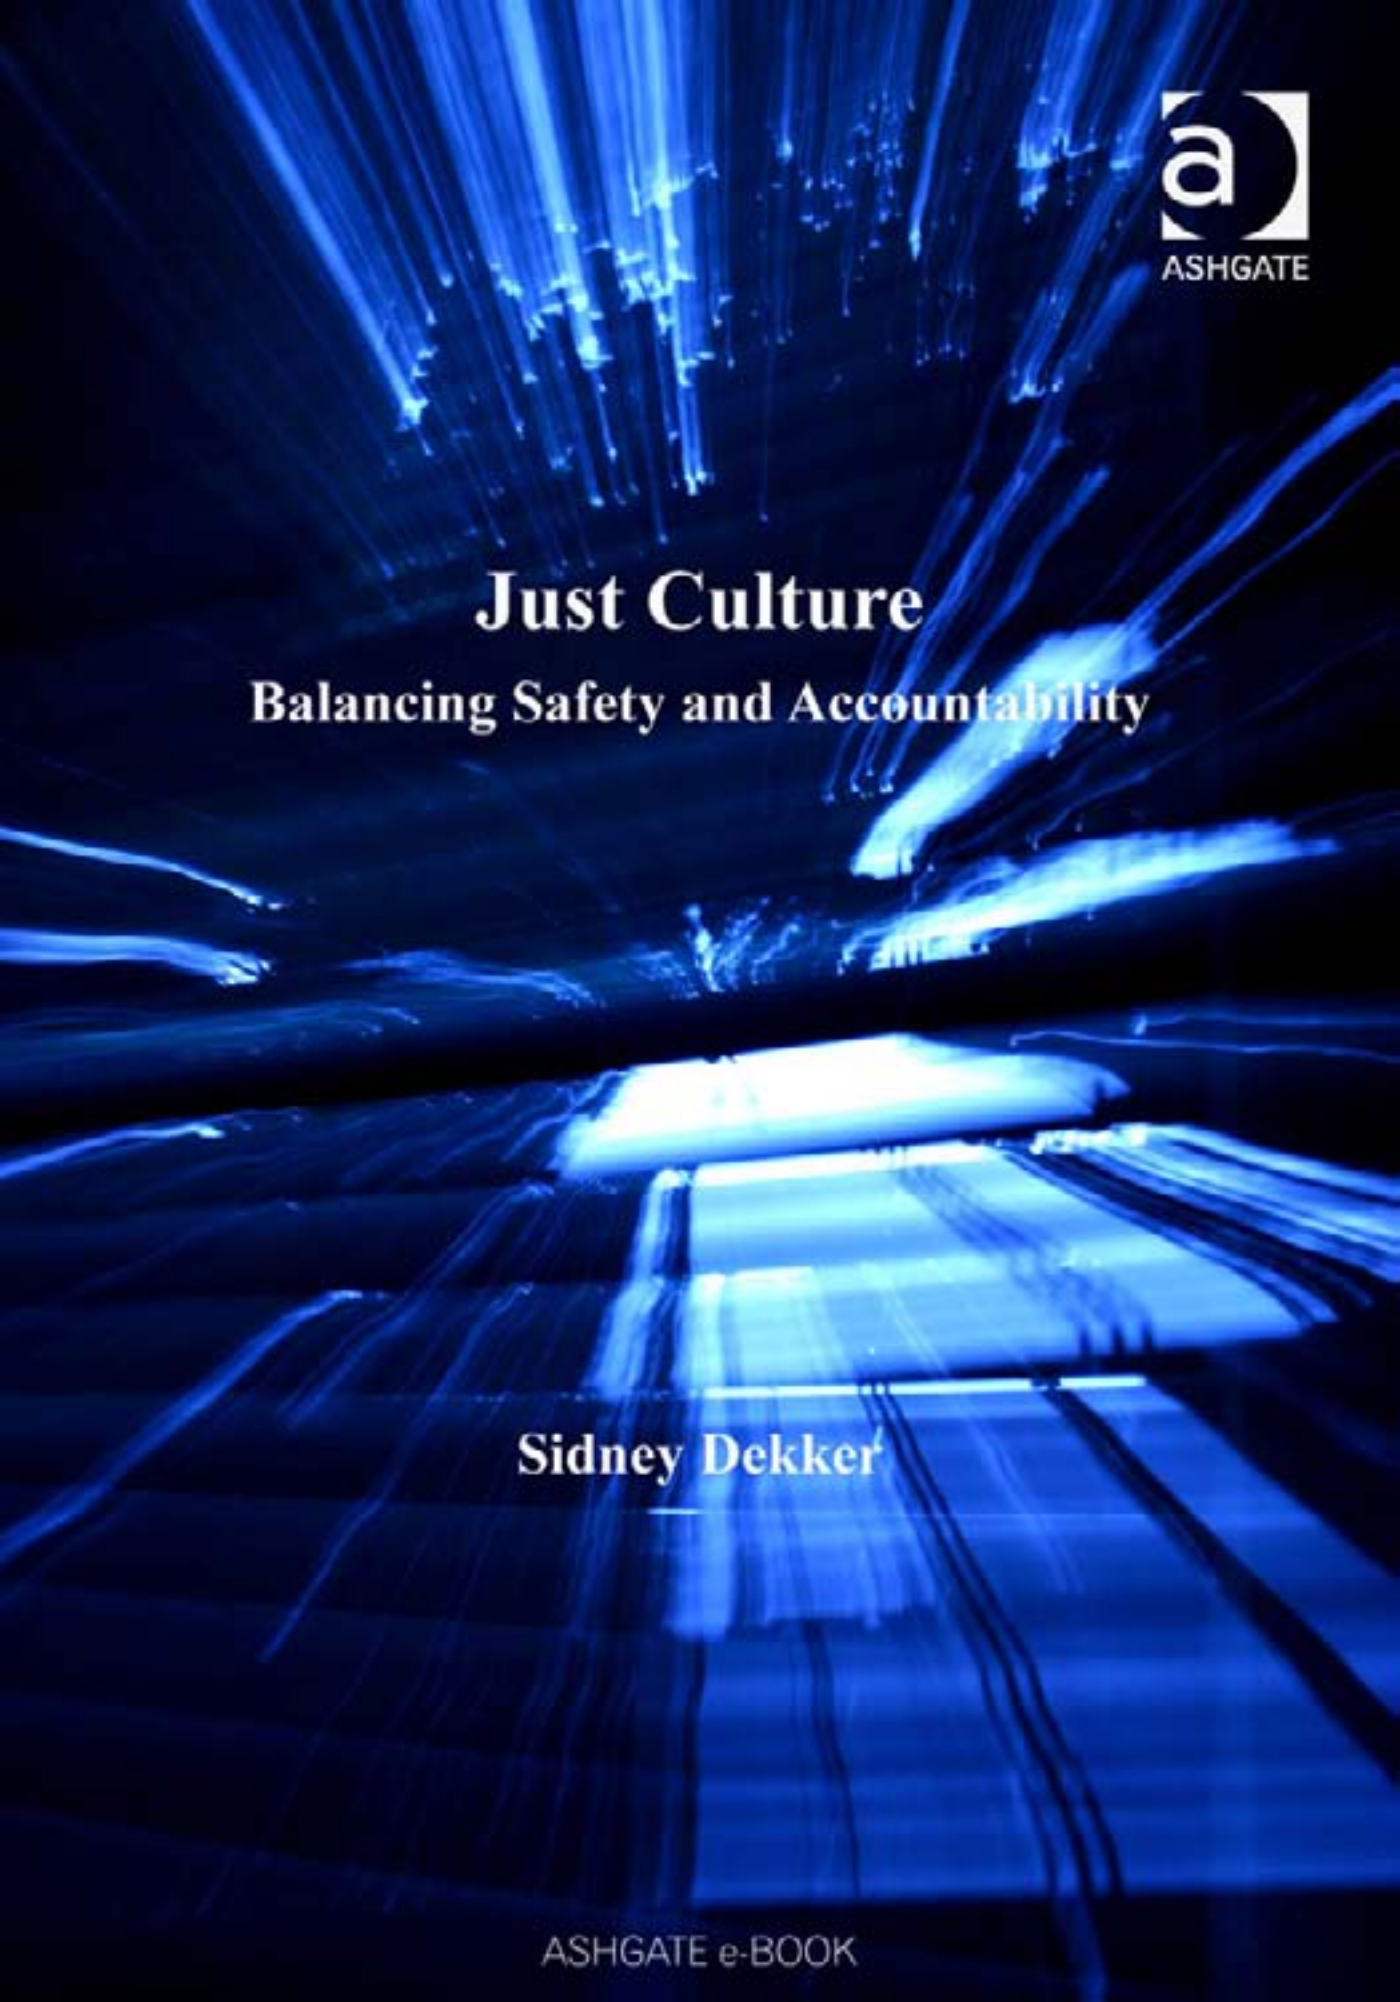
ASHGATE
Just Culture
Balancing Safety and Accountability
Sidney Dekker
ASHGATE e-BOOK

JUST CULTURE

To:
F.C.K., C.L., H.N., E.H. and P.N.

Just Culture
Balancing Safety and Accountability

SIDNEY DEKKER
Lund University, Sweden

ASHGATE

Published by
Ashgate Publishing Limited
Gower House
Croft Road
Aldershot
Hampshire GU11 3HR
England

Ashgate Publishing Company
Suite 420
101 Cherry Street
Burlington, VT 05401-4405
USA

Ashgate website: http://www.ashgate.com

British Library Cataloguing in Publication Data
Dekker, Sidney
Just culture : balancing safety and accountability
1. Professional ethics 2. Corporate culture
I. Title
174.4

Library of Congress Cataloging-in-Publication Data
Dekker, Sidney.
Just culture : balancing safety and accountability / by Sidney Dekker.
p. cm.
Includes bibliographical references and index.
ISBN: 978-0-7546-7266-1 (hbk) -- ISBN: 978-0-7546-7267-8 (pbk) 1. Professional ethics. I. Title.

BJ1725.D45 2007
174' .4--dc22

2007027819

ISBN: 978-0-7546-7266-1 (HBK)
ISBN: 978-0-7546-7267-8 (PBK)

Printed in Great Britain by MPG Books Ltd, Bodmin, Cornwall.

Contents

This page intentionally left blank

List of Figures and Tables

Figure

Table

Acknowledgements

A number of people have contributed to this book with ideas, questions and examples. A Eurocontrol grant to study the interface between the judiciary and operational organizations was helpful in creating time for writing, some of which ended up here, particularly in Chapters 10, 11 and 13. For this, I want to thank Tony Licu and his people in Brussels. The material in those chapters has been reproduced with kind permission of the European Organisation for the Safety of Air Navigation (Eurocontrol) ©2007 Eurocontrol. All rights reserved. Other people were instrumental in shaping ideas for Just Culture that ended up in this book, especially Tom Laursen in Zürich, Rob Robson in Winnipeg, Tom Hugh in Sydney, John Senders in Toronto, and Hans Houtman and Bert Ruitenberg in Amsterdam. I thank them all for their help and inspiration.

Preface

Keith Ramstead was a British cardiothoracic surgeon who moved to New Zealand. There, three patients died during or immediately after his operations, and he was charged with manslaughter.[1] 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 Ramstead.

From Acts of God to Culpable Mismanagement of Risk

We have not always looked at three dead patients as evidence of a possible crime, or as any form of reprehensible behavior. Turning to human error as explanation for an accident, and making it into a culpable act or a crime, is only a very recent way of dealing with failure. In fact, the whole idea of an "accident" is relatively modern.[2] Up until the scientific revolution in the seventeenth century, we apparently had no need for a concept like "accident." Religion and superstition supplied ample explanatory models for things that went wrong. We called it fate, predestination, God's will, witchcraft, taboo-breaking. Where misfortune was going to hit was—as far as mortal humans were concerned—random, uncontrollable.

It stayed this way for the next couple of hundred years, though ever fewer people bought into the idea that accidents had divine or demonic incitement behind them. In the early twentieth century, we began to see accidents as unfortunate but otherwise meaningless coincidences of space and time. As random physical events, though, we still did not consider accidents worthy of study. And we judged attempts to predict and prevent accidents as largely useless.

Over the last 30 years, however, this interpretation of accidents has shifted dramatically. Startling failures such as the Three Mile Island nuclear accident in 1973 and the collision of two jumbo jets at Tenerife in 1977 moved accidents back onto the centerstage of our societies.

No longer do we see accidents as meaningless, uncontrollable events. On the contrary: accidents are evidence that a particular risk was not managed well enough. And behind that mismanagement, there was a person, or multiple people. Today, even though we use the word quite freely, we have actually drifted from the idea of "accident" altogether. We spend huge amounts of resources on formally investigating all large accidents. Why would we do that if accidents are random events, if they are meaningless coincidences ("really" accidents)? We could investigate meaninglessness all we want, but there would be nothing to discover, nothing to change. No, we expect experts to make accidents comprehensible. We want them to explain which risk factors were not controlled, where and when and by whom. Accidents are no longer accidents at all. They are failures of risk management.

Failures of risk management invite us to look for somebody who was responsible. If misfortune hits today, we really don't see it as random or uncontrollable any longer. We often want to find out who didn't do her or his job. And then we want to put the "accident" on their account.

A Trend towards Criminalization

So the trend towards criminalizing human error, in a variety of fields of practice, is a relatively recent phenomenon. It has a lot of people worried, and understandably so. In Chapters 7, 8, 9, and 10 I will deal extensively with the problems of criminalization. For example, I will note how we delude ourselves that there should be consequences for operators or practitioners who "cross the line." I will explain how we don't realize that lines don't just exist "out there," ready to be crossed or obeyed, but that we—people—construct those lines, that we draw them differently every time, and that what matters is not *where* the line goes—but *who* gets to draw it.

Criminalization is only one bookend on a longer shelf of challenges with what we call "just culture." A just culture is something very difficult to define, as "justice" is one of those essentially contested categories. We will never agree with each other about what justice means, or what is just versus what is unjust. Essentially contested means that the very essence, the very nature, of the concept is infinitely negotiable. But that does not mean that we cannot agree, or make some progress on, some very practical problems related to what we could call a just culture.

A very daily challenge for many people in safety-critical domains, for example, is simply to get practitioners to talk about safety problems, to send

in reports, to honestly disclose. Building up trust, and giving people a sense of ownership and participation in system safety improvement is difficult enough. I start with that in the very next chapter and continue in Chapters 3 and 4, pausing to zoom out onto a case that takes various angles at just culture in Chapter 2. In Chapters 5 and 6, I look at the effects of hindsight in determining culpability and how different constructions of error (as technical or normative) have quite different ramifications for what we wish to do about it. In Chapter 12, I take the problem of the division between old view and new view of human error head-on. If human error is a symptom of trouble deeper inside a system, then we can simply blame the system. But what happens with people's accountability then? Chapter 13 gives you some concrete steps to go forward with building a just culture.

Different Interpretations, Different Countermeasures

To charge professionals like Keith Ramstead with a crime is just one possible response to failure. It is one possible interpretation of what went wrong and what should be done about it. As I try to indicate throughout the book, other ways are possible too, and not necessarily less valid:

- For example, one 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 six months?
- We could also see it as an organizational problem: there was a lack of quality control procedures at the hospital, and Ramstead 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 responsibility for implementing those changes?

The point is not that one interpretation is right and all the others wrong. The point is that multiple overlapping interpretations of the same act are always possible (and may even be necessary to capture its full complexity!). And all interpretations have different ramifications for what people and organizations think they should do to prevent recurrence.

Some interpretations, however, also have significant negative consequences for safety. They can eclipse or overshadow all other possible interpretations. The criminalization of human error seems to be doing exactly that. It creates many negative side-effects, while blotting out other possible ways forward. This is unfortunate and ultimately unnecessary. Unjust responses to failure, as I will argue in the last chapter, are not about bad performance. They are about bad relationships. And relationships can be managed. Just as nowadays we believe that risk can be managed.

Notes

1 Skegg, P.D.G. (1998). Criminal prosecutions of negligent health professionals: The New Zealand experience. *Medical Law Review*, 6, 220–46.

2 Green, J. (2003). The ultimate challenge for risk technologies: Controlling the accidental. In: J. Summerton and B. Berner (eds), *Constructing risk and safety in technological practice.* London, UK: Routledge.

Prologue
A Nurse's Error Became a Crime

Let me call her Mara.

It was on a Friday in March that I first met her. I had no idea what she would look like—an ICU nurse in her late forties, out of uniform. This could be anybody.

As I bounded up the stairs, away from the train platform, and swept around the corner of the overpass, there she was. It had to be her, at least from what I had been told. Late forties, an intensive care nurse of twenty-five years, a wife, a mother of three.

But now a criminal convict. An outcast. A black sheep. On sick leave, perversely with her license to practice still in the pocket.

We exchanged a glance, then embraced.

What else was I to do, to say? The telephone conversation from the night before fresh in my mind, here she was for real. Convicted twice of manslaughter in the medication death of a 3-month-old girl. Walking free now as her case before the Supreme Court was pending.

I stepped back and offered: "This sucks, doesn't it?"

She nodded, eyes glistening.

It was her all right. There hadn't been many other people around in any case.

"I recognized you from a video of a lecture you held," she explained as we turned to go down the stairs to go meet her lawyer.

"And how kind of you to travel up all this way."

"Well, it's the least I could do," I said.

Snow was everywhere. Unyielding, huge piles, blanketing the little town. The lawyer's address was distinguished. The most prominent address in the town, in fact. An imposing building in stately surroundings, with spacious offices, high ceilings, the quiet reverence and smell of an old library, its inhabitants

in archaic dress, using archaic language. For them, this was big: headed to the Supreme Court.

As I looked over to where Mara sat, I could not help finding her so out of place. A fish on the shore, gasping, trying to make sense of its surroundings as the burden of a final crawl for survival started sinking in. How on earth could a normal, diligent nurse who had practiced her entire adult life ever have expected to become the lead character in somebody's lofty law offices for a prelude to an appearance at the nation's highest court? Surreal. It must have felt surreal to her. She certainly looked as if it did.

As it turned out (how naïve I am), there is no substance to speak of in a defense before the Supreme Court, because it's all form. Mara began to discover this too, haltingly, stumblingly, and increasingly disgustedly.

"All I want is the truth to come out," she repeated.

"It won't," the lawyer found himself explaining over and over. "This is not about truth. It's about procedure and legal interpretation, and whether it has been correctly followed and applied. All we want is to get you off the hook. What we have to show is that the course of justice so far has been improper—the truth is secondary."

"But what about all the other people involved?" Mara appeared to become anguished. "The pediatrician, the prescription that magically disappeared days after the death, the nurse who administered the medication, the doctors who didn't really diagnose, the lousy routines at the hospital, what about them? The truth is that they are all part of this too!"

The lawyer turned to ice. "They are not on trial now, are they? This is about you. You are the only one. As soon as we bring them up in the Supreme Court, they will ask me 'So where are those co-defendants then, counselor? We thought this case was about the nurse, not all these others.' So don't bring it up, I plead with you, don't bring it up."

Mara seemed exasperated. If justice was like this—disinterested in truth, directed through dogmatic decisions by outsiders that limited what was relevant from the events that got her here, with people putatively helping her by telling her not to argue for her case—then why bother at all? Was it worth it? Justice was supposed to be about getting out the real story, what really happened. That would be just. Justice would be about righting what was wrong, and about avoiding that it would happen again. That would be just too. Yes, she made a mistake; yes, a baby died. She knew that. But she also knew that the entire system in which she worked was rotten, porous, and ready to kill again.

But it was plain to me that Mara knew why she was here. It wasn't just because of her, because of her role, because of her fate, or because everybody was suddenly gathering around invigorated efforts to make healthcare safer.

She knew who was paying her lawyer, and it wasn't her. Fewer than 1 per cent of cases presented actually get heard by the Supreme Court in my adopted country, and hers was among them. It must have mattered, somehow. The country had taken interest. The union certainly had too. Should medical practitioners involved in a patient's death be subject to the criminal justice system? Or should they be dealt with through the established professional channels: the medical disciplinary board? A lot was at stake, that much was obvious to Mara. Realizing that she may have used the stronger solution of the medicine, she had volunteered her possible contribution to the baby's death to her boss a few days after it had happened. Her boss duly reported the event to the relevant agency, but somebody also leaked it to the local press. Mara never found out who. Her role was played up, and a prosecutor happened to read it in the morning paper. After months of uncertainty—Mara even called up the prosecutor herself one day to get clarity about her intentions—charges were brought. A local court found her guilty of manslaughter. The conviction was upheld by a higher court, which toughened the punishment. Now the case was headed for the Supreme Court. Would people in healthcare ever volunteer information about incidents again? Was this the death knell for nascent medical event reporting systems? Was patient safety going to be dealt a serious setback?

We wandered back into town, in search of a cup of coffee.

When we had slipped into the warmth of a bakery, shaken the snow off our shoulders and sat down near a window, I cocked my head, glanced at her, and sighed, wonderingly. She must feel like a vehicle, sent out to test-drive the law, I mused. If the country and its healthcare system would get their day in court, if they were going to create clarity on the rules for dealing with medical error, then this was not going to help Mara. The black sheep would be herded through one more splendid spectacle of public judgment, but it was no longer about the sheep, if it ever was. It was about the principle. And she was merely its embodiment. When it was all over, whatever the outcome, she would have been used up. Her purpose to larger interests played out, expired. A mere piece of detritus mangled through a criminal justice system in its quest for new turf, disposed once the flag had been planted. She would be remembered only in faint echoes of *schadenfreude* ("thank God it wasn't me") and a waning trail of half-hearted collegial compassion ("we're so sorry for you, Mara"). The disillusionment with her work setting, her colleagues, her union, the justice

system—the world—was etched on her face. Vindication would remain elusive. The truth would not come out.

But is there truth in the aftermath of a little girl's medication death? Or are there only versions?

At the Supreme Court

A few weeks after the meeting with the lawyer, I saw Mara again, this time in the ornate halls of the Supreme Court. High ceilings soared up, up and away from two tables, one for the defense and one for the prosecution. They were arranged in front of a regal podium decked out with a row of seats. When the justices had filed in and sat down facing both teams, the prosecutor reached for his version of the truth. I remember his craftiness, his cultural conformity to the conflict-avoidance of my adopted country. He was sitting down, not standing up. He was reading from a prepared script, not ad-libbing or grandstanding in front of his audience. His tone was measured, quiet, reverential. This is, I suppose, what court proceedings are supposed to do: separate emotion from substance, sublimate conflict into negotiation, turn revenge into ritual.

Mara sat over at the other table, flanked by her lawyer. Hands in her lap, eyes cast downwards. As she sat there, the prosecutor's opening statement started rolling over her, his gentle voice reverberating around the hall unamplified.

Except it wasn't a statement. It was a story.

"The baby was born on February 24 in the regional hospital," he intoned. He recalled the happiness of the child's parents and mentioned details to paint a picture of family bliss, soon to be disrupted by a treatment gone awry.

> She had a normal birth weight, but showed some signs of seizures in her arm after delivery. Three days later, the seizures had become worse. She was given an anti-epileptic (containing phenobarbitone). After stabilizing on the March 5, she was discharged. But less than a month later, the seizures came back. The infant was rushed to the emergency room and taken in for observation. Her anti-epileptic dose got increased to 5 milligrams (mg) per milliliter (ml) and she was discharged again two days later. The day after, her mother called the hospital. After consultation, the dose was increased again—over the phone—to a twice daily 2 milliliter portion of the 5 milligram per milliliter mixture. On April 22 the baby was brought in as part of a routine checkup. Everything was normal.

He paused.

To recount his version of the truth, the prosecutor had created a narrative. Narratives are strong. He must have picked that up in class once, or in one of his many street fights. Or perhaps a story, or liking a story, understanding a story, is simply what makes us all human. Mara must have heard versions of the story hundreds of times now, I thought. She must have turned it over endlessly in her mind, picking away at her role, plaguing herself by retrospectively finding opportunities to not make a mistake, to not become the centerpiece of this imbroglio.

Act One was over. The justices looked at the prosecutor, silently. Spellbound or bored silly? It was difficult to tell. Time to set the stage for a plot twist. Act Two. A different ward: the intensive care unit (ICU). A new medication. And, of course, the introduction of the villain.

> On May12 the baby was admitted with a new bout of seizures, and sent to the ICU. Her anti-epileptic was increased to .5 ml, and she even received a bolus dose of anti-epileptic. But the seizures continued. The baby was then given Xylocard, a lidocaine-based medication, in a concentration of 2mg/ml. The seizures subsided. She was discharged again on May 16, off Xylocard, and back on the previous dose of anti-epileptic. But on May 18, her mother called the hospital to say that her baby was suffering a new onset of seizures, now lasting about 5 minutes each. In the evening, the child was taken to the hospital by ambulance and admitted to the pediatric ward. New seizures made that she was transferred to the ICU later that evening.

With the baby back on the scene of the crime-to-come, everything was ready for Mara to make her entry. The lines of the two lead roles could now converge.

> Early in the morning of Sunday May 19, Mara showed up for work. There were not many patients in the ICU, things were quiet. The baby was doing better now. In preparation for her transfer back to the pediatric ward, Mara went to the medication room to mix the Xylocard solution.

He paused and picked up the two little cartons in front of him on the table. Then he waved them around.

> There, in the cabinet, were two packages. One containing an intramuscular (IM) injection dose of 20 mg/ml Xylocard, and one with a 200 mg/ml Xylocard solution intended for intravenous, or IV, drop. Misreading the packages, Mara took the 200 mg/ml package to prepare the baby's drop, instead of the 20 mg/ml package, as was prescribed.

The chief justice motioned that she wished to see the packages. They were handed over. Passed from justice to justice, they got handled for what they were in that context: pieces of evidence in a manslaughter trial. The justices studied the packages with what looked like mild interest, but could just as well have been muffled puzzlement. What kind of evidence was this anyway? This was not just a common criminal instrument—a knife, a handgun, a fraudulent contract—this were pieces of highly-specialized medication, excised from their normal surroundings of thousands of normal, similar-looking packages that form the backdrop of a nurse's daily life. Now these two boxes looked quite out of place, floating along the court's elevated regal bench, examined by people with little idea of what it all meant. Questions must have mounted, one on top of the other. What was it with these peculiar Greek neologisms, and why were all these boxes white with green or light-blue lettering, and what were these befuddling volume-weight fusions?

The prosecutor continued. Not much longer now. Act Three. A rapid climax.

> That afternoon, back in the pediatric ward, the baby was hooked up to the new Xylocard drop, the one that Mara had mixed. But instead of subsiding, the infant's seizures quickly got worse. A pediatrician was called, and tried to intervene. But nothing helped. Not long after, the baby was declared dead. Post-mortem examination showed that she had died from lidocaine poisoning.

