

Just culture: who gets to draw the line?

Sidney W. A. Dekker

Received: 15 October 2007 / Accepted: 14 January 2008
© Springer-Verlag London Limited 2008

Abstract A just culture is meant to balance learning from incidents with accountability for their consequences. All the current proposals for just cultures argue for a clear line between acceptable and unacceptable behavior. This alone, however, cannot promote just culture as it falsely assumes that culpability inheres in the act, bearing immutable features independent of context, language or interpretation. The critical question is not where to draw the line, but who gets to draw it. Culpability is socially constructed: the result of deploying one language to describe an incident, and of enacting particular post-conditions. Different accounts of the same incident are always possible (e.g. educational, organizational, political). They generate different repertoires of countermeasures and can be more constructive for safety. The issue is not to exonerate individual practitioners but rather what kind of accountability promotes justice and safety: backward-looking and retributive, or forward-looking and change-oriented.

Keywords Incident reporting · Just culture · Human error · Accountability · Criminalization · Culpability

1 Drawing a line between legitimate and illegitimate behavior

The desire to balance learning from failure with appropriate accountability has motivated safety-critical industries and organizations to develop guidance on a so-called “just

culture”. In this paper I question whether such guidance can merely focus on a clear line between acceptable and unacceptable behavior—which all such guidance today does. This is based on the essentialist assumption that inherently culpable acts exist and should be dealt with as such. The counterproposition I advance in this paper is that culpable acts have no essentialist properties or immutable features, but that designations of acceptability or culpability are the result of processes of social construction steeped in context, language, history. After setting out the constructionist argument, I assess various alternatives of who gets the power to draw the line, and review the negative consequences for safety of leaving it in the hands of a judiciary alone. Then I try to clear up confusion between blame-free and accountability-free, suggesting that some forms of accountability, and accountability relationships between stakeholders, can be more constructive for safety than others. I conclude with a list of suggestions for organizations on building the basis for a just culture.

1.1 Balancing accountability and learning

Concern about just cultures has grown out of our changing interpretation of accidents since the 1970s (such as the nuclear incident at Three Mile Island, and the twin Boeing 747 disaster at Tenerife). We no longer see such accidents as meaningless, uncontrollable events, but rather as failures of risk management, and behind these failures are people and organizations (Green 2003). Today, almost every accident is followed by questions centering on “whose fault?” and “what damages, compensation?” It seems as if every death must be charged to somebody’s account (Douglas 1992). We have increasingly begun to see

S. W. A. Dekker (✉)
Lund University School of Aviation, 260 70 Ljungbyhed,
Sweden
e-mail: sidney.dekker@tfhs.lu.se

accidents as the result of people not doing their jobs properly, and the possibility of punishing them for that is no longer is remote. In 2006, for example, a nurse from Wisconsin was charged with criminal “neglect of a patient causing great bodily harm” in the medication death of a 16-year-old girl during labor. Instead of giving the intended penicillin intravenously, the nurse accidentally administered a bag of epidural analgesia. She lost her job, faced action on her nursing license and the threat of 6 years in jail as well as a 25,000\$ fine. Her predicament likened that of three nurses in Denver in 1998, who administered benzathine penicillin intravenously, causing the death of a neonate. The nurses were charged with criminally negligent homicide and faced 5 years in jail (Cook et al. 2000). This in turn was similar to a nurse in Sweden convicted of manslaughter in an order-of-magnitude medication error that led to an infant’s death (Dekker 2007).

Criminalization of any act is not just about retribution and explanation of misfortune, but also about putative deterrence, and so it is with the criminalization of human error. Responding to the 1996 ValuJet accident, where mechanics loaded oxygen generators into the cargo hold of a DC-9 airliner which subsequently caught fire, the editor of *Aviation Week and Space Technology* “strongly believed the failure of SabreTech employees to put caps on oxygen generators constituted willful negligence that led to the killing of 110 passengers and crew. Prosecutors were right to bring charges. There has to be some fear that not doing one’s job correctly could lead to prosecution” (North 2000, p. 66). The deterrence argument is problematic, however, as threats of prosecution do not deter people from making errors, but rather from reporting them (e.g. Merry et al. 2001; Cohen-Charash and Spector 2001; Sharpe 2003). Instead, anxiety created by such accountability leads for example to defensive medicine, not high-quality care, and even to a greater likelihood of subsequent incidents (e.g. Dauer 2004). The anxiety and stress generated by such accountability adds attentional burdens and distracts from conscientious discharge of the main safety-critical task (Lerner and Tetlock 1999).

