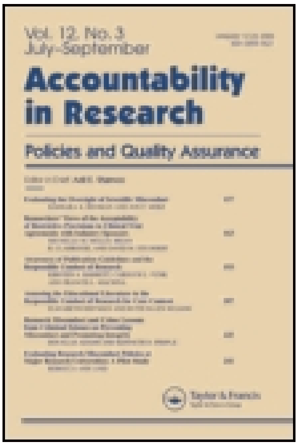

Accountability in Research: Policies and Quality Assurance
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/gacr20

The Epistemic Integrity of NASA Practices in the Space Shuttle Program

Jan De Winter M.A. [a] & Laszlo Kosolosky M.Phil [a]

[a] Centre for Logic and Philosophy of Science, Department of Philosophy and moral sciences, Ghent University, Ghent, Belgium

To cite this article: Jan De Winter M.A. & Laszlo Kosolosky M.Phil (2013): The Epistemic Integrity of NASA Practices in the Space Shuttle Program, Accountability in Research: Policies and Quality Assurance, 20:2, 72-92

To link to this article: http://dx.doi.org/10.1080/08989621.2013.767122

Accountability in Research, 20:72–92, 2013

ISSN: 0898-9621 print / 1545-5815 online
DOI: 10.1080/08989621.2013.767122

Taylor & Francis
Taylor & Francis Group

The Epistemic Integrity of NASA Practices in the Space Shuttle Program

Jan De Winter, M.A. and Laszlo Kosolosky, M.Phil

Centre for Logic and Philosophy of Science, Department of Philosophy and moral sciences, Ghent University, Ghent, Belgium

This article presents an account of epistemic integrity and uses it to demonstrate that the epistemic integrity of different kinds of practices in NASA's Space Shuttle Program was limited. We focus on the following kinds of practices: (1) research by working engineers, (2) review by middle-level managers, and (3) communication with the public. We argue that the epistemic integrity of these practices was undermined by production pressure at NASA, i.e., the pressure to launch an unreasonable amount of flights per year. Finally, our findings are used to develop some potential strategies to protect epistemic integrity in aerospace science.

Keywords: aerospace science, deception, epistemic integrity, NASA, production pressure, Space Shuttle

1. INTRODUCTION

In philosophy of science, there is an increasing awareness of the fact that scientific knowledge is usually not something that is constructed by one brilliant individual, but that it is typically a collective achievement, resulting from complex social processes in which several individuals with often diverse areas of expertise are involved (Nickles, 1980). This insight is the basis of social epistemology—a fairly recent branch of philosophy of science in which the social dimensions of knowledge are studied.[1] One of the main questions is how science should be socially organized. Several philosophers of science have tried to answer this question (see, e.g., Fuller, 1988, 2000; Kitcher, 1993, 2001; Longino, 1990, 2002). Their proposals are, however, often overly general and overly abstract. They are too general because they assume that the ideal organization applies to all different scientific domains—from sociology to quantum

Address correspondence to Jan De Winter, Centre for Logic and Philosophy of Science, Department of Philosophy and moral sciences, Ghent University, Blandijnberg 2, 9000 Ghent, Belgium. E-mail: Jan.DeWinter@UGent.be

physics—while in fact different scientific domains may ask for different kinds of social organization, as very different epistemic issues arise in different areas of research. The proposals are too abstract because they are not sufficiently informed by examinations of the actual organization of scientific communities, which makes them quite unhelpful for actual policymaking in science (Biddle, 2007, pp. 23–24).

We consider it important that social epistemologists look at actual scientific practice. If one wants to know which adjustments to the current regime in a certain scientific domain are optimal, one should know how the domain is currently socially structured, which problems arise within this specific domain, and what exactly caused these problems. So social epistemologists should examine which problems turn up in a certain area of research, and consider what caused them. Such reflections should form the basis for concrete recommendations on how to change the current social organization of the research area under consideration.

Some social epistemologists have already proceeded along these lines.[2] Biddle (2007), Brown (2008), De Winter (2012), Hollis and Pogge (2008), and Reiss (2010) all describe certain specific problems for the health sciences and discuss different concrete strategies to deal with them. Some other areas of research have, however, received considerably less attention from social epistemologists. This does not mean that these research areas do not have certain problems that fall within the field of social epistemology. In this article, we take a look at an area of research that, although it has received relatively little attention from social epistemologists, has some important epistemic issues to deal with: aerospace science. More specifically, we examine a particular case from this field: the Space Shuttle Program.

The Space Shuttle Program was conducted by National Aeronautics and Space Administration (NASA) from 1981 to 2011, and its aim was the development and exploitation of reusable spacecraft. Five Space Shuttle orbiters were developed as part of this program—Columbia, Challenger, Discovery, Atlantis, and Endeavour—and in total there were 135 Space Shuttle missions—Columbia had 28 flights, Challenger 10, Discovery 39, Atlantis 33, and Endeavour 25. Several different kinds of practices were involved in the program, ranging from the adjustment and testing of shuttle parts to informing the press about NASA achievements. What we will argue in this article is that the epistemic integrity of several of these practices was compromised by production pressure at NASA, i.e., the pressure to launch an unreasonable amount of flights per year. In order to do so, we first offer a clear account of epistemic integrity (Section 2). Next, we describe three different kinds of practices in the Space Shuttle Program (Section 3), and we give some specific examples of limited epistemic integrity for each of these practices (Sections 4–6). Section 7 shows that production pressure at NASA could explain why epistemic integrity was damaged in the examples considered. We hope this explanation can help social epistemologists in developing strategies to promote and protect the epistemic integrity of practices in aerospace science, and we draw some policy-related lessons from our findings ourselves in Section 8. We conclude in Section 9.

2. EPISTEMIC INTEGRITY

Before we can argue that the epistemic integrity of different practices in the Space Shuttle Program was limited, we should clarify what we mean by epistemic integrity. Elsewhere, we have explicated this concept and discussed it extensively (De Winter and Kosolosky, in press). Let us recapitulate the basics of our explication. Epistemic integrity is, as we understand it, a property that a practice, such as research, can have to a higher or lower degree.[3] More specifically, the epistemic integrity of a practice is a function of the degree to which the statements resulting from this practice are deceptive—the more deceptive these statements, the lower the epistemic integrity of the practice. Applied to research, we get the following: the more deceptive the (so-called) results or conclusions of research, the lower the epistemic integrity of this research.

