



# Studying organisational cultures and their effects on safety <sup>☆</sup>

Andrew Hopkins <sup>\*</sup>

*Department of Sociology, in the Faculty of Arts, and National Research Centre for OHS Regulation,  
Australian National University, Australia*

Received 23 January 2006; received in revised form 20 April 2006; accepted 10 May 2006

---

## Abstract

How do organisational cultures influence safety? To answer this question requires a strategy for investigating organisational culture. By far the most widely used research strategy is the perception survey. An alternative is for researchers is to immerse themselves in one or more organisations, making detailed observations about activities and drawing inferences about the nature of the organisation's culture (the ethnographic method). A third technique makes use of the wealth of material that is assembled by inquiries into major accidents. This paper describes how this material can be used to provide insights into organisational cultures. It draws on specific examples from the author's own work as well as the cultural analysis carried out by the Columbia Accident Investigation Board. It concludes with some additional suggestions for carrying out research on safety-relevant aspects of organisational culture.

© 2006 Elsevier Ltd. All rights reserved.

*Keywords:* Culture; Safety culture; Organisational culture; Major accident inquiries

---

## 1. Introduction

Despite all that has been written about safety culture, there is no agreement about just what this concept means. The titles of various journal articles allude to this uncertainty, for

---

<sup>☆</sup> Paper prepared for the International Conference on Occupational Risk Prevention, Seville, May 2006, and the 3rd International Conference, Working on Safety, The Netherlands, September 2006.

<sup>\*</sup> Tel.: +61 2 6494793; fax: +61 2 62490525.

*E-mail address:* [andrew.hopkins@anu.edu.au](mailto:andrew.hopkins@anu.edu.au)

example, “culture’s confusions” (Hale, 2000) and “safety culture: philosopher’s stone or man of straw?” (Cox and Flin, 1998). It is not the purpose of this article to outline these problems, but one crucial ambiguity should be noted. For some writers, every organisation has a safety culture of some sort, which can be described as strong or weak, positive or negative. For other writers, only an organisation which has an over-riding commitment to safety can be said to have a safety culture. On this view, relatively few organisations have safety cultures. As Reason puts it, “like a state of grace, a safety culture is something that is striven for but rarely attained” (Reason, 1997, p. 220; see also Parker et al., 2006). The fact that this very basic contradiction has not been resolved is but one indicator of the confusion which surrounds the use of this term.

This paper bypasses some of the confusion by shifting attention to the concept of organisational culture. This is a more general and in some respects a clearer concept. Every organisation has a culture (or perhaps a series of subcultures) and that culture can be expected to impact on safety. Understanding how this happens can provide insights into ways organisational cultures need to be modified to give a higher priority to safety. Distinguishing safety from culture in this way breaks the safety research task into two components – the study of organisational cultures and the study of the impact of these cultures on safety. This paper is largely concerned with the first of these – the study of organisational culture. This is in a sense a prior question; we must first identify and describe various elements of the culture of an organisation before we can evaluate their impact on safety.

Although clearer than safety culture, organisational culture has itself been defined in a variety of ways. Schein provides a useful summary of the way the concept of culture has been used by various writers: observed behavioural regularities, group norms, espoused values, formal philosophy, rules of the game, climate, embedded skills, habits of thinking, shared meanings and root metaphors (Schein, 1992, p. 8–9). It will be noted that some of these usages focus on values and attitudes as the key element of culture, while others stress behaviour. Cooper sees this as the crucial distinction. “The main difference between such definitions (he says) appears to reside in their focus on the way people think, or on the way people behave” (2000:112).

Perhaps the best known definition of organisational culture, “the way we do things around here” (Deal and Kenney, 1982, p. 4,59,60), is clearly behaviour focussed. Schein himself has at times referred to organisational culture simply as “the way we do things around here” (Schein, 1992, p. 8–9), although his formal definition is more complex. Moreover Hofstede, after discussing whether it is better to focus on values or practices in defining organisational culture, concludes that “shared perceptions of daily practices should be considered to be the core of an organisation’s culture” (Hofstede, 1997, p. 182–3).

It should be stressed that definitions in terms of values, on the one hand, and practices, on the other, are not necessarily in conflict in the way that the alternative definitions of safety culture are. It is simply a question of emphasis. “The way we do things around here” carries with it the connotation that this is the right, or appropriate or accepted way to do things. These judgments stem necessarily from shared assumptions or values. To reiterate, a definition in terms of practices does not deny the importance of values in any complete understanding of culture.

A further preliminary matter that must be mentioned is the relationship of culture to climate. Various writers argue that there are conceptual differences. For instance it is said that climate is a manifestation of culture (Mearns et al., 2003), that climate is directly measurable while culture is too abstract to be measured directly (Mearns et al., 2003), and that

climate refers to the situation at a particular point in time while culture refers to more enduring phenomena (Hale, 2000, p. 5). Some writers see climate as referring to attitudes and culture to behaviour (Guldenmund, 2000, p. 221).

However, even writers who argue for a conceptual distinction recognize that those carrying out empirical research on safety climate or safety culture have been unable or unwilling to make these distinctions and that in practice the terms appear to be used interchangeably. Guldenmund (2000, p. 220) suggests that the difference between climate and culture may be little more than terminological fashion. He notes that in 1970s people doing research on these topics tended to use the term climate but that gradually, during the 1980s, climate was replaced by culture.

