Memory & Cognition
2009, 37 (5), 624-631
doi:10.3758/MC.37.5.624

Changing your mind

Clare R. Walsh
University of Plymouth, Plymouth, England

and

P. N. Johnson-Laird
Princeton University, Princeton, New Jersey

When individuals detect an inconsistency in a set of propositions, they tend to change their minds about at least one proposition to resolve the inconsistency. The orthodox view from William James (1907) onward has been that a rational change should be *minimal*. We propose an alternative hypothesis according to which individuals seek to resolve inconsistencies by explaining their origins. We report four experiments corroborating the explanatory hypothesis. Experiment 1 showed that participants' explanations revised general conditional claims rather than specific categorical propositions. Experiment 2 showed that, when explanations did revise the categorical proposition, participants also tended to deny the consequences of a second generalization. Experiment 3 showed that this tendency persists when participants previously affirmed these consequences explicitly. Experiment 4 showed that, when participants could easily explain an inconsistency by revising a generalization, they were more likely to accept the consequences of a second generalization. All four results contravene minimalism but support the explanatory hypothesis.

If you encounter a fact that conflicts with your beliefs, you revise your beliefs. But what is a rational way in which to make such a revision? Logical thinking can detect inconsistencies, but it cannot tell you which propositions you should abandon in order to attain consistency. One view of rational change, however, is that it should be minimal. As William James (1907) wrote, "[The new fact] preserves the older stock of truths with a minimum of modification, stretching them just enough to make them admit the novelty" (p. 59). Such parsimony is sensible, and many cognitive scientists have advocated *minimalism*, as we refer to it, both for science and for everyday life (e.g., Gärdenfors, 1988; Harman, 1986). Likewise, computer programs for artificial intelligence have modeled such changes (de Kleer, 1986; Doyle, 1979), and measures have been developed to calculate the amount of change so that minimalism can be tested. These measures depend on counting all the beliefs that change their values (Elio & Pelletier, 1997; Harman, 1986; Hiddleston, 2005). For example, Elio and Pelletier wrote, "often this relies on counting the number of propositions whose truth values would change in one kind of revision versus another" (p. 426), and Harman's measure of change proposes that we "take the sum of the number of (explicit) new beliefs added plus the number of (explicit) old beliefs given up" (p. 59). Some proponents of minimalism might object to these measures on the grounds that they are simplistic, but, to the best of our knowledge, no other readily testable measure exists. Hence, we have adopted this measure too. As Elio and Pelletier point out, it implies that a change to a categorical belief is more minimal than is a change to a generalization. After a generalization is given up, inferences from it cannot occur, and so the overall number of changes is greater.

There are problems for minimalism. One is that there may be several ways to change your beliefs that are all equally minimal. But a more serious problem is that, when you encounter an inconsistency in daily life, your goal is seldom merely to edit your beliefs in order to restore consistency. A more important goal is to *explain* the occurrence of the inconsistency (Elio, 1998; Johnson-Laird, Girotto, & Legrenzi, 2004; Thagard, 1989). Research in decision making supports the view that, when individuals are faced with a number of conflicting options, they generate reasons to justify their choices (Shafir, Simonson, & Tversky, 1993). Likewise, we propose that, when individuals are faced with conflicting information, they generate reasons to resolve the inconsistency. A plausible explanation is bound to cause someone to revise at least some proposition that they held prior to their discovery of the discrepancy. It may do so in a minimal way, but there is no guarantee. Indeed, a plausible explanation, both in science and in daily life, may imply a wholesale rejection of propositions that individuals hitherto held. Accordingly, the aim of the present article is to compare these two theoretical accounts—minimalism and the explanatory hypothesis—and to report some experiments designed to test their predictions for cases in which they diverge.

C. R. Walsh, clare.walsh@plymouth.ac.uk

A crucial consequence of the explanatory hypothesis is that a plausible explanation may not be minimal. We illustrate this point by using examples of three sorts of inconsistency. The first sort occurs in the following problem based on a conditional generalization, a categorical proposition, and a fact that is inconsistent with what they imply:

Speaker A: If a drink contains sugar, then it gives you energy.
Speaker B: This drink contains sugar.
In fact, it doesn't give you energy. Why?