A story that makes sense, that is plausible, that has a powerful narrative arc and casts characters in recognizable roles of hero, victim, villain, bystander, can present a rather believable truth. And the prosecutor's story did. His plot painted a normal hospital, a normal, innocent little patient, attended to by normal physicians, suddenly all confronted by the sinister and totally unnecessary turn of events on a Sunday morning in May—the fatal denouement of Mara's mix-up. Quite impeccable. Quite logical.

A Calculation Gone Awry

But does that make it true? Consider another truth, the sort of "truth" that Mara had hoped in vain to bring out in the open on this day. After clocking on, the morning of May 19, she received a little briefing from the night ICU nurse. The original prescription had been unclear and not been signed by the doctor who wrote it. The hospital (even the ICU) was equipped with a computerized prescription system, but the physician had been sitting at a terminal that

happened to not be connected to the printer. Rather than moving to another terminal and print out a prescription, he wrote one by hand instead. Earlier that night the nurse had mixed a Xylocard solution with another physician's help, trying to divine the prescription. Now, in the morning, the doctor himself was asleep somewhere in the hospital and, given that it was a quiet Sunday, nurses would not become popular by waking him up to ask a simple clarification. The night nurse showed the unsigned prescription and her medication log entry to Mara:

"40 ml + Xylocard 200 mg = 10 ml = 4 mg/ml, total of 50 ml."

"Remember, 10 ml Xylocard," the doctor had said to the night nurse, who now relayed this to Mara. The infant, in other words, had received a total of 200 mg of Xylocard by mixing two 5 ml syringes each containing 20 mg/ml Xylocard, (the standard IM injection package, each with a total of 100mg Xylocard per syringe): into 40 ml of glucose solution. But in the ICU, syringes were never used for IV drops because they contain a weak solution. Syringes were for direct IM injection only. The ICU used vials, with a stronger solution, for IV drops. But pediatrics did not even have vials. They dealt with children, with little bodies, that needed no strong solutions. Pediatrics routinely discharged the prepackaged syringes into an IV drop instead. The ICU seldom had little infants, though, and no tradition of using syringes, rather than vials, for preparation of IV drops.

Later that day, when the night nurse had long gone home, Mara noticed that the infant's drop was running low and decided to prepare a new one. The baby would be transferred back to pediatrics, but the move had gotten delayed. Remembering the "10 ml" reference from the doctor, and reading 200 mg off the medication log (as the prescription was unclear), she took two boxes in each of which was a 5 ml vial containing 200 mg/ml Xylocard—10 ml total, and the figure of 200 mg—this was what the medication log said. She prepared the solution and wrote in the log:

"Xylocard 200 mg/ml = 10 ml = 4 mg/ml."

Mara showed her calculations to another nurse, and also the pediatrics personnel who came to collect the infant. The pediatrics staff did raise a question, but it focused only on the dose of 4 mg/ml, not the solution from which it supposedly would come. Five days earlier, when the infant had been with pediatrics too, she had been on 2 mg/ml, not 4 mg/ml. The ICU confirmed to

pediatrics that 4 mg/ml was now the prescribed dose. The baby was to receive 50 ml of the solution that was supposed to contain 4 milligrams of Xylocard for each milliliter.

But did it?

That night, Mara tossed in her bed. Her youngest son woke a couple of times, rendering his mother restless. In the darkened bedroom, the events of the day came back to her. As far as she knew, the baby lived, she had gone off shift before anything went awry. But something did not quite add up. Why had the night nurse, normally so assiduous, accepted such a messy and unsigned prescription? She had even had to call help from a physician to mix the thing. And what about that log entry of hers? It had read "Xylocard 200 mg," but did that make sense? Xylocard 200 mg was meaningless by itself. 200 mg per what? Per …?

Mara sat up with a start.

Could it be true that she had taken two vials, instead of two syringes? They both contained 5 ml of fluid each, so any combination of two would amount to the 10 ml the doctor had wanted. The two packages were side by side in the cabinet which was so neatly organized on alphabet. But two vials meant …

She quickly ran the numbers in her head, peering into the darkness. Two 5 ml vials both containing 200 mg/ml Xylocard would have amounted to 2000 mg Xylocard, or 40 mg/ml, not 4! This would add up to a lot for a little infant. Too much maybe. In that case her medication log entry didn't make sense either. Take 10 ml with each ml containing 200 mg, and you would not get 4 mg/ml. You'd get an order of magnitude more. Ten times more. Forty.

Why had nobody caught it? She had had people double-check! Pediatrics had checked! Also, an entry about the solution would have had to be made on the IV drop before it went into the child—another double-check. What had happened? She would try to figure this out as soon as she was at work again.

On her next shift, Mara asked about the little girl. "She has died," was the answer. Her heart must have sunk. But determined to figure out what had gone wrong, and if she may have had any role in it, Mara went to the binder with prescriptions and flipped back to Saturday night. Where was it? Where was the prescription, that messy, unsigned prescription, that her predecessor night-nurse had interpreted as "200 mg" Xylocard, setting her, Mara, up for a possible mistake?

The prescription was gone. It wasn't there. It had disappeared and would never be found again.

Years later, only a few weeks before the hearing at the Supreme Court, Mara would plead with her lawyer to bring up the missing prescription. He yielded not an inch.

"How can you bring up something that doesn't exist?" he asked.

"But," Mara countered, "we are not allowed to prepare medications without a prescription, there *has* to be a prescription, and in this case there was too. Somebody took it out!"

The lawyer sighed and was silent.

"Look," he said after a while. "This is not the time to introduce new evidence. And even if it was, you can't produce as evidence something that you don't have. It's that simple."

Mara's world must have spun around her. She was locked up inside a Kafkaesque entanglement that had erased any resemblance with the real world. Her mind must have cast around for anything stable, anything recognizable, anything sensible. Instead it was finding nothing to grab onto, no lifeline, no help. And no "truth."

Mea Culpa

What there was, and what had been introduced as evidence, of course, was her own medication log entry. The one that said that 10 ml of fluid, with each ml containing 200 mg of stuff, would amount to a measly four mgs of the stuff per ml in a 50 ml IV drop. It wouldn't. It would yield ten times as much. She had recorded her own miscalculation—putting the truth out there, for all to see.

Not long after learning of the baby's death, complying with reporting procedures in the hospital, she wrote to her superior:

> When I was going to mix Xylocard at around 10:45 that morning, I looked at the prescription and got Xylocard 20 mg/ml. I read both the package and the vial and recall that it said 20 mg/ml. I looked at what was prescribed and what I should prepare. So I got 20 mg/ml which I mixed with glucose 5% 40 ml.
>
> I asked another nurse to doublecheck but did not show her the empty vials. Then pediatrics came to get the infant, … and they took my prepared solution with them to hook it up in their ward. When the infant left us at 11:07, there was still about 3 ml in the previous drop, which had run through the night of 18–19 May.
>
> The following night, I awoke and suddenly realized that a vial normally contains 1000 mg/5 ml. And I had thought that I drew a solution of 20 mg/ml. When I was working the following Wednesday, I got to hear that the infant

had died. I then understood that it could have been my mistake in making the solution, as there are no vials of 20 mg/ml.[1]

Stories of mistake can be so simple. "My mistake," Mara had recorded. *Mea culpa*. To many others in the hospital, such an unprovoked admission must have been a godsend. Not that they would ever say, of course. They would not have to. The legal aftermath itself would prove them right. Mara was in the dock. Again and again. Nobody else.

Not that this would necessarily feel natural to anyone involved in the saga as it unfolded. Take a story as experienced from another nurse's point of view. When the infant started to show an increase in seizures and other problems after being hooked up to Mara's IV preparation in pediatrics, nurses called the attending physician.

He responded by phone: "Up the flow, give her more."

They did. The problems got worse.

They called again. "Give her more, give her a bolus dose," was the instruction again. They did.

But this did not seem to help at all—in fact, things were going from bad to worse very quickly now. The attending anesthetist was now called by phone, but nobody answered. Another was found by calling through the intercom, but nobody showed. Only minutes later did the attending pediatrician show up in person. He ordered another bolus dose of Xylocard, but this had no effect either. The baby now needed 100 per cent oxygen but she started vomiting into the mask, exacerbating her respiratory problems. The pediatrician ordered another bolus dose of Xylocard, thinking that this would finally stop the seizures. Then, during one attack, the girl presented respiratory failure. The pediatrician responded by intubating the baby, and cleaned the airways by suction. Then the anesthetist arrived. The baby got ventilated but the suction tube proved too narrow for her passages to be cleared. Another bolus dose of Xylocard got pumped into the IV port. Finally, a thicker tube was found and inserted, clearing her airway. It was all too late. The infant went into circulatory shock. Adrenalin, atropine and Tribonat were given, heart massage administered, the defibrillator was even pulled out. To no avail. The baby was declared dead not long after. A postmortem showed that the girl had ended up with 43 mg of lidocaine per gram of her blood. The therapeutic dose is less then 6 mg per gram of blood.

Even if Mara had mixed from the 20 mg/ml syringes and not the 200 mg/ml vials, the infant would still have ended up with twice the therapeutic dose due to the volley of bolus shots during her final moments. Yet that is

but one "truth" too. See the world from the pediatrician's perspective and another sensible story swims into view. The initial symptoms of lidocaine poisoning can include (yes) seizures. So the symptoms of too much Xylocard and too little Xylocard would have been similar, setting the physician onto a compelling plan to continue. Strong initial cues suggested his response was the right one. They had been confirmed before: this baby had responded well to the treatment of her seizures with lidocaine. The dose had been upped before, with good therapeutic consequences. He knew all this. And, for that matter, he never knew that the IV drop was administering the drug at ten times the ordered rate. The quality of his assessments and decisions could impossibly be rated against that knowledge—knowledge he did not possess at the time. That much would be "true."

But what did the doctors actually know? I remember Mara countering this even before the final trial. "Did they ever diagnose the source of the spasms?, Mara would ask. No, they didn't. Did they have any idea why the child responded better to Xylocard than to the previous anti-epileptic, even if Xylocard is not mainly intended to deal with seizures? Did anybody ever think to call in a neurologist? No. Did they ever ask themselves why the baby would suddenly develop such intense symptoms after getting back to pediatrics on Sunday afternoon? Not that Mara knew. Did they ever recognize their own role in the slippage of prescription routines? In taking a nap at work on a quiet Sunday morning and being really grumpy when awoken for no apparent good reason? In not bothering to get up and mosey ten feet to another computer to print out a prescription for Xylocard, rather settling for a bunch of handwritten squiggles instead? In not showing up for many, many critical minutes when a little baby was suffocating in her own vomit, wasting away on some drip? And then giving order after order after order of poisoning lidocaine? No, not that Mara would be aware. And who took that prescription away after the baby died? Where was it? And whose idea was it to start swapping a baby between pediatrics and the ICU, a ward designed in every way for taking care of big people, not little ones? Was any of those "truths" ever going to be brought out?

Criminal Law and Accidental Death

A legal system holds people accountable. But it does not allow people to hold their account. Mara had become a hostage of legal procedure and protocol, and she decried the shackles on what she was granted to say and when. At every turn in the legal plot, she went in to battle the limits, to break through

the constraints. She wanted permission to give her account. She just wanted the "truth" to come out. But at every end, she came out broken herself. Her account would still be inside of her—biting, festering. And increasingly bitter and partisan.

A legal system constructs an account from its own pick of the evidence. It makes its own story. It is interesting that society may turn increasingly to their legal systems to hand out that story, to provide accountability after a terrible outcome. There must be something in that account that we find terribly attractive; more enticing than what the people have to say who were actually there. Mara, for example.

Of course, we could dismiss their accounts as exculpatory, as subjective, biased, ulterior. Still struggling to understand her own performance, Mara had told a lower court that she may have misread the package labeling. By the time she got to the Supreme Court, however, she indicated that this was probably not the case: she mistakenly believed that 200 mg/ml was what she needed to have. This would certainly have made sense, given the prominence of the figure 200 in the medication log, and the reminder to end up with a volume of 10 ml Xylocard in total. But look at how the Supreme Court chose to interpret the various accounts that Mara had tried to provide. Put up as a last-gasp attempt to exonerate herself, to "find an explanation afterward," the Supreme Court painted Mara as ditzy when it came to getting an account of what had happened that Sunday in May:

> During the court proceedings, the ICU nurse described multiple ways how it could be that she mixed the IV drop with the wrong concentration of Xylocard. What she offered cannot therefore express what she really remembers. Rather, her accounts can be seen as attempts to find an explanation afterward. They are almost hypothetical and provide no certain conclusion as to why she did what she did.[2]

Whatever Mara offered, the sheer variety of her accounts had disqualified her as a purveyor of truth. In her stead, the Supreme Court was happy to provide the "certain conclusion" so sorely lacking from Mara's story. They speculated why Mara did what she did: she either "misread, miscalculated or took the wrong package" from the shelf—all because of "negligence." Mara did what she did (whatever it was), because she was careless. "She could have read the medication log more carefully, calculated more carefully or done any other double-check that would have revealed her error and its potentially fatal consequences."[3] But she did not. She was negligent. In the absence of a story from Mara that made sense, people turned to the legal system to serve them

a story with a cause and a culprit. The cause was misreading, miscalculating or grasping wrong due to negligence, and the culprit was Mara. Instead of listening to the protagonist, people legitimated a particular institution to get at the "truth" and mete out supposedly appropriate consequences. They may have thought, as many increasingly do, that this legitimated authority could deliver the veridical account—what really happened. For the one who was there could not be trusted to deliver an account that "expressed what she really remembered. After all, she could "provide no certain conclusion as to why she did what she did."

Of course, judicial proceedings do rely on the insider account as part of their evidence base. Mara was given a voice—here and there. But she never called the shots. She spoke when spoken to: merely proffering a hunch of answers to often inane questions gurgling from a tightly scripted ritual:

"So what did you read on this package, did you read anything at all, or did you take fluid directly from the vial?," the prosecutor in the Higher Court had insisted.

"I looked at both the package and the vial," Mara had replied.

"What did you see?"

"I don't know, I wrote 200 mg per ml, but I don't know"

"You don't know."

"No."

It sounded exasperated—feigned or real: "You don't know." If Mara did not know, then who would? She had been there, after all. Again, the incapability to give that final account, that deeper insight into the workings of her own mind that day, was taken as reticence, as foot-dragging. "You don't know" was taken, as it often is by the time adversarial positions are lined up in a criminal trial, not as "you really don't know," but as "you don't want to tell us."

I recall sitting in the lawyer's office with Mara when she offered the explanation in which she really believed she had taken the right package, the one she was supposed to take (as that was always the one she prepared IV drops from). There was no misreading, that had been a wrong explanation. But the Supreme Court justices would have none of that. They would not see the latest account as a genuine attempt of the insider to articulate what had happened, but as a ditch from the debris, as a ducking of responsibility.

Rational Systems that Produce Irrational Outcomes

And so we turn to our legal system to furnish us with the truth. The US Supreme Court put it most bluntly back in 1966: "The basic purpose of a trial is the determination of the truth."[4] Turning to a legal system to tease out accounts of failure would be an immaculate capping of the Enlightenment project: deploying rational techniques rather than institutional authority (such as that of the Church) to arrive at veridical accounts and the appropriate moral rules these suggested. People could now turn themselves, as individuals, to an objective, reasoned route to reality and morality. But intense attempts at deploying rationality, as the German sociologist Max Weber warned over a century ago, quickly deliver the opposite. The output of supposedly rational institutions is often—quite naturally, necessarily—irrational. There were many, both inside and outside the healthcare system, who thought just that about Mara's verdict. When a nurse herself reported a mistake in an honest effort to abide by the rules and perhaps help prevent recurrence, it made no sense at all to have her end up convicted of manslaughter for the very mistake she voluntarily divulged. This was irrational. Even more poignantly, why her? Singling out Mara for this adverse outcome of a discontinuous, wandering processes of care delivery that counted many contributions from many contributors, made no sense whatsoever. And then, this was not the first or only medication adverse event ever; not a uniquely egregious occurrence. In the same year that Mara was first charged, more than three hundred severe medication errors were reported to the country's health authority. Adverse medication events are "normal." They are the rule, or at least part of it, baked into the very fabric of delivering assorted compositions of volumes and weights and rates of substances through various means. This, moreover, is accomplished through a thoroughly discontinuous process, where gaps in the delivery of healthcare open up because of changes of medium (for example, from oral to written to oral prescriptions or dosage orders), handovers from one caregiver to another between shifts, movement of patients between wards, transferal of caretaking physician, or other interruptions in workflow. Patients, prescriptions, orders, medications, and healthcare workers all cross departments, shift responsibilities, flow through hierarchies and traverse levels of care as a matter of routine. It would be easy, then, and quite rational to show that Mara's adverse event was part of a systemic feature of healthcare delivery. So how a supposedly rational judicial process could come to the exact opposite conclusion is something that Weber would not have found surprising. The accounts of human error that a legal system produces can be so bizarre precisely because of its application of reason: the way judicial proceedings rationalize the

search for and consideration of evidence, closely script turn-taking in speech and form of expression, limit what is "relevant," are institutionally constrained in their deferral to domain expertise, and necessarily exclude the notion of an "accident" because there is no such legal concept.

When you come up close—close enough to grasp how case content becomes subjugated by judicial form, close enough to hear the doubts of the victims about the wisdom of having a trial in the first place, close enough to taste the torment of the accused, to feel the clap of manacles around the expression of their own account, to experience the world from the dock and sense the unforgiving glare it attracts—a more disturbing reality becomes discernable. In the view from below, there is a deep helplessness: an account is created by non-experts, who select bits and pieces in a process that runs its own course and over which there is very little, if any, external control. To those present when the controversial event happened, and who may now be in the dock (as well as to many of their co-practitioners), the resulting account may well be bizarre, irrational, absurd. And profoundly unfair.

The Shortest Straw

Mara had hoped that the process in the Supreme Court would end up bringing out a real version after the acrimony in lower courts. It did not. Instead of truth, she got an upheld conviction. Instead of getting vindication, she remained the villain.

Sitting in the twilight in her living room on a rainy day late in August, months after the hearing, I began to believe that her psychological devastation was due not just to the Supreme Court upholding the guilty verdict, including its heavier penalty. This may not even have been the chief source of her anguish. With her license to practice still intact, and the sentence turned into conditional time, it had few overt practical consequences (not that she could, or wanted to practice in the ICU ever again, by the way). No, I started to sense rather a resignation, a disillusionment, a dizzying realization that progress towards truth is not a movement from a less to a more objectively accurate description of the world. She might have hoped that we all could learn the truth behind the death of the little girl. But there is no such truth to find, to arrive at, to dig out. No final account, no last word—only versions, jostling for supremacy, media-light, popular appeal, legal sustainability. And her version had consistently drawn the shortest straw. Again and again.

Notes

1 The text of her report was given to me by the nurse herself. Translated from Swedish.
2 Swedish Supreme Court verdict B 2328-05, 19 April 2006, at 4–5.
3 Ibid.
4 Laudan, L. (2006). *Truth, error and criminal law: An essay in legal epistemology.* Cambridge, UK: Cambridge University Press, 2.

1 Why Bother with a Just Culture?

Building a just culture is hard.

I was in a conversation with two air traffic controllers recently. They were talking about an incident in their control center. They discussed what they thought had happened, and who had been involved. What should they do about it?

"Remember," said one controller to the other, "Omertà!"

The other nodded, and smiled with a frown.

I said nothing but wondered silently: "Omertà"?

Surely this had something to do with the mafia. Not with professional air traffic controllers.

Or any other professionals.

Indeed, a common definition of "omertà" is "code of silence." It seals people's lips. It also refers to a categorical prohibition to collaborate with authorities. These controllers were not going to talk about this incident. Not to anybody else, or anybody in a position of authority in any case. Nor were they going to voluntarily collaborate with supervisors, managers, investigators, regulators.

I live my professional life in occasional close contact with professional groups—firefighters, pilots, nurses, physicians, police, nuclear power plant operators, inspectors, air traffic controllers. I see a "code of silence" enforced and reproduced in various ways. A senior captain, who flies long-distance routes with a large, respectable airline told me that he does not easily volunteer information about incidents that happen on his watch. If only he and his crew know about the event, then they typically decide that that knowledge stays there. No reports are written, no "authorities" are informed.

"Why not?" I wanted to know.

"Because you get into trouble too easily," he replied. "The airline can give me no assurance that information will be safe from the prosecutor or anybody else. So I simply don't trust them with it. Just ask my colleagues. They will tell you the same thing."

I did. And they did.

Professionals under these circumstances seem to face two bad alternatives:

- either they report a mistake and get in some kind of trouble for it (they get stigmatized, they get a reprimand, or they get fired or even prosecuted);
- or they do not report the mistake and keep their fingers crossed that nobody else will do so either ("Remember: omertà!").

The professionals I talked to know that they can get into even worse trouble if they don't report and things come out anyway. But to not talk, and hope nobody else does either, often seems the safest bet. From the two bad alternatives, it is the least bad.

I once spoke at a meeting at a large teaching hospital, attended by hundreds of healthcare workers. The title of the meeting was "I got reported." The rules of the country where the meeting was held say that it is the nurse's or doctor's boss who determines whether an incident should be reported to the authorities. And the boss then does the reporting. "I got reported" suggests that the doctor or nurse is at the receiving end of the decision to report: a passive non-participant. A casualty, perhaps, of forces greater than themselves, and interests other than their own. The nurse or doctor may have to go their boss to report a mistake. But what motives have they to do so? The formal account of what happened, and what to do about it, ultimately rests in the hands of the boss.

A Few Bad Apples?

We could think that professionals who rely on "omertà" are simply a few bad apples. They are uncooperative, unprofessional exceptions. Most professions, after all, carry an obligation to report mistakes and problems. Otherwise their system cannot learn and improve. So if people do not want to create safety together, there must be something wrong with them.

This, of course, would be a convenient explanation. And many rely on it. They will say that all that people need to do is report their mistakes. They have nothing to fear. Report more! Report more! Then the system will learn and get better. And you will have a part in it. Indeed, every profession I have worked with complains about a lack of reporting.

"If only people would report more," supervisors or regulators sigh.

"Our biggest problem is under-reporting," a healthcare specialist commented. "And we don't even know how big that problem is," she added.

Yet often I am not surprised that people do not want to report. The consequences of disclosing mistakes can be quite dreadful. Organizations can

respond to mistake in many ways. In the aftermath of an incident or accident, pressures on the organization can be severe. The media wants to know what went wrong. Politicians may, too. They all want to know what the organization is going to do about it. Who made a mistake? Who should be held responsible? Even a prosecutor may get interested. National laws (especially those related to freedom of information) means that data that people voluntarily submit about mistakes and safety problems can easily fall into wrong hands. Reporting and disclosure can be dangerous.

Not reporting is hardly about a few bad apples. It is about structural arrangements and relationships between parties that either lay or deny the basis for trust. Trust is necessary if you want people to share their mistakes and problems with others. Trust is critical. But trust hard to build, and easy to break.

Responding to Failure: the Organization

Seeing stories about mishaps and their aftermath "from below," from the view of the operator or practitioner at the sharp end, is critical. And really interesting. But it is not enough. The view "from above"—that is, from an organizational leadership—can be equally knotty. How to respond to failure (once you learn about it) can be a wrenching question. If you are in an organization's management, you have to consider many different stakes and interests:

- What serves the organization best?
- What about the individual operator involved?
- What about the public (the consumers of your services or goods)?
- What about the regulator who is watching over you?
- What about your own position or survival as organizational manager?

The organization may even have invested in new technologies or routines, so as to improve services, or make them more efficient. Seeing those involved in a failure creates all kinds of trouble.

A colleague recently received a phone call from a hospital vice-president. A child had died a few days before from a tenfold chemotherapy overdose in their pediatric oncology unit. The hospital vice president led the investigation of this tragedy and had found a number of issues in processes, admixture formulation practices, and problematic new pharmacy technology that aligned to bring about the death of this child. The family was devastated. Everyone involved in the care of the child was devastated. Doing what they

usually did to create and administer chemotherapy admixtures, suddenly and lethally had not worked as intended. The introduction of the new pharmacy device was deemed a substantial factor—it had replaced familiar technology "on the fly" and this was one of the first uses.

The vice president said that he did not believe any of the personnel involved should be punished. Yet, despite his organization's publicly announced plan to develop a just culture, the CEO, CMO, and HR Director insisted on firing the two pharmacists involved in formulation of the admixture and the nurse who had administered the medication. There was absolutely no way the nurse could have known that the content of IV bag was not as labeled. The impetus for dismissal actually came from their consulting ethicist, who also happened to be a lawyer. He identified the child's death as evidence of a breach of 'duty ethic' and hence breach of legal duty—he deemed these three people unequivocally negligent.

How to respond to failure is, at its heart, an ethical question. We can wonder, then, whether it is smart to combine the function of ethicist and lawyer into one person, as was done in the example above. In fact, such a mix could be testimony to the confusion and difficulty of building a just culture:

- Is "just" something that meets legal criteria (for which you need a lawyer)?
- Or is "just" something that takes different perspectives, interests, duties and alternative consequences into evaluation (for which you might need an ethicist)?

Involving the Legal System

A final step into the muddled morass of responding to failure is to involve the legal system. Once this step is taken (or once the legal system starts involving itself), all bets about achieving "justice" are off. In fact, in all the cases I have seen up close, the outcome of a trial in the wake of failure was never "just." (Nor did it improve safety.)

- Victims would typically feel undercompensated and often started wondering about the wisdom of having a trial in the first place.
- The practitioner or professional on trial would definitely feel singled out as scapegoat, unfairly bearing the legal and moral load of the mishap.
- Proceedings would hardly be about the content of the case, and more about arcane legal protocol and procedure (and when they were about content,

they typically, and unjustly, ran roughshod over all kinds of operational subtleties and nuances).
- The organization would feel that it got unjust attention in the media, attention it would gladly do without.
- There would always be a losing side, even if the practitioner got off the hook.

So with a legal system, justice is hard to achieve in the wake of failure. But there is more. When a professional mistake is put on trial, safety almost always suffers. Rather than investing in safety improvements, people in the organization or profession invest in defensive posturing, so they themselves are better protected against prosecutorial attention. Rather than increasing the flow of safety-related information, legal action has a way of cutting off that flow. Safety reporting often gets a harsh blow when things go to court.

In 2006, Julie, 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. Julie accidentally administered a bag of epidural analgesia. Julie lost her job, faced action on her nursing license and the threat of six years in jail as well as a $25,000 fine. Julie's 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 five years in jail. One pleaded guilty to a reduced charge; another fought the charge and was eventually exonerated.