A just culture, then, is particularly concerned with the sustainability of learning from failure through the reporting of errors, adverse events, incidents. If operators and others perceive that their reports are treated unfairly or lead to negative consequences, the willingness to report will decline (e.g. Ruitenberg 2002 cited a 50% drop in incident reports after the prosecution of air traffic controllers involved a near-miss). Writings about just culture over the past decade (e.g. Reason 1997; Marx 2001; Dekker 2008) acknowledge this central paradox of accountability and learning: various stakeholders (e.g. employers, regulators) want to know everything that happened, but cannot accept everything that happened and will want to advertise their

position as such. Thus, rating certain behavior as culpable is not just about that behavior or its antecedent intentions, it performs a wider function of regulating a distinction between normal and abnormal, between order and disorder. “A ‘no-blame’ culture is neither feasible nor desirable. Most people desire some level of accountability when a mishap occurs” (GAIN 2004 p. viii). These are neo-Durkheimian ideas (see Durkheim 1950, 1895) about the boundary-maintaining function of the organizational rituals and languages that keep such a distinction in place:

“Confrontations in the form of criminal trials, excommunication hearings, courts-martial ... act as boundary-maintaining devices in the sense that they demonstrate ... where the line is drawn between behavior that belongs in the special universe of the group and behavior that does not” (Erikson 1966 p. 11).

Demonstrating a border between acceptable and unacceptable is deemed critical. After all, an environment of impunity, the argument holds, would neither move people to act prudently nor compel them to report errors or deviations. If there is no line, “anything goes”. So why report anything?

1.2 The line is a judgment, not a location

The essentialist assumption that animates current guidance on just culture is that some behavior is inherently culpable, and should be treated as such. The public must be protected against intentional misbehavior or criminal acts, and the application of justice is a prime vehicle for this (e.g. Reason 1997). As Marx (2001 p. 3) puts it, “It is the balancing of the need to learn from our mistakes and the need to take disciplinary action that (needs to be addressed). Ultimately, it will help you answer the question: ‘Where do you draw the disciplinary line?’” As another example (Eurocontrol 2006), a just culture is one in which “front-line operators or others are not punished for actions, omissions or decisions taken by them that are commensurate with their experience and training, but where gross negligence, willful violations and destructive acts are not tolerated”. Such proposals emphasize the establishment of, and consensus around, some kind of separation between legitimate and illegitimate behavior: “in a just culture, staff can differentiate between acceptable and unacceptable acts” (Ferguson and Fakelmann 2005 p. 34). Similarly, “in a Just Culture environment the culpability line is more clearly drawn” (GAIN 2004 p. viii).

But drawing an a priori line between the acts an organization will accept and those it will not is difficult. Culpability does not inhere in the act. Whether something is judged culpable is the outcome of processes of interpretation and attribution that follow the act, in which

assumptions of other people's volitional behavior and outcome control, as well as causal control, play a dominant role (Alicke 2000). Thus, to gauge whether behavior should fall on one side of the line or the other, variations on a basic decision tree are in circulation (e.g. Reason 1997). Yet its questions confirm the negotiability of the line rather than resolving its location:

- *Were the actions and consequences as intended?* This invokes the judicial idea of a mens rea ("guilty mind"), and seems a simple enough question. Few people in safety-critical industries intend to inflict harm, though that does not prevent them from being prosecuted for their "errors" (under charges of manslaughter, for example, or general risk statutes that hail from road traffic laws on "endangering other people," see e.g. Wilkinson 1994). Also, what exactly is intent and how do you prove it? And who gets to prove this, using what kind of expertise?
- *Did the person knowingly violate safe operating procedures?* People in operational worlds knowingly adapt written guidance, and have to do so to bridge the gap between prescriptive routines and actual work in worlds of imperfect knowledge, time constraints and infinite variation (Suchman 1987; Rochlin 1999; Vaughan 1999; Woods and Shattuck 2000; Smith 2001; Dekker 2003). Calling such adaptations "violations" (Reason 1997) already implies a moral judgment about who is wrong (the worker) and who is right (the rule). It is easy to show in hindsight which procedures would have been applicable, available, workable and correct for a particular task (says who, though?), but such overestimations of the role of procedural non-compliance in the wake of incidents conceals the real operational dilemmas faced by people (McDonald et al. 2002).
- *Were there deficiencies in training or selection?* "Deficiencies" seems unproblematic but what is a deficiency from one angle can be perfectly normal or even above industry standard from another.