We clarify what we mean by deceptiveness below, but first two remarks are in place. The first is that there is an important constraint on what we mean by "statements resulting from a practice." We assume that for a statement to be considered a result of a certain practice, it should be made within this practice, by persons or entities who are actually involved in the practice. The epistemic integrity of the practice only depends on the deceptiveness of such statements. So, for instance, if someone who is not involved in a certain research project draws false conclusions from it, and presents them as true, this does not entail that the research upon which he relies has low epistemic integrity. His analysis of the research and its results may have low epistemic integrity, but the epistemic integrity of the research project itself is not compromised as long as those involved in the project have not made any deceptive claims themselves. So in order to assess the epistemic integrity of NASA practices, we should look at the statements that those involved in these practices made, and more specifically, at how deceptive these statements were.

Secondly, it should be noted that for someone to make a statement, it is not always required that this person explicitly articulates this statement. The most obvious example is when a person asks someone else a question, say, "Is the Space Shuttle ready to fly?," and the latter answers with "Yes" or "No." Although the latter person did not literally say "the Space Shuttle is (not) ready to fly," his "Yes" or "No" answer is sufficient to conclude that he/she made this statement. A more subtle example is when a person P does not say anything, while there are certain generally acknowledged rules implying that in case event E happens, P should report this. We can then say that by not reporting anything, P implicitly made the statement that E did not occur.

We should also clarify what we mean by the deceptiveness of a statement, as our definition of epistemic integrity refers to this property. We consider a statement s, stated by person P_1, deceptive to a person P_2 if and only if (i) s is presented to P_2 as a true statement, and (ii) (a) s is false or (b) P_2 infers a false statement s' from s, and it is legitimate for P_2 to make this inference (given what is usually the case when a statement like s is true, what persons like P_1 usually mean by statements like s, what P_1 says about the interpretation of s, etc.). s is more deceptive to P_2 as P_2's inference from s to s' is more legitimate. The general degree of deceptiveness of a statement s is higher as s is more deceptive to a higher percentage of the people to whom it is communicated. We hope that the examples of deceptive claims (and hence, of limited epistemic integrity) we give in Sections 4 to 6 will make our accounts of deceptiveness and epistemic integrity sufficiently clear.[4]

But before we turn to the case of the Space Shuttle Program, let us briefly indicate why we consider the concept of epistemic integrity we have presented in this section valuable, and why we consider using it to evaluate practices in the Space Shuttle Program valuable. In De Winter and Kosolosky (in press), we argue that our concept of epistemic integrity is a good explication of the (epistemic) notion of research integrity. We do this by showing that it meets the following four requirements, which are based on Rudolf Carnap's (1950) characterization of the task of explication: (1) it can be used in most cases in which the term "research integrity" has so far been used, (2) it is exact, that is, clear and well-defined, (3) it is fruitful, that is, useful for the development and justification of norms for individuals and institutions involved in science, and (4) it is as simple as requirements (1), (2), and (3) permit. We show that while existing accounts of research integrity (Haack, 2007; National Academy of Sciences, 1992; Office of Research Integrity, 2007; Petrovečki and Scheetz, 2001; Steneck, 2006) do not meet these four requirements, our account does. Instead of repeating these arguments here, let us compare our concept of epistemic integrity to two concepts not considered in De Winter and Koslosky (in press): (1) integrity as the adherence to certain standards, such as the standard that one should disclose information that is disadvantageous to him/herself, and (2) the competence to evaluate and give relevant information.[5]

Firstly, consider the standard that one should disclose information that is disadvantageous to him/herself. Of course, a scientist should not necessarily disclose all information that is disadvantageous to him/herself, but only information that is relevant. A question that arises isthe following one: What counts as relevant? What one may consider relevant, may not be relevant according to someone else. Moreover, we can think of degrees of relevance, and ask *how* relevant a finding should be in order to consider it a finding that ought to be disclosed. Our account offers a criterion to settle such issues: If not including a finding in a report makes the report deceptive (because it then contains false statements, or because the audience will then legitimately infer false statements from the report), while it would not be deceptive if the finding *would* be included in the report, then the finding is relevant and should be disclosed. This implies that our concept of epistemic integrity is more fruitful than the standard that one should disclose information that is disadvantageous to him/herself: Contrary to the latter standard, our concept of epistemic integrity enables us to develop (adequate) norms that indicate which information is relevant and should be disclosed and which information isn't, and it enables us to (adequately) justify why certain findings were relevant and should have been disclosed.

A similar argument could be developed against the concept of competence to evaluate and give relevant information. This concept does not offer a criterion to determine what counts as relevant and what doesn't, nor does it specify how relevant a finding should be for this to be a finding that should be disclosed. Therefore, the concept is not exact, nor is it fruitful for the development and justification of norms on which information to disclose. We hope that this very brief defense of our concept of epistemic integrity is, at least for now, sufficient to convince the reader of the value of using this concept (rather than other concepts) to analyze the case of the Space Shuttle Program, but we will return to this later in this article (in Section 9, in which we explicitly mention some advantages of our analysis over earlier analyses of the Space Shuttle Program).

3. DIFFERENT KINDS OF PRACTICES IN THE SPACE SHUTTLE PROGRAM

If we want to assess the epistemic integrity of practices in the Space Shuttle Program, we should look at practices that resulted in certain statements, since it is on the basis of such statements, more specifically on the basis of their deceptiveness, that the epistemic integrity of a practice is determined. We can discern three general kinds of practices in the Space Shuttle Program. A first kind is research performed by working engineers. This included tests on shuttle components before launch, and postflight analysis of shuttle components (checking whether it was damaged during flight, whether it had worked properly, etc.). Involved in research on a shuttle component were: (1) working engineers from the NASA contractor that developed this component, and (2) working engineers from the NASA center that was responsible for the management of this component. They communicated their findings to each other and to middle-level managers. A formal channel for such communication were the Level IV and Level III meetings of Flight Readiness Review (FRR). FRR was a set of meetings (proceeding from Level IV to Level I) that preceded each launch, and at which the readiness of the shuttle to fly and to fly safely was determined. At Level IV meetings, contractor engineers presented their data analyses and conclusions to their own managers and to personnel from the relevant NASA center, and at Level III meetings, they informed NASA Project Managers of data and findings on risk acceptability (Vaughan, 1996, p. 84). There was also a lot of communication between contractor engineers, NASA center engineers, and their managers outside FRR. For instance, during flight, there was daily communication between contractor engineers and NASA center engineers (Vaughan, 1996, p. 85).