In other, similar cases where healthcare workers and other professionals were to stand trial on criminal charges, incident-reporting rates dropped. Sure, somebody may have been held "accountable." But the system did not get any wiser for it. Only dumber—literally. In the long run, it seems as if nobody benefits from this type of response to failure. Also, the things that get changed in response to legal action are not necessarily the things that make the operation or organization any safer.

In 2002, a 1-year-old girl had been taken from her home and placed in crisis care, as social workers had seen evidence of child abuse and malnutrition. After a few months of recuperation, the child was sent back home. The country's child protection council was not notified of the decision, and there was no record of whether the girl's family now fulfilled the conditions for return. A social worker visited three times, found little to report, and eventually went on sick leave. It was months before she was replaced. The

new social worker set out to find out more about the family. She drew up a plan for the mother, particularly to provide a steady daily routine for the little girl. The mother never managed. This, however, had no consequences, not even when authorities were notified that the girl looked blue in the face and was falling behind in her language development. She was not taken away from the mother again. A few months later, the girl was found dead—her mother had stuffed a rag in her mouth. She was three years old, and weighed less than twenty pounds.[1]

A prosecutor decided to get involved. Not by going after the mother, but the social worker—for second-degree manslaughter. She had been negligent, had ignored obvious signals of abuse, and yet not intervened. As usual, the prosecution of the one "bad apple" served to mask a whole host of systemic issues: cuts in funding for social services, increases in caseload per worker, time pressures, and similar cases in which the service decided not to intervene either. It glossed over the most difficult professional judgment, where only hindsight can tell right from wrong. Intervening in social work is never "right"—it is always too early or too late. Either children get placed in foster care early enough to avoid real problems (but then there may not be enough evidence for their placement and decisions can be appealed and overturned, hurting the child), or the decision to intervene comes when so much evidence has been gathered (and the case is so obvious) that the child has already been hurt. Prosecuting a social worker for second-degree manslaughter does nothing to alleviate this professional double bind. In the glare of hindsight, it in fact blinds wider society for its deep complexity. Instead, criminalization may cause social services to lower the threshold on what is considered enough evidence for placement in foster care, bloating caseloads and stretching demands on foster families (which in turn may lower the bar on their quality). It may encourage social services to become more bureaucratic, clogging the provision of critical social services with unnecessary red tape, bookkeeping and officialdom, ensuring a gradual crumbling of the quality of care.

Judicial proceedings can rudely interfere with an organization's priorities and policies. They can redirect resources into projects or protective measures that have little to do with the organization's original mandate, or with safety. What may be improved in the example above was all kinds of aspects of bureaucracy. Not again would this organization be "caught" by a prosecutor without an elaborate, auditable and defensible paperwork footprint of its actions and decisions. The other thing that likely happened was that the organization adjusted its criterion for intervention downward: it would now be satisfied with less evidence to step in, rather than get caught again by a prosecutor after (in hindsight) stepping in too late. This represents a dilemma, at many levels and in many ways, for various organizations—not just social services.

The things that get changed when a failure is met with an "unjust" response (the prosecution of an individual caregiver in the example above) are not typically the things that make the organization safer. It does not lead to improvement in primary processes. It can lead to "improvement" of all the stuff that swirls around those primary processes: bureaucracy, involvement of the organization's legal department, bookkeeping, micro-management. Paradoxically, many such measures can make the work of those at the sharp end, those whose main concern is the primary process, more difficult, lower in quality, more cumbersome, and perhaps even less safe.

A Just Culture: Balancing Safety *and* Accountability

Calls for accountability are important. And responding adequately to them is too. Calls for accountability themselves are, in essence, about trust. About people, regulators, the public, employees, trusting that you will take problems inside your organization seriously. That you will do something about them, and hold the people responsible for those problems to account. Accountability is fundamental to human relationships. If we cannot be asked to explain why we did what we did, then we somehow break the pact that *all* people are locked into. Being able to offer an account for our actions is the basis for a decent, open, functioning society.

But if we think calls for accountability are the same as holding people (criminally) responsible, we're way off.

I recall how one safety-critical industry was under intense media scrutiny in a country where I once lived. The newly-elected government had pledged to the public that it would let the industry continue to function if it were safe. Then reports started to leak out about operators drinking on the job, about an internal erosion in safety culture, about a lack of trust between management and employees. The regulator was under exceptional pressure to do something. To show that it, and the government, could be trusted.

So the regulator sent parts of the cases it had discovered to the prosecutor. The media loved it: now something was happening! Maybe the people entrusted with this safety-critical technology had committed crimes. Now somebody was finally going to be held accountable.

The regulator saw how some of the media spotlight on it got dimmed. It could breathe a little easier now. But it was a bittersweet lull. The relationship with the industry was dramatically disturbed. Regulators have to rely on open disclosure by people in the industry

they regulate, otherwise they have no accurate or truthful information to go and regulate on. Such disclosure was now going to be very unlikely. It would be, for years to come.

In addition, safety improvements, at least for the media (and thereby public opinion, and, by extension, the government's stance on the issue) could now be largely collapsed into the pursuit of a few bad apples in the industry's management. Now that these people would be held accountable, any other safety improvements could simply be assumed to be less important, or to follow automatically. Of course they would not. Publicly or legally reminding people of their responsibilities may have some effect in getting them or others to behave differently (though never for a long time). And the negative consequences of such accountability easily outweigh these effects.

Responding to calls for accountability by saying that crimes may have been committed is not likely to lead you to justice or improve safety. People will feel unfairly singled out, and disclosure of safety problems will suffer. A just culture, then, also pays attention to safety, so that people feel comfortable to:

- bring out information about what should be improved to levels or groups that can do something about it;
- allow the organization to invest resources in improvements that have a safety dividend, rather than deflecting resources into legal protection and limiting liability.

A just culture, then, means getting to an account of failure that can do two things at the same time:

- satisfy demands for accountability;
- contribute to learning and improvement.

Virginia Sharpe, a philosopher and clinical ethicist who has studied the problem of medical harm for many years, has captured these dual demands in what she calls "forward-looking accountability."[2] Accountability that is backward-looking (often the kind in trials or lawsuits) tries to find a scapegoat, to blame and shame an individual for messing up. But accountability is about looking ahead. Not only should accountability acknowledge the mistake and the harm resulting from it. It should lay out the opportunities (and responsibilities!) for making changes so that the probability of such harm happening again goes down. I will go into this more deeply in the final chapter.

For now, it may seem impossible to convert your profession or organization to forward-looking accountability. It may seem impossible to both hold people accountable and learn and improve at the same time. Which is why real just

cultures seem so elusive. If that is how you feel, you are not alone. The two seem impossible to reconcile:

- Set up ways that people can tell stories that contribute to learning and improvement (for example, confidential incident reporting) and some people will cry foul: your operators or managers should own up! They should take responsibility! I demand to know who messes up!
- But tell stories that satisfy such demands for accountability and you may find that there is very little learning or improvement leverage in them. In fact, you may find that the very act of forcing out such stories (for example, through a trial) makes learning very difficult.

Creating, and getting consensus around, an explanation of failure that both satisfies demands for accountability *and* contributes to learning and improvement is a wonderful challenge. It is the challenge at the heart of a just culture, and the chapters that follow will each somehow try to deal with it, leading you some way towards building a just culture yourself.

A Just Culture Has More Advantages

The main argument for building a just culture is that not having one is bad for both justice and safety. But there is more. Recent research[3] has shown that not having a just culture can be bad for people's:

- morale;
- commitment to the organization;
- job satisfaction;
- willingness to do that little extra, to step outside their role.

Indeed, the idea of justice seems basic to any social relation; basic to what it means to be human, and humans among each other. We tend to endow a just culture with benefits that extend beyond making an organization safer. Look at the hope expressed by a policy document from aviation, where a "just culture operates by design to encourage compliance with the appropriate regulations and procedures, foster safe operating practices, and promote the development of internal evaluation programs."[4] It illustrates the great expectations that people endow just cultures with: openness, compliance, fostering safer practices, critical self-evaluation. How all of this is supposed to happen is, of course, a more difficult question.

Let us focus on safety. That, after all, is the point of this book: how to reconcile accountability *for* failure with learning *from* that failure—with the aim to keep making progress on safety. Now it may seem obvious why employees may want a just culture. They may want to feel protected from capricious management actions, or from the (as they see it) malicious intentions of a prosecutor. But this oversimplifies and politicizes things. A just culture, in the long run, benefits everyone:

- **For those who run or regulate organizations, the incentive to have a just culture is very simple. Without it, you won't know what's going on.** A just culture is necessary if you want to monitor the safety of an operation. A just culture is necessary if you want to have any idea about the capability of your people, or regulated organization, to effectively meet the problems that will come their way.
- **For those who work inside an organization, the incentive of having a just culture is not "to get off the hook,"** but to feel free to concentrate on doing a quality job rather than on limiting personal liability; to feel involved and empowered to contribute to safety improvements by flagging for weak spots, errors and failures.
- **For those in society who consume the organization's product or service, just cultures are in their own long-term interest.** Without them, organizations and the people in them will focus on better documenting, hiding or defending decisions—rather than on making better decisions. They will prioritize short-term measures to limit legal or media exposure over long-term investments in safety.

Wanting Everything in the Open, but not Tolerating Everything

What is it that can make a just organization a safe organization, and an unjust one an unsafe one? People who write or think about just culture agree: it has to do with being open, with a willingness to share information about safety problems without the fear of being nailed for them. Most people also believe that the openness of a just culture is not the same as uncritical tolerance or generosity. If everything "goes," then in the end no problem may be seen anymore as safety-critical—and people will stop talking about them for that reason. It is precisely this tension between:

- wanting everything in the open;
- while not tolerating everything.

I will deal with both in this book. It will cover how the obligations to disclose are about wanting everything relevant in the open—and how a perceived lack of justice can mess that up really quickly. It will cover the problems with not tolerating everything—because the "everything" in there is not about a clear line or definition, but about who gets to decide. It will cover how a just culture is about the always uneasy, but exciting melding of the two. It is exactly the friction between wanting everything in the open so that you can learn, but not tolerating everything so that you can be "just," that makes building a just culture such an interesting venture.

Notes

1 Hilhorst, P. (2006, December 12). *Bij de dood van een peuter* (*On the death of a toddler*). De Volkskrant, 13.

2 Sharpe, V.A. (2003). Promoting patient safety: An ethical basis for policy deliberation. *Hastings Center Report Special Supplement*, 33(5), S1–S20.

3 See, for example, Cohen-Charash, Y. and Spector, P.E. (2001). The role of justice in organizations: A meta-analysis. *Organizational Behavior and Human Decision Processes*, 86, 278–321, or Colquitt, J.A., Conlon, D.E., Wesson, M.J., Porter, C.O.L.H., and Ng, K.Y. (2001). Justice at the millennium: A meta-analytic review of 25 years of organizational justice research. *Journal of Applied Psychology*, 86, 425–45.

4 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), 3.

This page intentionally left blank

2 Between Culpable and Blameless

When does a mistake stop being honest? Here I recount one such "mistake," that got turned into a crime. It happened on the morning of November 21, 1989, when a Boeing 747 on an instrument approach in heavy fog came very near to crashing at London Heathrow airport. The big airliner had been misaligned with the runway so that, when it started pulling up for a go-around, it was actually outside the airport's perimeter fence, and only about 75 feet off the ground. It narrowly missed a nearby hotel, setting off car alarms all over the parking lot, and fire sprinklers in the hotel. The second approach and landing were uneventful. And most passengers had no idea they had been so close to a possibly harmful outcome. For them, the flight was now over. But for the captain, a veteran with 15,000 flight hours, it was only the beginning of a greater drama. Two and a half years later, a divided jury (10 to 2) would find him guilty of negligently endangering his aircraft and passengers—a criminal offense. He would lose his job, and then some.

I re-analyze an account of the incident as masterfully told by Stephan Wilkinson[1] to illustrate the tension between different interpretations of the same event. Was it a mistake culpable enough to warrant prosecution? Or was it normal, to be expected, all in a day's work?

We can never achieve "objective" closure on these questions. You can only make up your own mind about them. Yet the case raises fundamental issues about a just culture: how can we produce explanations of failure that both satisfy demands for accountability *and* provide maximum opportunities for organizational learning?

A Bug, Weather and a Missed Approach

Wilkinson describes how the captain's problems began at a Chinese restaurant in Mauritius, an Island in the Indian Ocean off Africa. Together with his flight deck crew, a co-pilot and flight engineer, he dined there during a layover before

flying on to Bahrain and then to London. The leg from Bahrain to London would be the last portion of a trip that had begun in Brisbane, Australia.

Several days later, when the flight had gotten to Bahrain, both the co-pilot and flight engineer were racked with gastroenteritis (stomach flu). The captain, however, was unaffected. A Mauritian doctor had given the flight engineer's wife tranquilizers and painkillers. She also was on the trip and had dined with the crew. The doctor had advised the flight engineer to take some of his wife's pills if his symptoms got worse too. Flight crews, of course, can't just take advice or prescriptions from any doctor, but this man had been suggested by an airline-approved physician, who was too far away but had recommended the examining doctor to the crew. He would soon be added to the airline's list anyway. He did not, however, seem concerned that the crew had been scheduled to fly in a few days' time again. A colleague pilot commented afterward:

> This was apparently a doctor who didn't even understand the effects of self-medication in a pressurized aircraft on the performance of a complex task, and right there is a microcosm of everything that pressured the crew to get the job done. That doctor's vested interest is in sending flight crews out to fly. Certainly if he ever expects to work for the airline again, he isn't going to ground crews right and left. The company *wants* you to fly.[2]

(As part of the court case years later, however, the captain would be accused of violating the company's medical procedures.)

The subsequent flight to London was grim. Unexpected headwinds cut into the 747's fuel reserves, and the co-pilot had to leave the cockpit for several hours after taking some of the flight engineer's wife's medicines to control his symptoms. It left the captain to fly a stretch of five hours alone, much of it in the dark.

London Fog

Over Frankfurt, the crew heard that the weather at London Heathrow airport was bad. Thick fog meant that they probably would have to execute a so-called Category III instrument approach. In Category III conditions, a 747 is literally landing blind. While the wheels may just have poked out of the fog in the flare, the cockpit, considerably higher, is still in the clouds. Category III approaches are flown by the autopilot, with the crew monitoring the instruments and autopilot performance. The autopilot captures two radio beams (a localizer for lateral and a glideslope for vertical guidance). These are transmitted by the

instrument landing system on the ground, and the autopilot translates them into control commands to make the aircraft stay on track and on a gradual descent, exactly toward and onto the runway. At least, that is the idea.

The captain, like most airline pilots, had never flown a Category III approach down to minimums, despite his extensive instrument experience. The co-pilot, new with the airline, had neither. He had not even had the mandatory training for a Category III approach, and was not approved to fly one. But that was not going to stop anything. Still over Germany, the captain had gotten in touch with the airline and requested permission for the co-pilot to help out on this one approach into London to get them home. Dispensation was granted. It routinely is—it almost always is. The captain, however, never volunteered the information that his co-pilot was not in the best of states (in fact, he may not have been in the cockpit at that very moment). Nobody on the ground inquired either.

Later, the co-pilot testified that nobody had asked him if he wanted a dispensation. But even if he'd been asked, it would have been difficult to refuse.

> I accepted, with the airline's interest's at heart, the dispensation to operate to Category III autoland conditions," he later wrote to the court. "I personally would not mind if we had diverted. But what would the airline have said to the captain if he had diverted without asking for a dispensation? What would they have said to me if I had not accepted it?[3]

He really had been in a bind. Wanting to help the airline, wanting to get its passengers home, the co-pilot had agreed to go on with the flight. But he was sick, really. So if the flight would have had to divert because he was too poorly to do a Category III approach, this once, what was he doing on board anyway? And where had those medicines come from?

"This," Wilkinson, observed, "is the heart of the professional pilot's conflict. In one ear the airlines lecture, 'Never break regulations. Never take a chance. Never ignore written procedures. Never compromise safety.' Yet in the other they whisper, 'Don't cost us time. Don't waste our money. Get your passengers to their destination—don't find reasons why you can't.'"[4]

The Approach

Nearing London, the 747 was given a routine holding northeast of the airport. After some time of flying racetracks in the holding pattern, the flight engineer suggested, "Come on, we've got two minutes of holding fuel left, let's buzz off to

Manchester." The crew discussed the options—both Manchester and Gatwick (south of London) were diversion airports, though Manchester had better weather. But the captain "was a very determined man," as the flight engineer recalled. Just as he was deciding to head off to Manchester, Heathrow called and cleared the 747 for approach.

But a complication had arisen: instead of landing to the east (runway 09), as had been planned, they would now have to turn in shorter and land toward the west (runway 27), because the wind had changed. The approach became a hurried affair. The crew had to reshuffle charts, talk and think through the procedures, revise their mental pictures. A 10-knot tailwind at altitude meant that the 747 was motoring down the approach path toward the runway at an even greater groundspeed. Tightening their slack further still, the approach controller turned the 747 onto the localizer 10 miles from the runway, rather than the normal 12 miles or more. Halfway down, the tower radioed that some approach lights were apparently not working, requiring the flight engineer to take a quick look through his checklist to see how this affected their planned procedure, if at all. The tower controller also withheld clearance for the 747 to land until the last moment, as a preceding 747 was feeling its way through the fog, trying to find its turn-off from the runway.

But the autopilots really were about to become the final straw: they never seemed to settle onto the localizer, instead trundling back and forth through the beam, left to right. The two autopilots on this old "Classic" 747 may never have been able to capture the localizer: when the aircraft turned in to start its approach, the autopilots disconnected for some time and the airplane was flown manually. The autopilots, built by Sperry, were based on an earlier design. They were never really meant for this aircraft, but sort of "bolted on," and had to be nursed carefully.[5] On this flight the crew made a later attempt to re-engage the autopilots, though radar pictures showed that the 747 never settled on a stable approach path.

The flight engineer was getting worried about the captain, who had basically been flying solo through the night, and still was alone at the controls. The co-pilot was of little help. "I was not qualified to make this approach and could not make any suggestions as to what was wrong," he would later tell safety investigators. He stayed out of the way.

The captain was now technically illegal: trying to fly a Category III approach with autopilots that refused to settle down and function perfectly was not allowed. The right decision, to everybody, in hindsight, would have been to go around, to fly what's called a missed approach. And then to try again or go to

the alternative. "I'd have thrown away the approach, gone to my alternate or tried again. No question about it," one pilot questioned by Wilkinson said.

But other pilots, some with the same airline, believed the opposite.

> Look, he was concerned about fuel. He had a first officer who was no help. He knew a diversion to Manchester would cost the airline a minimum of 30,000 dollars. He realized he'd be sitting in the chief pilot's office trying to explain how he got himself into a position that required a missed approach in the first place. He figured the autopilots would settle down. And I'll bet he was convinced he'd break out at Category I limits (a higher cloud ceiling and better visibility than Category III) and could take over and hand-fly it the rest of the way. I can understand why he carried on.[6]

It might have worked, Wilkinson observed. And if it had, nobody would ever have heard of this case.

But it did not work. Ever concerned with passenger comfort, the captain waited with making a go-around. And then he made a gentle one. The 747 sank another 50 feet. The flight engineer glimpsed approach lights out the left window as they started pulling up, away.

As one 747 instructor said,

> This is a pilot who was critically low on fuel, which probably was one reason why he waited a second before going around. At decision height on a Category II approach, you look to see the slightest glow of approach lights, you wait "one-potato," see if anything comes into sight. Perhaps a thousand times before, he'd watched that same autopilot do strange things on the same approach to the same airport, and he'd break out at 200 or 500 feet and make a play for the runway. And on the crew bus everybody says, "Boy, that autopilot sucked again today."[7]

On climb-out after the first try, the co-pilot noticed how the captain's hands were shaking. He suggested that he fly the second approach instead, but the captain waved him away. The second approach was uneventful, and was followed by a landing that elicited applause in the passenger cabin.

No Disclosure, but a Trial

Back in the crew room after they had shut down the airplane, the captain found a note in his company letterbox. It requested that the crew see the chief pilot. The captain told the co-pilot and flight engineer to go home, and he would say that they had already left when he found the note.

But he did not go to the chief pilot either. Nor did he talk to an airline safety investigator about what had happened. Instead, he drove straight home and went to bed. That evening, a call came from the airline. The crew had been suspended.

An internal investigation was launched by the airline, who later issued a report chiding the co-pilot and flight engineer. The airline also demoted the captain to first officer. The aviation authority downgraded his license accordingly, and he was relegated to riding out the rest of his career in the right seat, no longer in command.

This was too much. Half a year after the incident, the captain resigned from the airline and began to appeal the authority's reduction of his license. Some did not see any problem. Recently, the pilot had been receiving grades of "average" on his half-yearly proficiency checks in the simulator, and instructors had taken note of his inability to perform well under pressure.

But why did the regulator take him to court? This "remains the subject of speculation," Wilkinson writes.

> There is considerable feeling that the airline was not sorry to see it happen, that the captain was a loose cannon who could have made things awkward for an airline that places great value on its public image. Some feel that the captain could have revealed some controversial company procedures. If the captain were branded a criminal, it would effectively negate whatever damage he might do … Others suspected empire building within the regulator's legal branch: this looked like a juicy case for an aspiring prosecutor to take public and demonstrate that even the flag carrier's jumbo jet captains dare not take on the aviation authority casually.[8]

Six weeks after the incident, the airline had announced that it was no longer granting bad-weather dispensations. But the fleet manager who had authorized the approach with the co-pilot's dispensation was not in the dock. Nor was the controller who turned the big 747 onto a tight approach, separated by what seemed like only five miles rather than the legal minim of six from the preceding 747. With traffic from all over the world converging onto London at 8:00 a.m., those rules were obviously allowed to be flexible.

It was the pilot who was in the dock. Seated next to a policeman. Why had he not filed a Mandatory Occurrence Report right after the flight? Because it did not constitute an occurrence, the pilot argued. After all, he had gone around, or at least initiated a go-around, and landed uneventfully the second time. Why had he gone around so slowly? Because the supposedly canonical technique was not described anywhere, he argued. At some point in the trial, the

pilot produced a transcript of every oral call-out, checklist response, and radio transmission that company and government regulations required the crew to accomplish during the approach. It showed that the entire routine took seven minutes. The approach had lasted only four, making it technically impossible to make an approach and follow all applicable rules at the same time.

Few cared. Jurors sometimes even napped. If the trial did not revolve around arcane legal points, it did so around finely-grained technical ones. The pilot was never called to testify on his own behalf.

The defense elaborated the fact that the old 747 was dispatched on its next leg out of London without a check of the autopilot, to see if it was somehow faulty. To this day, four crucial pages of the maintenance log, which might have told something about the autopilot, are missing (in a parallel to the prescription missing from the medication log in Mara's case—see Prologue).

"The regulator itself was at fault," a legal expert and airline pilot commented, "for permitting a situation to exist in which the airline's flight operations manual contained a provision that the captain would be expected to use, by which it could authorize him to make the approach without a qualified co-pilot. The approach was actually illegal at the fault of the airline, yet they were not charged. Had that provision not existed, the captain would have diverted to Frankfurt with cozy fuel reserves, to await better weather at London."[9]

A split jury found the pilot guilty. The judge fined him only £1,500, and rejected the regulator's demand that he pay £45,000 more to cover court costs. The pilot appealed the decision, but that was summarily rejected.

When he was young, the pilot lived near an airforce base where he would watch airplanes take off and land at the end of the war. That inspired him to become a pilot. "On December 1, 1992, three years and nine days after the incident, the pilot left home without a word to his wife. He drove some nine hours to a beach near the air force base. There he ran a hose from his car's exhaust pipe through a nearly closed window. In a matter of minutes he was dead. He left no letter or any explanation."[10]

What are the Implications for Just Culture?

It would be too easy to ask whether the prosecution and conviction of the captain was right. Or just. Because it is too difficult to answer. Was this a crime?

Multiple descriptions of the events are plausible. The disappearance of documents without a trace in these cases can always give people the chills.

Was it a conspiracy after all, a "cover-up", as some of Wilkinson's interviewees suggested? It could have been: turning one pilot into a highly visible scapegoat in order to silence him and others. This would save the reputation of both the airline and the regulator, who also happens to employ the aviation prosecutor in this country. But conspiracies take tight coordination and demand iron discipline from those involved in it.

Also, as a captain, this pilot had lately been "average," not stellar. He was stubborn and determined. He was ultimately responsible for getting himself and his crew into this jam. And then he apparently refused to cooperate, did not want to disclose or discuss the incident (it wasn't an occurrence to him, after all) until forced to do so in the adversarial setting of a trial.

Who is right? Whose version of event is true? The tension between multiple possible interpretations remains until the end of Wilkinson's story. But important traces about building a just culture do stand out:

- **A single account cannot do justice to the complexity of events**. We need multiple layers of description, partially overlapping and probably always somehow contradictory, to have any hope of approximating reality.
- **A just culture accepts nobody's account as "true" or "right"** and others wrong. This only leads to moral grandstanding, imperialism, and to losers like this pilot or Mara. Instead, it accepts the value of multiple perspectives, and uses them to encourage both accountability and learning.
- **A just culture is not about absolutes, but about compromise**. Achieving justice is not about black and white. Instead, it presumes compromise. Justice in a just culture cannot be enforced, it must be bargained. Such bargaining for justice is a process of discovery, a discovery that the best bargain may an outcome in which every party benefits,[11] for example an explanation of events that satisfies calls for accountability *and* helps an organization learn and improve.
- **A just culture pays attention to the "view from below"** among these multiple accounts, as that view (in this case from the person in the dock) may have little or no power to assert itself and is the easiest to quash. Silencing it can be organizationally or politically convenient. You may even see it as imperative. You may see putting others in an inferior position as a necessary, if sometimes annoying step in achieving other goals. But this makes it even more morally essential to give the view from below a voice.
- **A just culture is not about achieving power goals**, by using other people to deflect attention away from one's own flaws. This denies such

people their personhood, it makes them a mere instrument in the pursuit of protection of power, of existing structures or arrangements. Most people will see this as unethical,[12] and it violates the basic principles of Aristotelian justice that many of our societies still live by.[13]

- **Disclosure matters.** Not wanting to disclose can make a normal mistake look dishonest, with the result that it may be treated as such. Multiple examples in this book illustrate this. Disclosing is the practitioner's responsibility, or even duty.
- **Protecting those who disclose matters just as much.** The demand to disclose in the pilot's case above (a note in the letterbox) may not have given him confidence that honest disclosure would be treated fairly. Conditions at his airline may have been unfavorable for honest disclosure. Creating a climate in which disclosure is possible and acceptable is the organization's responsibility. And more protections are often necessary.
- **Proportionality and decency are crucial to a just culture.** People will see responses to a mistake as unfair and indecent when they are clearly disproportionate. "What was the guy found guilty of?" a pilot friend had asked Wilkinson in amazement. "Endangering his passengers," Wilkinson replied. "I do that every day I fly," the friend said with a laugh. "That's aviation."[14] The eventual punishment given to this pilot (a symbolic fine) may have indicated that the trial was seen as a disproportionate response to an event that perhaps should not have ended up in court. Proportionality means heeding Martin Buber's dictum: what is necessary is allowed, but what is not necessary is forbidden.

