were used (with titles included):

1. *Basic case.* 'The children who die from the side effects of the vaccination are not necessarily the same ones who would die from the flu.'
2. *Risk group for flu.* 'Suppose it were discovered that 100 out of every 10,000 children were susceptible to death from the flu. Children who were not susceptible do not experience any adverse effects. The test to determine who is susceptible is not generally available and cannot be given. The overall probability of death from the flu is still 10 out of 10,000 (with all 10 included in the 100 who are susceptible).'
3. *Risk group for vaccine.* 'Suppose it were discovered that 100 out of every 10,000 children were susceptible to death from the side effects of the vaccine (if any such deaths occur). Children who were not susceptible do not experience any adverse effects. The test to determine who is susceptible is not generally available and cannot be given. The overall probability of death from the flu is still 10 out of 10,000.'
4. *Same risk group for flu and vaccine.* 'Suppose it were discovered that 100 out of every 10,000 children were susceptible to death from the flu, and the same children were susceptible to death from side effects (if any such deaths occur). Children who were not susceptible do not experience any adverse effects. The test to determine who is susceptible is not generally available and cannot be given. This, of course, does not mean that a child who died from side effects of the vaccination would have died anyway, since he may not have contracted the flu at all. The overall probability of death from the flu is still 10 out of 10,000 (with all 10 included in the 100 who are susceptible).'
5. *Different risk groups for flu and vaccine.* 'Suppose it were discovered that 100 out of every 10,000 children were susceptible to death from the flu, and 100 out of every 10,000 children were susceptible to death from side effects, but they were not necessarily the same children. Children who were not susceptible to death from the flu never die from the flu, and children who are not susceptible to death from side effects never die from side effects. The tests to determine susceptibility are not generally available and cannot be given. The overall probability of death from the flu is still 10 out of 10,000 (with all 10 included in the 100 who are susceptible to death from the flu).'
6. *Chemical risk group for vaccine.* 'Suppose it were discovered that death from the vaccine is caused by the interaction of the vaccine with a certain chemical normally produced by the body. The interaction can occur when the level of this chemical goes above a certain point. 100 out of every 10,000 children have a chemical above this point. These children are at risk of death from the vaccine. No other children are at risk. The test for the chemical is not generally available and cannot be given. The overall probability of death from the flu is still 10 out of 10,000.'
7. *Vaccine might cause flu.* 'Suppose that the vaccine consists of a certain dose of weakened bacteria, which would encourage the body to produce specific antibodies to combat the flu. The dosage given in the vaccine is the minimum that would activate the production mechanism of the antibodies. What is considered death from 'side effects' in item A is actually death from the flu, caused by the vaccine. The overall probability of death from the flu is still 10 out of 10,000.'
8. *Vaccine failure.* 'Suppose that an alternative vaccine were developed. In this case, the vaccine causes no deaths, but it could be less effective. Of children who are given this vaccine, some could die from the flu, because the vaccine will fail. The death rate from the flu is still 10 out of 10,000, as in case A. In this case, of course, there are no side effects, but vaccinated children can die from the flu if the vaccine fails.'

Table 1. Mean maximal death rate (out of 10,000) for vaccinated children (N = 20)

Case	Personal decision	Support of law
1 Basic case	5.458	5.750
2 Risk group for flu	5.458	6.250
3 Risk group for vaccine	4.350	4.625
4 Same risk group	4.000	3.917
5 Different risk groups	3.667	3.500
6 Chemical risk group for vaccine	3.208	3.125
7 Vaccine might cause flu	5.167	5.542
8 Vaccine failure	7.125	6.833

RESULTS

Four subjects out of 28 were excluded from the analysis because their answers indicated that they have not understood the task (e.g., they were more willing to vaccinate when the death rate from the vaccine was high than when it was low). Eight subjects gave the same response for all cases. Table 1 shows the means across subjects of the maximal death rate for which subjects still decide to vaccinate and of the maximal death rate for which subjects still support the law requiring vaccination.

The results for policy decisions were the same as those for personal decisions. We found no significant difference in overall level of vaccination for policy vs. personal decisions ($t = .431$, $p = .67$). The ordering of the cases in the personal decisions did not differ from their ordering in the policy decisions either: Friedman analysis of variance of the differences in ranking of cases for the two types of decisions did not yield a significant result (Friedman test statistic = 6.698, $p = .461$). We therefore averaged the results for the two types of decisions.