A second kind of practice is the review of working engineers' research results by middle-level managers (from NASA contractors and NASA centers). These middle-level managers checked whether data were consistent, whether conclusions were adequately supported by data, whether arguments met NASA standards of quantification, etc. They communicated their findings to the working engineers, and after the required adjustments were made in the engineering analysis (by themselves and/or the working engineers), they passed the results up the hierarchy, to top-level managers. The former kind of communication (feedback to working engineers) occurred at Level IV and Level III meetings, and the latter kind of communication (passing the results up the hierarchy) occurred at Level II and Level I meetings (Vaughan, 1996, p. 84).

A third kind of practice is NASA's communication with the public. This includes the many press releases from NASA. For instance, in 1991, there were 203 press releases from NASA. Of the 195 press releases from 1991 we could consult on the website of NASA, 68 contained the word "shuttle" (which means that they related to the Space Shuttle Program). But this is not the only way in which NASA informed the public. It also published technical reports on the Space Shuttle, FRR reports, messages between NASA employees, etc. Several of these documents were released as a response to requests from the Freedom of Information Act.[6]

4. RESEARCH BY WORKING ENGINEERS

Now that we have a clear concept of epistemic integrity, and a picture of different kinds of practices in the Space Shuttle Program, we can argue that different kinds of practices in the Space Shuttle Program had limited epistemic integrity. We do this by giving, for each kind of practice, examples of deceptive claims that resulted from it. Given the account of epistemic integrity presented in Section 2, this is sufficient to conclude that the epistemic integrity of the practice under consideration was limited.

Let us start with research performed by working engineers. In order to demonstrate that the epistemic integrity of such research was damaged, we should show that those involved in it—the working engineers—made deceptive statements. Let us give some examples from the eve of the Challenger disaster. On January 27, 1986, the day before Challenger's fatal mission, two teleconferences were held, in which engineers and managers associated with the Solid Rocket Booster (SRB) Project located at Morton Thiokol Corporation (NASA contractor that built the SRBs), Marshall Space Flight Center (NASA center that bore managerial responsibility for the SRBs), and Kennedy Space Center (NASA center that assembled and tested the shuttle components and conducted launches) discussed the effect of the expected cold temperatures on the O-rings that were supposed to seal certain joints on the SRBs. Before the teleconferences, launch time was set at 9:38 A.M., and at that time, temperature was predicted to be 26°F. During the first teleconference, Thiokol suggested to delay launch until noon or after, when temperatures would be higher. A second teleconference was arranged in which Thiokol would take an official position and in which it would support this position with engineering data. At that teleconference, Thiokol initially recommended not to launch unless O-ring temperature was at least 53°F. One of the charts that was presented to support that position, by Thiokol engineer Arnie Thompson, concerned secondary O-ring resiliency. More specifically, it contained test data on how long it took for the secondary O-ring to re-establish contact after the field joint secondary seal lifted off the metal mating surfaces during motor pressurization. The chart revealed that this was a function of temperature, and it stated that at 50°F, the time to recover was 600 seconds (Vaughan, 1996, Chap. 8). But in fact, at 50°F, the O-ring did not re-establish contact at all in the test, and after 10 minutes, the test was terminated (Presidential Commission, 1986, Vol. 5, p. 1568). So 600 seconds was not the time to recover but the duration of the test. This means that conditions (i) and (ii)(a) of our account of deceptiveness are met: Thompson presented the statement that the time to recover was 600 seconds as a true statement, while it was in fact false. Hence, his chart was deceptive.

Arnie Thompson was not the only working engineer who made a deceptive statement during the second teleconference. During that teleconference, Thiokol engineer Roger Boisjoly was asked what evidence Thiokol had that O-ring erosion on the Space Shuttle's fifteenth mission was due to the cold. After all, there also was field joint blow-by on the Space Shuttle's twenty-second mission, when O-ring temperature was not that low. Boisjoly answered that he did not have the data to quantify the temperature concerns (Vaughan, 1996, pp. 303–304). But in fact, he did have such data. This was shown by two members of the investigative staff of the Presidential Commission that U.S. President Ronald Reagan established after the Challenger disaster to investigate the circumstances leading up to the accident. Alton Keel, the executive director of this investigative staff, and Randy Kehrli, a Department of Justice attorney, created a chart that included all missions, indicating for each mission the number of O-ring anomalies and calculated O-ring temperature. This chart showed a clear correlation between O-ring anomalies and temperature:

while only 17.6% of the flights with O-ring temperature above 65°F had O-ring anomalies, 100% of the flights with O-ring temperature below 65°F had O-ring anomalies (Vaughan, 1996, p. 382). So Boisjoly's claim that he did not have the data to quantify the temperature concerns was false: Boisjoly did have the data needed; he just wasn't able to translate them into a quantitative signal that cold temperatures could cause O-ring erosion. As Boisjoly's claim was false, it was deceptive [conditions (i) and (ii)(a) of our account of deceptiveness are met].

It is remarkable that none of the other working engineers participating in the second teleconference objected to Boisjoly's claim. Take for instance Marshall S&E engineer (engineer from Marshall's Science and Engineering Directorate) Ben Powers. The job of S&E engineers was to keep the contractor honest (Vaughan, 1996, p. 86). Hence, if Thiokol's temperature concerns could easily be quantified on the basis of the available data—and they could, since Keel and Kehrli, two non-engineers, were able to do so—and someone of Thiokol says that he/she does not have the data to quantify the temperature concerns, then Powers should have protested. This was his job. On the basis of the fact that he didn't, we can argue that he, too, was being deceptive. Because he did not protest, not even when George Hardy, Marshall's Deputy Director of Science and Engineering, asked at the end of the teleconference "Has anybody got anything to add to this?," he implicitly made the statement that he did not have anything to add (see Section 2: one can make a statement without explicitly articulating it). We can expect the audience to have inferred from this that Boisjoly was right and that the temperature concerns could not be quantified on the basis of the available data. It was legitimate for them to do so because in case Boisjoly was not telling the truth, it was Powers' job to make an objection. Therefore, the implicit statement that he did not have anything to add was deceptive [conditions (i) and (ii)(b) of our account of deceptiveness are met: Powers presented the statement that he did not have anything to add as a true statement, and the audience legitimately inferred the false statement that the temperature concerns could not be quantified on the basis of the available data from it].