Questions such as the ones above may form a good start, but they themselves cannot arbitrate between culpable or blameless behavior. Rather, they invoke new judgments and negotiations. This is also true for the very definition of negligence (a legal term, not a human performance concept):

"Negligence is a conduct that falls below the standard required as normal in the community. It applies to a person who fails to use the reasonable level of skill expected of a person engaged in that particular activity, whether by omitting to do something that a prudent and reasonable person would do in the circumstances or by doing something that no prudent or reasonable person would have done in the circumstances. To raise a question of negligence, there needs to be a duty of care on the person, and

harm must be caused by the negligent action. In other words, where there is a duty to exercise care, reasonable care must be taken to avoid acts or omissions which can reasonably be foreseen to be likely to cause harm to persons or property. If, as a result of a failure to act in this reasonably skillful way, harm/injury/damage is caused to a person or property, the person whose action caused the harm is negligent" (GAIN 2004 p. 6).

There is no definition that captures the essential properties of "negligence". Instead, definitions such as the one above open a new array of questions and judgments. What is "normal standard?" How far is "below?" What is "reasonably skillful?" What is "reasonable care?" What is "prudent?" Was harm indeed "caused by the negligent action?" Of course, making such judgments is not impossible. In fact, they remain judgments—made by somebody or some group in some time and place in the aftermath of an act—not objective features that stably inhabit the act itself. That judgments are required to figure out whether we deem an act culpable is not the problem. The problem is guidance that suggests that a just culture only needs to "clearly draw" a line between culpable and blameless behavior. Its problem lies in the false assumption that acceptable or unacceptable behavior form stable categories with immutable features that are independent of context, language or interpretation.

2 Different accounts and meanings of failure

2.1 The social construction of culpability

Just as the properties of "human error" are not objective and independently existing, so does culpability arise out of our ways of seeing and describing acts. What ends up being labeled as culpable does not inhere in the act or the person. It is constructed, or "constituted" as Christie (2004 p. 10) put it:

"The world comes to us as we constitute it. Crime is thus a product of cultural, social and mental processes. For all acts, including those seen as unwanted, there are dozens of possible alternatives to their understanding: bad, mad, evil, misplaced honour, youth bravado, political heroism—or crime. The same acts can thus be met within several parallel systems as judicial, psychiatric, pedagogical, theological".

It is tempting to think that culpability, of all things, must make up some essence behind a number of possible descriptions of an act, especially if that act has a bad outcome. We hope that the various descriptions can be sorted out by the rational process of an investigation, a hearing or a trial, and that it will expose Christie's "psychiatric, pedagogical, theological" explanations (I had failure anxiety! I was not trained enough! It was the Lord's will!) as false.

The application of reason will strip away the noise, the decoys, and the excuses to arrive at the essential story: whether culpability lay behind the incident or not. And if culpable behavior turns out not make up the essence, then there will be no retribution. But Christie argued that culpability is not an essence that we can discover behind the inconsistency and shifting nature of the world as it meets us. Culpability itself is that flux, that inconstancy, a negotiated arrangement, a tenuous, temporary stability achieved among shifting cultural, social, mental and political forces. Concluding that an unwanted act is culpable, is an accomplished project, a purely human achievement:

... deviance is created by society ... social groups create deviance by making the rules whose infraction constitutes deviance and by applying those rules to particular persons and labeling them as outsiders. From this point of view, deviance is not a quality of the act the person commits, but rather a consequence of the application by others of rules and sanctions to an “offender”. The deviant is the one to whom the label has successfully been applied; deviant behavior is behavior that people so label (Becker 1963 p. 9).

Becker argues that what counts as deviant or culpable is the result of processes of social construction. According to this, if an organization decides that a certain act constituted “negligence” or otherwise falls on the wrong side of the line, then this is the result of using a particular language and enacting a particular repertoire of post-conditions that turn the act into culpable behavior and the involved practitioner into an offender (e.g. Burr 2003).

2.2 Alternative readings of the same act

The social constructionist argument about culpability is that by seeing human error as a crime, we have evoked just one language for describing and explaining an event, relative to a multitude of other possibilities. If we subscribe to this one reading as true, it will blind us to alternative readings or framings that can frequently be more constructive. Take as an example a British cardiothoracic surgeon who moved to New Zealand (Skegg 1998). There, three patients died during or immediately after his operations, and he was charged with manslaughter. Not long before, a professional college had pointed to serious deficiencies in the surgeon’s work and found that seven of his cases had been managed incompetently. The report found its way to the police, which subsequently investigated the cases. This in turn led to the criminal prosecution against the surgeon. But the same unwanted act can be construed to be a lot of things at the same time, depending on what questions you asked to begin with. Ask Christie’s theological question and you may see in an error the manifestation of evil, or the weakness of the flesh. Ask

pedagogical questions and you may see in it the expression of underdeveloped skills. Ask judicial questions and you may begin to see a crime. Calling the surgical failures a crime is one possible interpretation of what went wrong and what should be done about it. Other ways are possible too, and not necessarily less valid:

- For example, we could see the three patients dying as an issue of cross-national transition: are procedures for doctors moving to Australia or New Zealand and integrating them in local practice adequate?
- And how are any cultural implications of practicing there systematically managed or monitored, if at all?
- We could see these deaths as a problem of access control to the profession: do different countries have different standards for who they would want as a surgeon, and who controls access, and how?
- It could also be seen as a problem of training or proficiency-checking: do surgeons submit to regular and systematic follow-up of critical skills, such as professional pilots do in a proficiency check every 6 months?
- We could also see it as an organizational problem: there was a lack of quality control procedures at the hospital, and the surgeon testified having no regular junior staff to help with operations, but was made to work with only medical students instead.
- Finally, we could interpret the problem as socio-political: what forces are behind the assignment of resources and oversight in care facilities outside the capital?

It may well be possible to write a compelling argument for each of these explanations of failure—each with a different repertoire of interpretations and countermeasures following from it. A crime gets punished away. Access and proficiency issues get controlled away. Training problems get educated away. Organizational issues get managed away. Political problems get elected or lobbied away. This also has different implications for what we mean by accountability. If we see an act as a crime, then accountability means blaming and punishing somebody for it. Accountability in that case is backward-looking, retributive. If, instead, we see the act as an indication of an organizational, operational, technical, educational or political issue, then accountability can become forward-looking. The question becomes: what should we do about the problem and who should bear liability for implementing those changes?

2.3 Overlapping and contradictory versions of history

The point is not that one interpretation of an incident is right and all the others are wrong. All the accounts are

inherently limited. Telling the story from one angle necessarily excludes the aspects from other angles. And all the interpretations have different ramifications for what people and organizations think they should do to prevent recurrence. Finding an act culpable, then, is settling onto one particular version or description of history. This version is not just produced for its own sake. It may serve a range of social functions, from emphasizing moral boundaries and enhancing solidarity (Erikson 1966), to sustaining subjugation or asymmetric power distribution within hierarchies (Foucault 1981), to protecting elite interests after an incident has exposed possibly expensive vulnerabilities in the whole industry (Perrow 1984; Byrne 2002), to mitigating public or internal apprehension about the system's ability to protect its safety-critical technologies against failure (Vaughan 1996; Galison 2000). This also denies the modernist objectification of history (captured, for example, in "probable cause" statements in incident reports) that considers the past to be an object; bygone, coagulated. Instead, the past is a dimension of our present experience. The past offers all kinds of opportunities to express and handle current issues, address current concerns, accommodate current agendas. This makes it critical to consider who owns the right to write history. Who has the power to tell a story of performance in such a way—to use a particular rhetoric to describe it, ensuring that certain subsequent actions are legitimate or even possible (e.g. pursuing a single culprit), and others not—so as to, in effect, own the right to draw the line?

3 Whom do we give the power to draw the line?

3.1 Judicial drawing of the line

People increasingly turn to the legal system to furnish them with an answer about the culpability of a practitioner's performance (Laudan 2006). For example, a directive from the European Union (2003/42/EC) says that a state must not institute legal proceedings against those who send in incident reports, apart from cases of gross negligence. But who gets to decide whether an act amounts to gross negligence? The same state, through its judiciary. Even so, we expect a court to apply reason, and objectivity. A disinterested party takes an evenhanded look at the case, the appropriate person gets to be held accountable and consequences are meted out. We tend to believe that an "objective" account (one produced by the rational processes of a court, or an independent investigation of the incident) delivers superior accuracy because it is well-researched and not as tainted by interests or a particular, partisan perspective. Many aspects of the justice system (and of formal accident investigation) are indeed designed

to impart an image of rationality, of consideration, objectivity and impartiality (e.g. Lady Justitia's blindfold, or the party system in certain investigations). But truths (or accounts that are taken as valid) are always brought into being by historically and culturally located groups of people, and as such open to the myriad influences that impact any social process.

Judicial involvement can consist of:

- The participation of law enforcement officials in investigations. There are countries in the developed world where the police is witness or participant in accident investigations (in for example road traffic or aviation). This can impede investigatory access to information sources, as pressures to protect oneself against criminal or civil liability can override a practitioner's willingness to cooperate in the accident probe.
- Judicial authorities stopping an investigation or taking it over when evidence of criminal wrong-doing emerges. This often restricts further access to evidence for safety investigators.
- Launching a criminal probe independent of a safety investigation or its status. Accident investigation boards typically say this retards their efforts to find out what went wrong and what to do to prevent recurrence (North 2002). For example, while the US National Transportation Safety Board was investigating a 1999 pipeline explosion near Bellingham, Washington, that killed three people, federal prosecutors launched their own criminal probe. They reportedly pressured employees of the pipeline operator to talk. Several invoked the US Constitution's Fifth Amendment, which protects against self-incrimination. They refused to answer questions from Safety Board investigators as well as from the police (McKenna 1999).
- Using a formal accident report in a court case. Even though using such reports as evidence in court is proscribed through various statutory arrangements (Eurocontrol 2006), these can get overridden or circumvented. And nobody can prevent a prosecutor or judge from simply reading a publicly-available accident report.
- Getting access to safety-related data (e.g. internal incident reports) because of freedom-of-information legislation in that country, under which any citizen (including the judicial system) has quite unfettered access to many kinds of organizational data. This access is particularly acute in organizations that are government-owned (such as many air traffic control providers, or hospitals).
- Taking the results of a safety inspection if these expose possibly criminal or otherwise liable acts. This does not have to take much: an inspection report listing