It is also observed that the terms have arisen in different academic disciplines: climate in social psychology and culture in anthropology. As a result they tend to be associated with the different research strategies of these disciplines, respectively, questionnaire research and ethnography. This, however, does not amount to a conceptual distinction between the two ideas, for there is no reason why a single concept cannot be studied using a variety of research strategies. For present purposes, therefore, I shall treat organisational climate and organisational culture as synonymous.

While the distinction between culture and climate remains elusive, what is clear is that there are real choices to be made in terms of research strategy. While a single concept can in principle be investigated using more than one research method, those methods have different strengths and weaknesses and can be expected to yield different insights. I turn, therefore, to this question of research strategies.

## 2. Culture surveys

The survey method appears to be the predominant strategy for studying organisational cultures and their effect on safety. There are numerous reports in the literature of safety culture or safety climate surveys. Questionnaires are administered to individuals and the data aggregated to the workplace or organisational level. This facilitates statements about organisations and comparisons of one organisation with another.

The survey method is well suited to studying individual attitudes and values and it might be thought that the method is thereby biased in favour of a definition of culture in these terms. However, the survey method is equally suited to studying practices, or “the way we do things around here”. The only qualification is that survey research of “the way we do things around” here necessarily measures people’s perceptions rather than what actually happens, which may not necessarily coincide. Hofstede (1997, p. 188) provides an excellent example of a study of organisational practices that uses the survey method. He asked respondents in twenty different organisations to indicate the extent to which they agreed with statements such as: “Where I work, meeting times are kept very punctually”; and “Where I work, quantity prevails over quality”. A total of 61 such questions were asked. The results were then factor analysed and six mutually independent dimensions were identified:

- process oriented vs results oriented;
- employee oriented vs job oriented;
- parochial vs professional;
- open system vs closed system;

loose control vs tight control;  
normative vs pragmatic.

The twenty organisational units studied were then able to be compared in terms of these six cultural dimensions. Hofstede does not correlate these variations with safety performance but one can easily envisage a study that does.

This example demonstrates a couple of related points. First, it shows that surveys can be used to study organisational practices, as well as attitudes. Second, as noted above, scholars seeking to identify a difference between organisational culture and organisational climate sometimes argue that the difference lies in the research methods used. [Guldenmund \(2000, p. 220–1\)](#) for instance argues that organisational climate studies typically use self-administered questionnaires that are analysed in quantitative fashion, while studies of organisational culture are typically carried out by participant observation and do not lead to quantitative results. Hofstede's work shows that this is not an invariable rule, for it is explicitly a study of organisational culture yet is just as quantitative and mathematically sophisticated as research which purports to investigate organisational climate.

The survey method does however have its drawbacks. It provides a relatively superficial description of the culture of an organisation. Many practices are too complex to be meaningfully described in the words of a survey question. Moreover the survey method tells us very little about dynamic processes – how the organisation goes about solving its problems. This is an important limitation. Consider, for a moment, [Westrum's](#) definition of organisational culture as “the organisation's pattern of response to the problems and opportunities it encounters” ([Westrum, 2004, p. ii22](#)). A moment's thought will reveal that this is not inconsistent with the definitions of culture provided above, but it does emphasise a dynamic element of culture that surveys have great difficulty capturing. [Schein \(1992, p. 187\)](#) makes a similar point when he notes that members of a culture are most likely to reveal themselves when they have problems to solve and he says that research findings “will often be analyses of how the problems were solved not simply pure description of cultures”. This is where the ethnographic method has much to contribute.

### 3. Ethnography

Ethnographic research originates in the discipline of anthropology where researchers immerse themselves for long periods in the culture of interest and provide detailed qualitative descriptions of what they observe ([Geertz, 1973](#)). Anthropologists have usually studied small scale non-industrial societies. Sociology is distinguished from anthropology principally by its focus on contemporary industrial society and it uses many of the same methods and theories as anthropology. In particular, sociology has used the ethnographic method to study organisations and their cultures (e.g. [Gouldner, 1955](#); [Kriegler, 1980](#); [Bourrier, 1998](#); [Westrum, 1999](#)). The ethnographic study of an organisation will not normally require the researcher to live within the organisation (except perhaps for what sociologists call total institutions, such as the military, prisons and mental hospitals, [Goffman, 1968](#)), but it does require the researcher to spend a great deal of time in the organisation, either as a participant observer, for instance an employee, or as non-participant observer of some sort.

Of course it is not only sociologists who do this kind of research. Some of the most important ethnographic research on organisational culture and its impact on safety, the

research on high reliability organisations, has been done by political scientists, among others. They have spent time on aircraft carriers, in nuclear power stations and in air traffic control rooms identifying the elements of the cultures of these organisations that contribute to reliable operations (LaPorte and Consolini, 1991; Weick and Sutcliffe, 2001).