5. REVIEW BY MIDDLE-LEVEL MANAGERS

Review by middle-level managers resulted in several deceptive statements as well. Consider some examples from Lawrence Mulloy, SRB Project Manager at Marshall from 1982 to 1986. On June 25, 1985, it was found that both the primary and secondary O-ring of one of the SRB joints had eroded on the seventeenth mission of the Space Shuttle (Winsor, 1988, p. 104). It was the first time that a primary O-ring had burned all the way through, allowing hot gases to erode a secondary O-ring. This finding resulted in a Launch Constraint, which meant that flight could only proceed if either the problem was resolved, or engineering analysis showed that the risk was acceptable (Vaughan, 1996, pp. 163–164). But the Launch Constraint only included the SRB nozzle joints, and not the SRB field joints, as it was a nozzle joint that failed on the Space Shuttle's seventeenth mission. The reason was that the nozzle joints were tested for leaks at 100 psi, while the other SRB joints were leak-tested at 200 psi, and Mulloy assumed that the damage in the nozzle joint was due to the less rigorous leak test for the nozzle joints. He assumed that a defective O-ring escaped notice in the 100-psi leak test, and that this caused the damage. As leak check pressure for the nozzle joints was raised to 200 psi for the subsequent flights, Mulloy waived the Launch Constraint (Winsor, 1988, p. 104). This was the result of Level IV and Level III meetings, and Mulloy reported a summary of the problem and its resolution to Levels II and I (Vaughan, 1996, p. 169).

The claim that the problem only concerned the nozzle joints raised the impression that there were good reasons to believe that the damage found did not have any safety implications for the field joints. But this is not correct. As we have seen, this view was based on the assumption that the damage in the nozzle joint was caused by a defective O-ring that escaped notice due to the 100-psi leak test. But this was only one possible hypothesis; another possibility was that the damage was caused by defective design. Wiley Bunn, Marshall Director of Reliability and Quality Assurance, justified the rejection of the latter hypothesis on the following grounds:

> We had six joints on that vehicle. If the design is that darn bad, all six of them should have leaked. We only had one leak. Therefore, if we only had one leak, it had to be a Quality escape. And so we just renewed our vigor to find that Quality escape. (Bunn, interview transcript, April 17, 1986, pp. 61–62, cited in Vaughan, 1996, pp. 164–165)

This is, however, not a good argument: the fact that five out of six joints did not leak does not imply that the design was fully adequate—even a design that is not fully adequate can work most of the time. In fact, the claim that only the safety of the nozzle joints, and not the safety of the field joints, had to be considered, was not supported by any good arguments, but only by the inadequately supported assumption that it was the leak test, and not the design, that was problematic.[7] Because we can expect that the people to whom this claim was communicated (Levels IV, III, II, and I), inferred that it *was* supported by good arguments, the claim was deceptive [conditions (i) and (ii)(b) are met].

This was not the first time Mulloy made a deceptive statement with respect to O-ring erosion. When the Space Shuttle's SRBs were examined after its fifteenth flight, Thiokol engineers found blow-by on two SRB field joints, which was "jet black" and intermixed with the grease between the O-rings. This alarmed them, because it indicated that there were two destructive forces at work on the O-rings, impingement erosion and blow-by erosion. In explaining this, Thiokol engineers referred to the cold: the fifteenth mission was preceded by the three coldest days in Florida history and O-ring temperature was 53°F, which was the lowest up to that time, and such low temperatures enhanced the probability of blow-by. This figured in their rationale for accepting risk: as Thiokol engineers did not expect to have the three coldest days in Florida history again, they accepted the risk for the next flight (Vaughan, 1996, pp. 154–163).

However, as the rationale for accepting risk was reported up the hierarchy, temperature considerations were omitted. When the risk associated with O-ring erosion was assessed for the sixteenth shuttle mission, the temperature concerns of Thiokol engineers were considered at lower levels of FRR, but when Mulloy presented the rationale for accepting risk to Level I, he did not mention these concerns. He stated that the risk was acceptable because (1) the amount of primary erosion (0.038") was within the experience base (the 0.053" erosion found after the Space Shuttle's second mission), (2) the amount of erosion was within the 0.090" safety margin, and (3) the effect of hot gas flow on the O-rings—impingement erosion and blow-by—was self-limiting. No reference was made to the temperature concerns (Vaughan, 1996, p. 161).

The reason why Mulloy did not mention the temperature considerations was that systematic data proving the association between temperature and O-ring erosion were not yet available (Vaughan, 1996, pp. 161–162). The appropriate response, however, would have been to collect such data, instead of omitting the temperature considerations, as if they did not play a role in the decision to accept risk. Because of this omission, NASA administrators at Level I have probably (legitimately) inferred from Mulloy's report that the three-factor rationale he presented (experience base, safety margin, and the self-limiting nature of the phenomenon) was the full rationale for accepting the risk associated with O-ring erosion, by those who were most familiar with this problem, that is, the engineers at Thiokol. This is false, since their rationale also included the expectation that the cold preceding the fifteenth mission would not recur, as we have seen. Because conditions (i) and (ii)(b) of our account of deceptiveness are met, we consider Mulloy's report deceptive.

Other middle-level managers made deceptive claims as well. An example is the report of the weather concerns by Stanley Reinartz, Manager of the Shuttle Projects Office at Marshall, to William Lucas, the Director of Marshall, on the morning of the Challenger accident. According to Lucas, Reinartz said "that an issue had been resolved, that there were some people at Thiokol who had a concern about the weather, that that had been discussed very thoroughly by the Thiokol people and by the Marshall Space Flight Center people, and it had been concluded agreeably that there was no problem, that he had a recommendation by Thiokol to launch and our most knowledgeable people and engineering talent agreed with that" (Presidential Commission, Vol. 1, pp. 100–101). If that was indeed what Reinartz said, then he made a deceptive claim, as some of Thiokol and Marshall's most knowledgeable people, among others Thiokol engineers Arnie Thompson and Roger Boisjoly, and Marshall engineers Ben Powers and Keith Coates, did in fact not agree with the recommendation to launch (Vaughan, 1996, Chap. 8) [conditions (i) and (ii)(a) are met].