“violations” (of regulations, which in turn are based in law) can be enough for a prosecutor to start converting those violations (which were discovered and discussed for the purpose of regulatory compliance and safety improvement) into prosecutable crimes.

In all these ways, judicial involvement (or the threat of it) can engender a climate of fear and silence (Ter Kulle 2004). So even as a court of law cannot bring the “truth” about human performance into necessarily sharper focus than any other social process (Nagel 1992), it has measurably negative consequences for practitioners’ (and sometimes even regulators’) inclination to share safety information (Ruitenbergh 2002; Dekker 2008). A recent European-wide Air Traffic Control survey confirms how the threat of judicial involvement after incidents (and certainly accidents) dampens people’s willingness to come forward with safety information (Eurocontrol 2006), and other examples are not hard to come by (e.g. Wilkinson 1994). In the wake of a June 1995 crash of an Ansett de Havilland Dash 8 near Palmerston North in New Zealand, accident investigators turned the aircraft’s cockpit voice recorder (CVR) over to criminal prosecutors. The crash killed four persons on the aircraft, but not the pilots, who faced charges of manslaughter. Pilots in New Zealand sued to block the police use of the CVR, arguing recorders should only be used for safety and educational purposes. But prosecutors prevailed and regained access to the CVR. Pilots soon began disabling CVR’s on their flights, prompting legislative changes that involved the country’s High Court and proscribing the public use of CVR information (McKenna 1999).

3.2 Alternatives to judicial drawing of the line

To mitigate the negative side-effects of judicial interference, some countries have moved ahead with installing a so-called judge of instruction, who functions as a go-between before a prosecutor can actually go ahead with a case. A judge of instruction gets to determine whether a case proposed by a prosecutor should be investigated (and later go to trial). The judge of instruction, in other words, can check the prosecutor’s homework and ambitions, do some investigation him- or herself, and weigh other stakeholders’ interests in making the decision to go ahead with a further investigation and possible prosecution or not. It is still the judge of instruction who gets to draw the line between acceptable and unacceptable (or: between worthy of further investigation and possible prosecution or not), but broader considerations can make it into the drawing of the line too (e.g. the interests of other industry stakeholders, as long as those are fairly represented).

Another adaptation is to make the prosecutor part of the regulator, as has been done in some countries (particularly in aviation). The prosecutor him- or her-self has a history in or affiliation with the domain, guaranteeing an understanding of and concern for its sources of safety. It is thus likely that the prosecutor is better able to balance the various interests in deciding whether to draw a line, and better able to consider subtle features of the professional’s performance that non-domain experts would overlook or misjudge. The risk in this solution, of course, is that the regulator itself can have played a role (e.g. insufficient oversight, or given dispensation) in the creation of an incident and can have a vested interest in the prosecution of an individual practitioner so as to downplay its own contribution. There is no immediate protection against this in this local solution, except for regulatory self-restraint and perhaps the possibility of appeals higher up in the judiciary.

Disciplinary rules within the profession are another alternative. Many professional groups (from accountants to physicians to hunters to professional sports players) have elaborate systems of disciplinary rules. These are meant foremost to protect the integrity of a profession. Usually, a judiciary delegates large amounts of legal authority to the boards that credibly administer these professional disciplinary rules. Professional sanctions can range from warning letters (which are not necessarily effective) to the revocation of licenses to practice. The judiciary will not normally interfere with the internal administration of justice according to these disciplinary rules. There is, however, great variation in the administration of internal professional justice and thus a variation in how much confidence a country can have in delegating it to an internal disciplinary board. And of course, it does not remove the problem of where the line goes: the judiciary will still have to judge whether a line has been crossed that prompts them to step in. This even raises a possible paradox in the justness of professional disciplinary rules. Because disciplinary rules aim to maintain the integrity of a profession, individual practitioners may still get “sacrificed” for that larger aim (especially to keep the system free from outside interference or unwelcome judicial scrutiny).