By the time a case reaches trial, much of the above has either been wasted or rendered impossible. A trial cannot do justice to the complexity of events, as it necessarily has to pick one account as the truest or most trustworthy one.

A seeming lack of honest disclosure is often a trigger for a trial. This could have been the case here. You can also see it in the literature on medical lawsuits. Patients or their families do not typically take a doctor to court *until* they feel that there is no longer any other way to get an account of what went wrong.[15] Stonewalling often leads to a trial. But a climate that engenders anxiety and uncertainty about how disclosure will be treated often leads to stonewalling. The more we take cases to trial, the more we could be creating a climate in which freely telling each other accounts is becoming more and more difficult.

Notes

1 Wilkinson, S. (1994). The November Oscar incident: Airline pilots are haunted by a missed approach that left only one casualty—the captain. *Smithsonian Air and Space*, February–March, 80–87.
2 Wilkinson, op. cit., 82.
3 Ibid., 84.
4 Ibid.
5 David Beaty, quoted ibid.
6 Ibid., 85.
7 Ibid.
8 Ibid., 86.
9 Ibid., 87.
10 Ibid.
11 Visotzky, B. (1996). *The genesis of ethics*. New York: Crown Publishers, Inc.
12 Ibid., 71.
13 Rawls, J. (2003). *A theory of justice*. Cambridge, MA: Harvard University Press.
14 Wilkinson, op. cit., 84.
15 Berlinger, N. (2005). *After harm: Medical error and the ethics of forgiveness*. Baltimore, MD: The Johns Hopkins University Press.

3 The Importance, Risk and Protection of Reporting

Many professions have codified the obligation to report. In air traffic control, for example, "all safety occurrences need to be reported and assessed, all relevant data collected and lessons disseminated."[1] There is an implicit understanding that reporting is critical for learning. And learning is critical for constantly improving safety (or, if anything, for staying just ahead of the constantly changing nature of risk).

Saying that all safety occurrences need to be reported is easy. But what counts as a "safety occurrence?" Recall from Chapter 2 that this can be open for interpretation: the missed approach of the 747 was, according to the pilot, not a safety occurrence. It was not worth reporting. But according to his bosses and regulators, it was. And the fact that he did not report it, made it all the more so.

Professional codes about reporting, then, should ideally be more specific than saying that "all safety occurrences" should be reported. What counts as a clear opportunity for organizational learning for one, perhaps constitutes a dull and unreportworthy event to somebody else. Something that could have gone terribly wrong, but did not, is not necessarily a clear indication of reportworthiness either. After all, in many professions things can go terribly wrong the whole time ("I endanger my passengers every day I fly!"). But that does not make reporting everything particularly meaningful.

Reporting is Important. But What to Report?

The point of reporting is to contribute to organizational learning. It is to help prevent recurrence by making systemic changes that aim to redress some of the basic circumstances in which work went awry. This means that any event that has the potential to shed some light on (and help improve the conditions for) safe practice is, in principle, worth reporting and investigating. But that still does not create very meaningful guidance.

Which event is worthy of reporting and investigating is, at its heart, a judgment. First, it is a judgment by those who perform safety-critical work at the sharp end. Their judgment about whether to report something is shaped foremost by experience—the ability to deploy years of practice into gauging the reasons and seriousness behind a mistake or adverse event.

To be sure, those years of experience can also have a way of blunting the judgment of what to report. If all has been seen before, why still report? What individuals and groups define as "normal" can glide, incorporating more and more non-conformity as time goes by and as experience mounts. In addition, the rhetoric used to talk about mistake can serve to "normalize" (or at least deflect) an event away from the professionals at that moment. A "complication" or "non-compliant patient" is not so compelling to report (though perhaps worth sharing with peers in some other forum), as when the same event were to be denoted as, for example, a diagnostic error.

Whether an event is worth reporting, in other words, can depend on what language is used to describe that event in the first instance. This has another interesting implication: in some cases a lack of experience (either because of a lack of seniority, or because of inexperience with a particular case, or in that particular department) can be immensely refreshing in questioning what is "normal" (and thus what should be reported or not).

Investing in a meeting where different stakeholders share their examples of what is worth reporting could be useful. It could result in a list of examples that can be handed to people as partial guidance on what to report.

But in the end, given the uncertainties about how things can be seen as valuable by other people, and how they could have developed, the ethical obligation should be "if in doubt, report."

But then, what delimits an "event"? The reporter needs to decide where the reported event begins and ends. She or he needs to decide how to describe the roles and actions of other participants who contributed to the event (and to what extent to identify other participants, if at all). Finally, the reporter needs to settle on a level of descriptive resolution that offers the organization a chance to understand the event and find leverage for change. Many of these things can be structured beforehand, for example by offering a reporting form that gives guidance and asks particular questions ("need-to-know" for the organization to make any sense of the event) as well as ample space for free-text description.

When Reporting becomes Dangerous

If people report their honest mistakes in a just culture, they will not be blamed for them. The reason is that an organization can benefit much more by learning from the mistakes that were made than from blaming the people who made them. So people should feel free to report their honest mistakes.

The problem is, often they don't.

Often they don't feel free, and they don't report.

This is because reporting can be risky. Many things can be unclear:

- How exactly will the supervisor, the manager, the organization respond?
- What are the rights and obligations of the reporter?
- Will the reported information stay inside of the organization? Or will other parties (media, prosecutor) have access to it as well?

The reason why most people fail to report is not because they want to be dishonest. Nor because they are dishonest. The reason is that they fear the consequences, or have no faith that anything meaningful will be done with what they tell. And often with good reason:

- either people simply don't know the consequences of reporting, so they fear the unknown, the uncertainty;
- or the consequences of reporting really can be bad, and people fear invoking such consequences when they report information themselves.
- or people know the consequences, but feel that there is no point in reporting because the organization won't do anything with the report anyway.

While the first reason may be more common, either reason means that there is serious work to do for your organization. In the first case, that work entails clarification. Make clear what the procedures and rules for reporting are, what people's rights and obligations are, and what they can expect in terms of protection when they report. In the second case it means trying to make different structural arrangements, for example with regulators or prosecutors, with supervisors or managers, about how to treat those who report. This is much more difficult, as it involves the meshing of a lot of different interests. I will talk more about the interests of the different stakeholders and possible relationships between them later in the book. In the third case, your organization clearly has work to do too. I will talk about that more below.

What if Reported Information Falls into the Wrong Hands?

The nurse in the Prologue honestly reported her contribution to the death of the infant to her supervisor. As a result, she was convicted as a criminal and is today without a job, or much of a life. Information about the incident was leaked to the media, and thereby into the hands of a prosecutor who happened to read about it in the local newspaper.

In many countries, it does not even have to go so haphazardly. Most democracies have strong freedom-of-information legislation. This allows all citizens from the outside access, in principle, to all non-confidential information. Such transparency is critical to democracy, and in some countries freedom-of-information is even enshrined in the constitution. But the citizen requesting information can easily be an investigating journalist, a police officer or a prosecutor. Freedom-of-information is really an issue when the organization itself is government-owned (and hospitals or air traffic control centers in many countries still are). Moreover, safety investigating bodies are also government organizations, and thus subject to such legislation. This can make people unwilling to collaborate with safety investigators.

The potential for such exposure can create enormous uncertainty. And uncertainty typically dampens people's willingness to report. People become anxious about leaving information in files with their organization. In fact, the organization itself can become anxious about even having such files. Having them creates the risk that names of professionals end up in the public domain. This, in turn, can subject safety information to oversimplification and distortion and misuse by those who do not understand the subtleties and nuances of the profession.

Some countries have succeeded in exempting safety data in very narrow cases from freedom-of-information legislation. The Air Law in Norway, for example, states, about the "Prohibition on use as evidence in criminal proceedings", that "Information received by the investigating authority may not be used as evidence in any subsequent criminal proceedings brought against the persons who provided the evidence."[2] Of course, this does not keep a prosecutor or judge from actually reading a final accident report (as that is accessible to all citizens), but it does prevent statements provided in good faith from being used as evidence. Similar legislation exists, though in other forms, in various countries. Many states in the US, for example, protect safety data collected through incident reporting against access by potential claimants. Most require a subpoena or court order for release of the information.[3]

One problem with this, of course, is that it locks information up even for those who can rightfully claim access, and who have no vindictive intentions. Imagine a patient, for example, or a victim of a transportation accident (or the family), whose main aim is to find out something specific about what happened to their relative. The protection of reporting, in other words, can make such disclosure (see the next chapter) more difficult. So when you contemplate formally protecting reported safety information, you should carefully consider these potential consequences.

Getting People to Report

Getting people to report is difficult. Keeping up the reporting rate once the system is running can be equally difficult, though often for different reasons. Getting people to report is about two major things:

- maximizing accessibility;
- minimizing anxiety.

The means for reporting must be accessible. If you have reporting forms, they need to be easily and ubiquitously available, and should not be cumbersome to fill in or send up.

Anxiety can initially be significant:

- What will happen to the report?
- Who else will see it?
- Do I jeopardize myself, my career, my colleagues?
- Does this make legal action against me easier?

As an organization you should ask yourself whether there is a written policy that explains to everybody in the organization what the reporting process looks like, what the consequences of reporting could be, what rights, privileges, protections and obligations people may expect. Without a written policy, ambiguity can persist. And ambiguity means that people will disclose less.

Getting people to report is about building trust: trust that the information provided in good faith will not be used against those who reported it. Such trust must be built in various ways. An important way is by structural (legal) arrangement. Making sure people have knowledge about the organizational and legal arrangements surrounding reporting is very important: disinclination to report is often related more to uncertainty about what *can* happen with a report,

than by any real fear about what *will* happen. One organization, for example, has handed out little credit-sized cards to its employees to inform them about their rights and duties around an incident.

Another way to build trust is by historical precedent: making sure there is a good record for people to lean on when considering whether to report an event or not. But trust is hard to build and easy to break: one organizational or legal response to a reported event that shows that divulged information can somehow be used against the reporter, can destroy months or years of building goodwill.

Keeping the Reports Coming In

Keeping up the reporting rate is also about trust. But it is even more about involvement, participation and empowerment. Building enough trust so that people do not feel put off from sending in reports in the first place is one thing. Building a relationship with participation and involvement that will actually get and people to send in reports and keep sending them in is quite another.

Many people come to work with a genuine concern for the safety and quality of their professional practice. If, through reporting, they have an opportunity to actually contribute to visible improvements, then few other motivations or exhortations to report are necessary. Making a reporter part of the change process can be a good way forward, but this implies that the reporter wants (or dares) to be identified as such, and that managers have no problems with integrating employees in their work for improved safety and quality.

Sending feedback into the department about any changes that result from reporting is also a good strategy. But it should not become the stand-in for doing anything else with the reports. Many organizations get captured by the belief that reporting is a virtue in itself: if only people report mistakes, and their self-confessions are distributed back to the operational community, then things will automatically improve and people will feel motivated to keep reporting. This does not work for long. Active engagement with that which is reported, and with those who report, is necessary. Active, demonstrable intervention that acts on reported information is too.

Reporting to Managers or to Safety Staff?

In many organizations, the line manager is the recipient of reports. This makes (some) sense: the line manager probably has responsibility for safety and quality

in the primary processes, and should have the latest information on what is or is not going well. But this practice has some side-effects:

- it hardly renders reporters anonymous (given the typical size of a department), even if no name is attached to the report;
- reporting can have immediate line consequences (an unhappy manager, consequences for one's own chances to progress in career);
- especially in cases where the line manager herself or himself is part of the problem the reporter wishes to identify, such reporting arrangements all but stop the flow of useful information.

I remember studying one organization that had shifted from a reporting system run by line managers, to one run by safety-quality staff. Before the transition, employees actually turned out very ready to confess an "error" or "violation" to their line manager. It was almost seen as an act of honor. Reporting it to a line organization—which would see an admission of error as a satisfactory conclusion to its incident investigation—produced rapid closure for all involved. Management would not have to probe deeper, as the operator had seen the error of his or her ways and had been reprimanded and told or trained to watch out better next time.

For the operator, simply and quickly admitting an error avoided even more or deeper questions from their line managers. Moreover, it could help avert career consequences, in part by preventing information from being passed to other agencies (for example, the industry's regulator). Fear of retribution, in other words, did not necessarily discourage reporting. In fact, it encouraged a particular kind of reporting: a mea culpa *with minimal disclosure that would get it over with quickly for everybody. "Human error" as cause seemed to benefit everyone—except organizational learning.*

As one employee told us: "I didn't tell the truth about what took place, and this was encouraged by the line manager. He had made an assumption that the incident was due to one factor, which was not the case. This helped me construct and maintain a version of the story which was more favorable for us (the frontline employees)."

Perhaps the most important reason to consider a reporting system that is not just run by line management is that it can drastically improve organizational learning. Here is what one line manager commented after having been given a report by an operator:

"The incident has been discussed with the concerned operator, pointing out that priorities have to be set according to their urgency. The operator should not be distracted by a single problem and neglect the rest of his working environment. He has been reminded

of applicable rules and allowable exceptions to them. The investigation report has been made available to other operators by posting it on the internal safety board."

Such countermeasures really do not represent the best in organizational learning. In fact, they sound like easy feel-good fixes that are ultimately illusory. Or simply very short-lived.

Opening up a parallel system (or an alternative one) can really help. The reports in this system should go to a staff officer, not a line manager (for example, a safety or quality official), who has no stakes in running the department. The difference between what gets reported to a line manager, and that is written in confidential reports can be significant. Also, the difference in understanding, involvement and empowerment that the reporter feels can be significant:

"It is very good that a colleague, who understands the job, performs the interviews. They asked me really useful questions and pointed me in directions that I hadn't noticed. It was very positive compared to before. Earlier you never had the chance to understand what went wrong. You only got a conclusion to the incident. Now it is very good that the report is not published before we have had the chance give our feedback. You are very involved in the process now and you have time to go through the occurrence. Before you were placed in the hot chair and you felt guilty. Now, during interviews with the safety staff, I never had the feeling that I was accused of anything."

Of course, keeping a line-reporting mechanism in place *can* be very productive for a department's continuous improvement work. Especially if things need to be brought to the attention of relevant managers immediately. But you should perhaps consider a separate, parallel confidential reporting system if you don't already have one. Both line-based and staff-based (or formal and confidential) reporting mechanisms offer several kinds of leverage for change. Not mining both data sources for improvement information could be a waste for your organization.[4]

Notes

1 Eurocontrol Safety Regulatory Requirement (*ESARR 2*). Brussels: Eurocontrol.

2 Norwegian Air Law, Article 12–24.

3 Sharpe, V.A. (2003). Promoting patient safety: An ethical basis for policy deliberation. *Hastings Center Report Special Supplement*, 33(5), S1–S20.

4 For more about this study, see Dekker, S.W.A. and Laursen, T. (2007). From punitive action to confidential reporting: A longitudinal study of organizational learning. *Patient Safety and Quality Healthcare*, 4(5), 50–56.

4 The Importance, Risk and Protection of Disclosure

The Difference between Disclosure and Reporting

Disclosure is different from reporting.[1] This chapter will deal with the difference, and with disclosure and honesty in particular.

- **Reporting is the provision of information to supervisors, oversight bodies or other agencies**. Reporting means given a spoken or written account of something that you have observed, participated in or done to an appointed party (supervisor, safety manager). Reporting is thought necessary because it contributes to organizational learning. Reporting is not primarily about helping customers or patients, but about helping the organization (for example, colleagues) understand what went wrong and how to prevent recurrence.
- **Disclosure is the provision of information to customers, clients, patients and families**. The ethical obligation to disclose your role in adverse events comes from a unique, trust-based relationship with the ones who rely on you for a product or service. Disclosure is seen as a marker of professionalism. Disclosure means making information known, especially information that was secret, or information that could be kept secret. Information about incidents that only one or a few people were involved in, or that only professionals with inside knowledge can really understand, could qualify as such.
- **Practitioners typically have an obligation to report** to their organization when something went wrong. As part of the profession, they have a duty to flag problems and mistakes. After all, they represent the leading edge, the sharp end of the system: they are in daily contact with the risky technology or business. Their experiences are critical to learning and continuous improvement of the organization and its work.
- **Many practitioners also have an obligation to disclose** information about things that went wrong to their customers, clients, patients. This

Table 1 The difference between disclosure and reporting for individuals and organizations

	Reporting	Disclosure
Individual	Providing written or spoken account about observation or action to supervisors, managers, safety/quality staff	Making information known to customers, clients, patients
Organization	Providing information about employees' actions to regulatory or other (for example, judiciary) authorities when required	Providing information to customers, clients, patients, or others affected by organization's or employee's actions

obligation stems from the relationship of trust that professionals have with those who make use of their services.

- **Organizations also have an obligation to disclose** information about things that went wrong. This obligation stems from the (perhaps implicit) agreement that companies have with those who make use of their services or are otherwise affected by their actions.
- **Organizations (employers), the judiciary and regulators have an obligation to be honest about the possible consequences** of failure, so that professionals are not left in the dark about what can happen to them when they do report or disclose.
- We could also think that **organizations have a (legal) obligation to report** certain things to other authorities (judiciary, regulatory).

Disclosure and reporting can clash. And different kinds of reporting can also clash. This can create serious ethical dilemmas that both individual professionals and their employing organizations need to think about:

- If an organization wants to encourage reporting, it may actually have to curtail disclosure. Reporters will step forward with information about honest mistakes only when they feel they have adequate protection against that information being misused or used against them. This can mean that reported information must somehow remain confidential, which rules out disclosure (at least of that exact information).
- Conversely, disclosure by individuals may lead to legal or other adverse actions (even against the organization), which in turn can dampen people's or the organization's willingness to either report or disclose.

- If organizations report about individual actions to regulatory or judicial authorities, this too can bring down the willingness to report (and perhaps even disclose) by individuals, as they feel exposed to unjust or unwelcome responses to events they have been involved in.

A representative of the regulator had been sent out for a field visit as a customer of an organization I once worked with. She had observed things in the performance of one practitioner that, according to the rules and regulations, weren't right. Afterwards, she contacted the relevant managers in the organization and let them know what she had seen. The managers, in turn, sent a severe reprimand to the practitioner.

It really strained trust and the relationship between practitioners and management: reporting was not encouraged by their reaction. The regulator would not have been happy either to find out that their visit was being used as a trigger to admonish an individual practitioner instead of resolving more systemic problems. It was as if the managers were offloading their responsibility for the problems observed onto the individual practitioner.

The difference between disclosure and reporting is not as obvious or problematic in all professions:

- where individual professional contact with clients is very close, such as in medicine, reporting and disclosure are two very different things;
- where the relationship is more distant, such as in air traffic control, the distinction blurs because for individual air traffic controllers there is not immediately a party to disclose to. The air traffic control organization, however, can be said to have an obligation to disclose.

If organizational disclosure or reporting does not occur, then the mistakes made by people inside that organization may no longer be seen as honest, and the organization can get in trouble as a result. This goes for individuals too. It may have played a role in the Heathrow 747 incident described in the previous chapter, as it plays a role in many other cases.

The Importance, Risk and Protection of Disclosure

Not providing an account of what happened may mean there is something to hide. And if there is something to hide, then what happened is probably not an "honest" mistake.

The killing of a British soldier in Iraq by a US pilot was a "criminal, unlawful act," tantamount to manslaughter, a British coroner ruled. The family of Lance Corporal of Horse Matty Hull, who died in March 2003, were told at the inquest in Oxford, England, that it was "an entirely avoidable tragedy." His widow Susan welcomed the verdict, saying it was what the family had been waiting four years for. Hull said she did not want to see the pilot prosecuted, but felt she been "badly let down" by the US government, which consistently refused to cooperate.

Susan Hull had also been told by the UK Ministry of Defence that no cockpit tape of the incident existed. This was proven untrue when a newspaper published the tape's contents and when it was later posted on the internet. It showed how Hull was killed when a convoy of British Household Cavalry vehicles got strafed by two US A10 jets. The British Ministry of Defence issued an apology over its handling of the cockpit video, while the US Department of Defense denied there had been a cover-up and remained adamant that the killing was an accident.

The coroner, Andrew Walker, was damning in his appraisal of the way the Hull family had been treated. "They, despite request after request, have been, as this court has been, denied access to evidence that would provide the fullest explanation to help understand the sequence of events that led to and caused the tragic loss of LCorp. Hull's life," he said. "I have no doubt of how much pain and suffering they have been put through during this inquisition process and to my mind that is inexcusable," he said.[2]

Not disclosing often means that a mistake will no longer be seen as honest. Once a mistake is seen as dishonest, people may no longer care as much about what happens to the person who made that mistake, or to the party (for example, the organization) responsible for withholding the information. This is where a mistake can get really costly—both financially and in terms of unfavorable media exposure, loss of trust and credibility, regulatory scrutiny or even legal action.

I recall one adverse event where the family was very upset, not only about what had happened, but about the organization not being seen as forthcoming. The family had been invited by the organization to stay in a nice hotel for some of the legal proceedings. Feeling injured and let down, they ordered as much expensive room service as possible, and then threw it all away. Having been hurt by the organization, they wanted to hurt the organization as much as possible in return.

Non-disclosure is often counterproductive and expensive. Silence can get interpreted as stone-walling, as evidence of "guilty knowledge." It is well known that lawsuits in healthcare are often more a tool for discovery than a mechanism

for making money. People don't generally sue (in fact, very few actually do). But when they do, it is almost always because all other options to find out what happened have been exhausted.[3] Paradoxically, however, lawsuits still do not guarantee that people will ever get to know the events surrounding a mishap. In fact, once a case goes to court, "truth" will likely be the first to suffer. The various parties may likely retreat into defensive positions from which they will offer only those accounts that offer them the greatest possible protection against the legal fallout.

The Ethical Obligation to Disclose

Not being honest, or not apologizing for a mistake, is what often causes relationships to break down, rather than the mistake or mishap itself. This makes honesty all the more important in cases where there is a prior professional relationship,[4] such as patient–doctor.

In a specific example, the Code of Medical Ethics of the American Medical Association says since 1981 that:

> It is a fundamental requirement that a physician should at all times deal honestly and openly with patients … Situations occasionally occur in which a patient suffers significant medical complications that may have resulted from the physician's mistake or judgment. In these situations, the physician is ethically required to inform the patient of all the facts necessary to ensure understanding of what has occurred … Concern regarding legal liability which might result following truthful disclosure should not affect the physician's honesty with a patient.[5]

This is unique, because few codes exist that specifically tell professionals to be honest. It also spells out in greater detail which situations ("that may have resulted from a physician's mistake or judgment") are likely to fall under the provisions of the Code. So this is a good start. But it still leaves a large problem. What does "honesty" mean? Being honest means telling the truth. But telling the truth can be reduced to "not lying." If it is, then there is still a long distance to full disclosure. What to say, how much to say, or how to say it, often hinges more on the risks that people see with disclosure, than with what a code or policy tells them to do.

The Risk with Disclosure

If structural arrangements and relationships inside an industry, or a profession, are such that all bets are off when you tell your story, then people will find

ways to not disclose, or to only disclose in ways that protect them against the vagaries and vicissitudes of the system.

Nancy Berlinger describes how medical education has a "hidden curriculum." This hidden curriculum can be seen as the sort of spontaneous system of informal mentoring and apprenticeship that springs up in parallel to any formal program. It teaches students and residents, mainly through example, how to think and talk about their own mistakes and those of their colleagues. They learn, for example, how to describe mistakes so that they no longer are mistakes. Instead, they can become:

- "complications;"
- the result of a "non-compliant" patient;
- a "significantly avoidable accident;"
- an "inevitable occasional untoward event;"
- an "unfortunate complication of a usually benign procedure."[6]

Medical professionals, in the hidden curriculum, also learn how to talk about the mistake among themselves, while reserving another version for the patient and family and others outside their immediate circle of professionals. A successful story about a mistake is one that not only (sort of) satisfies the patient and family, but one that also protects against disciplinary measures and litigation.

Many, if not all, professions, have a hidden curriculum. Perhaps it teaches professionals the rhetoric to make a mistake into something that no longer is a mistake. Perhaps it teaches them that there is a code of silence, an "ómerta," that proscribes collaborating truthfully with authorities or other outside parties.

The Protection of Disclosure

The protection of disclosure should first and foremost come from structural arrangements made by the organization or profession. One form of protecting disclosure is that of "I'm sorry laws." According to these laws (now implemented in for example the US States of Oregon and Colorado), doctors can say to patients that they are sorry for the mistake(s) they committed. This does not offer them immunity from lawsuits or prosecution, but it does protect the apology statement from being used as evidence in court. It also does not prevent negative consequences, but at least that which was disclosed cannot be used directly against the professional (as it was with the ICU nurse in the Prologue). Such protection is not uncontroversial, of course. If you make a mistake, you should not only own up to it but also face the consequences, some would say.

Which other professions have such cozy protections? This is where different ethical principles start to diverge.

What is Being Honest?

So what really is honesty, or telling the truth in reporting? We have seen a lot of different signals in the previous two chapters and before:

- Given an honest report of her role in the death of an infant got the nurse in the Xylocard case trouble.
- When managers are in charge of a reporting system, practitioners likely give them one story of what happened. Whether that is the "truth" or not is almost irrelevant: it is about making the aftermath of a mistake or incident as painless for everybody as possible.
- When taken to court, practitioners may tell yet another story. Again, whether that is the honest truth or not (something that legal systems often quite erroneously claim they can get out of people) is not the point: it is rather about minimizing the spiraling negative consequences of being put on trial.
- Honestly disclosing to a family what happened to a patient is often under pressure from what a caregiver learned in the hidden curriculum.

The question that comes up is this: is honesty a goal in itself? Perhaps honesty should fulfill the larger goals of:

- learning from a mistake to improve safety; and
- achieving justice in its aftermath.

These are two goals that serve the common good. Supposedly pure honesty can sometimes weaken that common good. For example, both justice and safety were hurt when the nurse from the prologue was put on trial as a result of her honest reporting. Honesty, or truth-telling—should we always pursue it because it is the "right" thing to do, no matter what the consequences could be?