Finally, we would also like to offer an example of deception by NASA contractor management (in order to show that not only Marshall managers were being deceptive). Richard Feynman, who was a member of the Presidential Commission on the Space Shuttle Challenger Accident, exposed a strange use of the term "safety factor." When the risk associated with O-ring erosion was assessed for the Space Shuttle's twenty-fifth mission, it was asserted that there was a safety factor of three. The idea was that on the Space Shuttle's fifteenth mission, erosion depth was one-third of the radius, while it had to be at least one radius before the O-ring failed (Feynman, 1986, pp. F1–F2). This use of the term "safety factor" was adopted by Thiokol management, as the following quote from Allan McDonald, Director of the Solid Rocket Motor Project at Thiokol, indicates:

> If you took our worst measured erosion on the O-ring relative to what it took to really fail it, it was nearly a factor of three to one. Recognizing the fidelity of the math model is not real good, we did not feel it was that bad at three to one, and as long as we could retain the secondary seal during a good portion of the erosion time period, we felt good. (Presidential Commission, 1986, Vol. 5, p. 1591)

But to speak in such a context of a factor of three to one, from which safety could be inferred, is deceptive. This is because the audience has probably inferred from such claims that the O-rings were more or less safe, or more specifically, that their safety was similar to safety in other technological contexts in which there is a safety factor of three to one. But this is false [conditions (i) and (ii)(b) are met], as Feynman shows on the basis of the following example:

> If a bridge is built to withstand a certain load without the beams permanently deforming, cracking, or breaking, it may be designed for the materials used to actually stand up under three times the load. This "safety factor" is to allow for uncertain excesses of load, or unknown extra loads, or weaknesses in the material that might have unexpected flaws, etc. If now the expected load comes on to the new bridge and a crack appears in a beam, this is a failure of the design. There was no safety factor at all; even though the bridge did not actually collapse because the crack went only one-third of the way through the beam. The O-rings of the Solid Rocket Boosters were not designed to erode. Erosion was a clue that something was wrong. Erosion was not something from which safety can be inferred. (Feynman, 1986, p. F2)

6. COMMUNICATION WITH THE PUBLIC

The third kind of practice we identified is communication with the public. An example is NASA's public announcement that the Space Shuttle Program was operational on July 4, 1982, the day that the fourth test flight was finished successfully. We can expect the audience, including potential payload customers, holders of the congressional purse strings, voters, and foreign nations, to have (legitimately) inferred that this meant that the Space Shuttle had attained an airline-like degree of routine operation (Presidential Commission 1986, Vol. 1, p. 5; Vaughan 1996, p. 125). But this was false; the Space Shuttle was, at that time, and years afterwards, a developmental craft with constantly changing technology and mysterious problems that were not predicted from design, features that are not characteristic of the operational phase in other areas of aviation (Hall, 2003, p. 240). Since conditions (i) and (ii)(b) of our account of deceptiveness are met, we consider NASA's public declaration that the Space Shuttle Program was operational deceptive.

Another example is the estimation that the probability of mission failure was 1 in 100,000, a figure that was published in "Space Shuttle Data for Planetary Mission Radioisotope Thermoelectric Generator (RTG) Safety Analysis" on February 15, 1985 (Salmon, 2005, p. 127). But this figure was deceptive. An estimate of the reliability of the SRBs was made by the range safety officer on the basis of past performance. As 121 out of 2,900 flights failed, the probability of failure was approximately 1 in 25. But this included rockets that were flown for the first few times, and for the mature rockets, a figure of 1 in 50 might have been more reasonable. Furthermore, if parts would be carefully selected and inspected, even a figure of 1 in 100 might have been achievable. A figure of 1 in 1,000 was probably not achievable at the time. But this conflicted with NASA management's opinion on the matter. NASA officials estimated that the probability of failure with loss of vehicle and of human life was much lower than 1 in 100—the lowest estimate being 1 in 100,000. They argued that the high figures were based on past performance of unmanned rockets, and that a distinction had to be made between manned space flight programs and unmanned programs. The argument was that because manned programs had an extremely high degree of mission success, standard statistical methods could not be used to determine the probability of mission failure, and therefore they had to rely on engineering judgment, and not on numerical probability usage. An obvious inference was then that the figure of 1 in 100,000 was supported by the judgment of working engineers. But this is false; in fact, working engineers estimated the probability of mission failure much higher (Feynman, 1986, p. F1). Hence, NASA's report on the probability of mission failure was deceptive [conditions (i) and (ii)(b) are met].

7. AN EXPLANATION

As we were able to identify some deceptive statements resulting from different practices in the Space Shuttle Program, we can conclude that the epistemic integrity of these practices was limited—epistemic integrity could, at least in theory, be higher, that is, if the practices did not result in any deceptive statements. But why was epistemic integrity limited? In this section, we try to explain this.

Let us start with a very brief sketch of the historical context in which the Space Shuttle Program was situated. It was a context of competition for scarce resources. The Space Shuttle Program was born in the aftermath of the Vietnam War, when spaceflight was no longer national priority and when NASA lost the budgetary certainty it had before. NASA personnel was reduced by 1,000 employees per year, and of the three projects that NASA planned—a mission to Mars, a space station in earth orbit, and a space shuttle to transport people and materials in space—it could only execute the space shuttle project, which aimed at the development of a reusable space shuttle that should reduce the cost of putting objects into orbit (Vaughan, 1996, pp. 18–19).

While the Apollo Program was justified by the desire to respond to the Soviet launch of Sputnik, the Space Shuttle Program became justified on the basis of cost-effectiveness. The Space Shuttle Program gained approval on the basis of a study by Mathematica, Inc., a think tank that NASA called in in 1971 to assess the program's cost-effectiveness. Mathematica reported that, given the Space Shuttle's payload capacity, it would pay for itself provided that there would be more than 30 flights each year. In light of this economic justification for the Space Shuttle Program, NASA had to maintain before Congress and the general public that the program was a good investment on economic grounds. But the estimates by Mathematica were overly optimistic—among others because it was based on data furnished by contractors hoping to receive shuttle contracts—and given workforce reductions at NASA, it became increasingly difficult for NASA to meet performance expectations. The gap between what NASA could do and what it was expected to do was widening, and this resulted in production pressure, i.e., the pressure to launch a certain amount of flights each year, more than was possible given NASA's means (Vaughan, 1996, pp. 19–32). Of course, this pressure also had an impact on NASA contractors, since they had to meet NASA requirements in order to maintain contracts.

Several analysts of the Challenger disaster refer to production pressure in explaining this event. For instance, the U.S. House Committee on Science and Technology made the following statements:

> The Committee found that NASA's drive to achieve a launch schedule of 24 flights per year created pressure throughout the agency that directly contributed to unsafe launch operations. (House Committee, 1986, p. 3)

> There is no doubt that operating pressures created an atmosphere which allowed the accident on 51-L to happen. Without operating pressures the program might have been stopped months before the accident to redesign or at least understand the SRB joint. Without operating pressure the flight could have been stopped the night of January 27. (House Committee, 1986, p. 123)

But, as we mentioned earlier in this section, our aim is not to explain the Challenger disaster, but to explain why the epistemic integrity of practices in the Space Shuttle Program was limited. We think this can be explained along similar lines: production pressure compromised epistemic integrity.