Such a problem has a single inconsistency arising from a single generalization and a categorical proposition. A putative explanation such as *the drink doesn't contain sugar* rejects the categorical proposition that Speaker B asserts. And, as Elio and Pelletier (1997) have pointed out, it is a minimal resolution of the inconsistency, because it merely blocks the inconsistent consequence of the two propositions but leaves the conditional generalization intact. Yet individuals tend to revise the generalization, at least when they select from a choice of possible revisions (Dieussaert, Schaeken, De Neys, & d'Ydewalle, 2000; Elio & Pelletier, 1997; Politzer & Carles, 2001). The explanatory hypothesis predicts this tendency. To deny the categorical proposition, as in the claim that the drink does not contain sugar, seems arbitrary and unprincipled. Explanations should have general power, and almost all generalizations about events in daily life are susceptible to disabling conditions. Is it really true that a drink containing sugar necessarily provides energy? You can readily think of disabling conditions.

In some cases, disabling conditions may be retrieved from long-term memory. For example, you may know that a certain amount of sugar is required to provide an increase in energy. But you can also generate novel possibilities on the fly. For example, you may imagine that a person has taken a sleeping pill, and so sugar will not increase the person's energy. Belief revision theorists have not distinguished between these two sorts of possibility, perhaps because it is difficult to determine whether a disabling condition is novel, if only because introspections are an unreliable source of evidence (Nisbett & Wilson, 1977). Hence, we will not distinguish between existing and novel disabling conditions in our measures. Because disabling conditions provide a reason why the consequent of a generalization may fail to occur, the explanatory hypothesis predicts that generalizations are more likely to be revised than are categorical propositions about a particular individual.

One possibility is that people revise generalizations, because generalizations are less believable than categorical propositions are. Indeed, less believable generalizations often have more potential disabling conditions than highly believable ones do, and so they are more likely to be revised (Dieussaert et al., 2000; Politzer & Carles, 2001). However, the relative believability of general and categorical propositions is unlikely to explain which is chosen for revision. One reason is that generalizations are often revised even when individuals had been completely certain that they were true (Markovits & Schmeltzer, 2007). A second reason is that judgments of the believability of a conditional are context dependent. In the absence of an inconsistency, categorical propositions tend to be judged as more believable than are generalizations. For example, in one study, participants judged the categorical assertion *the experiment was conducted according to procedure* as more believable than the generalization *if the experiment was conducted according to procedure, then the helium is in liquid form*. But, when the two assertions were presented together with an inconsistent fact, *the helium is not in liquid form*, participants switched their judgment (Hasson & Johnson-Laird, 2003).

The second sort of problem illustrating the contrast between minimalism and the explanatory hypothesis consists of two generalizations, a categorical proposition, and a fact that is inconsistent with the consequences of *one* of the generalizations and the categorical proposition:

Speaker A: If you follow this diet, then you lose weight.
Speaker B: If you follow this diet, then you have a good supply of iron.
Speaker C: John followed this diet.
In fact, John did not lose weight. Why?

Such problems have single inconsistencies arising from one of the two generalizations. Minimalism predicts that you should give up the categorical proposition but that you should be less likely to do so than in the first sort of problem because there are two generalizations from which it yields a conclusion as opposed to only one generalization. In other words, the categorical proposition now has more inferential power, and so you should be more likely to retain it (Elio & Pelletier, 1997). If, for example, you deny that John followed the diet, you can no longer infer that he has a good supply of iron. If you have already drawn this inference, then, according to minimalism, you should retain it. A minimal change calls for you to retain all beliefs, even if they no longer have a satisfactory justification, provided that they are consistent with your other beliefs (Gärdenfors, 1992). In contrast, the explanatory hypothesis predicts that individuals should still resolve inconsistencies such as the one above by revising the generalization yielding the inconsistency (i.e., the first generalization in the example). Once again, this is not a minimal change.

The third sort of problem is also based on two generalizations, but the facts are inconsistent with the consequences of both of them:

Speaker A: If you follow this diet, then you lose weight.
Speaker B: If you follow this diet, then you have a good supply of iron.
Speaker C: John followed this diet.
In fact, John did not lose weight, and he did not have a good supply of iron. Why?

Such problems have double inconsistencies arising from two generalizations. In this case, minimalism still predicts that individuals' explanations should revise the categorical proposition. They should be more likely to do so than in the first problem above, because, if they do not, then they must revise both generalizations as opposed to a single generalization in the first problem. The explanatory hypothesis, however, makes a different prediction for this problem than for the previous problems. It is quite difficult to explain the failure of two generalizations, and so individuals should be more likely to revise the categorical proposition than in the case of an inconsistency with a single generalization or with only one of two generalizations.