This approach has also been championed by Schein, originally trained as a psychologist. He makes an interesting distinction between the ethnographic model and what he calls the clinician/consultant model. Ethnographers seek to minimize their impact on the culture under investigation. Organisational clinicians or consultants, on the other hand, are called in to solve a problem and to bring about some change. Schein, an organisational consultant, is understandably an advocate of the second model. The fundamental assumption, he says, is that “one can understand a system best by trying to change it ... Some of the most important things I learned about (company) cultures surfaced only as reactions to some of my interventionists efforts” (1992, p. 30). Schein admits however that the distinction between his two models breaks down in practice. Even the most scrupulous ethnographer cannot help but affect the culture under study. “One’s very presence is an intervention”, he says (1992, p. 186). Both models therefore are interventionist. For this reason I shall ignore Schein’s distinction here and describe his research as ethnographic (see also Cox and Flin, 1998, p. 194; Glendon and Stanton, 2000; Mearns and Flin, 1999).<sup>1</sup>

While ethnographic research provides a much richer account of the culture of an organisation than surveys can, a question arises as to the validity of the description provided. How can we be sure that the elements of culture identified by the researcher are correct? Schein provides an answer to this question. “The accuracy of the depiction is ... judged by the credibility of that description to insiders who live in the culture and, at the same time, to outsiders who are trying to understand it” (1992, p. 187). Most importantly, if the members of the culture recognize the description and agree that that is how it is, then the researcher can feel reasonably confident in the findings.<sup>2</sup> This does not mean, of course, that the ethnographer’s description will be the same as that provided by members of the culture. The ethnographer’s description is designed to speak to the wider concerns of his or her discipline.<sup>3</sup> However, if members of the culture fail to recognize the ethnographer’s description of their culture, that description must be called into question.<sup>4</sup>

One draw back of the ethnographic method is the commitment of time it requires from the researcher. It may well be that the only piece of ethnographic work ever carried out by

---

<sup>1</sup> In a sense the two models mentioned are at opposite ends of a continuum. Participant observation of an organisation (e.g. as an employee) lies somewhere in between. As a participant, the researcher inevitably influences the organisation, but the organisational culture is usually too embedded to be significantly distorted by influences of this sort. Studies at any point along this continuum can reasonably be described as ethnographies, and indeed are. They are all based on detailed observation by someone who is immersed in the organisation for long periods.

<sup>2</sup> Schein (1992, p. 187) also takes a swipe at the traditional social psychological paradigm which he says “leads to more of an illusion of objectivity than actual objectivity. It feels safe to argue that culture is simply that which has been operationally defined as culture but that approach may lead to conclusions that have very little to do with what actually goes on in organisations. To satisfy certain canons of ‘normal science’, we end up sacrificing credibility and validity”.

<sup>3</sup> An anonymous reviewer provided the following vivid comment: “The ethnographic description is not a simple mirror image. It is a back and forth effort to take into account how people see things, and how they can actually be interpreted in a larger framework. One needs at the same time to be ready to “go native” and then back! This is the ambivalent posture of the researcher using ethnographic method.”

<sup>4</sup> Margaret Mead’s informants subsequently denied the reality of her description of coming of age in Samoa, and later anthropological research showed that she was indeed wrong (Freeman, 1983).

a researcher is a PhD. Fortunately, there is another way in which researchers can gain access to rich bodies of material about the way organisations operate, without engaging in ethnographic field work.

#### 4. Major accident inquiries

Major accidents, such as rail crashes, space shuttle disasters and petrochemical plant explosions where many lives are lost, often give rise to multi-million dollar inquiries. These are a priceless source of information about organisational cultures and the way they impact on safety. These inquiries may sit for many days taking evidence from a large number of people. Inquiry panel members or counsel assisting the inquiry may question individual witnesses for hours. Questioners may pursue numerous lines of inquiry, probing, looking for things that might have been overlooked, exploring inconsistencies and conflicts of evidence, day after day. Proceedings are taped and many thousands of pages of transcript evidence are generated. This is far more material than an individual researcher engaged in an intensive interview process could ever produce. Moreover, the fact that witnesses can be required to give evidence to these inquiries and that witnesses can be interrogated in quite hostile fashion means that inquiries can gain access to information that no interviewer could ever hope to uncover.

Inquiry reports normally make use of a relatively small amount of this material. Moreover, their primary concern is to explain the event in question and only secondarily, if at all, to describe the culture of the organisation. As a result the transcript provides the researcher with a relatively unexplored data source, a ready made set of interview transcripts which can be mined for insights about the culture of the organisation. The material is often so rich and so diverse that researchers reading it are effectively immersed in the world of that organisation without ever setting foot on its premises. Their research amounts to armchair ethnography.

It is of course not possible to provide a complete picture of the culture of an organisation using this method. Schein (1992, p. 28) argues that it is never possible to describe an entire culture. What is possible is to identify what he describes as “elements of the culture”, in more concrete terms, groups of practices that hang together in some way.