4 Blame-free or accountability-free?

4.1 A discretionary space for accountability

Moves to redirect the power to draw the line away from the judiciary can be met with suspicions that operators want to blame “the system” when things go wrong, and that they do not want to be held liable in the same way as other citizens (Merry et al. 2001; Pellegrino 2004). Yet perhaps

the choice is not between blaming people or systems. Instead, we may reconsider the accountability relationships of people in systems (Berlinger 2005). All safety-critical work is ultimately channeled through relationships between human beings (such as in medicine), or through direct contact of some people with the risky technology. At this sharp end, there is almost always a discretionary space into which no system improvement can completely reach. This is in part a space in the almost literal sense of “room for maneuvering” that operators enjoy while executing their work relatively unsupervised (in the examination room, the operating theatre, cockpit). It is also a space in a metaphorical sense, of course, as its outlines are not stipulated by decree or regulation, but drawn by actions of individual operators and the responses to them. It is, however, a final kind of space filled with ambiguity, uncertainty and moral choices. And a space that is typically devoid of relevant or applicable guidance from the surrounding organization, leaving the difficult calls up to the individual operator or crews. Systems cannot substitute the responsibility borne by individuals within that space. Individuals who work in those systems would not even want that. The freedom (and concomitant responsibility) that is left for them is what makes them and their work human, meaningful, a source of pride.

But organizations can do a number of things. One is to be clear about where that discretionary space begins and ends. Not giving practitioners sufficient authority to decide on courses of action, but demanding that they be held accountable for the consequences anyway, creates impossible and unfair goal conflicts (for which managers may sometimes be held accountable, but they too could have been the recipients of similar goal conflicts). It effectively shrinks the discretionary space before action, but opens it wide after any bad consequences of action become apparent. Second, an organization must deliberate how it will motivate people to conscientiously carry out their duties inside the discretionary space. Is the source for that motivation fear or empowerment? There is evidence that empowering people to affect their work conditions, to involve them in the outlines and content of that discretionary space, most actively promotes their willingness to shoulder their responsibilities inside of it (Kohn 1999; Wiegmann et al. 2002; Dekker and Laursen 2007). For example, during surgery, an anesthetist reached into a drawer that contained two vials that were side by side, both with yellow labels and yellow caps. One, however, had a paralytic agent, the other a reversal agent for when paralysis was no longer needed. At the beginning of the procedure, the anesthetist administered the paralyzing agent. But toward the end, he grabbed the wrong vial, administering additional paralytic instead of its reversal agent. There was no bad outcome in this case.

But when he discussed the event with his colleagues, he found that this had happened to them too, and that they were all quite aware of the potential risks. Yet none had spoken out about it, which could raise questions about the empowerment anesthetists may have felt to influence their work conditions, their discretionary space (Morreim 2004).

4.2 Blame-free is not accountability-free

Equating blame-free systems with an absence of personal accountability, as some do (e.g. Pellegrino 2004) is wrong. The kind of accountability wrung out of practitioners in a trial is not likely to contribute to future safety in their field, and in fact may hamper it. We can create such accountability not by blaming people, but by getting people actively involved in the creation of a better system to work in. Holding people accountable and blaming people are two quite different things. Blaming people may in fact make them less accountable: they will tell fewer accounts, they may feel less compelled to have their voice heard, to participate in improvement efforts. Blame-free or no-fault systems are not accountability-free systems. On the contrary: such systems want to open up the ability for people to hold their account, so that everybody can respond and take responsibility for doing something about the problem. This also has different implications for what we mean by accountability. If we see an act as a crime, then accountability means blaming and punishing somebody for it. Accountability in that case is backward-looking, retributive. If, instead, we see the act as an indication of an organizational, operational, technical, educational or political issue, then accountability can become forward-looking (Sharpe 2003). The question becomes what should we do about the problem and who should bear responsibility for implementing those changes.

5 Creating the basis for a just culture

Whereas the judicial climate in a country can discourage open reporting and honest disclosure (e.g. Berlinger 2005), this does not mean that an organization charged with running a safety-critical operation (in e.g. healthcare, aviation, nuclear power generation) cannot try to build a basis for a just culture. The first steps involve a normalization of incidents, so that they become a legitimate, acceptable part of organizational development. Then, the organization must consider what to do about the question “who gets to draw the line?” both inside its own operation and in influencing the judicial climate surrounding it. Here are some suggestions:

First, normalize and try to legitimize incidents:

- An incident must not be seen as a failure or a crisis, neither by management, nor by colleagues. An incident is a free lesson, a great opportunity to focus attention and to learn collectively.
- Abolish financial and professional penalties (e.g. suspension) in the wake of an occurrence. These measures render incidents as something shameful, to be kept concealed, leading to the loss of much potential safety information and lack of trust.
- Monitor and try to prevent stigmatization of practitioners involved in an incident. They should not be seen as a failure, or as a liability to work with by their colleagues.
- Implement, or review the effectiveness of, any debriefing programs or critical incident/stress management programs the organization may have in place to help practitioners after incidents. Such debriefings and support form a crucial ingredient in helping practitioners see that incidents are “normal”, that they can help the organization get better, and that they can happen to everybody.
- Build a staff safety department, not part of the line organization that deals with incidents. The direct manager (supervisor) of the practitioner should not necessarily be the one who is the first to handle the practitioner in the wake of an incident. Aim to decouple an incident from what may look like a performance review or punitive retraining of the practitioner involved.
- Start with building a just culture at the very beginning: during basic education and training of the profession. Make trainees aware of the importance of reporting incidents for a learning culture, and get them to see that incidents are not something individual or shameful but a good piece of systemic information for the entire organization. Convince new practitioners that the difference between a safe and an unsafe organization lies not how many incidents it has, but in how it deals with the incidents that it has its people report.
- Ensure that practitioners know their rights and duties in relation to incidents. Make very clear what can (and typically does) happen in the wake of an incident (e.g. to whom practitioners were obliged to speak, and to whom not). A reduction in such uncertainty can prevent practitioners from withholding valuable incident information because of misguided fears or anxieties.

Second, the important discussion for an organization is who draws the line between acceptable and unacceptable inside the organization? This means not only who gets to handle the immediate aftermath of an incident (the line

organization: supervisor/manager, or a staff organization such as safety department), but how to integrate practitioner peer expertise in the decision on how to handle this aftermath, particularly decisions that relate to the individual practitioner’s stature. Empowering and involving the practitioner him- or her-self in the aftermath of an incident is the best way to maintain morale, maximize learning, and reinforce the basis for a just culture.

Third, think about how to protect the organization’s data from undue outside probing (e.g. by a prosecutor). The consequences of this step must be thought through. One problem is that better protection for incident reporters can lock information up even for those who rightfully want access to it, and who have no vindictive intentions (e.g. patients or their families). The protection of reporting can make disclosure to such parties more difficult.

Fourth, it could be profitable to start a discussion with the prosecuting authority in the country on how to help them integrate domain expertise (to support them in making better judgments about whether something is worthy of further investigation and prosecution). This may require that previous mistrust is overcome and may seem difficult in the beginning. In the end, however, it may tremendously benefit all parties, as it may also create a better understanding of each other’s point of view and interests.

Uncertainty about, and perceived unfairness of, who gets to draw the line is likely to overrule any guidance in use today on where that line goes. The socially constructed judgment of that line means that its location will forever be more unpredictable than relatively stable arrangements among stakeholders about who gets to draw the line, with or without help from others.

Acknowledgment Part of the work for this paper was conducted on a grant from The European Organisation for the Safety of Air Navigation (Eurocontrol) to study the interface between the judiciary and air navigation service providers. With thanks to Tony Licu and his people in Brussels.

References

- Alicke MD (2000) Culpable control and the psychology of blame. *Psychol Bull* 126:556–74
- Becker HS (1963) *Outsiders: studies in the sociology of deviance*. Free Press, New York
- Berlinger N (2005) *After harm: medical error and the ethics of forgiveness*. The Johns Hopkins University Press, Baltimore
- Burr V (2003) *Social constructionism*. Routledge, London
- Byrne G (2002) *Flight 427: anatomy of an air disaster*. Copernicus books, New York
- Christie N (2004) *A suitable amount of crime*. Routledge, London
- Cohen-Charash Y, Spector PE (2001) The role of justice in organizations: a meta-analysis. *Organ Behav Hum Decis Processes* 86:278–321
- Cook RI, Render M, Woods DD (2000) Gaps in the continuity of care and progress on patient safety. *BMJ* 320:791–794