To test the hypothesis that presence of risk group for vaccinated children affects the decision, we compared the cases with a risk group for the vaccine (case 3 and case 5) to the otherwise equivalent cases without a risk group for vaccine (cases 1 and 2). Indeed, the presence of a risk group for the vaccine significantly reduced subjects' willingness to vaccinate: 12 subjects were less willing to vaccinate in case 3 than in case 1, and only 2 subjects went the other way ($p < .02$); 15 subjects were less willing to vaccinate in case 5 than in case 2, and none went the other way ($p < .001$).

Our description of the risk group might have aroused suspicion. Subjects might have thought that sufficient effort would yield the missing information. To test this possibility, we ran an additional experiment in which we made clear why the information was unavailable. Twenty-three subjects were given the basic case plus a new vaccine risk-group version, identical to version 3, except that subjects were told: 'A test to determine who is susceptible is available now, but it must be done when the child is born, and it was not available when your child was born.' Results were essentially unchanged. Nine subjects were less willing to vaccinate in the risk-group condition than in the basic case versus one who went the other way ($p < .02$), six gave the maximum rating (9 or 10) to both versions, and seven gave equal ratings, less than 9, to the two versions.

Returning to the main experiment, the decisions in cases with risk group for the flu (case 2 and case 5) did not significantly differ from the decisions in the matched cases without it (case 1 and case 3). Four subjects were more willing to vaccinate in case 2 than in case 1, while 4 other subjects were less willing to vaccinate in this case ($p = 1$); 2 subjects were more willing to vaccinate in case 5 than in case 3. 9 subjects went the other way ($p = .065$, but here the risk group increases rather than decreases willingness to vaccinate). In sum, a risk group for the vaccine makes subjects reluctant to vaccinate, but a risk group for the flu has no significant effect.

The difference between case 4 (same risk group for flu and vaccine) and case 5 (different risk groups) is not significant: 7 subjects were less inclined to vaccinate in case 5 than in case 4, 2 subjects went the other way ($p = .180$). This is consistent with the general lack of effect of the risk group for flu.

Making the missing information more salient caused subjects to be still less willing to vaccinate. Comparing case 3 (risk group for the vaccine) with case 6 (chemical risk group for the vaccine), we find that the additional information given in case 6 caused 10 subjects to give a lower answer in this case than in case 3. None of the subjects went the other way ($p < .01$). Information concerning the mechanism by which the vaccine causes death does not seem to matter (7 subjects gave a higher response in case 1 than in case 7, 3 subjects went the other way, $p = .344$).

Finally, subjects are willing to accept a higher death rate of vaccinated children when the death is a result of vaccine failure rather than vaccine side effects (10 subjects were more willing to vaccinate in case 8 than in case 1, only one subject went the other way, $p < .02$). When death results from vaccine failure, the decision to vaccinate does not cause death: this condition is analogous to the 'same children' in Experiment 1. This finding is consistent with the finding of Spranca *et al.* (1989) that omission bias is determined by belief that the commission itself causes bad outcomes that would not have occurred if the decision maker were unaware of the possibility of making a decision.

EXPERIMENT 3

The lack of flu risk group effect in the previous experiment suggests that missing information is weighed differently when it concerns the effects of omission or the effects of commission. However, it could also result from the correspondence between the missing information and the response mode: subjects responded in terms of number of deaths from the vaccine, so they may have been more sensitive to information concerning death from the vaccine. To test this alternative hypothesis, we designed the following experiment.

As before, subjects were presented with a hypothetical situation, in which they had to decide whether to vaccinate their child against an epidemic flu. They were first asked how willing they were to vaccinate their child in the basic case, in which no risk group was mentioned. Then they were asked whether they would be more (or less) inclined to vaccinate their child if they knew of a risk group for the vaccine, and whether they would be more inclined to vaccinate if they knew of a risk group for the flu.

Two basic cases were used. In one case the death rate from the vaccine was given, and the subjects were asked to give the minimal death rate from the flu that will make them decide to vaccinate their child. In the second case the death rate from the flu was given, and subjects were asked for the maximal death rate from the vaccine that would still make them decide to vaccinate. For each of these basic cases two versions of risk group questions, for the vaccine and for the flu, followed.