In summary, minimalism predicts that explanations should tend to revise categorical propositions rather than generalizations, and it yields a predicted trend for explanations that rule out the categorical proposition: The highest proportion should occur for inconsistencies with both of two generalizations, an intermediate proportion should occur for inconsistencies with a single generalization, and the lowest proportion should occur for inconsistencies with one of two generalizations. The explanatory hypothesis predicts instead that explanations should tend to revise generalizations. It also predicts that this tendency should be greater for an inconsistency arising from a single generalization or from only one of two generalizations than for inconsistencies arising from both of two generalizations. We carried out four experiments designed to test the differing predictions of the two competing hypotheses.

EXPERIMENT 1

Experiment 1 examined the three sorts of problem described in the introduction (see also Table 1). The participants' task was to explain the inconsistencies, and we examined the explanations' implications for the maintenance or revision of the propositions in the problems.

Method

Design. The participants acted as their own controls and carried out four different problems of each of the three sorts of inconsistencies: with single generalizations, with one of two generalizations, and with both of two generalizations. We used 12 different sorts of subject matter from the following categories: physiology, physics, economics, and psychology (see the Appendix). The generalizations were all highly plausible and similar to those in the high-plausibility category used by Politzer and Carles (2001). The contents were rotated so that they were presented equally often with each of the 3 sorts of problem. The problems were presented to each participant in a different random order.

Participants and Procedure. We tested 20 volunteers employed by Educational Testing Service, Princeton. The key instructions were "You will be presented with information from two or three different speakers. . . . You will then be given some additional information that you know, for a fact, to be true. Your task, in essence, is to explain what is going on." The participants were tested individually and wrote their responses to each problem.

Results and Discussion

The participants had no difficulty in generating explanations to answer the question "why?" about the inconsistencies. Table 1 presents the percentages of generalizations and categorical propositions that the participants' explanations revised. The explanations were classified as implying revisions to the generalization if they were of one of three forms ("It is not the case that if A then B," "A is not sufficient for B," or "B does not always follow A") or if they described a disabling condition that would prevent B. They were classified as implying revisions to the categorical proposition if they were of the forms "not A" or "perhaps not A." This coding scheme was used to classify explanations in earlier studies (e.g., Byrne & Walsh, 2002), and it classified 91% of the responses. The remaining responses affirmed or denied the *new* information or were too vague to refute either of the propositions (e.g., "I don't know why," "One of the speakers is wrong," and "The speakers interpreted the phrases differently").

Table 1
Percentages of Explanations That Implied Revisions to the Generalizations or to the Categorical Propositions in Experiment 1

Type of Inconsistency	Revised the Generalization: If A then B	Revised the Categorical Proposition: A
With a single generalization: If A then B. A. Fact: Not B.	72	14
With one of two generalizations: If A then B. If A then C. A. Fact: Not B.	78	15
With both of two generalizations: If A then B. If A then C. A. Fact: Not B and Not C	63	30

Note—The balances of percentages are cases in which the explanations did not imply revisions to either proposition.

As Table 1 shows, the overwhelming majority of the explanations for all three sorts of problem implied revision of the generalization rather than the categorical proposition (17 out of the 20 participants showed this pattern, 2 showed the opposite pattern, and there was one tie; binomial test, $p < .01$). Likewise, the one reliable difference among the three sorts of problem was the increase in explanations revising the categorical propositions in the inconsistency with both of the two generalizations (30%) as opposed to those in which they occurred with one of the two generalizations (15%; Wilcoxon test, $z = 2.07$, $p < .02$, one-tailed). There was no reliable difference in the proportions of explanations ruling out the categorical proposition between the problems in which inconsistencies occurred with single generalizations (14%) and those in which inconsistencies occurred with one of two generalizations (15%; Wilcoxon test, $z = 0.07$, $p > .9$). These results all corroborate the explanatory hypothesis, but are contrary to minimalism.

When explanations implied a revision to a generalization, they rarely (5% of the explanations) were simple negations of the generalization. Instead, they typically described a disabling condition that would prevent the antecedent from producing the consequent (35% of all explanations) or proposed that the antecedent of the generalization was not always sufficient to produce the consequent (27% of all explanations). There was no reliable difference in the sorts of implied revisions of the generalizations among the three sorts of problem.