Dietrich Bonhoeffer, writing from his cell in the Tegel prison in Berlin in 1943 Nazi Germany, drafted a powerful essay on this question. He was being held on suspicion of a plot to overthrow Hitler, a plot in which he and his family were actually deeply involved.[7] If he were to tell the truth, he would have let murderers into his family. If he were to tell the truth, he would have to disclose where other conspirators were hidden. So would not telling this make him a

liar? Did it make him, in the face of Nazi demands and extortions, immoral, unethical? Bonhoeffer engaged in nondisclosure, and outright deception.

The circumstances surrounding truth-telling in professions today is not likely as desperate and grave as Bonhoeffer's (he was executed in a concentration camp just before the end of the war, in April 1945). But his thoughts have a vague reflection in the fears of those who consider disclosing or reporting today. What if they tell the whole truth—rather than a version that keeps the system happy and them protected? Bonhoeffer makes a distinction between the morality and epistemology of truth-telling that may offer some help here:

- the **epistemology** of truth-telling refers to the validity or scope of the knowledge offered. In that sense, Bonhoeffer did not tell the truth (but perhaps the nurse in the Xylocard case did);
- the **morality** of truth-telling refers to the correct appreciation of the real situation in which that truth is demanded. The more complex that situation, the more troublesome the issue of truth-telling becomes (the nurse in the Xylocard case may not have done this, but perhaps should have).

Bonhoeffer's goal in not disclosing was not self-preservation but the protection of the conspiracy's efforts to jam the Nazi death machine, thereby honoring the perspective of the most vulnerable. That his tormentors wanted to know the truth was unethical, much more so than Bonhoeffer's concealment of it.

Translate this into the situations faced by the nurse in the Xylocard case (or the pilot in the story of Chapter 2, or the nurse Julie in Chapter 1). Here it may be less ethical for prosecutors or judges in positions of power to demand the full truth, than it would have been for these professionals to offer only a version of that truth.

Asking for honesty initially, as the airline did, and as the hospital's procedures proscribed, is reasonable. It is here that honesty can contribute to the larger goals of accountability and learning. Responding to this, as the nurse did, was reasonable too, and an honest attempt to give account and perhaps help the hospital learn. Not responding to it, as the captain did, was perhaps foolhardy and unreasonable (but we do not know what history the airline had, or what signals it had sent out earlier about its tolerance for reporters and their mistakes).

But going to court, and demanding honesty there, became a different issue altogether in both these cases. Once adversarial positions were lined up against each other in a trial, where one party had the capacity to wreak devastating consequences on another, the ethics of honesty got a whole new

dynamic. Demanding honesty in these cases ended up serving only very narrow interests, such as the preservation of a company's or hospital's reputation, or *their* protection from judicial pressure. Or it deflected responsibility from the regulator (who employed the prosecutor) after allowing the airline to routinely hand out dispensations from existing rules—something that played a role in the incident in Chapter 2. Quite unlike Bonhoeffer, who must have been under tremendous pressure, self-preservation did become the overriding aim of the parties in these trials.

Wringing honesty out of people in vulnerable positions is neither just nor safe. It does not bring out a story that serves the dual goal of satisfying calls for accountability and helping with learning. It really cannot contribute to just culture.

Notes

1 Sharpe, V.A. (2003). Promoting patient safety: An ethical basis for policy deliberation. *The Hastings Center Report*, 4, S1–S19.
2 *The Observer* (Oxford, UK), Sunday, March 18, 2007, 1.
3 Berlinger, N. (2005). *After harm: Medical error and the ethics of forgiveness*. Baltimore, MD: The Johns Hopkins University Press.
4 Cohen, J.R. (2002). Legislating apology: The pros and cons. *University of Cincinnati Law Review*, 70, 819–72.
5 American Medical Association (1981). *Code of Medical Ethics, Ethical opinions E-8.12*. Chicago, IL: AMA.
6 Berlinger, op. cit.
7 Nancy Berlinger has written extensively on disclosure and apology. See: Berlinger, ibid., 42–62.

This page intentionally left blank

5 Are All Mistakes Equal?

If you report or disclose a mistake, or if your organization discloses information about an incident, what are the risks? We are back at that question from the previous chapters.

The risk depends not only on the culture in which you disclose or report. Of course, if that culture is "just," according to the ideas in this book, it will have found a balance between responding to calls for accountability and making improvements in safety. Mistakes will be seen as an opportunity for learning, and responsibility will be designated to make sure that lessons get implemented.

But still, not all mistakes or incidents will be seen the same way. Not all mistakes, in other words, are equally "forgivable." In this chapter, I look at one basic distinction that can be made in many professions: that between technical and normative errors.

Technical vs Normative Errors

Most professions make distinctions in how they treat mistakes and incidents. When studying the way surgeons treat errors that can (or have) hurt patients, Charles Bosk, a sociologist, saw a remarkable pattern. Surgeons and other physicians made a distinction between what he began to call *technical* and *normative* errors.[1] To be sure, it was not the error that is either technical or normative. It became technical or normative because of the way people looked at the error, because of what they saw in it, talked about it and how they responded to it.

The distinction can have powerful consequences for how your organization (or surrounding society) is prepared to deal with an error that occurred. Whether you construct an error as normative or technical has far-reaching consequences for exacting accountability and encouraging learning.

Technical Errors: Errors in a Role

When a practitioner makes a technical mistake, she or he is performing her or his role diligently, but the present skills fall short of what the task requires.

For example, when a pilot makes a rough landing, this could likely be the effect of skills that have yet to be developed or refined. Technical errors will be seen as opportunities for instructors or colleagues to pass on "tricks of the trade" (for example, "start shifting your gaze ahead when flaring").

People can be quite forgiving of even serious lapses in technique, as they see these as a natural by-product of learning-by-doing. Technical errors do not just have to be connected to the physical handling of a process or its systems; they can also involve interaction with others in the system, for instance, air traffic control, nurses, night staff, physicians. The person in question, for example, may have seen the need to coordinate (and may even be doing just that), but does not have the experience or finely-developed skills to recognize how to be sensitive to the constraints or opportunities of the other members of the system.

For an error to be constructed as technical, however, it has to meet two conditions:

- One is obviously that the frequency or seriousness should decrease as experience goes up. When a person keeps making the same mistakes over and over, it may be difficult to keep seeing them as purely technical. As long, however, as the person making the errors shows a dedication to learning and to his or her part in creating safety, that person is still conscientiously filling his or her role.
- The other condition for a technical error is that it should not be denied, by the pilot involved, as an opportunity for learning and improvement. If a practitioner is not prepared to align discrepant outcomes and expectations by looking at him- or herself, but rather turns onto the one who revealed the discrepancy, trainers or supervisor or managers (or courts) will no longer see the error as purely technical.

A flight instructor reports: We had been cleared inbound to a diversion airport due to weather. We were on downwind when an airliner came on the frequency, and was cleared for the ILS (instrument landing system) toward the opposite runway. The student proceeded to extend his downwind to the entry point he had chosen, even though the field was now fully visible. He was entirely oblivious to hints from air traffic control to turn us onto base so we could make it in before the airliner from the opposite side. When the student still did not respond, I took control and steered our aircraft onto base. We completed the landing without incident before the airliner came in. Upon debriefing, however, the student berated me for taking control, and refused to accept the event as an

opportunity to learn about "fitting in" with other traffic at a dynamic, busy airport. He felt violated that I had taken control.

Professionals with limited experience may not be very sensitive to the unfolding context in which they work. They have simply not yet learned which cues to pick up on and how to respond to them. People will see such insensitivity as a technical issue consistent with the role of student, and a due opening for enlightenment. Sticking to the plan, or behaving strictly in the box, even though a situation has unfolded differently, has been known to lead to problems and even accidents. So valuable lessons are those that demonstrate how textbook principles or dogged elegance sometimes have to be compromised in order to accommodate a changing array of goals. Professionals can otherwise end up in a corner. Surgery has a corollary here: "excellent surgery can make dead patients."

The benefits of technical errors almost always outweigh the disadvantages. Of course, this is so in part because the division of labor between senior and junior practitioners in most operating worlds (or between instructors and students), is staggered so that no-one advances to more complex tasks until they have demonstrated their proficiency at basic ones.

Bosk tells how Carl, a surgical intern, was closing an incision, while Mark, the chief resident, was assisting. Carl was ill at ease. He turned to Mark and said "I can't do it." Mark said, "What do you mean, you can't? Don't ever say you can't. Of course you can." "No, I just can't seem to get it right." Carl had been forced to put in and remove stitches a number of times, unable to draw the skin closed with the proper tension. Mark replied, "Really, there is nothing to it;" and, taking Carl's hand in his own, he said, "The trick is to keep the needle at this angle and put the stitch through like this," all the while leading Carl through the task. "Now, go on." Mark then let Carl struggle through the rest of the closure on his own.[2]

If aid is necessary, there are almost always only two responses:

- verbal guidance is offered, with hints and pointers;
- or the superordinate takes over altogether.

The latter option is taken when time constraints demand quick performance, or when the task turns out to be more complex than people initially assumed. This division of labor can also mean that subordinates feel held back, with not enough opportunity to exercise their own technical judgment. The example

above could be an instance of this, where the division of labor is seen by the student as stacked in favor of the instructor. For instructors, supervisors, managers and others, the challenge is always to judge whether the learning return from letting the practitioner make the mistake is larger than from helping her or him avert it and clearly demonstrating how to do so.

In another example, Bosk tells of the difficulty of performing a myleogram (a diagnostic procedure involving the removal of spinal fluid and the injection of dye in the spinal column) that had been ordered for a patient named Mr Eckhardt. A senior student was to instruct a junior student in the procedure. They tried without any success to get the needle in the proper space. After some fumbling and a few sticks at Eckhardt, the senior student instructed the junior student to go "get Paul" (a second-year resident). Paul came in and surveyed the situation. After examining Eckhardt's back he told the students, who were profusely apologizing for their failure, not to worry; that the problem was in Mr Eckhardt's anatomy and not in their skills. He then proceeded with some difficulty to complete the procedure, instructing the students all the while.[3]

As for professionals, they should not be afraid to make mistakes. They should be afraid of not learning from the ones that they do make. Bosk's study showed how self-criticism is strongly encouraged and expected of surgeons in the learning role (which is to say, almost every surgeon). Everybody can make mistakes, and they can generally be managed.

Denial or defensive posturing instead discourages such learning. It allows the trainee or subordinate to delegitimize mistake by turning it into something shameful that should be brushed aside, or into something irrelevant that should be ignored. Denying that a technical error has occurred is not only inconsistent with the idea that they are the inevitable by-product of training. It also truncates an opportunity for learning. Work that gets learned-by-doing lives by this pact: technical errors and their consequences are to be acknowledged and transformed into an occasion for positive experience, learning, improvement. Not going along with that implicit pact is no longer a technical error, it is a normative one.

Normative Errors: Errors in Assuming a Role

Technical errors say something about the professional's level of training or experience. Normative errors say something about the professional him- or herself relative to the profession. Normative errors are about professionals not discharging their role obligations diligently.

- Technical errors create extra work, both for superordinate and subordinate. That, however, is seen as legitimate: it is part of the game, the inevitable part of learning by doing, of continuous improvement.
- The extra work of normative errors, however, is considered unnecessary.

In some cases, it shows up when a crewmember asserts more than his or her role allows:

A senior airline captain told me about one case which he constructed as a normative error. "It was my turn to go rest," he said, "and, as I always do, I told the first and second officer 'If anything happens, I want to know about it. Don't act on your own, don't try to be a hero. Just freeze the situation and call me. Even if it's in the middle of my break, and I'm asleep, call me. Most likely I'll tell you it's nothing and I'll go right back to sleep. I may even forget you called. But call me.' When I came back from my break, it turned out that a mechanical problem had developed. The first officer, in my seat, was quite comfortable that he had handled the situation well. I was irate. Why hadn't he called me? How can I trust him next time? I am ultimately responsible, so I have to know what's going on."

The situation was left less resilient than it could (and, in the eyes of the captain, should) have been: leaving only two more-junior crewmembers, with no formal responsibility, in charge of managing a developing problem. Of course, there are potential losses associated with calling:

- the superordinate could think the call was superfluous and foolish, and get cranky because of it (which the first officer in the example above may have expected and, as it turned out, misjudged);
- the subordinate foregoes the learning opportunity and gratification of solving a problem her- or himself.

But the safe option when in doubt is always to call, despite the pressures not to. That is, in many cases, how a subordinate crewmember is expected to discharge her or his role obligations. In other cases, fulfilling those obligations is possible *only* by breaking out of the subordinate role, as a chief pilot once told me:

My problem is with first officers who do not take over when the situation calls for it. Why do we have so many unstabilized approaches to runways in (a particular area of our network)? If the captain is flying, first officers should first point out to him or her that he or she is out of bounds, and if that does not work, they should take over. Why don't they, what makes it so difficult?

The chief pilot here flagged the absence of what may turn out critical for the creation of safety in complex systems: the breaking-out of roles and power structures that were formally designed into the system. Roles and power structures often go hand-in-glove (such as captain and first officer, doctor and nurse), and various programs (for example, crew resource management training in aviation and healthcare, team resource management in air traffic control, bridge resource management in shipping, and so forth) aim to soften role boundaries and flatten hierarchies. These programs want to increase opportunities for coordinating viewpoints and sharing information. Where people do not do this, they fail to discharge their role obligations too—in this case by not acknowledging and deploying the flexibility inherent in any role.

Errors and the Importance of Reporting and Disclosure

Let us now revisit the theme of the previous two chapters.

If a practitioner is not "honest" in openly discussing and wanting to learn from a mistake, then we will tend to see this as a normative breach of professional responsibility.

This is not so strange. In the kinds of operating worlds where we believe a just culture is important, it is very difficult to know and anticipate all the problems that may occur during a lifetime of practice. There will always be things that practitioners remain inexperienced with, simply because that kind of problem, in that kind of way, has not appeared before. Indeed, in complex and dynamic work, where resource limitations and uncertainty reign, failure is going to be a lasting statistical reality.

The possibility of suffering technical errors will consequently never go away entirely. In such worlds, where the knowledge base on how to create safety is inherently and permanently incomplete, many believe firmly in the importance of disclosing, discussing and learning from error. When that does not happen, even an honest, technical error can become seen as a dishonest normative one.

Covering up is never really excusable, Bosk quotes an attending physician as saying. You have to remember that each time a resident hides information, he is affecting someone's life. Now in this business it takes a lot of self-confidence, a lot of maturity, to admit errors. But that's not the issue. No mistakes are minor. All have a mortality and a morbidity. Say I have a patient who comes back from the operating room and he doesn't urinate. And say my intern doesn't notice or he decides it's nothing serious and he doesn't catheterize the

guy and he doesn't tell me. Well, this guy's bladder fills up. There's a foreign body and foreign bodies can cause infections; infection can become sepsis; sepsis can cause death. So the intern's mistake here can cause this guy hundreds of dollars in extra hospitalization and it could cost him his life. All mistakes have costs attached to them. Now a certain amount is inevitable. But it is the obligation of everyone involved in patient care to minimize mistakes. The way to do that is by full and total disclosure.[4]

The obligation to report or disclose, discuss and learn seems to be a critical hinge in how we believe a just culture should work. But, as we saw in the previous chapter, honest and open accounting can seem dangerous to many practitioners. How an error might be interpreted after-the-fact is sometimes entirely up for grabs. A technical one (missing an approach, or supplying the wrong drug because of inexperience with that particular drug or procedure or kind of patient) can easily be converted into a normative error—with much more serious consequences for accountability (such as a criminal trial).

Knowing the Outcome

Another aspect that often enters this deliberation is hindsight. If the outcome of a mistake is really bad, we are actually more likely to see that mistake as more culpable than if the outcome is positive. There is more to account for, in other words. This can be strange, because the same mistake can be looked at in a completely different light (for example, without knowledge of outcome) and then it does not look as bad or culpable at all. There is not much to account for. So hindsight plays a huge role in how we handle the aftermath of mistake. In the next chapter, we briefly pause to look at one such case: it looks normal, professional, plausible and reasonable from one angle. And culpable from another. The hinge between the two is hindsight: knowing how things turned out.

Notes

1 Bosk, C.L. (2003). *Forgive and remember: Managing medical failure* (2nd edition). Chicago: University of Chicago Press.
2 Ibid., 45.
3 Ibid., 44.
4 Ibid., 60–61.

This page intentionally left blank

6 Hindsight and Determining Culpability

There is a children's song in Swedish about a fire breaking out in a store. The firemen show up, do their job and succeed in putting out the fire. The song then ends: "Think about it—nobody got hurt: you did a great job!" Similarly, a saying in Dutch proclaims: "Eind goed, al goed." If it ended well, all is well.

We assume that if an outcome is good, then the process leading up to it must have been good too—that people did a good job. The inverse is true too: we often conclude that people may not have done a good job when the outcome is bad.

This is reflected, for example, in the compensation that patients get when they sue their doctors. The severity of their injury is the most powerful predictor of the amount that will be awarded. The more severe the injury, the more the compensation. As a result, physicians believe that liability correlates not with the quality of the care they provide, but with outcomes over which they have little control.[1]

Here is the common reflex: the worse the outcome, the more we feel there is to account for. This is strange: the process that led up to the outcome may not have been very different from when things would have turned out right.

Consider this: "If catnapping while administering anesthesia is negligent and wrongful, it is a behavior that is negligent and wrongful whether harm results or not."[2]

Also, quality processes can still lead to bad outcomes because of the complexity, uncertainty and dynamic nature of work.

Hindsight Bias

If we know that an outcome is really bad, then this influences how we see the behavior that led up to it. We will be more likely to look for mistakes. Or even negligence. We will be less inclined to see the behavior as "forgivable." The worse the outcome, the more likely we are to see mistakes, and the more things we discover which people have to account for. Here is why:

- after an incident, and especially after an accident (with a dead patient, or wreckage on a runway), it is easy to see where people went wrong, what they should have done or avoided;
- with hindsight, it is easy to judge people for missing a piece of data that turned out to be critical;
- with hindsight, it is easy to see exactly the kind of harm that people should have foreseen and prevented. That harm, after all, has already occurred. This makes it easier for behavior to reach the standard of "negligence."

The reflex is counterproductive: like physicians, other professionals and entire organizations may invest in ways that enable them to account for a bad outcome (more bureaucracy, stricter bookkeeping, practicing defensive medicine). These investments may have little to do with actually providing a safe process.

Yet hindsight bias is one of the most consistent and well-demonstrated biases in psychology. But incident reporting systems or legal proceedings—systems that somehow have to deal with accountability—have essentially no protections against it.

Lord Anthony Hidden, the Chairman of the investigation into the devastating Clapham Junction railway accident in England wrote, "There is almost no human action or decision that cannot be made to look flawed and less sensible in the misleading light of hindsight. It is essential that the critic should keep himself constantly aware of that fact."[3]

If we don't heed Anthony Hidden's warning, hindsight bias can have a profound influence on how we judge past events: Hindsight means that we:

- oversimplify causality ("this led to that") because we can start from the outcome and reason backwards to presumed or plausible "causes;"
- overestimate the likelihood of the outcome (and people's ability to foresee it), because we already have the outcome in our hands;
- overrate the role of rule or procedure "violations." While there is always a gap between written guidance and actual practice (and this almost never

leads to trouble), that gap takes on causal significance once we have a bad outcome to look at and reason back from;
- misjudge the prominence or relevance of data presented to people at the time;
- match outcome with the actions that went before it. If the outcome was bad, then the actions leading up to it must have been bad too—missed opportunities, bad assessments, wrong decisions and misperceptions.

Shooting Down an Airliner

Zvi Lanir, a researcher of decision-making, tells a story of a 1973 encounter between Israeli fighter jets and a Libyan airliner. He tells the story from two angles: from that of the Israeli Airforce, and then of the Libyan airliner.[4] This works so well that I do that here too. Not knowing the real outcome, Israeli actions make good sense, and there would be little to account for. The decision-making process has few if any bugs:

- the incident occurred during daylight, unfolded during 15 minutes, and happened less than 300 kilometers away from Israeli headquarters;
- the people involved knew each other personally, were well acquainted with the terrain, and had a long history of working together through crises that demanded quick decisions;
- there was no evidence of discontinuities or gaps in communication or the chain of command;
- the Israeli Air Force Commander happened to be in the Central Air Operations Center, getting a first-hand impression as events advanced;
- the Israeli Chief of Staff was on hand by telephone for the entire incident duration too.

The process that led up to the outcome, in other words, reveals few problems. It is even outstanding: the Chief of Staff was on hand to help with the difficult strategic implications; the Air Force Commander was in the Operations Center where decisions were taken, no gaps in communication or chain of command occurred. Had the outcome been as the Israelis may have suspected, then there had been little or nothing to account for. Things went as planned, trained, expected, and a good outcome resulted.

But then, once we find out the real outcome (the real nature of the Libyan plane), we suddenly may find reason to question all kinds of aspects of that very same process. Was it right or smart to have the Chief of Staff involved? What

about the presence of the Air Force Commander? Did the lack of discontinuities in communication and command chain actually contribute to nobody saying "wait a minute, what are we doing here?" The same process—due to our learning about its real outcome—gets a different hue. A different accountability. As with the doctors that get sued: the worse the outcome, the more there is to account for. Forget the process that led up to it.

A Normal, Technical Professional Error

At the beginning of 1973, Israeli intelligence received reports on a possible suicide mission by Arab terrorists. The suggestion was that they would commandeer a civilian aircraft and try to penetrate over the Sinai desert with it, to self-destruct on the Israeli nuclear installation at Dimona or other targets in Beer Sheva. On February 21, the scenario seemed to be set in motion. A sandstorm covered much of Egypt and the Sinai desert that day.

At 13:54 (1:54 pm), Israeli radar picked up an aircraft flying at 20,000 feet in a northeasterly direction from the Suez bay. Its route seemed to match that used by Egyptian fighters for their intrusions into Israeli airspace, known to the Israelis as a "hostile route." None of the Egyptian war machinery on the ground below, supposedly on full alert and known to the Israelis as highly sensitive, came into action to do anything about the aircraft. It suggested collusion or active collaboration.

Two minutes later, the Israelis sent two F-4 Phantom fighter jets to identify the intruder and intercept it if necessary. After only a minute, they had found the jet. It turned out to be a Libyan airliner. The Israeli pilots radioed down that they could see the Libyan crew in the cockpit, and that they were certain that the Libyans could see and identify them (the Shield of King David being prominently displayed on all Israeli fighter jets).

At the time, Libya was known to the Israelis for abetting Arab terrorism, so the Phantoms were instructed to order the intruding airliner to descend and land on the nearby Refidim airbase in the south of Israel. There are international rules for interception, meant to prevent confusion in tense moments where opportunities for communication may be minimal, and opportunities for misunderstanding huge. The intercepting plane is supposed to signal by radio and wing-rocking, while the intercepted aircraft must respond with similar signals, call the air traffic control unit it is in contact with and try to establish radio communication with the interceptor.

The Libyan airliner did none of that. It continued to fly straight ahead, toward the northeast, at the same altitude. One of the Israeli pilots then sided up to the jet, flying only a few meters beside its right cockpit window. The co-pilot was looking right at him. He then appeared to signal, indicating that the Libyan crew had understood what was going on and that they were going to comply with the interceptors. But it did not change course, nor did it descend.

At 14:01, the Israelis decided to fire highly luminescent tracer shells in front of the airliner's nose, to force it to respond. It did. The airliner descended and turned toward the Refidim airbase. But then, when it had reached 5,000 feet and lowered its landing gear, the airliner's crew seemed to change its mind. Suddenly it broke off the approach, started climbing again, putting away the landing gear, and turned west. It looked like an escape.

The Israelis were bewildered: a Captain's main priority is the safety of his or her passengers: doing what this Libyan crew was doing showed none of that concern. So maybe the aircraft had been commandeered and the passengers (and crew) were along for the ride, or perhaps there were no passengers onboard at all. Still, these were only assumptions. It would be professional, the right thing to do, to double-check. The Israeli Air Force commander decided that the Phantoms should take a closer look, again.

At 14:05, one of the Phantoms flew by the airliner within a few meters and reported that all the window blinds were drawn. The Air Force Commander became more and more convinced that it may have been an attempted, but foiled, terrorist attack. Letting the aircraft get away now would only leave it to have another go later.

At 14:08, he gave the order for the Israeli pilots to fire at the edges of the wings of the airliner, so as to force it to land. The order was executed. But even with the tip of its right wing hit, the airliner still did not obey the orders and continued to fly westward. The Israelis opened all international radio channels, but could not identify any communication related to this airliner. Two minutes later, the Israeli jets were ordered to fire at the base of the wings. This made the airliner descend and aim, as best it could, for a flat sandy area to land on. The landing was not successful. At 14:11, the airliner crashed and burned.

A Normative, Culpable Mistake

Had the wreckage on the ground revealed no passengers, and a crew intent on doing damage to Israeli targets, the decisions of the relevant people within the Israeli Air Force would have proven just and reasonable. There would be no

basis for asserting negligence. As it turned out, however, the airliner was carrying passengers. Out of 116 passengers and crew, 110 were killed in the crash.

The cockpit voice recorder revealed a completely different reality, a different "truth." There had been three crew members in the cockpit: a French captain, a Libyan co-pilot and a French flight engineer (sitting behind the two pilots). The captain and the flight engineer had been having a conversation in French, while enjoying a glass of wine. The co-pilot evidently had no idea what they were talking about, lacking sufficient proficiency in French. It was clear that the crew had no idea that they were deviating more than seventy miles from the planned route, first flying over Egyptian and later Israeli war zones.

At 13:44, the captain first became uncertain of his position. Instead of consulting with his co-pilot, he checked with his flight engineer (whose station has no navigational instruments), but did not report his doubts to Cairo Approach. At 13:52, he got Cairo's permission to start a descent toward Cairo international airport. At 13:56, still uncertain about his position, the captain tried to receive Cairo's radio navigation beacon, but got directions that were contrary to those he had expected on the basis of his flight plan (as the airport was now gliding away further and further behind him).

Wanting to sort out things further, and hearing nothing else from Cairo approach, the crew continued on their present course. Then, at 13:59, Cairo came on the radio to tell the crew that they were deviating from the airway. They should "stick to beacon and report position." The Libyan co-pilot now reported for the first time that they were having difficulties in getting the beacon.

At 14:00, Cairo approach asked the crew to switch to Cairo control: a sign that they believed the airliner was now within range to land, close to the airport. Two minutes later the crew told Cairo control that they were having difficulties receiving another beacon (Cairo NDB, or Non-Directional Beacon, with a certified range of about 50 kilometers), but did not say they were uncertain of their position. Cairo control asked the aircraft to descend to 4,000 feet.