METHOD

Thirty students were solicited as in previous experiments.

The instructions to the questionnaire repeated almost exactly the description of the hypothetical situation given in Experiment 2, except that the probability of death from the flu was not given. Then subjects were told, 'In each of the following cases you are given some information concerning the vaccine or the flu.' The cases were as follows:

Case 1: Basic case, flu response mode. Subjects were given the information that 10 out of 100,000

Table 2. Number of subjects affected by presence of risk groups in Experiment 3 (N = 30)

Risk group	Response mode in basic case	Less likely to vaccinate	More likely to vaccinate
Vaccine	Death rate from flu (Case 2)	19	5
	Death rate from vaccine (Case 6)	21	5
Flu	Death rate from flu (Case 3)	12	13
	Death rate from vaccine (Case 5)	13	13

children will die from the vaccine, and they were asked to complete the following sentence: 'I will vaccinate my child if more than _____ out of 100,000 children will die from the flu.'

Case 2: Vaccine risk group, flu response mode. Subjects were asked to assume that a risk group for the vaccine has been discovered, although the tests to determine susceptibility are not generally available and cannot be given. They were asked whether, in this case, they would be more or less likely to vaccinate their child than in Case 1.

Case 3: Flu risk group, flu response mode. This case was parallel to Case 2 with a risk group for the flu instead of the vaccine.

Case 4: Basic case, vaccine response mode. This was like Case 1, except that the available information was death rate from the flu (40 out of 100,000 children will die from the flu). Subjects were asked to complete the following sentence: 'I will vaccinate my child if no more than _____ out of 100,000 children will die from the vaccine.'

Case 5: Flu risk group, vaccine response mode. This was like Case 3 except that the subjects were asked to refer their decision to Case 4.

Case 6: Vaccine risk group, vaccine response mode. This was like Case 2 except that subjects were asked to refer their decision to Case 4.

RESULTS

Table 2 shows, for each relevant case, the number of subjects who would be affected by the knowledge of a risk group and the direction of change in their decision (relative to the corresponding basic case). Note first that knowledge of the risk group had nearly identical effects in the two response modes. This was true for both the risk group for vaccine and the risk group for flu.

To examine the effect of risk group, we averaged the numerical responses across the two cases with the risk group and the two cases without it. For the vaccine risk group (Cases 2 and 6), we find that 18 subjects were less likely to vaccinate when they knew of a risk group for the vaccine (relative to the basic case, without mention of a risk group), and only 3 subjects were more likely to vaccinate ($p < .001$). The presence of risk group for the flu did not have a systematic effect on subjects' decisions: 8 subjects were less inclined to vaccinate when they knew of a risk group for the flu (Cases 3 and 5), and 10 subjects were more inclined to vaccinate ($p = .815$).

We have therefore replicated the result of Experiment 2 and eliminated an alternative explanation of this result. Subjects are less inclined to vaccinate their child when they know of a risk group for vaccinated children but are not similarly affected by existence of a risk group for children who get the flu. This finding does not seem to result from having the information concerning death rates from the vaccine made more salient by the response mode.

EXPERIMENT 4

It is evident from the previous experiments that presence of a risk group for vaccinated children decreases the willingness to vaccinate when there is no way of knowing who belongs to the risk group, even when the overall risk from the vaccine is kept constant. One possible reason for this effect may be that when faced with ambiguity, people tend to assume the worst. That is, if they know of the existence of a risk group, and have no way of finding out whether they belong to this group, they will be inclined to assume that they do. If this is true then we should expect people to be less inclined to vaccinate as the death rate within the risk group increases.

Experiment 4 tested this hypothesis by varying the size of the risk group relative to the population and the vaccine caused death rates of children who are in the risk group. Specifically, the risk group was either small (1000 out of 100,000) or large (10,000 out of 100,000), and the death rate for children in the risk group was either high (10%) or low (1%). Accordingly, we had three levels of overall death rate in the population: high (1%), medium (.1%), or low (.01%). Obviously, the probability of death from the vaccine in the population as a whole is equal to the product of the relative size of the risk group and the death rate within this group.