At least two classic studies of major accidents come close to the research method advocated here, but differ in important respects. Turner’s study of several man-made disasters allowed him to identify a number of practices which amount to a culture of risk denial (Pidgeon, 1998, p. 206; Turner and Pidgeon, 1997; Hopkins, 2005, p. 19–22). But his study is based on official reports alone and not on the more detailed proceedings of the inquiries. A second classic is Vaughan’s study (1996) of the Challenger space shuttle disaster in 1986. This was quite explicitly a study of organisational culture and how it contributed to disaster. She identifies various cultural elements including: the normalization of deviance, the culture of production and structural secrecy, and she shows how these contributed to the outcome. Her study makes use of inquiry transcripts, but more than this, it draws on numerous documents created before the event, as well as her own interviews. She describes it as an historical ethnography (1996, p. 61). “My work (she says) was in harmony with the work of many social historians and anthropologists who use documents to examine how cultures shape ways of thinking” (1996, p. 61). Her very extensive sources meant that she was able to study NASA’s culture with an intensity that is normally only possible when the researcher works or even lives within the culture of interest. Valuable though her method

is, it comes at a price: Vaughan's study took her nine years to complete! This is not a price that many researchers can afford to pay.

## 5. The accident analysis method in action

The potential of the research method advocated above is illustrated by two studies of organisational cultures that I have carried out (Hopkins, 2005, p. 23–109). The first of these was a study of the culture of the rail organisations in the Australian state of New South Wales. It was based on the official inquiry into the Glenbrook rail crash. After reading the transcript and inquiry report I identified four main cultural themes that had contributed to the accident in fairly obvious ways. First, the railways were obsessively rule focused, in ways that hindered safety. Second, the railway system was organisationally and occupationally fragmented, resulting in a culture of silos.<sup>5</sup> Third, there was a powerful culture of punctuality – on-time-running. Finally, the culture was risk-blind, even risk-denying. The last of these elements is very similar to an aspect of organisational culture identified by both Turner and Vaughan. Turner's work was largely about how it is that organisations develop cultures of risk denial, while Vaughan talks about NASA's culture as a "way of seeing that is simultaneously a way of not seeing" (1996, p. 394). These are alternative ways of talking about risk-blindness. A diagrammatic representation of the way these cultural elements gave rise to the accident is provided in Fig. 1. The diagram shows just how these four cultural elements led to the outcome. The details are not relevant here; the diagram is included to indicate the importance of the four cultural elements identified in the analysis in explaining this accident.

The second study was an official inquiry in which I was a member of the Board. The Inquiry was set up to investigate why the Australian Air Force had allowed hundreds of aircraft maintenance workers to be poisoned by exposure to toxic chemicals over a 20 year period. I was in effect the research member of the Board and I wrote the report so as to highlight the way Air Force values had contributed to this outcome. The following five values were identified. First, higher priority was afforded to operations than to logistical matters such as maintenance. This was the Air Force equivalent of valuing production ahead of safety (Vaughan's culture of production). Second, there was a "can do" attitude, according to which nothing must be allowed to stand in the way of getting the job done. (Vaughan talks about the can-do culture at NASA, 1996, p. 209) Third, the Air Force valued its aircraft more than its people. Fourth, like many airlines, the Air Force valued air safety more highly than ground safety. Finally, the command and discipline system, to which all defence forces attach great importance, meant that affected workers had no way of objecting to their fate. Fig. 2 represents the causal analysis carried out by the Board of Inquiry. Again, it is included here to demonstrate the centrality of Air Force values in the analysis.

The question arises as to how the various cultural themes were identified in these two pieces of work. In some cases they were themes that were stressed by witnesses to such an extent that they could not be ignored. In the rail case, the focus on making rules and blaming rule violators for accidents was referred to constantly, as was the fact that on-time-running took precedence over safety. In the Air Force matter, only one of the five values was

---

<sup>5</sup> Grain silos in rural areas are isolated from one another. The metaphor refers to the way parts of an organisation can be similarly disconnected.

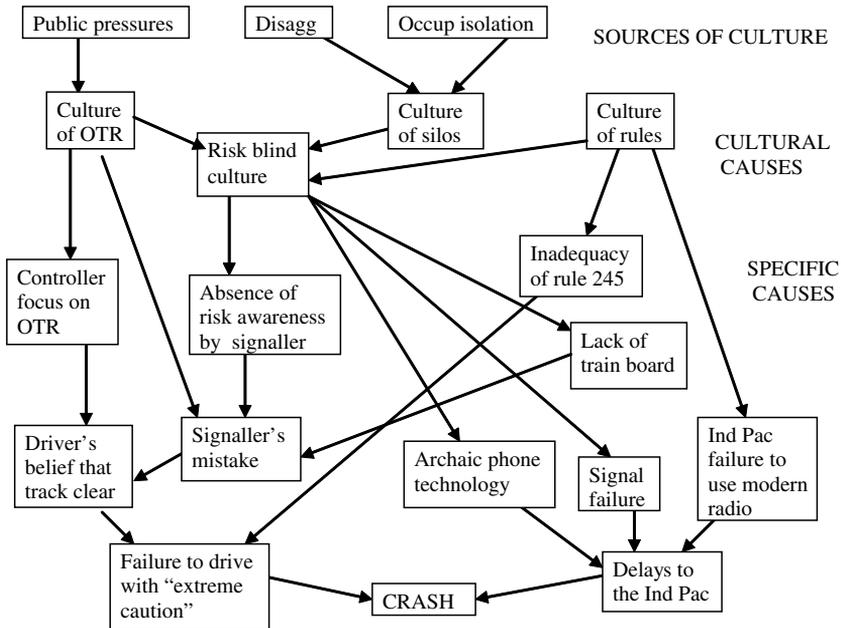


Fig. 1. Causes of the Glenbrook crash (Source: Hopkins, 2005, p. 77).

explicit in this way – the “can-do” approach to work. Where cultural elements were not referred to explicitly by participants, they were inferred from the stories told and the practices described in testimony. For instance, in the rail accident there was no explicit mention by witnesses of a culture of silos, but there were numerous comments about organisational fragmentation and occupational isolation and their effect on safety. The culture of silos seemed a natural way to draw these ideas together, especially given that the phrase “silo mentality” has quite wide currency.