EXPERIMENT 2

This experiment compared the two hypotheses using a different task. The problems contained facts that were inconsistent with the consequences of one of two generalizations. On half of the trials, the participants had to explain the inconsistency and then to decide whether the consequences of the *second* generalization still held. The following is an example:

> Speaker A: If you follow this diet, then you lose weight.
> Speaker B: If you follow this diet, then you have a good supply of iron.
> Speaker C: John followed this diet.
> In fact, John did not lose weight. Why?

After the participants had offered an explanation, they were asked, "Did John have a good supply of iron?"

It is possible that requesting participants to generate explanations causes them to change the inferences that they draw from the second generalization. To assess this possibility, on the other half of the trials, the participants did not have to explain the inconsistency but had merely to answer the question about the second generalization.

Logically, the resolution of the inconsistency calls for individuals to abandon either the categorical proposition that John followed the diet or the generalization that the diet yields a loss of weight, but it does not call for them to deny their conclusion that John has a good supply of iron, which, according to Harman (1986), would be an unnecessary and nonminimal change. Hence, it would be a minimal change to abandon the categorical proposition that John followed the diet and to maintain everything else, including the conclusion that he has a good supply of iron. Similarly, coherence theories of belief revision propose that individuals should retain conclusions even if they withdraw their justifications, provided that the conclusions are consistent with their other beliefs (Gärdenfors, 1992). Hence, individuals should retain conclusions that they have drawn from the second generalization, regardless of the sort of explanation that they create. But, according to the explanatory hypothesis, neither of these phenomena should tend to occur. An explanation should tend to imply a revision to the generalization yielding the inconsistency. And individuals should often give up the second generalization if their explanation is a potential disabler for it too (see Walsh & Sloman, 2007). When individuals do revise the categorical proposition, they also should tend to abandon the consequences of the second generalization, because they have no ready explanation for why it should hold.

Method

Participants. We tested individually 31 Princeton University undergraduates, who were paid for their participation.

Materials, Design, and Procedure. The format of the problems was the same as that in the previous experiment, but every problem was based on two conditional generalizations. On half the trials (in separate blocks, presented in a counterbalanced order), the participants answered the "why" question and then answered a question about the inference from the second generalization. On half of the trials, the participants answered only the question about the second generalization. Likewise, we counterbalanced the order of the two generalizations; on half of the trials, the fact denied the consequences of the first generalization, and on half of the trials, it denied the consequences of the second generalization. We used eight pairs of conditional generalizations using four types of subject matter: physiology, physics, economics, and psychology (see the Appendix). The contents were rotated so that they were presented equally often in the explanation condition, in which the participants wrote explanations for the inconsistencies, and in the no-explanation condition. Each participant completed the eight problems, which were presented in a different random order.

Results and Discussion

Table 2 presents the percentages of responses according to whether the participants' explanations implied a revision to the relevant generalization (*if A then B*) or to the categorical proposition (*A*), and according to whether they accepted, rejected, or were uncertain about the inference from the other generalization (*if A then C*). On 84% of trials, the participants violated minimalism, because they either implied a revision to the generalization or rejected the inference from the other generalization

Table 2
Percentages of Responses in Experiment 2 According to Whether the Participants' Explanations Implied a Revision to the Generalization or to the Categorical Proposition and Whether They Accepted, Were Uncertain About, or Rejected the Inference From the Other Generalization

Was the Inference of C From the Second Generalization, if A Then C, Accepted?	Explanations That Revised the Generalization: If A Then B.	Explanations That Revised the Categorical Proposition: A.	Totals
Yes	27	4	31
Maybe	27	10	37
No	4	26	30
Total	58	40	98*

Note—Both a revision to the generalization and a rejection of the inference from the other generalization are more than minimal changes. *Explanations for 2% of responses could not be classified.

(Wilcoxon test, $z = 3.62$, $p < .0002$). The probability of a response corroborating the violation of minimalism is 2/3, because responses in four of the six categories in Table 2 do so; hence, the Wilcoxon test compared frequencies of actual violations with this probability. Likewise, if the participants were minimalists, they should accept the inference from the other generalization. Yet they were just as likely to reject it as to accept it, and the majority of participants rejected it on at least one occasion. In the first block of problems, the participants were less likely to reject the inference when they had created an explanation (23%) than when they had not (44%, Mann–Whitney U test, $z = 1.94$, $p < .05$). The creation of an explanation seemed to have sensitized them to alternative possibilities, and so they gave more "maybe" answers.