- Dauer EA (2004) Ethical misfits: mediation and medical malpractice litigation. In: Sharpe VA (ed) *Accountability: patient safety and policy reform*. Georgetown University Press, Washington, pp 185–202
- Dekker SWA (2003) Failing to adapt or adaptations that fail: contrasting models on procedures and safety. *Appl Ergon* 34:233–238
- Dekker SWA (2007) Discontinuity and disaster: gaps and the negotiation of culpability in medication delivery. *J Law Med Ethics* 35:463–470
- Dekker SWA (2008) *Just culture: balancing safety and accountability*. Ashgate Publishing Co, Aldershot
- Dekker SWA, Laursen T (2007) From punitive action to confidential reporting. *Patient Saf Qual Healthc* 5:50–56
- Douglas M (1992) *Risk and blame: essays in cultural theory*. Routledge, London
- Durkheim E (1950/1895) *The rules of the sociological method*. Free Press, New York
- Erikson K (1966) *Wayward puritans*. Wiley, New York
- Eurocontrol Performance Review Commission (2006) *Report on legal and cultural issues in relation to ATM safety occurrence reporting in Europe: outcome of a survey conducted by the Performance Review Unit in 2005–2006*. Eurocontrol, Brussels
- Ferguson J, Fakelmann R (2005) The culture factor. *Front Health Serv Manage* 22:33–40
- Foucault M (1981) The order of discourse. In: Young R (ed) *Untying the text: a post-structuralist reader*. Routledge, London, pp 48–79
- GAIN (2004) *Roadmap to a just culture: Enhancing the safety environment*. Global Aviation Information Network (Group E: Flight Ops/ATC Ops Safety Information Sharing Working Group)
- Galison P (2000) An accident of history. In: Galison P, Roland A (eds) *Atmospheric flight in the twentieth century*. Kluwer, Dordrecht, NL, pp 3–44
- Green J (2003) The ultimate challenge for risk technologies: controlling the accidental. In: Summerton J, Berner B (eds) *Constructing risk and safety in technological practice*. Routledge, London
- Kohn LT, Corrigan JMK, Donaldson M (eds) (1999) *To err is human*. Institute of Medicine, Washington DC
- Laudan L (2006) *Truth, error and criminal law: an essay in legal epistemology*. Cambridge University Press, Cambridge
- Lerner JS, Tetlock PE (1999) Accounting for the effects of accountability. *Psychol Bull* 125:255–75
- Marx D (2001) Patient safety and the “just culture”: a primer for health care executives. Columbia University, New York
- McDonald N, Corrigan S, Ward M (2002 June). Well-intentioned people in dysfunctional systems. Keynote paper presented at the 5th Workshop on human error, safety and systems development, Newcastle
- McKenna JT (1999) Criminal and safety probes at odds. *Aviat Week Space Technol* 47–48 (13 December)
- Merry AF, McCall Smith A (2001) *Errors, medicine and the law*. Cambridge University Press, Cambridge
- Morreim EH (2004) Medical errors: pinning the blame versus blaming the system. In: Sharpe VA (ed) *Accountability: patient safety and policy reform*. Georgetown University Press, Washington, pp 213–232
- Nagel T (1992) *The view from nowhere*. Oxford University Press, Oxford
- North DM (2000 May 15) Let judicial system run its course in crash cases. *Aviation Week Space Technol*, p 66
- North DM (2002, February 4) Oil and water, cats and dogs. *Aviation Week Space Technol*, p 70
- Pellegrino ED (2004) Prevention of medical error: where professional and organizational ethics meet. In: Sharpe VA (ed) *Accountability: patient safety and policy reform*. Georgetown University Press, Washington, pp 83–98
- Perrow C (1984) *Normal accidents: living with high-risk technologies*. Basic Books, New York
- Reason JT (1997) *Managing the risks of organizational accidents*. Ashgate Publishing Co, Aldershot
- Rochlin GI (1999) Safe operation as a social construct. *Ergonomics* 42:1549–1560
- Ruitenbergh B (2002) Court case against Dutch controllers. *The Controller* 41:22–5
- Sharpe VA (2003) Promoting patient safety: an ethical basis for policy deliberation. *Hastings Center Report Special Suppl* 33(5):S1–S20
- Skegg PDG (1998) Criminal prosecutions of negligent health professionals: the New Zealand experience. *Med Law Rev* 6:220–46
- Smith K (2001) Incompatible goals, uncertain information and conflicting incentives: the dispatch dilemma. *Human Factors Aerospace Safety* 1:361–380
- Suchman LA (1987) *Plans and situated actions: the problem of human-machine communication*. Cambridge University Press, Cambridge
- Ter Kulle A (2004) Safety versus justice. *Canso News* 18:1–2
- Vaughan D (1996) *The challenger launch decision: risky technology, culture and deviance at NASA*. University of Chicago Press, Chicago
- Vaughan D (1999) The dark side of organizations: mistake, misconduct, and disaster. *Annu Rev Sociol* 25:271–305
- Wiegman D, Zhang H, von Thaden T, Sharma G, Mitchell A (2002) A synthesis of safety culture and safety climate research (Technical report ARL-02-3/FAA-02-2). Aviation Research Lab, University of Illinois, Urbana-Champaign
- Wilkinson, S. (1994) The Oscar November incident. *Air Space*, pp 80–87 (February–March)
- Woods DD, Shattuck LG (2000) Distant supervision: local action given the potential for surprise. *Cogn Technol Work* 2:242–245