Not much later, the co-pilot reported that they had "four MiGs" behind them, mistaking the Israeli Phantoms for Soviet-built Egyptian fighter jets. The captain added that he guessed they were having some problems with their heading and that they now had four MiGs behind them. He asked Cairo for help in getting a fix on his position. Cairo responded that their ground-based beacons were working normally, and that they would help find the airliner by radar.

Around that same time, one of the Phantoms had hovered next to the co-pilot's window. The co-pilot had signaled back, and turned to his fellow crewmembers to tell them. The captain and flight engineer once again engaged

in French about what was going on, with the captain angrily complaining about the Phantom's signals that that was not the way to talk to him. The co-pilot did not understand.

At 14:06 Cairo control advised the airliner to climb to 10,000 feet again, as they were not successful in getting a radar fix on the airplane (it was way out of their area and probably not anywhere near where they expected it to be). Cairo had two airfields: an international airport on the west side, and a military airbase on the east. The crew likely interpreted the signals from the "MiGs" as them having overshot the Cairo international airport, and that the fighter jets had come to guide them back. This would explain why they suddenly climbed back up after approaching the Refidim airbase. Suspecting that they had lined up for Cairo East (the military field), now with fighters on their tail, the crew decided to turn west and find the international airport instead.

At 14:09, the captain snapped at Cairo control that they were "now shot by *your* fighter," upon which Cairo said they were going to tell the military that they had an unreported aircraft somewhere out there but did not know where it was. When they were shot at again, the crew became panicked, accelerating their talk in French. Were these Egyptians crazy? Then, suddenly, the co-pilot identified the fighters as Israeli warplanes. It was too late, with devastating consequences.

Hindsight and Culpability

The same actions and assessments that represent a conscientious discharge of professional responsibility can, with knowledge of outcome, become seen as a culpable, normative mistake.

With knowledge of outcome, we know what the commanders or pilots should have checked better (because we now know what they missed: for example that there were passengers on board and that the jet was not hijacked). After the fact, there are always opportunities to remind professionals what they could have done better (could you not have checked with the airline? Could your fighters not have made another few passes on either side to see the faces of passengers?). Again, had the airliner not contained passengers, nobody would have asked those questions. The professional discharge of duty would have been sufficient if that had been the outcome. And, conversely, had the Israelis known that the airliner contained passengers, and was not hijacked but simply lost, they would never have shot it down.

Few in positions to judge the culpability of a professional mistake have as much (or any) awareness of the debilitating effects of hindsight. Judicial proceedings, for example, will stress how somebody's behavior did not make sense, how it violated narrow standards of practice, rules or laws.

Jens Rasmussen once pointed out that if we find ourselves (or a prosecutor) asking "how could they have been so negligent, so reckless, so irresponsible?," then this is not because the people in question were behaving bizarrely. It is because we have chosen the wrong frame of reference for understanding their behavior. The frame of reference for understanding people's behavior, and judging whether it made sense, is their own normal work context, the context they were embedded in. This is the point of view from where decisions and assessments are sensible, normal, daily, unremarkable, expected. The challenge, if we really want to know whether people anticipated risks correctly, is to see the world through their eyes, *without* knowledge of outcome, without knowing exactly which piece of data will turn out critical afterward.

The Worse the Outcome, the More to Account For

If an outcome is worse, then we may well believe that there is more to account for. That is probably fundamental to the social nature of accountability. We may easily believe that the consequences should be proportional to the outcome of somebody's actions. Again, this may not be seen as fair: recall the example from the beginning of the chapter. Physicians believe that liability is connected to outcomes that they have little control over, not to the quality of care they provided. To avoid liability, in other words, you don't need to invest in greater quality of care. Instead, you invest in defensive medicine: more tests, covering your back at every turn.

The main question for a just culture is not about matching consequences with outcome. It is this: did the assessments and actions of the professionals at the time make sense, given their knowledge, their goals, their attentional demands, their organizational context? Satisfying calls for accountability here would not be a matching of bad outcome with bad consequences for the professionals involved. Instead, accountability could come in the form of reporting or disclosing how an assessment or action made sense at the time, and how changes can be implemented so that the likelihood of it turning into a mistake goes down.

Notes

1 Brennan, T.A., Sox, C.A., and Burstin, H.R. (1996). Relation between negligent adverse events and the outcomes of medical malpractice litigation. *New England Journal of Medicine*, 335, 1963–7.

2 Dauer, E.A. (2004). Ethical misfits: Mediation and medical malpractice litigation. In V.A. Sharpe (ed.), *Accountability: Patient safety and policy reform*, 185–201. Washington, DC: Georgetown University Press, 194.

3 Hidden, A. (1989). *Clapham Junction Accident Investigation Report*. London: HM Stationery Office, 147.

4 A version of the story was published in Lanir, Z. (1989). The reasonable choice of disaster: The shooting down of the Libyan airliner on 21 February 1973. *Journal of Strategic Studies*, 12, 479–93.

This page intentionally left blank

7 "You Have Nothing to Fear if You've Done Nothing Wrong"

"You have nothing to fear if you have done nothing wrong." That was said by a prosecutor in a European country who was responding to concerns from the aviation sector that human errors—normal, honest mistakes—were being converted into criminal behavior by his office. Some pilots and air traffic controllers were being fined or charged for rule infractions that were part and parcel of getting the job done. And they and their colleagues were getting anxious. Was data supplied in good faith, for example through incident reports, going to be used against them? Were there enough protections against the prying of a prosecutorial office?

Don't worry, said the prosecutor. Trust me. There is nothing to fear if you have done nothing wrong. I can judge right from wrong. I know a willful violation, or negligence, or a destructive act when I see it.

But does he? Does anybody?

Just Culture and the Concern with a Line

All existing definitions of just culture draw a line between acceptable and unacceptable behavior. A willful violation is not acceptable. An honest mistake is. And if what you have done is acceptable—if you have done nothing wrong—you have nothing to fear.

For example, says a proposal for air traffic control, 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."

The definition invokes the distinction between technical and normative errors that we looked at in Chapter 5 (technical errors would be those that are commensurate with a person's level of experience and training; normative ones include negligence and willful violations). But, we saw in Chapter 4, the line between technical and normative errors isn't all that clear. In fact, it is often not until after the fact, in the process of interpreting the error and considering its aftermath, that we know whether to call it one or the other. And then the choice may still be contested.

Why the Idea of a Line makes Sense

The idea of a line makes sense. If just cultures are to protect people against being persecuted for honest mistakes (when they've done nothing wrong), then some space must be reserved for mistakes that are not "honest" (in case they *have* done something wrong). Consequently, all proposals for a just culture emphasize the establishment of, and consensus around, some kind of line between legitimate and illegitimate behavior: "in a just culture, staff can differentiate between acceptable and unacceptable acts."[1] An environment of impunity, the argument holds, would neither move people to act prudently nor compel them to report errors or deviations. After all, if there is no line, "anything goes". So why report anything? This is not good for people's morale or for the credibility of management or for learning from failure.

So calls for some kind of border that separates tolerable from culpable behavior make intuitive sense, and ideas on just culture often center on its embrace and clarity: "A 'no-blame' culture is neither feasible nor desirable. Most people desire some level of accountability when a mishap occurs. In a Just Culture environment the culpability line is more clearly drawn."[2] Another argument for the line is that the public must be protected against intentional misbehavior or criminal acts, and that the application of justice is a prime vehicle for such protection.

A recent directive from the European Union (2003/42/EC) governs occurrence reporting in civil aviation. This directive has a qualification: a state must not institute legal proceedings against those who send in incident reports, apart from cases of gross negligence. But guess who decides what counts as "gross negligence?" The state, of course. Via its prosecutors or investigating magistrates.

The directive, as all guidance on just culture today, assumes that cases of "gross negligence" jump out by themselves. That "willful violations" represent a non-

problematic category, distinct from violations that are somehow not "willful." It assumes that a prosecutor or other authority can recognize—objectively, unarguably—willful violations, or negligence or destructive acts when they show up. There is not a single proposal for just cultures—indeed, not a single appeal to the need to learn from failure in aviation—that does not build in some kind of escape clause into the realm of essentially negligent, unwanted, illegitimate behavior.

Why the Idea of a Line is an Illusion

If we want to draw a line, we have to be clear about what falls on either side of it. Otherwise there is no point in a line—then the distinction between acceptable and unacceptable behavior would be one big blur. Willful violations, say many people, clearly fall on the "unacceptable" side of the line. Negligence does too. But what is negligence then? Look at a definition of negligence below:

> Negligence is 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.[3]

First, the definition is long, very long. Second, it does not capture the essential properties of "negligence," so that you can grab negligent behavior and put it on the unacceptable side of the line. Instead, the definition lays out a whole array of other questions and judgments that we should make. Rather than this definition solving the problem of what is "negligence" for you, you now have to solve a larger number of equally intractable problems instead:

- 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?"

Rather than clarifying which operational behavior is "negligent," such a characterization shows just how complex the issue is. There is an amazing array of judgment calls to be made. Just see if you, for your own work, can (objectively, unarguably) define things like "normal in the community," "a reasonable level of skill," "a prudent person," "a foresight that harm may likely result." What, really, is normal (objectively, unarguably)? Or prudent, or reasonable (objectively, unarguably)? And don't we all want to improve safety precisely because the activity we are engaged in can result in harm?

Of course, it is not that making such judgments is impossible. In fact, we probably do this quite a lot every day. It is, however, important to remember that judgment is exactly what they are. They are not objective and not unarguable. To think that there comes a clear, uncontested point at which everybody says, "yes, now the line has been crossed, this is negligence," is probably an illusion. What is "normal" versus "negligence" in a community, or "a reasonable level of skill," versus "recklessness" is infinitely negotiable. You can never really close the debate on this.[4]

What is interesting is not whether some acts are so essentially negligent as to warrant serious consequences. What matters is which processes and authorities we in society (or you in your organization) rely on to decide whether acts should be seen as negligent or not.

Recall also that all of these judgments can become significantly clouded by the effects of hindsight. With knowledge of outcome, it becomes almost impossible for us to go back and understand the world as it looked to somebody who did not yet have that knowledge of outcome. The so-called substitution test (which the definition of negligence above already applies) can be of some help. But even here, whether another reasonably prudent person would have done the same thing in the same circumstances becomes a whole different matter once you have a dead body as the outcome. Or multiple dead bodies. All research on the hindsight bias shows that it turns out to be very difficult for us not to take this into account, somehow, when we apply the substitution test.

The Social Construction of an Offense

A few months ago, my wife and I went for dinner in a neighboring city. We parked the car along the street, amongst a line of other cars. On the other side of the street, I saw a ticket machine, so I duly went over, put some cash in the machine, got my ticket and displayed it in the car windshield. When we returned from dinner, we were aghast to find a parking ticket the size of a half manila envelope proudly protruding from under one

of the wipers. I yanked it away and ripped it open. Together we pored over the fine print to figure out what on earth we had violated. Wasn't there a ticket in our car windshield? It had not expired yet, so what was going on? It took another day of decoding arcane ciphers buried in the fine print, to find the one pointing to the exact category of violation. It turned out that it had somehow ceased to be legal to park on that side of that particular piece of that street sometime during our dinner on that particular evening. I called a friend who lives in this city to get some type of explanation (the parking police only allowed us to listen to a taped recording, of course). My friend must have shaken his head in blend of disgust and amusement.

"Oh, they do this all the time in this town," he acknowledged. "If it hasn't been vandalized yet, you may find a sign the size of a pillowcase suspended somewhere in the neighborhood, announcing that parking on the left or right side of the street is not permitted from like 6 o'clock until midnight on the third Tuesday of every month except the second month of the fifth year after the rule went into effect. Or something."

I felt genuinely defeated (and yes, we paid our fine). A few weeks later, I was in this city again (no, I did not park, I no longer dared to), and indeed found one of the infamous exception statements, black letters on a yellow background, hovering over the parking bays in a sidewalk. "No parking 14–17 every third day of the second month," or some such abstruse decree.

This city, I decided, was a profile in the construction of offense. Parking somewhere was perfectly legal one moment, and absolutely illegal the next. The very same behavior—which had appeared so entirely legitimate at the beginning of the evening (there was a whole line of cars on that side of the street, after all, and I did buy my ticket—had morphed into a violation, a transgression, an offense inside the space of a dinner. The legitimacy, or culpability of an act, then, is not inherent in the act. It merely depends on where we draw the line. In this city, on one day (or one minute), the line is here. The next day or minute, it is there. Such capriciousness must be highly profitable, evidently. We were not the only car left on the wrong side of the road when the rules changed that evening. The whole line that had made our selection of a parking spot so legitimate was still there—all of them bedecked with happily fluttering tickets. The only ones who could really decrypt the pillowcase-sized signs, I thought, were the ones who created them. And they probably planned their ambushes in close synchronicity with whatever the signs declared.

What we see as a crime, then, and how much retribution we believe it deserves, is hardly a function of the behavior. It is a function of our interpretation of that behavior. And that can differ not only from day to day or minute to minute. This can slide over time, and differ per culture, per country.

I did not argue with the city in the example above. I allowed them to declare victory. They had made the rules and had evolved a finely-tuned

game of phasing them in and out as per their intentions announced on those abstruse traffic signs. They had spent more resources at figuring out how to make money off of this than I was willing to invest in learning how to beat them at their own game and avoid being fined (I will take public transport next time). Their construction of an offense got to reign supreme. Not because my parking suddenly had become "offensive" to any citizen of that city during that evening (the square and surrounding streets were darker and emptier than ever before), but because the city decided that it should be so. An offense does not exist by itself "out there," as some objective reality. We (or prosecutors, or city officials) are the ones who construct the offense—the willful violation, the negligence, the recklessness.

Decision Trees for Determining Culpability

In making those judgments, however, there is some help to be had. A number of tools, most in the form of decision trees, for determining culpability are currently in circulation. They are another form into which to put huge definitions such as the one above. And they are a great start. But that is exactly what they are: a start. They still leave the analytic heavy lifting to you. They actually leave the problem of whether an action falls on this or that side of the line for you. All they will help you do is break down the problem. But are the resulting, smaller components more manageable?

One popular decision tree is that one that appears in Reason's *Managing the risks of organizational accidents*.[5] Here are some of the questions that it presents, and some of the problems that they create:

- **Were the actions and consequences as intended?** This seems a simple enough question, but what, exactly, is intent? Philosophers and judicial experts alike still cannot really agree, so why would this suddenly be simple for you to decide? Asking the person whose actions they are may not be of much help either. Yes, the nurse would say, I intended to mix 20 mg/ml Xylocard. And as far as I know and can recount, that is exactly what I did. And no, I did not intend to poison a little baby girl. That the intended actions and consequences did not match up in this case did little to protect the nurse from prosecution or conviction. Other factors than "I did not mean to" play a role judgments of culpability.
- **Did the person knowingly violate safe operating procedures?** People in all kinds of operational worlds knowingly violate safe operating procedures all the time. Even procedures that can be shown to have been available, workable and correct (though here, of course, the question once

again pops up: workable? says who? (objectively, unarguably)). Following all applicable procedures means not getting the job done in most cases. Hindsight is great for laying out exactly which procedures were relevant (and available and workable and correct) for a particular task, even if the person doing the task would be the last in the world to think so.

- **Were there deficiencies in training or selection?** "Deficiencies" seems like an unproblematic word, but what exactly does it mean? Again, what looks like a deficiency to one seems perfectly normal or even above standard to another. The question here is: who gets to decide? Most people maintain that doctors who intentionally murder their patients are criminals, even if somebody could argue that this clearly shows problems with physician selection and proficiency checking (which, according to Reason's decision tree, would be a mitigating factor), or could raise issues of cultural standards related to end-of-life care and euthanasia in that particular country.

These are, once again, good questions to start with. But they do not solve the problem of determining culpability. They only redefine the problem. Another question in Reason's decision tree is whether there was a matter of inexperience. This is a great question too: it recalls the difference between technical and normative errors from Chapter 4. But again, whether something is judged a technical error (due to a lack of experience) or a normative error (due to a lack of discharging professional responsibility) is the outcome of the processes of interpretation and attribution that *follow* the error. It is much less determined by the behavior that led up to the error.

Psychological Research: Culpability and Control

Decision trees are not just born out of practice. Psychological research shows that even if we are not prompted, we will evaluate actions and their consequences along various criteria. It points out that the criminal culpability of an act is likely to be constructed as a function of three things:

- the amount of volitional behavior control the person had (was the act freely chosen or compelled?);
- volitional outcome control (did the actor know what was going to happen?); and
- the actor's causal control (his or her unique impact on the outcome).[6]

In this triad, "factors that establish personal control intensify blame attributions, whereas constraints on personal control potentially mitigate blame."[7]

When we apply these criteria to the case of the previous chapter, we can see how answers to these questions are difficult. They are not really answers, they are judgments. When it comes to volitional behavior control, did the nurse from the Prologue (Mara) act on purpose or by accident? (This is like asking whether consequences and actions were intended and whether they matched.) We will read into her act more control if she had behaved purposely and knowingly. Still struggling to understand her own performance, Mara had told a lower court that she might have misread the package labeling. By the time she got to the Supreme Court, however, she indicated that this was probably not the case: she mistakenly believed that 200 mg/ml was what she needed to have. This would certainly have made sense, given the prominence of the figure 200 in the medication log, and the reminder to end up with a volume of 10 ml Xylocard in total. As a result, the Supreme Court observed how,

> during the court proceedings, the ICU nurse described multiple ways how it could be that she mixed the IV drop with the wrong concentration of Xylocard. What she offered cannot therefore express what she really remembers. Rather, her accounts can be seen as attempts to find an explanation afterward. They are almost hypothetical and provide no certain conclusion as to why she did what she did.[8]

In other words, the inability to know or remember how an "error" occurred (which is quite normal, even when it comes to our own immediately past behavior), was converted into an inability to disprove volitional behavior control. To the court, whether the nurse acted knowingly or purposely could be ruled neither out nor in.

Additionally, volitional behavior control was amplified by the absence of what are called capacity and situational constraints. The Supreme Court emphasized how Mara had 25 years of experience and ample time to prepare the mixture. There was no lack of knowledge or experience (though Mara had never prepared this particular drug for an infant). She had just come on shift, there was no stress or manpower shortage during that morning.

These conditions would also have helped the nurse foresee the consequences of her actions: "Whether the nurse's negligence stemmed from misreading, miscalculating or taking the wrong package, it is obvious that she could have read the medication log more carefully, calculated more carefully or done any other double-check that would have revealed her error and its potentially fateful consequences."[9] In other words, volitional outcome control could also

be established: the nurse was experienced enough and had time enough to find out what could, or would, be the consequences of her actions.

Then to causal control. With the various truths swirling around the case, there would be ample opportunity to find other contributors to the outcome that would reduce Mara's unique impact. Yet Mara's initial *mea culpa* corrupted later appeals to additional, and necessary, actors. Recall how the pediatrician who ended up giving the infant its overdose, for example, successfully asserted that his administrations would not have had the fatal effect they did if the drug solution had been correct—which he could only believe it was. Mitigating circumstances related to long-eroded practices in drug management in the hospital were dismissed as playing no serious role in exerting causal force on the outcome: the court admitted that "there were serious shortcomings in routines and procedures at the regional hospital, this did not remove the nurse's own responsibility for checking that her mixture was correct."[10] But however lousy the workplace, its organization or traditions, that still did not relieve an individual actor of the responsibility to not err. At least not in how the Supreme Court drew the line.

It is Not Where to Draw the Line, but Who Draws it

What matters in creating a just culture is not to come up with a definition that leaves a number of supposedly self-evident labels ("willful violation," "negligence," or people that are not "prudent," or "normal," or "reasonably skilled") on the wrong side of the law and the rest on the right side. For those labels are far from self-evident. Almost any act can be constructed into willful disregard or negligence, if only that construction comes in the right rhetoric, from the legitimated authority. Drawing a line does not solve any problem, it simply displaces it. *What matters instead is to consider very carefully, and preferably make structural arrangements about, who gets to draw the line.* This gets us to the next chapter:

- Who has the authority to draw the line?
- Who in your organization, and in your society, has the right language, and the official legitimacy to say that the line has been crossed?
- Do these people rely on a claim to a "view from nowhere," an objective, unarguable, neutral point of view from which they can separate right from wrong?

Rather than a prosecutor saying, "you have nothing to fear if you have done nothing wrong," a more accurate portrayal would be "if I decide you have done nothing wrong, you have nothing to fear." Which in itself might mean there could be something to fear.

This is why a just culture should not give anybody the illusion that it is simply about drawing a line. Instead, it should give people clarity about who draws the line, and what rules, values, traditions, language and legitimacy this person uses. Whether this person is a prosecutor or a manager, or even a committee of peers, is not really the point (though all have different stakes and biases in drawing their lines). The point of a just culture is to get clarity and agreement about it.

Notes

1 Ferguson, J. and Fakelmann, R. (2005). The culture factor. *Frontiers of Health Services Management*, 22(1), 33–40.

2 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), viii.

3 Ibid., 6

4 Of course, we have evolved as civilization in large part by outlawing certain practices. Acts such as rape are beyond any discussion: we see them as universally reprehensible. If you were to argue that rape is a crime only because we label it one, you would be deemed so ridiculous that nobody would take you seriously any longer.

Yet beyond this basic and sustained bedrock of what a civilization considers criminal, it is easy to show that the goal posts for what *else* counts as a crime shift with time and with culture. The "crimes" I deal with in this book are a good example. They are acts in the course of doing normal work, "committed" by professionals—nurses, doctors, pilots, air traffic controllers, policemen, managers. These people have no motive to kill or maim or otherwise hurt, though we as society have given them the means and their work offers them plenty of opportunities. Somehow, we, as society, or as employing organization, manage to convert acts with bad outcomes into culpable or even criminal behavior that we believe should be punished as such. This, though, hinges not on the essence of the acts, but on our interpretation of them.

5 Reason, J.T. (1997). *Managing the risks of organizational accidents*. Aldershot, UK: Ashgate Publishing Co., 209.

6 Alicke, M.D. (2000). Culpable control and the psychology of blame. *Psychological Bulletin*, 126(4), 556–74.

7 Ibid., 557.

8 Swedish Supreme Court verdict B 2328-05, 19 April 2006, 3–4.

9 Ibid., 4–5.

10 Ibid., 5.

8 Without Prosecutors, There Would be No Crime

It is great if you try to define "negligence," or other categories of offense (willful violation, destructive act, deliberate rule-breaking). You can even try to reason your way to a definition by using one of the culpability decision trees that are in circulation. But these definitions and decision trees are really only tools to get a discussion started about whether something can be seen by a community of peers or judges as negligence (or as a normative error as opposed to a technical one). Definitions will not supply you with an answer: you will still have to do the hard analytic work.

As concluded in the previous chapter: even more important for the outcome of these discussions is the question of who eventually gets to decide. Who in your society, or in your organization, has the authority and the legitimacy to claim that behavior falls on the other side of the line, the wrong side? If you do not think critically about this question, or if you have not made any structural arrangements around it, you could find that what counts as negligence in your organization or country may have very little to do with who knows best. Or with justice. Instead, it may have more to do with power.

Why We Do Not Ask the People Themselves

We do not normally ask professionals themselves whether they believe that their behavior "crossed the line." But they were there, perhaps they know more about their own intentions than we can ever hope to gather. Perhaps they are in a better position to say whether substance abuse played a role, or whether the procedures that they violated were workable or correct or available. And whether they knowingly violated them or not. Yet we don't rely on insiders to give us the truth. After all:

- we suspect that those people are too biased for that;
- we reckon they may try to put themselves in the most positive light possible;
- we will see their account as one-sided, distorted, skewed, partial—as a skirting of accountability rather than embracing it.

To get a truthful account of what happened, we do not typically listen to the people who were there, even if we do sometimes give them a voice (like we do in a trial, for example).

The View from Nowhere

So who gets to decide instead? Is there a perspective that is not biased? A perspective that is impartial, neutral? We often turn to our legal systems for this. We believe that it will furnish us with the truth—unvarnished and unbiased. This expectation is not new. The US Supreme Court put it most bluntly back in 1966: "The basic purpose of a trial is the determination of the truth."[1]

We expect a court to apply reason, and objectivity, and come up with the real story, with the truth. And then hand out consequences for those responsible for the outcome. From a distance, it may well come across this way. A disinterested party takes an even-handed look at the case. The appropriate person gets to be held accountable. Appropriate consequences are meted out. Truth and justice are served.

Is the Legal System Impartial?

The legal system certainly goes to great pains to show as if it is impartial. Many of the trappings of the justice system are designed to impart an image of rationality, of consideration, of objectivity and impartiality:

- think, for starters, of Lady Justitia's blindfold—the very profile of neutrality.
- the pace of judicial proceedings is measured, the tone solemn;
- the rules of proceedings are tight and tightly controlled;
- the uniforms and settings and language invoke a kind of otherworldliness, of not exactly belonging to the daily, messy hubbub of the real world out there;

- even the buildings are often designed so as to be set apart from the rest of the world: just imagine your own local courthouse. It is probably separated from the sidewalk by gates, lawns, forecourts, high steps;
- the judges are often behind enormous doors, seated at a distance from other people, on podia, behind solid desks, under high ceilings.

Does this symbolism and imagery, this elevation and separation—meant to offer the assurance of rationality and impartiality—really give a court a better, more neutral view of the truth?

There is No View from Nowhere

Telling the story from an objective angle is impossible, no matter how objective, disinterested, unbiased you may think you are, or how neutral we make Lady Justitia look with her blindfold. Just ask yourself, if you were to take an objective look at the world, from where would you look? An objective view is a "view from nowhere."[2] And there is no view from nowhere, as there would be nobody to form the view.

So no view can be neutral, or objective. Because no view can be taken from nowhere. This means that all views somehow have values and interests and stakes wrapped into them. Of course, we can try to control the influences of those values and interests. And the legal system has great traditions and symbols and rituals to do just that. But in the end, nobody can discover or generate a value-free truth. Judges are stakeholders in the healthcare system too. They may be consumers of it, after all. And they have a larger role: helping maintain stability, and confidence in a society's institutions:

In the Xylocard case from the Prologue, the Supreme Court admitted that its agenda was in part to reassure any disquiet about the safety of the healthcare system: "Concern for patients' safety and their confidence in the healthcare system, demand that the nurse's actions be seen as so clumsy that they imply culpable negligence. She therefore cannot avoid being responsible for manslaughter."[3] *The maintenance of "confidence in the healthcare system" demanded the construction of a version where one anti-hero could be singled out to receive the blame, to bear the explanatory and moral weight of the infant's death.*

For a court to find an offense, and to call it criminal, is not the product of blind arbitration. It is not the clearest view on things from an objective stance. It is not the cleanest, truest rendering of a story. Instead, it is the negotiated outcome of a social process. And as such it is not much different (if at all) from

any other social process, in how it is influenced by history, tradition, institutions, personal interactions, hopes, fears, desires.