Table 2 also shows a correlation between the nature of the explanation and the answer to the question about the second generalization. When the participants' explanations implied a revision to the relevant generalization, they tended to accept rather than reject the inference from the second generalization; but when their explanations implied a revision to the categorical proposition, they tended to reject rather than accept the inference from the second generalization. This interaction was highly reliable (Wilcoxon test, $z = 4.12$, $p < .00005$). Hence, contrary to minimalism but in accordance with the explanatory hypothesis, when the participants' explanations implied a revision to the categorical proposition, they then denied the consequence from the second generalization.

EXPERIMENT 3

Experiment 2 showed that, when participants received problems of the following form, they frequently inferred *not C*:

If A then B.
If A then C.
A.
In fact, not B.

Minimalism predicts that individuals should maintain C only if they drew this inference prior to encountering the inconsistent fact *not B*. Hence, Experiment 3 tested whether participants would abandon C if they explicitly drew this inference before they discovered the inconsistency.

Method

Participants. We tested 40 undergraduate and graduate students from the University of Plymouth, who were paid for their participation.

Materials, Design, and Procedure. We used the same eight pairs of conditional generalizations as those used in Experiment 2 and four new pairs (see the Appendix). The format of the problems was similar, but participants were asked "what follows?" before they were presented with the inconsistent fact:

Speaker A: If A then B.
Speaker B: If A then C.
Speaker C: A.
What follows?

After the participants wrote their responses, they turned the page. In eight of the problems, they were presented with an inconsistent fact and were asked to evaluate the inference from the second generalization:

In fact, not B. Why?
Does C follow?

In the remaining four problems, the new fact confirmed the consequent of one of the generalizations:

In fact, B. Why?
Does C follow?

The contents were rotated so that they were presented with each sort of problem. As before, the new fact, regardless of whether it was consistent, was relevant to the first generalization on half of the trials and relevant to the second generalization on half of the trials.

Results and Discussion

Three participants were eliminated because they failed to follow the instructions. Table 3 presents the percentages of responses for trials with an inconsistency, according to whether the participants' explanations implied a revision to the relevant generalization or to the categorical proposition and according to whether they accepted, rejected, or were uncertain about the inference from the other generalization. On 91% of trials, the participants violated minimalism, because they either implied a revision to the generalization or rejected the inference from the other generalization (Wilcoxon test, $z = 4.88$, $p < .00005$). As in Experiment 2, the probability of a response corroborating the violation of minimalism is 2/3, because responses in

Table 3
Percentages of Responses in Experiment 3 According to Whether the Participants' Explanations Implied a Revision to the Generalization or to the Categorical Proposition and Whether They Accepted, Were Uncertain About, or Rejected the Inference From the Other Generalization

Was the Inference From the Second Generalization Accepted?	Explanations That Revised the Generalization: If A Then B.	Explanations That Revised the Categorical Proposition: A.	Totals
Yes	15	0	15
Maybe	25	7	33
No	12	39	51
Total	52	46	98*

Note—Both a revision to the generalization and a rejection of the inference from the other generalization are more than minimal changes. *Explanations for 2% of responses could not be classified.

four of the six categories in Table 3 do so; hence, again the Wilcoxon test compared frequencies of actual violations with this probability. Furthermore, the difference in the proportions of rejections of the inference from the other generalization following a consistent and inconsistent fact was highly significant (1% vs. 51%, Wilcoxon test, $z = 4.21$, $p < .00005$). If the participants were minimalists, then, after an inconsistency, they should accept the inference from the other generalization. Yet, once again, they rejected it (51% of trials) more often than they accepted it (15% of trials, Wilcoxon test, $z = 3.22$, $p < .001$), and the balance of responses was uncertainty ("maybe"). After a consistent fact, only 1% of responses rejected the inference from the other generalization.

Table 3 also shows a correlation between the nature of the explanation and the answer to the question about the inference from the second generalization. When the participants' explanations implied a revision to the categorical proposition, they tended to reject the inference from the other generalization. In contrast, when their explanations implied a revision to the relevant generalization, they were more likely to accept the inference from the other generalization. This interaction, which the explanatory hypothesis predicts, was highly reliable (Wilcoxon test, $z = 4.59$, $p < .000005$). In sum, the participants often went beyond minimalism to deny a consequence that they had explicitly inferred in the first part of the trial, and they were particularly likely to do so when their explanations of an inconsistency ruled out the categorical proposition.