As in the previous experiments subjects were told that it is not possible to determine in advance who belongs to the risk group. Subjects were asked to indicate the minimal number of deaths from the flu in the population that will make them decide to vaccinate their child. As the death rate from the vaccine was not kept constant across the different conditions, we introduced three control conditions with the corresponding overall vaccine-caused death rates, but without risk groups. Thus, the risk group effect can be determined at each level of overall death rate.

METHOD

Forty subjects were solicited as in previous experiments.

Subjects went through all the cases. To control for order effects, two different orders of the cases were used, with about half of the subjects assigned to each of them. However, in both orders the control conditions preceded the risk group conditions.

The instructions to the questionnaire were identical to the instructions in Experiment 2, except for the description of the task, which read: 'In each of the following cases you are given some information concerning the vaccine of the flu, and you are asked to decide what is the minimal death rate from the flu (how many children out of every 100,000 will die) that will make you decide to vaccinate.' The subjects had to fill in the blank in the following sentence: 'I will vaccinate my child if more than ____ out of 100,000 will die from the flu.'

Seven different versions of the vaccination problem were used. In the three control cases, 10 (in Case 1), 100 (Case 2), or 1000 (Case 3) 'out of 100,000 children will die from the vaccine.' The risk group cases were:

Case 4. 1000 children are at risk, 10 of them will die from the vaccine.

Case 5. 1000 children are at risk, 100 of them will die.

Case 6. 10,000 children are at risk, 100 of them will die.

Case 7. 10,000 children are at risk, 1000 of them will die.

The exact wording of the risk group cases was (we give case 4 as an example): 'Suppose it were discovered that 1000 out of 100,000 children are susceptible to death from the vaccine. This means that outside of these 1000 children no one is in danger of death from the vaccine. Of the children

Table 3. Mean minimal death rates from flu for giving the vaccine (N = 40)

Case	Risk group size (out of 100,000)	Vaccine deaths (out of 100,000)	Mean minimal flu death rate (out of 100,000)
1	—	10	2,584
2	—	100	4,316
3	—	1,000	10,693
4	1,000	10	8,139
5	1,000	100	7,739
6	10,000	100	7,287
7	10,000	1,000	14,435

Table 4. Mean log of the ratio: response to risk group case over response to corresponding control case (N = 40)

Case	Risk group size (out of 100,000)	Vaccine deaths (out of 100,000)	Control case	Log of ratio
4	1,000	10	1	1.465
5	1,000	100	2	0.848
6	10,000	100	2	1.073
7	10,000	1,000	3	0.702

in the risk group, 10 out of 1000 will die from the vaccine. The tests to determine susceptibility are not generally available and cannot be given.'

RESULTS

Table 3 shows the mean response for each version. Examining the control conditions first, we find an immense omission bias: the minimal number of deaths from the flu that would cause subjects to vaccinate their child is at least ten times the number of deaths from the vaccine. This result extends the results of previous experiments through the use of a free-response format. A comparison between the three control conditions shows that as the number of deaths caused by vaccination increases, the ratio between the number of deaths from the flu and the number of deaths from the vaccine decreases.

For the risk group versions we computed the log of the ratio between the response to each version and the response to the corresponding control version. (Positive logs indicate less willingness to vaccinate with the risk group.) The means of those logs, across all subjects are reported in Table 4. Each of the cells in Table 4 is significantly larger than zero. Thus, we replicated here the 'vaccine risk group' effect found in earlier experiments: the presence of a risk group for the vaccine results in an augmentation of the omission bias. Other things being equal, subjects are even less inclined to vaccinate when they know the risk is 'unevenly' spread across the population.

A multivariate test of the differences between the cells yielded a significant result ($F(4,36) = 4.61$, $p < .01$). However, the only significant pairwise comparison is the comparison between the two extreme cases, Case 4 and Case 7 ($F(1,39) = 11.34$, $p < .01$): Willingness to vaccinate was affected more by

a small risk group with low risk than by a large risk group with high risk. It is likely that this distinction is due to the difference in overall death rates than the independent effect of risk group size or death rate within this group. Indeed, a comparison of the large risk group with low risk and the small risk group with high risk cases did not show a significant difference ($F(1,39) = .25$, $p = .61$).