To Get to the "Truth" You Need Multiple Stories

Recall from the prologue how the justices were struggling to divine what the medicine cartons were all about, what the strange names and figures meant. And recall from Chapter 2 how the pilot tried to show to the court that configuring the airplane for approach took more time than had been available—and nobody cared. While the professionals on trial doggedly searched for ways to get "the truth" out, it never would.

Multiple versions competed and contradicted each other, but many of them seemed equally valid. All illuminated different aspects of the case. In the Xylocard trial, the pediatrician had a point: his repeated bolus doses of Xylocard into the baby could not be judged in light of the fact that the baby was already getting ten times the prescribed dose through her drip. He could not have known that, after all. The nurses had a point too: ordering bolus dose after bolus dose, with only worsening effects, and without ever having established a diagnosis for the baby's condition, did not make perfect sense. Settling for only one version amounts to an injustice to the complexity of the adverse event that the nurse was on trial for.

Similarly, the captain had a point: it was the airline, its image, production pressures and routine dispensations to as yet unlisted doctors and unqualified copilots that helped box him in. But the other side had a point too: why had this pilot not voluntarily contributed to learning and improvement after the incident?

This implies that forcing one story onto other people as if it were the true and only one (like the justice system sometimes does) is actually quite unjust. A just culture always takes multiple stories into account, because:

- telling the story from one angle necessarily excludes aspects from other angles;
- no single account can claim that it, and it alone, depicts the world as it is;
- innumerable stories are possible, and, if you want to be "just," or approximate the "truth," a number are even necessary;
- also, if you want to explore as many opportunities for safety improvement as possible, you probably want to listen to as many stories or angles as possible. The world is complex—live with it. And learn from it what you can.

A colleague in healthcare told me how he believed that some acts are objective, self-evident, or even unarguably criminal—murder of a patient for example, or adverse events

involving substance abuse by the provider (a doctor being drunk on duty), or deliberately unsafe acts. He told me the story of some nurses who substituted diuretic tablets for pain relieving tablets as a prank to make patients demand urine bottles from the night staff. These were egregious acts, he said. Criminal acts. That could only be dealt with through discipline or other legal forms.

I am in no position to say that these things are not crimes. But what I find interesting is how we come to give the acts meanings as crimes, committed by these individuals at that moment. Seeing these acts as criminal can rule out or obscure a host of other factors that may once again trigger other people to behave similarly "criminal." When it comes to doctors deliberately murdering patients, for example, this raises a host of questions about access control to the profession (is there a psychiatric evaluation to become a doctor? To become an airline pilot there is. Are there regular proficiency checks for doctors practicing on their own? For pilots there are). Drunk or stoned doctors raise questions about working hours (36-hour shifts, 80+ hour weeks) and the effects on their personal lives. Playing a prank on the night staff at the cost of patients raises questions about organizational staff disputes that are left unaddressed, and about the ethical awareness of the staff involved.

Yes, through the eyes of a lawyer or prosecutor, these acts may well look like crimes. The language of "crimes" is one that would seem to fit the acts above quite well. But that is not necessarily the only language in which we can talk about things like the ones above. Or do something about them. Even these "crimes" can be constructed as different things:

- as societal or professional trade-offs (we make our doctors work long hours in part because healthcare is hugely expensive already, and we trust them to remain healthy, alert, and self-responsible once we license them);
- as managerial issues (simmering inter-departmental or cross-shift conflicts are not resolved early enough through higher-level intervention);
- as pedagogical ones (ethical training for staff).

Again, I am not saying (because I can't) that one interpretation is better or more "right" than another. I am only saying that different interpretations are possible. And all interpretations have a logical repertoire of action appended to them:

- a crime gets punished away;
- a pedagogical problem gets taught away;
- an organizational problem gets managed away.

See only one interpretation and you may miss other important possibilities for progress on safety.

There is Never One "True" Story

Perhaps we should give up trying to dig out the "true" account of a failure altogether. As soon as you make such a claim, somebody will come around and point to "untrue" elements in your story. Or missing parts. Or misconstrued parts, or mischaracterized ones, or underemphasized parts. Trying to tell a story from the perspective "from nowhere" is impossible. As soon as anybody starts describing what happened and what went right or wrong in that story, that person is already using his or her own language, thereby inevitably importing his or her own values, interests, background, culture, traditions, judgments. The courts may have laid a claim on an objective account of a professional's actions. But from the professional's perspective (and that of almost all their colleagues) that account was incomplete, unfair, biased, partial.

Remember, in trying to build a just culture, what matters is not getting to a true or objective account of what happened. That is not where the criterion for success lies. And however hard you try, you won't, anyway. Instead, to achieve a just culture, you need to get to an account of failure that can do two things at the same time:

- satisfy demands for accountability;
- contribute to learning and improvement.

Notes

1 Laudan, L. (2006). *Truth, error and criminal law: An essay in legal epistemology*. Cambridge, UK: Cambridge University Press, 2.
2 Nagel, T. (1992). *The view from nowhere.* Oxford, UK: Oxford University Press.
3 Swedish Supreme Court verdict B 2328-05, 19 April 2006, 5–6.

9 Are Judicial Proceedings Bad for Safety?

The great psychologist William James (1842–1910) said that if you want to study religion, you need to study the most religious person at his most religious moment. So let's do that here for a moment. What happens to a just culture when the judicial system gets involved? What happens when the pursuit of justice goes all the way—when human error is converted into a crime by the legal system, which then metes out supposedly appropriate consequences?

Paradoxically, when the legal system gets involved, things seem to get neither more just, nor safer. In fact, with the evidence in hand, you could argue that the opposite happens. This also holds important lessons for "less religious men during less religious moments." Lessons for those who want to pursue a just culture in their own department, their own organization, their own industry or country—without quite going all the way and calling an error a crime.

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 possible charges of manslaughter. Pilots in New Zealand sued to block the police use of the CVR, saying recorders should only be used for safety and educational purposes. Prosecutors prevailed and regained access to the CVR, but pilots soon began disabling CVRs on their flights. Officials have crafted a plan that would permit police use of CVRs in future cases, provided New Zealand's High Court deemed it necessary. In any case, the plan calls for the CVR information not to be made public.[1]

Is Criminalizing Human Error Bad for Safety?

The sheer threat of judicial involvement is enough to make people think twice about coming forward with information about an incident that they were involved in.[2]

Just imagine how the colleagues of the nurse in the Prologue may have felt about this. The nurse, after all, stepped forward voluntarily with her view on the death of the infant. As long as there is fear that information provided in good faith can end up being used by a legal system, practitioners are not likely to engage in open reporting.

Many admit that they will only file a report when there is the chance that other parties will disclose the incident (for example, an air traffic controller may think that a pilot will report a close call if s/he does not do it; a nurse may feel the same way with respect to a resident physician present during the same event, or vice versa), which would make the event known in any case.

This puts practitioners in a catch-22: either report facts and risk being persecuted for them, or not report facts and risk being persecuted for not reporting them (if they do end up coming out along a different route). Many seem to place their bet on the latter: rather not report and cross your fingers that nobody else will find out either.

Practitioners in many industries, the world over, are anxious of inappropriate involvement of judicial authorities in safety investigations that, according to them, have nothing to do with unlawful actions, misbehavior, gross negligence or violations.[3]

They are not alone in their anxiety. Operational organizations, and even regulatory authorities (which fall under departments or ministries other than justice—for example, transportation) are concerned that their safety efforts, such as encouraging incident reporting, are undermined.[4] But what is it, exactly, that people are afraid of? Judicial involvement can consist of:

- **The participation of law enforcement officials in investigations**. There are countries in the developed world where the police are witnesses or even participants 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 any practitioner's willingness to cooperate in the accident probe.
- **Judicial authorities stopping an investigation altogether** 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 in many countries say that this severely retards their efforts to find out what went wrong and what to do to prevent recurrence.[5]

- **Using a formal accident report in a court case**. Even though using such reports as evidence in court is proscribed through various arrangements, these routinely get overridden or circumvented. And, in any case, nobody can prevent a prosecutor or judge from reading a publicly-available accident report.
- **Getting access to safety-related data** (for example, 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.

The safety manager for one organization told me how the person involved in an incident flatly refused that the incident be used for recurrent training, precisely because of the perceived risk of persecution. Even assurances of complete anonymity and de-identification of incident data were not enough to sway the practitioner. While understandable, this denied colleagues an opportunity to engage in a meaningful lesson from their own operation. Normal, structural processes of organizational learning are thus eviscerated; frustrated by the mere possibility of judicial proceedings against individual people.

In all of these ways, judicial involvement (or the threat of it) can engender a climate of fear and silence. In such a climate it can be difficult, if not impossible, to get access to information that may be critical to finding out what went wrong, or what to do to not have it happen again. Here is another example of what that can lead to.

A prosecutor responsible for aviation decided to launch what she termed a "test case."[6] *On take-off the crew of a large passenger jet had suddenly seen another aircraft, being towed by a truck, cross the runway in front of them. Immediately they aborted their take-off and stopped before reaching the intersection. Nobody was hurt. The air traffic control organization, as well as the country's independent transportation safety board, both launched investigations and arrived at pretty much the same conclusions. After unclear radio transmissions to the tow truck driver, an assistant controller had passed her*

interpretation of the tow's position to the trainee controller responsible for the runway. The assistant controller did not have a screen that could show ground-radar pictures. The trainee controller did, and took the position of the tow at the edge of the runway to mean that the crossing had been completed. Buttons on a newly-added panel in the tower for controlling lighted stop-bars at runway intersections proved ambiguous, but at the time all looked in order, and he cleared the other jet for take-off. Meanwhile, the coach of the trainee controller was performing supervisory duties in the tower. The account, in other words, was straightforward in its complexity: mixing elements of interface design, production pressure, weather conditions, handovers, short-staffing, screen layouts, and communication and teamwork—among many other factors. This, the safety community knows, is what organizational incidents and accidents are made of. Many factors, all necessary and only jointly sufficient, are required to push a system over the edge of breakdown. And all of those factors are connected to normal people doing normal work in what seems a perfectly normal organization. These factors, then, are also the stuff of which recommendations for improvement are made. And they were, also in this case. The Air Traffic Control organization issued no fewer than 23 recommendations, all of them aimed at rectifying systemic arrangements in for example design, layout, staffing, coaching, communications and handovers. The independent safety investigation board issued nine, quite similar, recommendations. This, as far as the community was (and is) concerned, is how the incident cycle was supposed to work. A free lesson, in which nobody got hurt, was milked for its maximum improvement potential. The people involved had felt free to disclose their accounts of what had happened and why. And they had felt empowered to help find ways to improve their system. Which they then did, for everybody's benefit.

But two years after the incident, the aviation prosecutor of the country decided to formally charge the coach/supervisor, the trainee and the assistant controller with "the provision of air traffic control in a dangerous manner, or in a manner that could be dangerous, to persons or properties" (the country's law actually contains such provisions). Each of the three controllers was offered a settlement: they could either pay a fine or face further prosecution. Had they paid the fine, the prosecutor would have won her "test" and the door for future prosecutions would have stood wide open. The controllers collectively balked. A first criminal court case was held a year and a half after the incident. The judge ruled that the assistant controller was not guilty, but that both the trainee and the coach/supervisor were. They were sentenced to a fine of about US$450 or 20 days in jail. The trainee and the coach/supervisor decided to appeal the decision, and the prosecutor in turn appealed against the assistant controller's acquittal.

More than a year later, the case appeared before a higher court. As part of the proceedings, the judges, prosecutor and their legal coterie were shown the airport's tower (the "scene of the crime"), to get a first-hand look at the place where safety-critical work was created. It was to no avail. The court found all three suspects guilty of their crime. It

did not, however, impose a sentence. No fine, no jail time, no probation. After all, none of the suspects had criminal records and indeed: the air traffic control tower had had its share of design and organizational problems. The court had found legal wiggle room by treating the case as an infringement of the law, as opposed to an offense. It was as if they were proving themselves right and wrong at the same time. The court was wrong to bring and prosecute the case because there was no offense, but did not waste tax money after all because they managed to find an infringement. This was actually a no-brainer, as an infringement means "guilt in the sense that blame is supposed to be present and does not need to be proven." The only admissible defense against this is being devoid of all blame. This would work only if the air traffic controller was off-duty and therefore not in the tower to begin with. It was a celebration of perverse formalism (to use judge Thomas's words): a decorous nod to the prosecutor who had gone out to test the waters, and a measly but still unsettling warning to air traffic controllers and other professionals that they were not above the law. And it stopped all appeals: appealing an infringement is not possible as there is no conviction of an offense, and no punishment. The real punishment, however, had already been meted out. It was suffered by the safety efforts launched earlier by the air traffic control organization, particularly its incident reporting system. Over the two years that the legal proceedings dragged on, incident reports submitted by controllers dropped by 50 per cent.

Many people, especially from the various professional communities, are duly concerned. The Secretary-General of the Worldwide Association of Air Traffic Control Providers warned of "grave and undesirable consequences for safety" when judicial systems get involved.[7]

But Isn't There Anything Positive about Involving the Legal System?

Some in the legal community see the criminalization of error as a long-overdue judicial colonization of rogue areas of professional practice. It is, they say, a clamp-down on closed, self-serving and mutually protective professional "brotherhoods" that somehow assert a special status and hold themselves to be above the law. Law is seen as authoritative, neutral and fair, and it should reign equitably over everybody (hence Lady Justitia's blindfold): there should be no exception or discrimination either way.[8]

An increasingly vocal consumer movement, wanting greater control over safety in a variety of products and services, has been seen as sponsoring this view.[9] Pilots, doctors, air traffic controllers—already adequately compensated monetarily for the responsibility bestowed upon them—should be treated like

everybody else. If they commit a culpable act, they should be held accountable for it. Exceptionalism is anti-democratic.

There is no evidence, however, that the original purposes of a judicial system (such as prevention, retribution, or rehabilitation—not to mention getting a "true" account of what happened or actually serving "justice") are furthered by criminalizing human error.

- The idea that a charged or convicted practitioner will serve as an example to scare others into behaving more prudently is probably misguided: instead, practitioners will become more careful only in not disclosing what they have done.
- The rehabilitative purpose of justice is not applicable either, as there is usually little or nothing to rehabilitate in a pilot or a nurse or air traffic controller who was basically just doing her or his job.
- Also, correctional systems are not equipped to rehabilitate the kind of professional behaviors (mixing medicines, clearing an aircraft for take-off) for which people were convicted.

Not only is the criminalization of human error by justice systems a possible misuse of tax money—money that could be spent in better ways to improve safety—it can actually end up hurting the interests of the society that the justice system is supposed to serve. Indeed, other ways of preventing recurrence can be much more effective:

Alan Merry dryly remarked: "The addition of anti-hypoxic devices to anesthetic machines and the widespread adoption of pulse oximetry have been much more effective in reducing accidents in relation to the administration of adequate concentrations of oxygen to anesthetized patients than has the conviction for manslaughter of an anesthetist who omitted to give oxygen to a child in 1982."[10]

If you want a people in a system to account for their mistakes in ways that can help the system learn and improve, then charging and convicting a practitioner is unlikely to do that.

Judicial Proceedings and Justice

But wait, you may say, doesn't the legal system help society understand what went wrong and why, and what we can do about it? The chances that a legal system

will tease out a meaningful and just account of what happened are actually remote. It is not its charter and even if it were, it is not particularly good at it.

Go back again to the nurse's case from the Prologue: heaping all responsibility for the baby's death on her shoulders made no historical sense whatsoever, and it was really hard to see this as fair or just. Lots of other people had been involved, and she had not even administered the drug in question. The judicial proceedings in the aftermath of the baby's death, through sheer design and rules of relevancy, played down or ignored these other contributions. It ended up with an account of a complex system failure that contradicted decades of research into how such accidents actually happen.

The potential for bad outcomes lies baked into the very activity that we ask practitioners to do for us. The criminal trial of the airline captain from Chapter 2, for example, found him guilty of "endangering his passengers" while flying an approach to a runway in fog. "I do that every day I fly," a colleague pilot had responded. "That's aviation."[11]

Pilots and nurses and doctors and similar practitioners endanger the lives of other people every day as a matter of course. How something in those activities slides from normal to culpable, then, is a hugely difficult assessment, for which a judicial system often lacks the data, the education and the expertise. The decision whether to prosecute a practitioner, then, can turn out to be quite haphazard, and the practitioner on the receiving end will likely see this as quite unjust.

In the same year that Mara, the nurse from the Prologue, was first charged, more than 300 severe medication errors were reported to the country's health authority. In another study, a full 89 per cent of responding anesthetists reported having made drug administration errors at some stage in their careers. Most had done so more than once, and 12.5 per cent reported having actually harmed patients in this way.[12] *So why the nurse in the Xylocard case, and not one of scores of other medical practitioners who fall victim to similar medication misadventures—all the time, all over the place?*

It is the whole point of legal proceedings to narrow in on a few acts by a few individuals or even a single individual. By its very nature, however, this clashes with what we know about accident causation in complex, dynamic systems today. Many factors, all necessary and only jointly sufficient, are needed to push a basically safe system over the edge into breakdown. Single acts by single culprits are neither necessary nor sufficient. This, logically, does not make judicial proceedings about complex events "just."

The accounts of an accident that a legal system produces can be so limited in many ways because of the way it conducts its business—among other things through:

- the way judicial proceedings rationalize the search for and consideration of evidence;
- how they closely script turn-taking in speech and form of expression;
- how they limit what is "relevant," are institutionally constrained in their deferral to domain expertise;
- how they necessarily exclude the notion of an "accident" or "human error" because there are typically no such legal concepts.

This is not to deny the relevance or even authority of a legal tradition, at least not on principle. It is, rather, to see it as that: one tradition, one perspective on a case of failure. One way for which prosecutors and judges have received the power to enforce it on others, one language for describing and explaining an event, relative to a multitude of other possibilities.

Another consequence of the accountability demanded by legal systems is that it is easily perceived as illegitimate, intrusive and ignorant. If you are held "accountable" by somebody who really does not understand the first thing about what it means to be a professional in a particular setting (a ward, a cockpit, a control room, a police beat), then you will likely see their calls for accountability as unfair, coarse and uninformed. Indeed, as unjust. Research shows that this results in less disclosure and a polarization of positions, rather than an openness and willingness to learn for the common good.[13]

Judicial Proceedings and Safety

If judicial processes in the wake of accidents can be bad for justice, what about their effects on safety? Here is a summary of some of the adverse effects:

- Judicial proceedings after an incident can **make that people stop reporting incidents**. The air traffic control provider in the example in this chapter reported a 50 per cent drop in incidents reported in the year following criminal prosecution of controllers involved in a runway incursion incident. Interestingly, the criminal prosecution does not even have to be started, let alone lead to a conviction: the threat of criminal prosecution can make people hesitant about coming forward with safety information.

- Judicial proceedings, or their possibility, can **create a climate of fear**, making people reluctant to share information. It can hamper an organization's potential to learn from its own incidents. People may even begin to tamper with safety recording devices, switching them off.
- **Judicial proceedings can interfere with regulatory work**. Some regulators, for example, have become more careful in using language such as "deviation" in their inspection reports. If it is a "deviation" that a regulator takes notice of, it is very likely a deviation from some regulation. And regulations have their basis in law. A "deviation" can then easily become a breaking of the law—a crime, rendering sources at the operator silent as a result). Regulators can become much less direct about what is wrong and needs to be done about it.
- Judicial proceedings can help **stigmatize an incident as something shameful**. Criminalizing an incident can send the message to everybody in the operational community that incidents are something professionally embarrassing, something to be avoided, and if that is not possible, to be denied, muffled, hidden.
- The stress and isolation that practitioners can feel when subject to legal charges or a trial typically **makes them perform less well in their jobs**. And investing cognitive effort in considering how actions can get you in legal trouble detracts attention from performing quality work.[14]
- Finally, judicial proceedings in the aftermath of an accident **can impede investigatory access to information** sources, as people may become less willing to cooperate in the accident probe.[15] This could make it more difficult for investigators to get valuable information, particularly when judicial proceedings are launched at the same time as the safety investigation. There is, however, a suggestion (at least from one organization) that criminal prosecution in the aftermath of an accident does not dampen people's willingness to report regarding incidents. This could point to a subtlety in how employees calibrate their defensive posture: accidents, and becoming criminally liable for one, are somehow judged to be qualitatively different from liability for incidents.

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.[16]

Tort Liability

So far, I have basically talked about criminal legal proceedings (and will do so again in the next chapter). This has a reason: there may be a trend towards criminalizing human error. So it is useful to assess whether or not that is a reasonable way to achieve the dual goals of a just culture: explanations of failure that satisfy calls for accountability and offer opportunities for change and progress on safety. So far, the evidence suggests that criminal law does not contribute to the achievement of these goals.

But another kind, called tort (or civil) liability, has been in use to deal with human error for quite a while, particularly in healthcare. Tort is a legal term that means a civil (as opposed to a criminal) wrong. To be liable under tort law, you do not have to have a formal contract with the other party, as it covers duties for all citizens under a particular jurisdiction (which is true of criminal law too of course). If a court concludes that an action is a crime, then the state can impose punishment (such as imprisonment or fines). If an action is a tort, however, the consequence is usually the payment of damages to the party injured or disadvantaged by the action.[17]

Tort law too, has come under criticism for neither contributing to safety nor to justice when it comes to human error:[18 19]

- Tort law is a very irregular mechanism to compensate victims of error. According to one study, only one in seven patients who can be said to have been "negligently" harmed ever gain access to the malpractice system. Those who are older and poorer are disproportionately excluded from access.[20]
- Tort law also delivers compensation inefficiently. Administrative costs account
- for more than 50 per cent of total system costs, and a successful plaintiff recoups only one dollar of every $2.50 spent in legal and processing costs.[21]
- Malpractice claims offer only the chance of financial compensation. They do not have as a goal to encourage corrective action or safety improvements, they do not help people get an apology or any other expression of regret or concern.
- Tort law includes practices such as pre-trial discovery and all kinds of rules that govern disclosure and the protection of information. And, of course, a trial is in itself adversarial, lining up people against each other in competitive positions. The upshot is that tort law makes it *more* difficult to

get facts out, rather than helping people find out what went wrong and what to do about it so it does not happen again.

- Also, the adversarial process is based on the idea that the presentation of relentless, one-sided arguments to an impartial judge or jury is the best way to get to the "truth." The previous two chapters did acknowledge that multiple stories are necessary if we want to learn anything of value about complex events, but that does not mean only two, necessarily opposing stories, where what is true in one is almost automatically false in the other. Those who tell these stories are often not the ones who know them best (the physician or the patient), but rather their lawyers, who will have to abstract away from the details and cast things in a legal language that can get far removed from the actual meaning of people's actions and intentions at the time.

  As with criminal trials (which do not deter people from making mistakes but *do* deter people from talking about their mistakes), tort law promotes defensive practice rather than high quality care.[22]

Summing Up the Evidence

The cases of human error that have gone to trial so far suggest that legal proceedings—tort or ciminal—in the wake of incidents or accidents could be bad for safety, and may not help in creating a just culture.

Many inside and outside professional circles see a trend towards criminalization of human error as troublesome. If justice exists to serve society, then prosecuting human error may work against that very principle. The long-term consequence for society of turning errors into crimes or culpable malpractice could be less safe systems. Criminalizing error, or pursuing tort claims, can:

- erode independent safety investigations;
- promote fear rather than mindfulness in people practicing safety-critical work;
- make organizations more careful in creating a paper trail, not more careful in doing their work;
- make work of safety regulators more difficult by stifling primary sources of information and having to package regulatory findings in a language that does not attract prosecutorial attention;
- waste money on legal processes that do not really end up contributing to justice or to safety;
- ignore needs of victims other than mere financial ones, such as apology or the recognition of having been harmed;

- discourage truth-telling and instead cultivate professional secrecy, evasion, and self-protection.

If they become the main purveyor of accountability, legal systems could help create a climate in which freely telling accounts of what happened (and what to do about it) becomes difficult. There is a risk of a vicious cycle. As I warned before, we may end up turning increasingly to the legal system because the legal system has increasingly created a climate in which telling each other accounts openly is less and less possible. By taking over the dispensing of accountability, legal systems may slowly strangle it.

Notes

1 McKenna, J.T. (1999). Criminal and safety probes at odds. *Aviation Week and Space Technology*, 47–8.
2 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*. Brussels: Eurocontrol.
3 See, for example, Ter Kuile, A. (2004). Safety versus justice. *Canso News*, 18, 1–2.
4 Eurocontrol Performance Review Commission (2006), op. cit.
5 See, for example, North, D.M. (2002, February 4). Oil and water, cats and dogs. *Aviation Week and Space Technology*, 70.
6 This case is described in Ruitenberg, B. (2002). Court case against Dutch controllers. *The Controller*, 4(41), 22–5.
7 Ter Kuile, op. cit.
8 Bittle, S. and Snider, L. (2006). From manslaughter to preventable accident: Shaping corporate criminal liability. *Law and Policy*, 28(4), 470–96.
9 Merry, A.F. and McCall Smith, A. (2001). *Errors, medicine and the law*. Cambridge, UK: Cambridge University Press, 127.
10 Ibid., 51.
11 Wilkinson, S. (1994). The November Oscar incident. *Air and Space*, February–March, 82.
12 Merry, A.F. and Peck, D.J. (1995). Anaesthetists, errors in drug administration and the law. *New Zealand Medical Journal*, 108, 185–7.
13 Lerner, J.S. and Tetlock, P.E. (1999). Accounting for the effects of accountability. *Psychological Bulletin*, 125, 255–75; and see also Thomas, E.W. (2005). *The judicial process: Realism, pragmatism, practical reasoning and principles.* Cambridge, UK: Cambridge University Press.
14 Lerner and Tetlock, op. cit.
15 See, for example, North, D.M. (2002). Oil and water, cats and dogs. *Aviation Week and Space Technology*, 70.
16 McKenna, op. cit.
17 Tort law is applicable particularly in legal systems that stem from English common law, but even Napoleonic and other legal systems have ways of compensating victims through civil legal procedures. The technical variations are of course both subtle and many. Also, there can be overlap between crime and tort in some countries: the same action can be prosecuted

as a crime (possibly resulting in the state imposing penalties) *and* as a civil tort (possibly resulting in damages to the victim).

18 Dauer, E.A. (2004). Ethical misfits: Mediation and medical malpractice litigation. In V.A. Sharpe (ed.), *Accountability: Patient safety and policy reform*. Washington, DC: Georgetown University Press, 185–201.