EXPERIMENT 4

If the principal task in resolving an inconsistency is to explain its origins, then individuals should be influenced by how easy it is to find an exception to a generalization. But this factor has no effect on a minimal reconciliation with the fact. The present experiment compared the two hypotheses using problems similar to those used in Experiment 3, except that the task was solely to decide whether the inference from a second generalization followed:

Speaker 1: If A then B.
Speaker 2: If A then C.
Speaker 3: A.
But, as a matter of fact: not B.
Does C follow?

Suppose that the generalization yielding the inconsistency has many disablers. According to the explanatory hypothesis, individuals can then envisage a disabler rendering the generalization false. Its falsity in turn explains the inconsistency, and so they should accept the inference from the other generalization. But, if the generalization yielding the inconsistency has few disablers, individuals should tend instead to reject the categorical proposition in order to resolve the inconsistency, and so they should also tend to reject the inference from the other generalization. Minimalism, of course, makes neither of these predictions: Individuals should have a universal tendency to accept the inference from the other generalization.

Method

Materials and Design. We constructed six pairs of problems, each having the form illustrated above. One generalization in each pair had few disablers (e.g., Speaker A: *If the TV is on, then energy is consumed*) and the other had many disablers (e.g., Speaker B: *If the TV is on, then there is a picture on the screen*), in line with previous normative studies (e.g., Cummins, 1995). A third speaker affirmed the antecedent common to both generalizations (Speaker C: *This TV was on*).

There were two groups of participants. The few disablers group received problems with a fact that denied the consequent of the generalization with few disablers:

In fact, no energy was consumed.
Was there a picture on the screen?

The many disablers group received problems with a fact that denied the consequent of the generalization with many disablers:

In fact, there was no picture on the screen.
Was energy consumed?

We also counterbalanced the order of the two generalizations within the problems. There were accordingly four versions of each problem. For both groups, the task was to write an answer to the question about the inference from the other generalization.

Participants. We tested individually 25 Princeton University undergraduates, who were paid for their participation. The participants were assigned at random to the few disablers group or to the many disablers group.

Results and Discussion

Table 4 presents the results for the two groups. The group with problems in which the facts conflicted with generalizations having many disablers tended to accept the inference from the other generalization (57%), whereas the group with problems in which the facts conflicted with generalizations with few disablers tended not to accept the inference from the second generalization (21%; Mann–Whitney $U = 22.5$, $p < .01$). This pattern of responses is inexplicable on a minimalist account, but the explanatory hypothesis predicts it. When a fact conflicted with a generalization with many disablers, the participants could explain the inconsistency in terms of the falsity of the generalization. The other generalization and the categorical proposition remained intact, and so the participants tended to accept their consequences. But when a fact conflicted with a generalization with few disablers, individuals could not so easily explain the inconsistency in terms of the falsity of the generalization, but they could explain it in terms of the falsity of the categorical proposition. They then

Table 4
Percentages of Responses in Experiment 4 Accepting, in Doubt of, or Rejecting the Inference From the Other Generalization for the Two Groups, Depending on Whether the Generalization Yielding the Inconsistency Had Many or Few Disablers

Was the Inference From the Second Generalization Accepted?	Many Disablers Group	Few Disablers Group
Yes	57	21
Maybe	11	32
No	32	47

Note—The rejection of the inference is a more than minimal change.

withdrew its consequences when it was coupled with the other generalization.

GENERAL DISCUSSION

From James (1907) to Gärdenfors (1988), a long-standing theory of how individuals cope with inconsistencies is that they should resolve them with the loss of as little information as possible: They should make minimal changes to accommodate the facts. This hypothesis is plausible, but, as our experiments show, it is not always true. The crux is that individuals do not just revise propositions to resolve inconsistencies, but instead they seek to explain the origins of the inconsistency (Johnson-Laird et al., 2004). This explanatory hypothesis leads to predictions that diverge from minimalism.

First, explanations of inconsistencies should tend to revise conditional generalizations more often than they do categorical propositions. Generalizations such as *if you strike a match, then it lights* have disabling conditions (e.g., the match was soaking wet). Propositions such as *Pat struck a match* are indeed categorical, so it is harder to envisage a plausible and principled reason why they might be false. All the experiments yielded such a tendency, and it was highly robust in Experiment 1, in which the participants' explanations tended to rule out the generalizations. This tendency is not a minimal change, because, logically speaking, it also removes support for other conclusions apart from the one giving rise to the conflict (Harman, 1986).