These results suggest that as the overall risk gets higher, subjects are less sensitive to the presence of risk group. The results do not support the 'worst case' hypothesis, which holds that subjects who do not know whether they belong to the risk group tend to assume that they do.

GENERAL DISCUSSION

Subjects are reluctant to vaccinate when the vaccine can cause bad outcomes, even if the outcomes of not vaccinating are worse. This is true regardless of whether the outcomes concern an individual child, in which case the difference is expressed in probabilities, or whether they concern a large population, in which case the outcomes differ in the number of children affected. Some subjects make an absolute rule, and will accept no risk whatsoever that they will 'cause' a death even in return for complete elimination of the risk of death from other causes (see Baron, 1986). These findings show a strong bias toward omissions of the sort found by Spranca *et al.* (1990), who discuss at greater length the determinants of this bias.

In a pilot study we asked subjects to explain their reasons for deciding not to vaccinate at the optimal level. Many subjects did not write any arguments, hence we cannot subject the list of arguments to a quantitative analysis. However, it is worth noting that many of the arguments that were given revolved around the issue of responsibility.

One subject wrote, 'I feel that if I vaccinated my kid and he died I would be more responsible for his death than if I hadn't vaccinated him and he died — sounds strange, I know. So I would not be willing to take as high a risk with the vaccine as I would with the flu.' Another subject wrote, 'I'd rather take my chance that the child will not catch the flu than to be responsible for giving my child a vaccine which could be fatal.' A third wrote, '... I did not want to risk killing the child with a vaccine that is optional. It would have been my fault if the child died from the vaccine.' These arguments illustrate the main concern of subjects regarding the vaccine: One is perceived to be more responsible for outcomes of commissions than for outcomes of omissions.

Reluctance increases when we call to subjects' attention a piece of missing information about the existence of a risk group for death from the vaccine. This finding supports the proposal of Frisch and Baron (1988) that the perception of missing information can make people reluctant to act, even when the information is unobtainable. In the present experiments, unlike previous experiments on ambiguity effects, the salience of missing information is varied independently of subjects' knowledge of the final probability. Frisch (1988) has found that subjects in experiments such as Ellsberg's often do not know that the expected probability is constant across the conditions being compared. We tell subjects this explicitly.

It is interesting to note that the missing information is, in a sense, nothing new. If subjects thought about it, they could easily imagine that death from vaccine is predicted by a great many factors. Subjects have typically heard about risk factors for most diseases. If the information is unavailable, knowing of its potential existence cannot affect action. We therefore suggest that the ambiguity effect we have found is a kind of framing effect (Frisch & Baron, 1988), a result of our calling the information to subjects' attention rather than a result of the existence of the information itself (which subjects might well imagine).

The effect of ambiguity does not appear to result from a tendency to assume that the missing

information is necessarily bad. In the vaccination problems we used, a tendency to assume the worst would mean presuming that one's child belongs to the relevant risk group. This would imply greater willingness to vaccinate when subjects know of a risk group for the flu, and a decreasing willingness to vaccinate as the risk for the children in the vaccine risk group increases. We found no support for either of those predictions.

Ambiguity (salient missing information) is considered relevant only in the case of commissions. Consistent with the view that one feels more responsible for results of commission than for results of omission, subjects seem to think of the effect of missing information on the consequences of their action (vaccinating), not the consequences of their inaction (not vaccinating). Ambiguity concerning the consequences of action increases the reluctance to act, but there is no corresponding effect of the omission option. A possible explanation of this result is that ambiguity increases the feeling of responsibility for a bad outcome that a decision maker causes. Ambiguity therefore has no effect on commissions because those subjects who are affected by feelings of responsibility do not feel responsible for the results of omissions.

Our findings were obtained in certain hypothetical situations, so their generality is unclear. They do show that the patterns of inference we have found are fairly easy to detect, and it therefore seems likely that these patterns are found elsewhere, including some real situations. Parallels with real cases, such as pertussis vaccine, reinforce our findings. We and our colleagues have begun to study real decisions about vaccination as well, and these results will be reported separately.