Production pressure operated in two ways.[8] Firstly, it made it more difficult to communicate claims that would slow down the process of getting a shuttle ready for launch, while it promoted the production of claims that would accelerate this process. Take, for instance, the incomplete rationale for accepting risk associated with O-ring erosion, presented by Mulloy to Level I, which was deceptive because it did not include Thiokol engineers' temperature concerns. We have seen that the appropriate response to these concerns would have been to collect systematic data on the association between temperature and O-ring erosion. This would, however, have taken time, and a more time-saving option was to simply omit these concerns. This could explain why temperature concerns were omitted, and why a deceptive rationale was presented to Level I. Similar explanations could be offered for other deceptive statements discussed in Sections 4 to 6: the claim that the problem of O-ring erosion on the seventeenth shuttle mission only concerned the nozzle joints (and not the field joints) was time-saving because it meant that the possibility of field joint failure did not have to be addressed, telling William Lucas that Thiokol and Marshall's most knowledgeable people agreed with the recommendation to launch was time-saving because it made launch delay less likely, declaring the Space Shuttle Program operational was time-saving because an operational system requires less testing and fewer procedural constraints (Vaughan, 1996, p. 125), and so on.

It should be remarked that we do not claim that deceptive statements were produced *intentionally*, that those who produced deceptive statements knowingly deceived their audiences in order to meet the launch schedule. This is possible, but deception could also be unintended. Consider Roger Boisjoly's claim on the eve of the Challenger accident that he did not have the data to quantify Thiokol's temperature concerns. It is possible that Boisjoly knew that he *did* have such data and that he lied, but a more plausible assumption is that Boisjoly did not know he had such data, that he simply failed to see how the data that were available to him could be used to construct a quantitative argument for the temperature concerns. But even in the latter case, production pressure could explain why Boisjoly made a deceptive claim. Due to the pressure to launch, contractors, such as Thiokol, and their employees did not want to be the organization/person that was responsible for postponing launch. This applied especially to Thiokol on the eve of the Challenger launch, since a teleconference to discuss Thiokol's one-billion-dollar contract was scheduled for January 28, after the Challenger lift-off (Charles, 1989, p. 118). Now, if Boisjoly would have been able to develop a quantitative argument for the temperature concerns, launch would have probably been delayed, and he and Thiokol would be responsible for this. Of course, being responsible for launch delay was still preferable to being responsible for mission failure. But this was not what was expected; even Boisjoly seemed to expect that the mission would return (even though he was aware of the risk involved) (Vaughan, 1996, p. 380). Moreover, if the mission would fail, as it did, Boisjoly would not be seen as the person who was responsible for this, but rather as the hero who tried to prevent the tragedy.[9] These circumstances might have inhibited Boisjoly's creativity in developing arguments that would actually convince managers to delay launch. If there would be no production pressure, so that the expected consequences of causing launch delay would not be that bad, Boisjoly might have been more inventive in arguing that cold temperatures could cause O-ring erosion, and he might have seen how the available data could be translated into a quantitative argument.

A second way in which production pressure operated was by implying that those involved did not have the time necessary to make sure that their claims were correct and that their audiences were not deceived. If Boisjoly had more time to prepare his argument for the temperature concerns—after the first teleconference, Thiokol only had about one hour and a half to prepare a formal presentation of conclusions and launch recommendation (Vaughan, 1996, pp. 287–288)—he would have probably been able to construct a quantitative argument on the basis of the data available, and he would not have made the deceptive claim that he did not have the data to construct such an argument. Another example is the chart that stated that at 50°F the time to recover was 600 seconds, while in fact the O-ring did not re-establish contact at all at 50°F. This error could be explained by time pressure: the chart was created by Arnie Thompson between the two teleconferences, when there was, as we have seen, pressure by an unreasonable deadline. Such pressure made errors such as the one under consideration more likely, while detection and correction of errors became less likely.

8. POLICY SUGGESTIONS

Now, let us draw some policy-related lessons from this case study. How could the epistemic integrity of practices in aerospace science be promoted? In light of the above considerations, an obvious strategy is to reduce production pressure. One way to do this is for the government to base funding decisions on how valuable it considers aerospace science, and on how much resources it is willing to allocate to it, without focusing too strictly on cost-effectiveness; the justification for funding space agencies such as NASA should not be primarily based on their cost-effectiveness or the assumption that their activities will pay for themselves. Space agencies, too, should not focus too strictly on cost-effectiveness in decisions on the allocation and renewal of contracts. Those involved in research can then communicate statements that slow down the process of getting a shuttle ready for launch without the risk of losing funding, and they can take the time needed to adequately and accurately perform their tasks.

However, cost-effectiveness considerations should and can not be entirely omitted from aerospace science. Resources should not be wasted, and therefore, it is important that space agencies and their contractors proceed efficiently. When they do not, this should not be without consequences; those who are responsible for the waste of resources could, for instance, be replaced. If inefficiency would not have any undesirable consequences, the danger exists that no goals are accomplished, while huge amounts of resources are drained. Also note that the more efficiently an agency such as NASA proceeds, the easier it is for the government to justify government funding of that agency to the public.

With respect to cost-effectiveness considerations, we have two suggestions. First, in considering cost-effectiveness, one should not rely too heavily on data furnished by contractors hoping to receive shuttle contracts—such data should always be taken with a grain of salt. We have seen that performance expectations of the Space Shuttle Program were shaped by a cost-effectiveness study that was based on such data. As a consequence, performance expectations were overly optimistic and could not be met, resulting in a pressure to launch more flights than was reasonable given NASA's means. If the data under consideration would have been less crucial in the justification of the Space Shuttle Program, this pressure, which undermined epistemic integrity (see Section 7), might have been avoided.

Secondly, we suggest that cost-effectiveness should not be defined solely in terms of flights per year, but also in terms of safety. The goal is to have as much flights as possible, *and that are safe as possible*. Therefore, we consider it crucial that both the government and NASA leadership stress the importance of flight safety,[10] and that at all levels of decision-making, cost-effectiveness assessments take safety into account as well. Fewer flights would then not necessarily imply lower cost-effectiveness; if the flights are significantly safer, cost-effectiveness could even be higher. Making pro-launch statements that undermine flight safety would no longer be promoted—such statements would not imply higher cost-effectiveness—and making contra-launch statements that result in higher flight safety would probably be a lot easier. The current system does not stimulate scientists and other actors to push through on safety concerns and say, e.g., "I think we should not launch," or "more safety testing is required" (see note 9): this would make them responsible for launch delay and hence reduced cost-effectiveness, and possibly loss of funding/contract due to low cost-effectiveness. But if cost-effectiveness is, at all levels of decision-making, understood (partly) in terms of flight safety, delaying launch in order to improve flight safety would no longer endanger further funding or contract renewal, as it would not reduce cost-effectiveness. Note that including safety considerations in cost-effectiveness assessments would also stimulate scientists to take their time in order to avoid errors, as fewer errors usually implies higher safety.