Second, when the facts conflict with the consequences of one of two generalizations, individuals should still tend to think of explanations that revise the conflicting generalization. For example, consider the following problem:

Speaker A: If you follow this diet, then you lose weight.
Speaker B: If you follow this diet, then you have a good supply of iron.
Speaker C: John followed this diet.
In fact, John did not lose weight.

Explanations should tend to be based on a disabling condition for the first generalization (e.g., a failure to exercise). There was indeed no reliable decline in such explanations in Experiment 1 in comparison with problems that contained no other generalization. But, as Experiments 2 and 3 showed, when explanations did revise the categorical proposition in cases such as the example above, the participants tended also to reject the consequences of the second generalization, because they had no explanation for why it should hold. This result contravenes minimalism, because a minimal change would be to abandon the categorical proposition in the example above and nothing else.

Third, the tendency to create explanations that revise a generalization should be biased by whether the generalization itself has many potential disablers (e.g., *If coal is burnt, then the room is warm*) or only a few (e.g., *If coal is burnt, then carbon dioxide is produced*). It is easier to think of disabling conditions for generalizations with many potential disablers than for those with few of them. When both generalizations occur in a problem yielding an inconsistency, a critical variable is which of the two yields a consequence in conflict with the facts. As Experiment 4 showed, when participants could easily resolve an inconsistency by rejecting the generalization giving rise to it, they were more likely to accept the consequences of the other generalization. In this case, the categorical proposition and the other generalization remained intact, and so the participants accepted their consequences. When the participants could not resolve the inconsistency so easily by rejecting the generalization, because it had few disablers, they presumably rejected the categorical proposition instead. They now had no reason to accept its consequences from the other generalization, and so they rejected the inference. This interaction is inexplicable in terms of the minimalist hypothesis, which in both cases predicts that individuals should reject the categorical proposition but nothing else.

Minimalism is a putative solution to the problem of maintaining a consistent set of propositions. It ensures that individuals do as little as possible in order to accommodate new information that conflicts with old. In contrast, our results imply that human reasoners use a different strategy. They are concerned about explaining inconsistencies, and their quest for plausible explanations leads them to imagine how the events might have unfolded. This approach can lead them to make greater than minimal changes to the information that they have.

AUTHOR NOTE

The research was supported by a postdoctoral fellowship to C.R.W. from the Educational Testing Service, Princeton, and by grants to P.N.J.-L. from the National Science Foundation (BCS-0076287 and SES 0844851) to investigate strategies in reasoning and deductive and probabilistic reasoning. We thank Sam Glucksberg and Uri Hasson for helpful discussions and Clare Partridge for her help in running one of the experiments. Address correspondence to C. R. Walsh, Centre for Thinking and Language, School of Psychology, University of Plymouth, Drake Circus, Plymouth PL4 8AA, England (e-mail: clare.walsh@plymouth.ac.uk).

REFERENCES

Byrne, R. M. J., & Walsh, C. R. (2002). Contradictions and counterfactuals: Generating belief revisions in conditional inference. In W. Gray & C. Schunn (Eds.), *Proceedings of the 24th Annual Conference of the Cognitive Science Society* (pp. 160-165). Mahwah, NJ: Erlbaum.

Cummins, D. D. (1995). Naive theories and causal deduction. *Memory & Cognition*, **23**, 646-658.

de Kleer, J. (1986). An assumption-based TMS. *Artificial Intelligence*, **28**, 127-162.

Dieussaert, K., Schaeken, W., De Neys, W., & d'Ydewalle, G. (2000). Initial belief state as predictor of belief revision. *Current Psychology of Cognition*, **19**, 277-288.

Doyle, J. (1979). A truth maintenance system. *Artificial Intelligence*, **12**, 231-272.

Elio, R. (1998). How to disbelieve p→q: Resolving contradictions. *Proceedings of the 20th Annual Meeting of the Cognitive Science Society* (pp. 315-320). Mahwah, NJ: Erlbaum.

Elio, R., & Pelletier, F. J. (1997). Belief change as prepositional update. *Cognitive Science*, **21**, 419-460.

Gärdenfors, P. (1988). *Knowledge in flux: Modeling the dynamics of epistemic states*. Cambridge, MA: MIT Press.