ACKNOWLEDGMENTS

This work was supported by grants from the National Institute of Mental Health (MH-37241) and from the National Science Foundation (SES-8509807 and SES-8809299). We thank John C. Hershey and Howard Kunreuther for comments. Address correspondence to J. Baron, Department of Psychology, University of Pennsylvania, Philadelphia, PA 19104-6196, U.S.A.

REFERENCES

- Baron, J. *Rationality and intelligence*, New York: Cambridge University Press, 1985.
- Baron, J. 'Tradeoffs among reasons for action,' *Journal for the Theory of Social Behavior*, **16** (1986), 173-195.
- Baron, J. *Thinking and deciding*, New York: Cambridge University Press, 1988.
- Bennett, J. 'Whatever the consequences,' *Analysis*, **26** (1966), 83-102 (reprinted in B. Steinbock, Ed., *Killing and letting die*, pp. 109-127. Englewood Cliffs, NJ: Prentice Hall).
- Bennett, J. 'Morality and consequences,' in S. M. McMurrin (Ed.), *The Tanner Lectures on human values* (Vol. 2, pp. 45-116). Salt Lake City: University of Utah Press, 1981.
- Brun, W. and Teigen, K. H. 'Prediction and postdiction preferences in guessing,' *Journal of Behavioral Decision Making*, **3** (1990), 17-28.
- Ellsberg, D. 'Risk, ambiguity, and the Savage axioms,' *Quarterly Journal of Economics*, **75** (1961), 643-669.
- Frisch, D. *The effect of ambiguity on judgment and choice*. Unpublished doctoral dissertation, Department of Psychology, University of Pennsylvania, 1988.
- Frisch, D. and Baron, J. 'Ambiguity and rationality,' *Journal of Behavioral Decision Making*, **1** (1988), 149-157.
- Howard, R. A., Matheson, J. E. and North, D. W. 'The decision to seed hurricanes,' *Science*, **176** (1972), 1191-1202.
- Huber, P. W. *Liability: The legal revolution and its consequences*, New York: Basic Books, 1988.
- Inglehart, J. K. 'Compensating children with vaccine-related injuries,' *New England Journal of Medicine*, **316** (1987), 1283-1288.
- Kagan, S. 'The additive fallacy,' *Ethics*, **99** (1988), 5-31.
- Kahneman, D. and Miller, D. T. 'Norm theory: Comparing reality to its alternatives,' *Psychological Review*, **93** (1986), 136-153.

- Kahneman, D. and Tversky, A. 'Choices, values, and frames,' *American Psychologist*, **39** (1984), 341–350.
- Kamm, F. M. 'Harming, not aiding, and positive rights,' *Philosophy and Public Affairs*, **15** (1986), 3–32.
- Knetsch, J. L., Thaler, R. H. and Kahneman, D. *Experimental tests of the endowment effect and the Coase theorem*. Manuscript, Department of Economics, Simon Fraser University, 1988.
- Kunreuther, H. and Hogarth, R. M. 'Risk, ambiguity, and insurance,' *Journal of Risk and Uncertainty*, **2** (1989), 5–35.
- Ritov, I. and Baron, J. 'Status quo and omission bias'. *Journal of Risk and Uncertainty* (in press).
- Ritov, I., Hodes, J. and Baron, J. *Biases in decisions about compensation for misfortune*. Manuscript, Department of Psychology, University of Pennsylvania, 1989.
- Samuelson, W. and Zeckhauser, R. 'Status quo bias in decision making,' *Journal of Risk and Uncertainty*, **1** (1988), 7–59.
- Spranca, M., Minsk, E. and Baron, J. 'Omission and commission in judgment and choice,' *Journal of Experimental Social Psychology*, in press.
- Steinbock, B. (Ed.), *Killing and letting die*, Englewood Cliffs, NJ: Prentice Hall, 1980.
- Viscusi, W. K., Magat, W. A. and Huber, J. 'An investigation of the rationality of consumer valuation of multiple health risks,' *Rand Journal of Economics*, **18** (1987), 465–479.

Author biographies:

Ilana Ritov received her Ph.D. in psychology from the Hebrew University, Jerusalem, in 1988. She is now Lecturer in the Department of Industrial and Management Engineering, Ben-Gurion University, Beer-Sheva, Israel.

Jonathan Baron received his Ph.D. in psychology from the University of Michigan in 1970. He is Professor of Psychology, University of Pennsylvania.