A final suggestion is to involve independent outsiders, who are completely free from any pressure to launch, in aerospace research. Their task could, for instance, be to detect as much deceptive statements as they can find in the relevant research, and to replace them by non-deceptive statements. Outsiders may be able to see deception where this is difficult to see from an insider's point of view. This claim is supported by the fact that Alton Keel and Randy Kehrli, two non-engineers, were able to show that Thiokol's temperature concerns on the eve of the Challenger launch could easily be quantified on the basis of the data that were available at that time, and hence, that Boisjoly's claim that he did not have the data to quantify these concerns, was mistaken (see Section 4).

9. CONCLUSION

We have argued that production pressure at NASA, i.e., the pressure to launch an unreasonable amount of flights per year, compromised the epistemic integrity of the following practices in the Space Shuttle Program: (1) research performed by working engineers, (2) the review of the results of this research by middle-level managers, and (3) NASA's communication with the public. More specifically, production pressure caused these practices to result in deceptive statements by (1) making it easier for those involved to communicate pro-launch statements than to communicate contra-launch statements, and (2) causing those involved to have insufficient time to make sure that their claims were non-deceptive. Furthermore, we have proposed some potential strategies to protect the epistemic integrity of practices in aerospace science.

Our analysis of practices in the Space Shuttle Program differs from earlier analyses in two important ways. A first is that we were able to identify some deceptive claims that were not recognized as such in earlier analyses (e.g., no one seemed to realize that Arnie Thompson's chart, which stated that the time to recover was 600 seconds, was deceptive), and for those statements that were recognized as deceptive or misleading in earlier analyses, we clarified why exactly they were deceptive. For instance, Feynman showed that the term "safety factor" was used in a strange way. But those who stated that there was a safety factor of three might respond that they specified what they meant by this and that therefore, they did not mislead anyone. We were able to show why they did: they raised the impression that safety was similar to safety in other technological contexts in which there is a safety factor of three, and this was not the case.

A second difference is that our analysis is not based on the fact that an accident occurred. Most earlier analyses focus on a certain accident (the Challenger disaster or the Columbia disaster) and search for factors that contributed to this accident, which are then regarded as problematic. But things might have been different. If some environmental factors (e.g., temperature) were a bit different, the Challenger or Columbia disaster might not have occurred, even if NASA and its contractors would have proceeded exactly the same as they currently have. Perhaps NASA just had bad luck. The problem with accounts that evaluate practices in light of events such as the Challenger or Columbia disaster is then that the outcomes of such studies are heavily dependent upon contingent circumstances. If things had been different, and no accidents had occurred, such studies might have very different outcomes; they might conclude that NASA should be praised for its efficiency at getting shuttles in space and back on earth safely. What we did, is show that certain practices in the Space Shuttle Program were problematic—they had limited epistemic integrity—in a way that is independent of whether or not these practices resulted in certain accidents.

Finally, we would also like to suggest some routes for further research. Previous philosophical research reveals that the pressure on pharmaceutical companies to make a profit could compromise the epistemic integrity of biomedical research (Biddle, 2007, De Winter and Kosolosky, in press). We think this case study complements such research very well: it shows that the pressure on a space agency to launch could compromise the epistemic integrity of practices in aerospace science. But to obtain an accurate overall picture of how different kinds of pressure and performance goals affect the epistemic integrity of scientific practices, a lot more research is needed. We should investigate more cases from biomedical and aerospace science, and we should also scrutinize how commercial goals affect epistemic integrity in other domains of science than biomedical science. Another interesting question for further research is how the pressure to publish influences the epistemic integrity of academic research.[11] Such studies would in our opinion be very useful for policymaking in science.

ACKNOWLEDGMENTS

Jan De Winter is a Ph.D. fellow of the Research Foundation (FWO)—Flanders. Research for this paper by Laszlo Kosolosky was supported by subventions from the Research Foundation (FWO)—Flanders through research project G.0122.10. We are very grateful to Erik Weber and two anonymous reviewers for reviewing earlier versions of this paper.

NOTES

1. In detail, we conceive of this "new" discipline in the following manner: social epistemological research arose as an epistemological and philosophy of science reaction against the grown division between, on the one hand, analytic philosophy of science (as it became present after World War II) and, on the other hand, sociology of science (from the early 1970s). The difference between both disciplines is often perceived as a difference between a mainly normative and a mainly empirical discipline. Whereas philosophers of science (and epistemologists) focused more on grasping the right methodological rules that a single, rational scientist is to pursue, sociologists of science (and researchers within science studies in general) focus more on the description and explanation of the social history of science. Social epistemologists aspire to transcend this grown division. Social epistemology can be described by a number of characteristics, of which we align two here. First, social epistemologists emphasize the social or collective aspect of science and knowledge in general, as opposed to an individualistic approach in the traditional philosophy of science and epistemology; scientists accept claims as a result of interaction with, and mutual dependence of, others (and society in general). Methodological rules are always to comprise rules on how the social interaction between scientists should look like and how institutions should be shaped accordingly. Second, social epistemologists do not conclude that the social character of knowledge gaining is a source of bias or irrationality that would undermine or negatively influence the acceptance of true (or truth conforming) statements, this in contrast to a large audience within sociology of science. Social epistemologists regard the social dimension as constitutive for good knowledge and see it as their duty to sort out how the quest for knowledge should be organized—including its social (and institutional) dimensions. Social epistemologists do normative research, without hereby losing grip on the social dimension of knowledge. (This note is based on an unpublished research project proposal written by Jeroen Van Bouwel.)

2. David Resnik has been pursuing this type of research for a number of years. He has explored applied ethical and political issues in scientific practice from a social epistemology viewpoint (Resnik, 1996, 1998, 2007, 2009).

3. This notion should be distinguished from other notions of integrity, such as epistemic integrity of individuals, ethical integrity of practices, and ethical integrity of individuals. We leave the question of what the latter notions exactly stand for, and how the different notions exactly relate to each other to further research. Also see Resnik (1996, 1998).