Gärdenfors, P. (1992). Belief revision: An introduction. In P. Gärdenfors (Ed.), *Belief revision* (pp. 1-28). Cambridge: Cambridge University Press.

HARMAN, G. (1986). *Change in view: Principles of reasoning.* Cambridge, MA: MIT Press.
HASSON, U., & JOHNSON-LAIRD, P. N. (2003). Why believability cannot explain belief revision. In R. Alterman & D. Kirsh (Eds.), *Proceedings of the 25th Annual Conference of the Cognitive Science Society* (pp. 516-521). Mahwah, NJ: Erlbaum.
HIDDLESTON, E. (2005). A causal theory of counterfactuals. *Nous*, **39**, 632-657.
JAMES, W. (1907). *Pragmatism: A new name for some old ways of thinking*. New York: Longmans, Green.
JOHNSON-LAIRD, P. N., GIROTTO, V., & LEGRENZI, P. (2004). Reasoning from inconsistency to consistency. *Psychological Review*, **111**, 640-661.
MARKOVITS, H., & SCHMELTZER, C. (2007). What makes people revise their beliefs following contradictory anecdotal evidence: The role of systemic variability and direct experience. *Cognitive Science*, **31**, 535-547.
NISBETT, R. E., & WILSON, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, **84**, 231-259.
POLITZER, G., & CARLES, L. (2001). Belief revision and uncertain reasoning. *Thinking & Reasoning*, **7**, 217-234.
SHAFIR, E., SIMONSON, I., & TVERSKY, A. (1993). Reason-based choice. *Cognition*, **49**, 11-36.
THAGARD, P. (1989). Explanatory coherence. *Behavioral & Brain Sciences*, **12**, 435-502.
WALSH, C. R., & SLOMAN, S. A. (2007). Updating beliefs with causal models: Violations of screening off. In M. A. Gluck, J. R. Anderson, & S. M. Kosslyn (Eds.), *Memory and mind: A festschrift for Gordon H. Bower* (pp. 345-357). Mahwah, NJ: Erlbaum.

APPENDIX
Materials Used in the Experiments

Conditional Generalizations Used in Experiments 1–3

If you follow this diet, then you lose weight. If you follow this diet, then you have a good supply of iron.

If the drink contains sugar, then it tastes sweet. If the drink contains sugar, then it gives you energy.

If people are worried, then they find it difficult to concentrate. If people are worried, then they have insomnia.

If people are nervous, then their hands shake. If people are nervous, then they get butterflies in their stomach.

If coal is burnt, then heat is produced. If coal is burnt, then it yields carbon monoxide.

If a match is struck, then it produces light. If a match is struck, then it gives off smoke.

If the economy is doing well, then share prices rise. If the economy is doing well, then unemployment is low.

If the banks cut interest rates, then the economy grows. If the banks cut interest rates, then there is an increase in inflation.

Conditional Generalizations Used in Experiment 1 Only

If sales go up, then profits improve. If sales go up, then their production costs are reduced.

If someone is very kind to you, then you like that person. If someone is very kind to you, then you are kind in return.

If people have a fever, then they have a high temperature. If people have a fever, then they lack an appetite.

If there is high pressure, then the night is cold. If there is high pressure, then the skies are clear.

Conditional Generalizations Used in Experiment 3 Only

If there is heavy rain, then the football pitch becomes muddy. If there is heavy rain, then there are puddles on the footpath.

If there is very loud music, then it is difficult to have a conversation. If there is very loud music, then the neighbours complain.

If it is windy, then the wind chimes jingle. If it is windy, then the clothes on the line move.

If the manufacturing process is operating, then smoke is emitted. If the manufacturing process is operating, then chemical by-products are emitted.

Conditional Generalizations Used in Experiment 4

Few Disablers	Many Disablers
If coal is burnt, then carbon dioxide is produced.	If coal is burnt, then the room is warm.
If the window is opened, then air escapes.	If the window is opened, then the room becomes cold.
If the dishes are washed, then they get wet.	If the dishes are washed, then they are clean.
If a sponge is placed in water, then it absorbs water.	If a sponge is placed in water, then it floats.
If a person sprints, then their heartbeat is raised.	If a person sprints, then they lose weight.
If the noise in the room is extremely loud, it is difficult to have a conversation.	If the noise in the room is extremely loud, the neighbors complain.

(Manuscript received March 4, 2008;
revision accepted for publication February 8, 2009.)