Given that many of these cultural elements are research constructs, not explicitly identified by witnesses, how is their validity to be established? How confident can we be in these particular ways of grouping and describing the practices revealed in the inquiry? Schein addresses this question in relation to elements of culture identified using the ethnographic method, and his approach is applicable here. One of his criteria is the credibility of the description to those who live within the culture. The Air Force Board of Inquiry had a built-in guarantee of credibility from this point of view. One of the Board members was a senior Air Force officer. In putting his signature to the report he was effectively certifying the validity of the cultural elements identified. Here is an example of how it worked. A number of witnesses at the Inquiry spoke about the attention paid to air safety. A systematic pattern appeared to be emerging of a culture focused on air safety but lacking any focus on the safety of maintenance workers. I proposed this to the Air Force member of the Board and he recognized it immediately. The priority of air safety over ground safety was so ingrained in the Air Force, he said, that he was barely conscious of it. Schein notes that deep, taken-for-granted assumptions of this nature are the bedrock of culture (1992, p. 21–7). The priority of air over ground safety was clearly a bedrock element of air force culture in this sense. All other elements of the Air Force culture identified above were similarly validated.

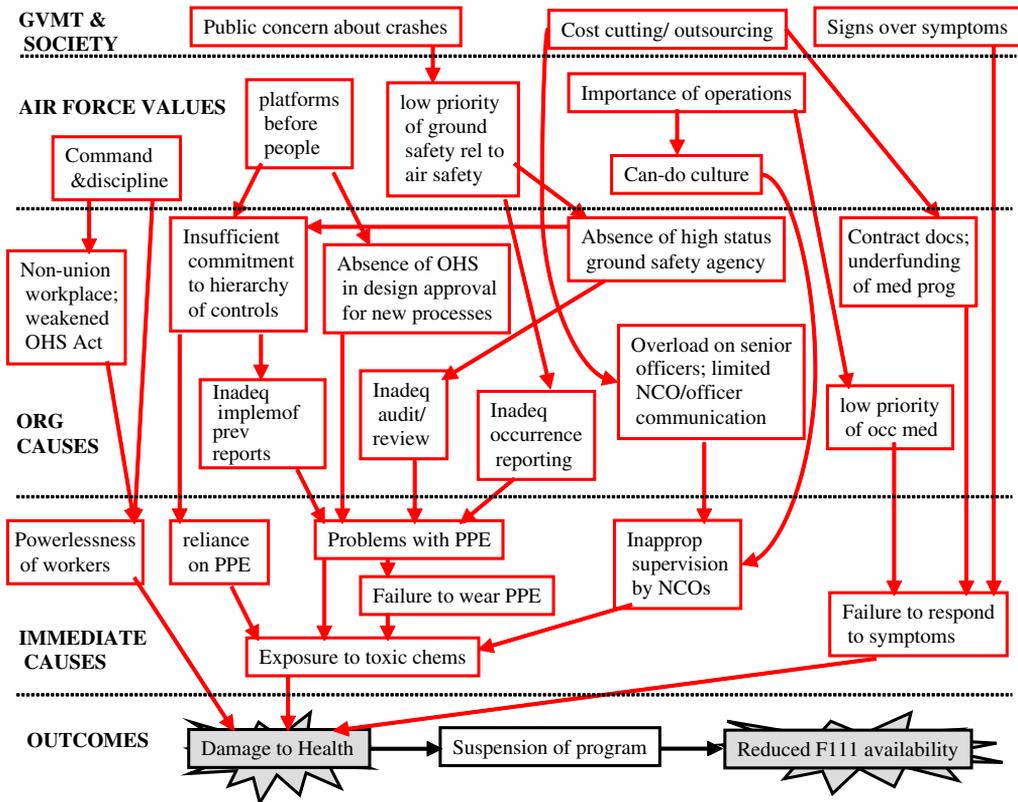


Fig. 2. Causes of damage to Air Force workers' health (Source: Hopkins, 2005, p. 109).

My rail research was not integral to the Inquiry in this way. It was done quite independently and with no access to insiders who might validate the cultural elements that were not explicitly identified by witnesses but that I had inferred. Schein's second criterion of validity is relevant in these circumstances. If the description makes sense to outsiders trying to understand the culture, this can be taken as evidence of the validity of the description. Outsiders are likely to see the cultural elements as making sense if the descriptions draw together the organisational practices identified in the inquiry in a natural way and if these descriptions are recognisable or at least plausible based on other organisations of which outsiders have some knowledge. From this point of view it is ultimately the reader of the research account who makes the judgment as to validity. For instance, readers familiar with Vaughan's comment about the NASA culture as providing a way of not seeing, and with other research on cultures of denial, will find the description of the rail culture as risk-blind, or risk-denying, to be quite credible. The idea of a culture of silos will also strike a chord with many readers familiar with management literature. Indeed, the various cultural elements identified were named as far as possible so as to facilitate this recognition. As it happens, one of the features of high reliability organisations, sensitivity to operations, is the inverse of the silo phenomenon. Sensitivity to operations means that front line workers strive to maintain as complete a picture as possible of what is going on at any one time. They do not focus merely on one aspect of their job, perhaps the most immediate aspect, but seek

to be aware of all aspects of the operation. This is just the opposite of the limited awareness demonstrated by rail workers. The fact that the silo culture turns out to be the converse of one of the elements of the culture of high reliability adds further credence to this description.