4. For illustrations from other areas of research, see De Winter and Kosolosky (in press).

5. These two concepts were suggested by an anonymous reviewer of an earlier version of this paper.

6. See http://www.nasa.gov/

7. Thiokol's Boisjoly seemed to be aware of this. In an internal memo he sent to Robert Lund on July 31, 1985, he stated that the same scenario that resulted in the failure of the nozzle joint could also occur in a field joint (Presidential Commission, 1986, Vol. 1, pp. 249–250).

8. It should be noted that the pressure to produce is not unique to aerospace science. Similar pressures feature other areas of research. Think, for instance, of the pressure to produce results and the pressure to publish (Shamoo and Resnik, 2009).

9. Also see the report of the U.S. House Committee on Science and Technology: "[T]he present system permits [contractors] to "express concern" without actually saying, "stop the flight, it is unsafe". If the odds favor a successful flight they do not have to be responsible for cancelling, yet if the mission fails they are on record as having warned about potential dangers" (House Committee, 1986, p. 152).

10. The recommendation that the importance of safety should be stressed within the organization has also been offered in the field of health care; see Institute of Medicine (2000).

11. We should mention that there already is a lot of research on the effects of different kinds of pressure in science (e.g., Shamoo and Resnik, 2009). What we propose here more specifically is to use our concept of epistemic integrity to clarify whether, how, and why different kinds of pressure compromise epistemic integrity.

REFERENCES

Biddle, J. (2007). Lessons from the Vioxx debacle: What the privatization of science can teach us about social epistemology. *Social Epistemology* 21: 21–39.

Brown, J. R. (2008). The community of science®. In Carrier, M., and Howard, D. (Eds.) *The Challenge of the Social and the Pressure of Practice: Science and Values Revisited*. pp. 189–216. Pittsburgh: University of Pittsburgh Press.

Carnap, R. (1950). *Logical Foundation of Probability*. London: Routledge and Keegan Paul.

Charles, M. T. (1989). The last flight of Space Shuttle Challenger. In Rosenthal, U., Charles, M. T., and Hart, P. T. (Eds.) *Coping with Crises: The Management of Disaster, Riots and Terrorism*. pp. 110–122. Springfield: Charles C. Thomas Publisher, Ltd.

De Winter, J. (2012). How to make the research agenda in the health sciences less distorted. *Theoria* 27: 75–93.

De Winter, J., and Kosolosky, L. The epistemic integrity of scientific research. *Science and Engineering Ethics* (2012). doi:10.1007/s11948-012-9394-3.

Feynman, R. (1986). Personal observations on the reliability of the Shuttle. In Presidential Commission, *Report of the Presidential Commission on the Space Shuttle Challenger accident*. Vol. 2, Appendix F. Washington, DC: United States Government Printing Office.

Fuller, S. (1988). *Social Epistemology*. Bloomington, IN: Indiana University Press.

Fuller, S. (2000). *The Governance of Science: Ideology and the Future of the Open Society*. Buckingham: Open University Press.

Haack, S. (2007). The integrity of science: What it means, why it matters. In Conselho Nacional de Ética para as Ciências da Vida (Ed.), *Ética e Investigação nas Ciências da Vida – Actas do 10° Seminário do CNECV* (pp. 9–28). Lisbon: Presidência do Conselho de Ministros. Available at http://www.as.miami.edu/phi/haack/PORTUGAL.pdf. Last accessed June 4, 2012.

Hall, J. L. (2003). Columbia and Challenger: Organizational failure at NASA. *Space Policy*, 19:239–247.

Hollis, A., and Pogge, T. (2008). *The Health Impact Fund: Making new medicines accessible for all*. New Haven, CT: Incentives for Global Health.

House Committee. (1986). *Investigation of the Challenger accident: Report of the Committee on Science and Technology, House of Representatives, Ninety-ninth Congress, second session*. Washington, DC: United States Government Printing Office. Available at http://www.gpo.gov/fdsys/pkg/GPO-CRPT-99hrpt1016/pdf/GPO-CRPT-99hrpt1016.pdf. Last accessed September 25, 2012.

Institute of Medicine. (2000). To err is human: Building a safer health system. Washington, DC: National Academy Press.

Kitcher, P. (1993). *The Advancement of Science*. New York: Oxford University Press.

Kitcher, P. (2001). *Science, Truth, and Democracy*. New York: Oxford University Press.

Longino, H. (1990). *Science as Social Knowledge*. Princeton: Princeton University Press.

Longino, H. (2002). *The Fate of Knowledge*. Princeton: Princeton University Press.

National Academy of Sciences. (1992). *Responsible Science, Volume I: Ensuring the Integrity of the Research Process*. Washington, DC: National Academies Press.

Nickles, T. (1980). *Scientific Discovery, Logic, and Rationality*. Dordrecht, the Netherlands: Reidel.

Office of Research Integrity. (2007). *Research on research integrity*. Available at http://grants.nih.gov/grants/guide/rfa-files/RFA-RR-07-004.html. Last accessed October 25, 2011.

Petrovečki, M., and Scheetz, M. D. (2001). *Croatian Medical Journal* introduces culture, control, and the study of research integrity. *Croatian Medical Journal*, 42:7–13.

Presidential Commission. (1986). *Report of the Presidential Commission on the Space Shuttle Challenger accident*. Washington, DC: United States Government Printing Office. Available at http://history.nasa.gov/rogersrep/genindex.htm. Last accessed September 25, 2012.

Reiss, J. (2010). In favour of a Millian proposal to reform biomedical research. *Synthese*, 177:427–447.

Resnik, D. B. (1996). Social epistemology and the ethics of research. *Studies in the History and Philosophy of Science* 27:565–586.

Resnik, D. B. (1998). *The Ethics of Science*. New York: Routledge.

Resnik, D. B. (2007). *The Price of Truth*. New York: Oxford University Press.

Resnik, D. B. (2009). *Playing Politics with Science*. New York: Oxford University Press.

Salmon, W. C. (2005). *Reality and Rationality*. New York: Oxford University Press.

Shamoo, A., and Resnik, D. B. (2009). *Responsible Conduct of Research*, 2nd ed. New York: Oxford University Press.

Steneck, N. H. (2006). Fostering integrity in research: Definitions, current knowledge, and future directions. *Science and Engineering Ethics*, 12:53–74.

Vaughan, D. (1996). *The Challenger launch decision: Risky technology, culture, and deviance at NASA*. Chicago, IL: University of Chicago Press.

Winsor, D. A. (1988). Communication failures contributing to the Challenger accident: An example for technical communicators. *IEEE Transactions on Professional Communication*, 31:101–107.