It must be acknowledged that cultural descriptions are inevitably subjective, to some extent, and that validations by outsiders, or even insiders, cannot entirely remove that element of subjectivity. Leading anthropologist, Clifford Geertz, takes the view that ethnographic descriptions are “essentially contestable”. He comments:

“Anthropology, or at least interpretative anthropology, is a science whose progress is marked less by a perfection of consensus than by a refinement of debate. What gets better is the precision with which we vex each other” (Geertz, 1973, p. 29)

Geertz provides a powerful, if rhetorical, justification of this approach:

“It is not against a body of uninterpreted data, radically thinned descriptions, that we must measure the cogency of our explications, but against the power of the scientific imagination to bring us into touch with the lives of strangers. It is not worth it, as Thoreau said, to go round the world to count cats in Zanzibar” (1973, p. 16)

It should be noted that identifying cultural elements that have contributed to an accident involves no automatic judgments about their desirability; there is no presumption that these values are necessarily a bad thing. For instance, a culture of on-time-running is clearly desirable for rail organisations; it is only when it takes priority over safety that it becomes problematic. Similarly, both a can-do attitude and a command system in which orders are obeyed are appropriate for a military organisation, particularly during military operations. It is obviously less appropriate in peace-time, non-operational settings. In short, identifying elements of organisational culture as contributing to certain undesired outcomes does not automatically imply that these values need to be changed. In some situations it may be enough to put in place mechanisms to ensure that such features of the organisational culture do not operate at the expense of safety.

## **6. The Columbia space shuttle accident report**

The Australia Air Force Inquiry was simultaneously an investigation into what went wrong and at the same time an explicit inquiry into an organisational culture. The cultural elements emerged from the evidence in a relatively unstructured way. An alternative strategy in identifying the culture of the Air Force would have been to compare it in some more systematic way against one or more existing models of organisational cultures. The investigation of the Columbia space shuttle tragedy of 2003 provides an excellent example (CAIB, 2003).

The Columbia Board included academics with a research interest in organisational cultures and its inquiry was consequently a research project on the organisational culture of NASA. The research design involved comparing the NASA culture revealed in the inquiry process with other theoretical models. In the Board's words,

“To develop a thorough understanding of accident causes and risk, and to better understand the chain of events that led to the Columbia accident, the Board turned to the contemporary social science literature on accidents and risk and sought insights from experts in High Reliability, Normal Accident, and Organisational Theory...

Insight from each figured prominently in the Board's deliberations... The Board selected certain well-known traits from these models to use as a yardstick to assess the Space Shuttle Program, and found them particularly useful in shaping its views on whether NASA's current organisation ... is appropriate" (CAIB, 2003, p. 180).

In fact the insights on which the Board ultimately relied came almost exclusively from High Reliability Theory. It selected the following features of high reliability organisations (HROs) to be used as a standard of comparison:

- Commitment to a safety culture
- Ability to operate in both a centralized and decentralized manner
- Recognition of the importance of communication
- Avoidance of over-simplification
- A wariness of success
- A commitment to redundant systems

There is no need to describe these features of high reliability organisations here. Suffice it to say that the Board found NASA lacking with respect to each of these cultural elements.

The Board also studied some real examples of HROs, specifically the US Navy's submarine and nuclear reactor safety programs (CAIB, 2003, p. 182–4). It identified how these manifested the theoretical characteristics of HROs and again showed how NASA fell short.

The Columbia inquiry was thus a theoretically driven inquiry into the culture of NASA. It is a very persuasive account, but it is not the only possible account of NASA's organisational culture. Independent researchers who study the same material, but use different organisational models, could come up with different descriptions. Consider, for example, the models used by Westrum. He argues that "the most critical issue for organisational safety is the flow of information" (Westrum, 2004, p. 22). He therefore divides organisational cultures into three types, on the basis of how they process information: the pathological, the bureaucratic and the generative (2004:23. Hudson extends this into a five-category classification: pathological, calculative, bureaucratic, proactive, generative – Parker et al., 2006). Without going into details here, one can imagine that NASA might turn out to be a combination of the pathological and bureaucratic. Interestingly, Vaughan describes the culture of NASA at the time to the Challenger disaster seventeen years earlier, as "bureaupathological" (1996:210). The point is simply that an inquiry driven by different theoretical concerns would generate a different cultural description. However, given that they are based on the same data, these various descriptions will presumably be consistent.

## 7. The impact of organisational culture on safety

The earlier mentioned methods of investigating organisational culture – the survey method and the ethnographic method – are carried out independently of any accident or accident investigation. The question of whether or how the cultures so identified impact on safety is a separate question. Mearns and co-workers argue that there is some, though rather limited, evidence that organisations which do well in safety climate surveys actually have fewer accidents (Mearns et al., 2003). The ethnographic method is not well suited to producing evidence about the impact on safety of the cultural elements it identifies. In the

absence of an accident, the ethnographic method can only speculate or hypothesise about the impact of organisational culture on safety.

Studying organisational cultures in the context of a major accident overcomes this difficulty. The links between the cultural elements and the outcome can usually be demonstrated. The formal logic involves counterfactual comparisons. If it can be concluded that in the absence of the cultural element concerned, the accident would not have occurred, or probably would not have occurred, then the cultural element can be counted as one of the causes of the accident. Figs. 1 and 2 represent this logic. Each arrow is a necessary condition, or “but-for” cause, in the following sense: but for the matter identified at the beginning of the arrow, the outcome at the end of the arrow would not have occurred. The resultant chains of causation show exactly how the cultural elements identified in the analysis gave rise to the unwanted outcome (For a development of this approach see Ladkin, 2001). The style of representation used here has been adapted from Rasmussen, who describes such diagrams as *accimaps* (Rasmussen, 1997).

## 8. The sources of organisational culture

It is sometimes said that the key to cultural change is leadership and that safety cultures or generative cultures can most easily be brought about by installing leaders who have the appropriate vision (Westrum, 2004, p. ii26). But inquiries into organisational culture reveal that there are often less personal sources of culture which need to be understood and counteracted if there is to be any hope of changing the organisational culture itself. For example, some aspects of the culture of the organisation may be rooted in wider societal factors, such as the public demand that trains run on time. It may require strong pressure from other external sources, such as safety regulators, to ensure that this pressure does not override safety.

Alternatively, culture change may depend on internal organisational change. For instance, the Columbia Board believed that NASA’s single-minded focus on getting on with the mission, regardless of safety concerns, would only be counteracted if a powerful Technical Engineering Authority was set up within NASA. The Authority should be concerned solely with safety, it should be in no way responsible for launch schedules and program costs, and it should have the capacity to withhold authorization for launches (CAIB, 2003, p. 193). The Board took this model from the US Navy which it regarded as a high reliability organisation.

The Australian Air Force Board of Inquiry came to a very similar conclusion. It noted that air safety was underpinned by a powerful technical engineering authority and that the only way ground safety would ever achieve the same status was if there was a similarly powerful ground safety authority within the Air Force.

The important point is that major accident research that identifies the impact of organisational culture on safety is likely also to provide insights into the source of that culture and hence ways in which it might be changed or modified. Neither the questionnaire technique nor the ethnographic method provide such ready insights.

## 9. Some practical suggestions

The previous discussion envisages the researcher drawing on the evidence of a major accident investigation to explore the culture of an organisation and its impact on safety.

An obvious limitation of this method is that many organisations have not had a major accident or, if they have, it has not been investigated by a public inquiry in the manner described above. Suppose, however, we are interested in an organisation that has not been the subject of a public inquiry and we wish to study how aspects of the culture of that organisation may be affecting safety. Researchers may easily find themselves in such a situation, especially if they engage in consultancy work. Does the preceding discussion of major accident inquiries provide any indication of how we might proceed?

The data gathering process in an official inquiry is quite eclectic and relatively untargeted. As a result much of the information collected may turn out to be of relatively little use. However, since such a mass of information is collected there is still much that is useful. Individual researchers with limited resources, particularly limited time, cannot afford such an inefficient data collection strategy. They need a more focused strategy, one that is guided by some theory about what to look for. Moreover, in major accident inquiries, the impact of the culture on safety is relatively easy to demonstrate. Where research on organisational culture does not take place in the context of an accident investigation, the impact of elements of organisational culture on safety will be more difficult to demonstrate. Here again the most efficient research strategy is to draw on theory with a ready-made connection between culture and safety.

An obvious candidate is the theory of high reliability organisations. The researcher can set out to examine the extent to which the organisation in question approaches the model of an HRO, just as the Columbia inquiry did. Weick and Sutcliffe argue that HROs have the following characteristics: a preoccupation with failure; a reluctance to simplify; a sensitivity to operations; a commitment to resilience; and deference to expertise (Weick and Sutcliffe, 2001, p. 10).<sup>6</sup> It is clear, however, that studying such an array of issues is still time consuming and will require researchers to spend quite long periods of time within the organisation.

An even more streamlined approach starts with the observation that the way an organisation handles information about potentially safety relevant matters is crucial to safety. We have already noted Westrum's claim that "the most critical issue for organisational safety is the flow of information". Pidgeon agrees. "The heart of a safety culture (he says) is the way in which organisational intelligence and safety imagination regarding risk and danger are deployed" (1998:206). Reason defines a safety culture as consisting of constellations of practices, most importantly, concerning reporting and learning. A safety culture, he argues, is both a reporting culture and a learning culture. Where safety is an organisation's top priority, the organisation will aim to assemble as much relevant information as possible, circulate it, analyse it, and apply it. Hence, researchers with limited time can restrict themselves to investigating the organisation's reporting practices (who reports, what gets reported and what is done about these reports) and its strategies for learning from accidents, both its own and others'. Such research is not just about one aspect of organisational culture. As Westrum notes, "information flow is also a type of marker for organisational culture" more generally. Finally, findings about how reporting and accident investigation systems are working are likely to be sufficiently precise to enable comparisons to be made among organisations. Research focused in this way therefore has the potential to move beyond the case study approach to comparative analysis.

---

<sup>6</sup> NB This is not the same as the list drawn up by the Columbia Board.

## 10. Conclusion

This paper has been about strategies for studying organisational cultures and their effects on safety. It has argued that the debates in the safety literature about the meaning of terms such as culture, climate and safety culture are somewhat sterile, and that the more important issue is to explore the kinds of methods that are available for studying such phenomena. The most commonly mentioned methods are the survey questionnaire and the ethnography. This paper has sought to draw attention to another method which involves researchers immersing themselves in the mass of material generated by public inquiries into major accidents. Such material provides the researcher with such deep insights into the way an organisation thinks and acts that it can reasonably be described as a form of armchair ethnography. Moreover, using the method of counter-factual thinking, that is, imagining what would have happened had things been otherwise, the researcher can demonstrate quite convincingly the way various elements of the culture of the organisation contributed to the accident. In other words, taken as a whole the method allows the researcher to clarify the impact of various elements of organisational culture on safety. The paper concludes with some comments about more time-efficient strategies available for studying safety-relevant aspects of organisational culture.

## References

- Bourrier, M., 1998. Elements for designing a self-correcting organisation: examples from nuclear plants. In: Hale, A., Baram, M. (Eds.), *Safety Management: The Challenge of Change*. Pergamon, Oxford, pp. 133–146.
- CAIB (Columbia Accident Investigation Board), 2003. Report, Vol. 1. Washington NASA.
- Cox, S., Flin, R., 1998. Safety culture: philosopher's stone or man of straw? *Work and Stress* 12 (3), 189–201.
- Deal, T., Kenney, A., 1982. *Corporate Culture: The Rites and Rituals of Corporate Life*. Addison-Wesley, Reading MA.
- Freeman, D., 1983. *Margaret Mead and Samoa: The Making and Unmaking of an Anthropological Myth*. Harvard University Press, Cambridge.
- Geertz, C., 1973. *The Interpretation of Cultures*. Basic New York. See especially chapter 1, "Thick description: towards an interpretive theory of culture".
- Glendon, A., Stanton, N., 2000. Perspectives on safety culture. *Safety Science* 23 (1–3), 193–214.
- Goffman, E., 1968. *Asylums*. Penguin, Harmondsworth.
- Gouldner, A., 1955. *Patterns of Industrial Bureaucracy*. RKP, London.
- Guldenmund, F., 2000. The nature of safety culture: a review of theory and research. *Safety Science* 34, 215–257.
- Hale, A., 2000. Culture's confusions. *Safety Science* 34, 1–14.
- Hofstede, G., 1997. *Cultures and Organisations*. McGraw-Hill, New York.
- Hopkins, A., 2005. *Safety, Culture and Risk*. CCH, Sydney.
- Kriegler, R., 1980. *Working for the Company*. Oxford UP, Melbourne.
- Ladkin, P., 2001. *Causal System Analysis*. <http://www.rvs.uni-bielefeld.de/publications/books/ComputerSafety-Book/index.html>.
- LaPorte, T., Consolini, P., 1991. Working in practice but not in theory: the theoretical challenges of "high-reliability organisations". *Journal of Public Administration Research and Theory* 1 (1), 19–47.
- Mearns, K., Flin, R., 1999. Assessing the state of organisational safety – culture or climate? *Current Psychology*. Spring 18 (1).
- Mearns, K., Whitaker, S., Flin, R., 2003. Safety climate, safety management practices and safety performance in offshore environments. *Safety Science* 41 (8), 641–680.
- Parker, D., Lawrie, M., Hudson, P., 2006. A framework for understanding the development of organisational safety culture. *Safety Science* 44 (6), 551–562.
- Pidgeon, N., 1998. Safety culture: theoretical issues. *Work and Stress* 12 (3), 202–216.
- Rasmussen, J., 1997. Risk management in a dynamic society: a modelling problem. *Safety Science* 27 (2/3), 183–213.

- Reason, J., 1997. *Managing the Risks of Organisational Accidents*. Ashgate, Aldershot.
- Schein, E., 1992. *Organisational Culture and Leadership*, 2nd ed. Jossey-Bass, San Francisco.
- Turner, B., Pidgeon, N., 1997. *Man-Made Disasters*, 2nd ed. Wykeham, Oxford.
- Vaughan, D., 1996. *The Challenger Launch Decision: Risky Technology, Culture and Deviance at NASA*. Univ of Chicago Press, Chicago.
- Weick, K., Sutcliffe, K., 2001. *Managing the Unexpected: Assuring High Performance in an Age of Complexity*. Jossey-Bass, San Francisco.
- Westrum, R., 1999. *Sidewinder: Creative Missile Design at China Lake*. Naval Institute Press, Annapolis, MD.
- Westrum, R., 2004. A typology of organisational cultures. *Quality and Safety in Health Care* 13 (Suppl II), ii22–ii27.