19 From Sharpe, V.A. (2003). Promoting patient safety: An ethical basis for policy deliberation. *Hastings Center Report Special Supplement*, 33(5), S1–S20.

20 See, for example, Sloan, F.A. and Hsieh, C.R. (1990). Variability in medical malpractice payments: Is the compensation fair?" *Law and Society Review*, 24, 997–1039, and Burstin, H.R., Johnson, W.G., Lipsitz, S.R., and Brennan, T.A. (1993). Do the poor sue more? A case-control study of malpractice claims and socioeconomic status. *Journal of the American Medical Association*, 13, 1697–701.

21 Weiler, P. et al. (1993). *A measure of malpractice*. Cambridge, MA: Harvard University Press.

22 Kessler, D. and McClellan, M. (1996). Do doctors practice defensive medicine? *Quarterly Journal of Economics*, 111, 353–90.

This page intentionally left blank

10 Stakeholders in the Legal Pursuit of Justice

There are many different ways in which we can hold people accountable. And many of these forms of accountability do not produce the legal system's rivalry, the competition for "truth," the defensiveness that can be destructive for trust, and corrosive for attempts to retain or improve safety.

While a legal system may produce one form of accountability, we should not believe that it delivers anywhere near "full" accountability, or even a meaningful form. The accounts that get wrung out of practitioners in a judicial process are not necessarily completely forthcoming, full, unbiased, open, or honest. Practitioners on trial have reason to be defensive, adversarial, and ultimately limited in their disclosure.

So who are the parties that make it so?[1] Whose interests are at stake, and why do they collide? How can a society keep supporting a way of dealing with risk that perhaps works against its own long-term interests?

The Victims

Those who should benefit most from any legal action in the wake of an incident or accident are those who were affected most by it: the victims. Take the parents of the infant in the case from the Epilogue, or the family of a passenger in an airliner crash.

Most countries give victims the role of factual witness in a trial (in addition, a so-called witness-impact statement has become popular in the US, allowing witnesses to reveal to juries even the emotional and other toll taken by the accident or "crime." There are suggestions that this can unfairly sway juries against the suspect).

Witness testimony can highlight all kinds of angles on the case, from the emotional and practical consequences of the loss they suffered to their own observations of the behavior of, for example, doctors or other caregivers before, during and after the occurrence. This can bring aspects to light that would not

otherwise have been known, as victims may sometimes have had a close-up view of the unfolding incident.

This in turn can put judges (or juries) in a tricky position. If court cases are conducted on the presumption of innocence, and scrupulously avoid the word "perpetrator," instead using a more tentative "defendant" or "suspect," then what does this do to the status of the victim? And what does the victim do to the status of the defendant or suspect?

Without a proven crime, there can be no victim of a crime. A strong validation of a victim's account in court, then, could perhaps make people lose sight of the difference between suspect and guilty. It can become difficult to remain unbiased and retain the presumption of innocence.

But having victim's testimony in court serves other important functions that may sometimes outweigh these risks. Giving testimony in court offers victims an opportunity to get their voices heard. This is very important. They often want to tell their stories, or a part at least. A trial may be the first time that anybody bothers to seriously listen to the victim.

Such blocked yearning may be one of the grounds for going to court in the first place. Michael Rowe, a sociologist at Yale, captured this in an essay about his son's death after two failed liver transplants: "Many of those who sue doctors … have no place else to hand their grief when that grief—and seemingly their loved one's life—is being ignored, even declared, in the space left by silence, a thing of no value."[2]

Yet not everybody turns to courts to get their voices heard. Victims can turn to the media instead. These can provide an outlet for what could end up a rather one-sided account of the incident or accident.

What matters for an organization involved in a tragic incident, then, is to validate victims' concerns and wants, and to do it quickly. Not all organizations have well-developed response mechanisms in place that deal respectfully and timely with the needs of victims. A basic desire of the victim is simply to be recognized, to get a chance to tell their side of the story. And to not have to wait for months or to force the organization to listen.

Do Victims Believe that Justice is Served by Putting Error on Trial?

If a trial is the first opportunity for victims to tell their story, then that part of justice will likely be seen as served. But what about the consequences for the accused? Do victims see those as "just?" The record may be surprising. In the nurse's case from the Prologue, the mother of the infant began to doubt both

the point and the fairness of the trials against one ICU nurse well before it was all over. And there are similar examples.

In one case of criminalization, air traffic controllers in Yugoslavia were charged with murder and were jailed in the wake of a mid-air collision between two passenger aircraft. 176 lives were lost. It was 1976, and Zagreb was one of the busiest air traffic control centers in the Europe. Its navigation beacon formed a crossroads of airways heavily used by traffic to and from southeastern Europe, the Middle East, the Far East and beyond. The center, however, had been structurally understaffed for years. At the time of the accident, the radar system was undergoing testing and the center's radio transmitters often failed to work properly. Through a combination of different languages and flawed data presentation to the controller, one of the aircraft managed to level off exactly at the altitude of the other. Three seconds later, its left wing smashed through the other's cockpit and both aircraft plummeted to the ground. "Improper ATC operation," the accident investigation concluded. One controller, however, was singled out and sentenced to a prison term of seven years, despite officials from the aviation authority offering testimony that the Zagreb center was understaffed by at least 30 controllers. Significantly, the father of one of the victims of the collision led an unsuccessful campaign to prevent the controller's jailing. He then joined the efforts of other controllers to have him released after serving two years.[3] *It was not until the early 1990s that the whole air traffic control system around Zagreb was revamped.*

This is one reason why victims can have doubts about putting practitioners on trial for their alleged errors. The organizations that helped produce the problem are often left untouched; the norms, values, policies and regulations that drive their business are not critically examined.

Putting the front-end operator on trial is an example of single-loop learning, which focuses on the first part (possibly a human) that can be connected to the failure and replacing or otherwise dealing with just that part.[4]

For some victims this can seem too easy, too quick, too convenient. And it does not get at the heart of the issue: making sure that there is no next time. This is often one of the few recourses that victims have left. They have already been bereaved or injured by the incident or accident, and putting somebody in jail is not going to give them back what they lost. What uplifts instead is getting some confidence that it will not happen again, that somebody else will not have to go through what they had to suffer.

This confidence can perhaps evaporate when victims realize how a trial confines the remedy to a judgment about the right or wrong of only one person's actions. It does not get much better if the person in the dock is a manager or

a director-general instead of an operator. Accused of deficient management or insufficient oversight, these individuals get to bear the full brunt of the diffuse failings of an entire system. Not many, not even victims, would see this as either reasonable or fair.

Are Victims in it for the Money?

What about financial compensation? Are victims interested in monetary compensation and is that why they will pursue or help pursue a trial? Not really. A recurring finding from lawsuits against physicians, for example, is not only that they are surprisingly rare, but also that patients or their families do not primarily engage in legal action because of money. They sue primarily to get the story out.[5] Patients and families do not typically engage in legal action until they have found that they are being stonewalled, that no "account" is forthcoming from the practitioners or organization involved in the adverse event. They want to hold the practitioners or organization involved accountable—literally, and initially often even without prejudice or reprisal. They want to hear the story from the side of the involved practitioners and their employing organizations: what went wrong? Where? And why? How can other patients or passengers or spouses of soldiers be protected from the same kinds of failures? These are often among the most pressing questions.

If there is no other route to such disclosure, people turn to the legal system as their final address for forcing out "accountability." Again, the "accounts" produced under such duress, of course, may have little to do with what happened and much more with protecting vested personal or organizational interests (see the previous chapter).

The Suspect or Defendant

The person on trial (typically a "suspect" in a criminal trial, and a "defendant" in a civil trial) really suffers two kinds of consequences:

- **Psychological**. The suspect or defendant may experience stigmatization and excessive stress, and feel humiliation, isolation, shame, and depression. Judicial proceedings occur essentially in a foreign language for practitioners of other professions, and they may feel very little control over what is going on or what the outcome may be.
- **Practical**. Practical consequences can include jail time or significant financial costs (fines, court costs, lawyers' fees). Such costs can be borne

by insurance (in case of malpractice suits) and otherwise by professional associations (and sometimes by employers), because few practitioners (currently) have insurance that covers the cost of criminal prosecution. One other real consequence of criminal prosecution is the risk of losing the license to practice. A criminal or otherwise judicially tainted record is enough for some organizations to avoid a practitioner altogether. Loss of license often means loss of income, livelihood. It can mean loss of colleagues, context, familiarity and perhaps loss of meaning itself. Some organizations that have the resources redeploy the practitioner, but not all have the wherewithal to do so.

Where licenses were not lost, the employing organization may still not dare to have the practitioner work operationally any longer, or the practitioner him- or herself may elect not to. The nurse in the Xylocard case actually did not lose her license to practice as nurse. In a bizarre twist of legal protocol, the medical licensing board lost its access to and control over the investigation once the judiciary stepped in. They were never able to form their own judgment about the case or the nurse's ability to practice. She still has her license in her pocket today. But what does that mean? The nurse won't practice anymore. It is not likely that she could face a prescription from any physician that was even remotely unclear. It is unlikely that she could hook a patient up to any drip without asking herself a thousand questions. It is not likely that she could be effective anymore. Or safe.

In most countries, testimony of the suspect can be used as evidence in court. Interestingly, courts are mostly, or exclusively, interested in confessions. Denials are generally not seen as convincing. But if a suspect confesses the "crime," then this can be adequate for a conviction. No other evidence may be necessary.

What this means is that the police, or other investigating authorities, may sometimes have an interest in "helping" the suspect remember certain things, or stating them in a certain way. Add to this that courts in some countries are content to review only a summary of the interrogation transcript, which may have been drawn up months after the actual encounter with the suspect, and the distance between what was intended and what can get interpreted by a judge or jury becomes huge.

It is not strange that, also for this reason, suspects may feel as if they are caught up in a Kafkaesque process. They get accused of things they do not know or understand, because these are cast in a language profoundly foreign from that which makes up their own world, their own expertise.

The Prosecutor

Prosecutors are on the front line of defending and upholding the law. They have to decide which acts should be prosecuted. The role of a prosecutor is to launch a prosecution on behalf of the state.

What to Prosecute?

In the wake of an incident, whether to prosecute or not is often a very difficult call to make. In making this call, prosecutors could benefit from some guidance and perhaps even domain expertise. But access to objective domain expertise can be very hard. Whether to go ahead with prosecution or not is mostly at the prosecutor's discretion—in principle. In practice, there can be pressure from various directions:

- There may be political pressure to prosecute. Where prosecutors are elected, their constituencies could demand that they go ahead with prosecution. Where they are appointed, politicians could make clear in various ways that prosecution is desired (because politicians want to be seen as "doing something" about the problem).
- The role of the media can be significant here too: it could be that when the media calls for holding people accountable, then politicians may too.
- There can also be political pressure in the other direction (that is, to not prosecute): some organizations and professional associations have lobbied successfully for agreements between politicians and other stakeholders, so that prosecutors leave professional incidents in particular industries alone.

What to prosecute is clear—in principle. Just look in the law or jurisprudence. Yet in practice, and particularly in cases of "human error," it appears more random and unsystematic. One important reason is the sorts of laws used for such prosecution. Most stem from what could be called general risk statutes, which proscribe, for example, "endangering the lives" of other people. In many countries such statutes have their roots in road traffic law or laws governing damage to third parties in the normal course of daily life. Such laws are deliberately vague, and their jurisprudence predictably diverse, because of the infinite variation of situations that judges or juries may have to handle. But consider what happens when such general notions of danger or risk slide into considerations of culpability of practitioners' performance in a high-risk, safety-critical profession. Their very jobs involve the endangerment of the lives of other people.

Although safety data in many countries is unprotected because of freedom-of-information acts, prosecutors normally do not look into an organization's database in the hope of finding evidence of prosecutable acts. Something else must often rouse interest. A prosecutor may get a cue about the presumed seriousness of an error from the media, for example. This can be entirely coincidental, as in Xylocard case, where the prosecutor stumbled upon her story in the local newspaper. Errors can sometimes be portrayed in the media as sufficiently culpable (even before any investigation) so as to capture a prosecutor's imagination.

Safety Investigations that Sound like Prosecutors

Prosecutors and judges are not supposed to use official investigation reports in their judicial proceedings—at least this is the rule in many countries. There, the official investigation report cannot be used as evidence in court. But there is nothing in those laws that forbids prosecutors or judges from reading publicly available reports, just like any other citizen.

Over the past few years, I have counseled various investigative bodies about the language they used to describe people's actions in an incident or accident. All of these cases had attracted judicial interest. The people involved knew that judges and prosecutors were waiting for the formal report to come out (even though they were not supposed to use it formally in their judicial work).

If a trend towards criminalization is indeed happening, then recent safety board conclusions such as the ones about an aircraft accident that happened to two pilots on a repositioning flight, could be counterproductive. Unforeseen effects of high altitude flying, for which the crew was not trained, made that they entered a stall and suffered a dual engine failure as well as other unfamiliar problems in their attempts to re-start the engines (which had a history of in-flight re-start problems). Cockpit procedures did not contain specific guidance on how to recover from the situation they had gotten into.

The transportation board, however, thought that "the pilots' unprofessional operation of the flight was intentional and causal to this accident … the pilots' actions led directly to the upset and their improper reaction to the resulting in-flight emergency exacerbated the situation to the point that they were unable to recover the airplane … the probable causes were the pilots deviation from standard operating procedures, and poor airmanship."[6] *While such responses can be understandable (and may even be seen as justified), they are a little difficult to reconcile with the typical mandate of a safety investigation (which is not to find people to blame but to help prevent recurrence). Also, a focus on people's putative lack of professionalism and a direct link between their actions and the bad*

outcome can overshadow the more diffuse contributions, from inadequacies in training to general unfamiliarity with high-altitude operations, a history of engine re-start problems, incomplete flight manuals and a host of deeper organizational issues.

Perhaps language in investigation reports should be oriented towards explaining why it made sense for people to do what they did, rather then judging them for what they allegedly did wrong before a bad outcome.[7] An investigation board should not do the job of a prosecutor.

The Prosecutor as Truth-finder

Countries whose laws stem from the Napoleonic tradition (sometimes called inquisition law) typically offer their prosecutors or investigating magistrates the role of "truth-finder." This means that they and their offices are tasked with finding all facts about the case, including those that acquit the suspect or mitigate his or her contribution. Just like a judge or jury, they have to presume that the suspect is innocent until the opposite has been proven.

Combining a prosecutorial and (neutral) investigative role in this way can be difficult: a magistrate or prosecutor may be inclined to highlight certain facts over others. Accusatory law (that stems from common law tradition), in contrast, actually assumes that a prosecutor is partisan. As shown in the previous chapter, however, putting two versions of the "truth" opposite each other in an adversarial setting may still not be the best way to get to a meaningful, let alone honest, story of what happened and what to do about it. Also, the resources available to the two opposing parties may be quite asymmetric, with the prosecutor often in a better position.

Prosecutors can actually get access to evidence collected in safety investigations quite easily and use it in criminal cases. In the US, National Transportation Safety Board (NTSB) investigators can be called to testify in civil cases, but only on factual information. They cannot be forced to offer their analysis or opinions about information collected in an accident investigation. There is no such restriction, however, when they are called to testify in a criminal court case.

There is also no legal restriction in the US against the use of the actual tape from a cockpit voice recorder (CVR) in a criminal trial. This despite the Board's own extensive limits on CVR usage: it does not extend to other agencies. The NTSB, for example, strictly limits who is permitted to listen to the actual CVR tape, and these people cannot make notes of its contents. The NTSB does not release the recording or any copy of it and only makes public a transcript of the recording that is limited to details pertinent

to the safety investigation. But those restrictions end at the NTSB's doors. There is no prohibition against criminal prosecutors issuing a subpoena for the CVR tape and using it in court.[8]

The Defense Lawyer

The defense lawyer has an important role in laying out the defense strategy of the suspect. He or she can, for instance, recommend that the suspect not answer certain questions, or not testify at all. Judges or juries are not supposed to draw conclusions about suspects' culpability if they choose to remain silent. But, consistent with the fundamental nature of social relations and accountability, such silence can get interpreted as a desire to skirt responsibility.

A real and practical problem faced by most defense lawyers is that they are unlikely to understand the subtleties of practicing a particular safety-critical profession. Nor may they really have to understand. Contesting that a particular action is culpable or not is grounded in legal interpretation, rather than a deeper understanding of the meanings of risk, normative boundaries and acceptable performance as the insider would have seen them in that operational world at the time. Indeed, the legal terms that get people in trouble in court (like "negligence") are not human performance terms. These things are worlds apart.

Defense lawyers can also be limited—in budget, in human resources, and in their authorizations to investigate—to dig up their own facts about the case. In contrast, prosecutors can for example deploy the police to force facts into the open (though even there, prosecutors often face competition for limited resources: others may want or need to deploy the police elsewhere). Prosecutors can sometimes draw on the resources of government crime labs, witnesses or forensic institutes. Defense lawyers instead often have to rely on voluntary disclosure of facts by parties that think it is their duty, or in their interest to help the suspect (the employing organization often does not, by the way). This is another reason why cases can get argued on legal rather than substantive grounds. Finding minor procedural or formal flaws that scuttle the prosecution's case can be a cheaper and more effective defense than trying to match the investment in lining up facts that prosecutors can usually make.

The Judge

A judge in inquisition law generally has three tasks:

- establishing the facts;
- determining whether the facts imply that laws were broken;
- if there were, decide adequate retribution or other consequences.

Establishing the Facts

The first task, establishing the facts, is a really hard one. Facts, after all, get assembled and then brought to the bench by different parties, foremost the prosecutor. Here the border between facts on the one hand and interpretations or values on the other can begin to blur. Of course facts are disputed during a court case, this is the whole point of having a trial. But what a fact means in the world from which it came (for example, a rule "violation") can easily get lost. Neither judges, nor many of the other participants in a trial, necessarily possess the expertise to understand the language and practice from a particular domain such as nursing or air traffic control. They do not know how that world looks from the inside and, were they given a chance for such a look, they may still not really understand what they saw (the legal teams in the air traffic control case from the previous chapter were given such an opportunity, and the nurse's judges were given the Xylocard packages to look at). What the facts meant in context can remain hazy.

For this reason, judges sometimes rely on outside experts to help them decode the facts that are delivered to the bench. This is where expert witnesses come in: other practitioners or perhaps scientists whose field is relevant to the issue at hand. But judges and prosecutors and lawyers often want to ask questions that lie outside the actual expertise of the witness. Either the expert witness must decline answering, or indicate that she or he is not really confident about the answer. Neither is likely to bolster their credibility or usefulness in a courtroom. Expert witnesses are supposed to be friends of the court, that is, help the judge understand the facts from an unbiased point of view. But witnesses are selected by one of the parties, and neither party is obliged to disclose how long they looked around to find an expert witness whose opinion was favorable to their side of the story.

Determining Whether Laws were Broken

Determining whether the facts imply that laws were broken is at least as difficult as establishing the facts. How does a judge move from the facts to this judgment?

Scientists are required to leave a detailed trace that shows how their facts produced or supported particular conclusions. Such a trace typically involves

multiple stages of analysis. The researcher shows, for instance, how he or she moved from the context-specific empirical encounter (the "facts") to a concept-dependent conclusion. What scientists know, in other words, cannot be taken on faith: they have to show how they got to know what they know. This is hammered into the rules of the game, it is part of the prerequisites for publication.

For judges, however, such burden of proof does not seem to exist to the same extent. How they believe that the facts motivate a particular conclusion (and thereby judgment) can be expressed in a few lines of text.

Is a jury any better at this than a judge? Law based on Napoleonic principles does not use a jury to move from fact to judgment (nor to decide on punishment), but common law typically does. A jury also has a few of the problems that the judge faces (they are not likely to be trained in the practitioner's domain either; establishing facts and basing a judgment of unlawfulness on them is probably difficult for them too), and also introduces new problems.[9] One is the peculiarities of group behavior, from groupthink to the emergence of a dominant jury member. Jury selection is another, especially where jury members get selected on how they are likely to vote on particular aspects of the case. And the resulting group is unlikely to be a "jury of peers" where the "peer" to be judged is somebody who exercised a complex safety-critical profession that required many years of specialist education and training.

Deciding Adequate Punishment

Professionals convicted of wrongdoing often do not end up in jail, or not for a long time. Judges do seem to conclude that this is not going to be rehabilitative. Fines or conditional sentences may be given instead. Of course, neither is likely to help improve safety in the domain from which the practitioner came, and they may not even be seen as "just" either.

Lawmakers

Lawmakers do not have a direct stake in legal proceedings or what it does to the creation of just cultures—other than the stakes they represent for their constituencies (voters). But legislators do play an important role, as they are eventually the ones who help sketch out the lines in laws that will then be drawn more clearly and applied by prosecutors and judges. They may also have a stake in aligning national laws with those of international bodies. Employing organizations or professional organizations may find that without some type of

access to relevant legislators, making changes in the direction of a just culture could be difficult.

The Employing Organization

At first sight, employing organizations would not seem to benefit from the prosecution of one of their practitioners. It often generates bad press, the brand name can get tarnished, and management can be made to look bad or incompetent in the media.

On the other hand, employers can sometimes feel that they have to protect vested organizational interests, which may involve a degree of defensive posturing and shifting of blame.

What can get lost in the struggle to handle the immediate stress and challenges of legal proceedings is the organization's ethical mandate. This is, for example, to create safety (such as air traffic control) or to care for people (a hospital). Creating safety means not relying on simple, individual explanations for failure. Implicitly or explicitly supporting simplistic accounts of a bad apple could be seen as violating the very mandate the organization has. And why would that mandate not extend to the period after an accident that exposed the opposite? Caring for people means not discarding a nurse or doctor during or after he or she has been made to carry the blame for failure.

Some professions have come quite far with the development of so-called crisis intervention, peer support, or stress management programs that are intended to help practitioners in the aftermath of an incident. The importance of such programs cannot be overestimated: they help incidents become less of a stigma, that they can happen to everybody, and that they can help the organization get better if the aftermath is managed well.

Most Professionals do not Come to Work to Commit Crimes

In considering the stakes of the various parties involved in the legal pursuit of justice, it is important to remember that most professionals do not come to work to commit a tort or a crime. They do not come to work to do a bad job at all. Their actions make sense given their pressures and goals at the time. Their actions are produced by and within a complex technological system, and are part and parcel of a normal workday. Professionals come to work to do a job, to do a good job. They do not have a motive to kill or cause damage. On the

contrary: professionals' work in the domains that this book talks about focuses on the creation of care, of quality, of safety.

Notes

1 Inspiration and material for this section comes from the first chapter of Wagenaar, W.A. (2006). *Vincent plast op de grond: Nachtmerries in het Nederlands recht (Vincent urinates on the ground: Nightmares in Dutch law)*. Amsterdam, NL: Uitgeverij Bert Bakker.
2 Rowe, M. (2002). The rest is silence. *Health Affairs*, 21(4), 232–6, 236.
3 Thomas, G. (2002). Aviation on trial. *Air Transport World*, 9, 31–3.
4 Argyris, C. and Schön, D.A. (1995). *Organizatonal learning II: Theory, method and practice.* New York: Addison Wesley.
5 Berlinger, N. (2005). *After harm: Medical error and the ethics of forgiveness*. Baltimore, MD: The Johns Hopkins University Press.
6 National Transportation Safety Board (2007). *Report of Aviation Accident: Crash of repositioning flight, Pinnacle Airlines flight 3701, Bombardier CL-600-2B19, N8396A, Jefferson City, Missouri, October 14, 2004 (NTSB/AAR-07/01)*. Washington, DC: Author.
7 See particularly Dekker, S.W.A. (2006). *The field guide to understanding human error*. Aldershot, UK: Ashgate Publishing Co.
8 McKenna, J.T. (1999, 13 December). Criminal and safety probes at odds. *Aviation Week and Space Technology*, 47–8.
9 Wagenaar, op. cit.

This page intentionally left blank

11 Three Questions for a Just Culture

"What is just?" ask colleagues in the aftermath of an incident "caused" by one of them. "How do we protect ourselves against disproportionate responses?" they add. "What is wise?" ask the supervisors. "What do people—other employees, customers, the public—expect me to do?" ask managers. And then other parties (such as prosecutors) ask, "Should we get involved?" The confusion about how to respond justly and still maintain a sense of organizational cohesion, loyalty and safety can be considerable.

Three Questions

At the same time, many organizations (whether they know it or not) seem to settle on pragmatic solutions that at least allow them to regain some balance in the wake of a difficult incident. When you look at these "solutions" a little more closely, you can see that they really boil down to answers to three central questions:

1 Who in the organization or society gets to draw the line between acceptable and unacceptable behavior?
2 What and where should the role of domain expertise be in judging whether behavior is acceptable or unacceptable?
3 How protected against judicial interference are safety data (either the safety data from incidents inside of the organization or the safety data that come from formal accident investigations)?

The differences in the directions that countries or organizations or professions are taking towards just cultures come down to variations in the answers to these three questions. Some work very well, in some contexts, others less so. Also, the list of solutions is far from exhaustive, but it could inspire others to think

more critically about where they or their organization may have settled (and whether that is good or bad).

In general, though, we can already see this for the three questions.

On Question 1

The more a society, industry, profession or organization has made clear, agreed arrangements about who gets to draw the line, the more predictable the managerial or judicial consequences of an occurrence are likely to be. That is, practitioners will suffer less anxiety and uncertainty about what may happen in the wake of an occurrence, as arrangements have been agreed on and are in place.

Anxiety that results from uncertainty about what may happen is really bad, and in many ways unnecessary. For example, while the real risk of being sued for medical malpractice is actually low, doctors perceive it to be quite high. And it is perception, not reality, that creates anxiety. The anxiety leads to defensive medicine, unnecessary tests, rather than high-quality care. Even if you, as an organization, have to tell your people that the aftermath of an incident may be unpleasant for them, that is still better than not telling them anything at all.

On Question 2

The greater the role of domain expertise in drawing the line, the less practitioners and organizations may be likely to get exposed to unfair or inappropriate judicial proceedings.

There is actually no research that suggests that domain experts automatically prevent the biases of hindsight slipping into their judgments of past performance. Hindsight is too pervasive a bias. It takes active reconstructive work, for everyone, to even begin to circumvent its effects. Domain experts, however, do have an easier time forming an understanding of the situation as it looked to the person at the time, as they probably know such situations from their own experience. Here is how that may influence their ability to make a fairer judgment of the controversial action:

- It is easier for domain experts to understand where somebody's attention was directed. This is one area where domain experts may have an easier time avoiding the hindsight bias: even though the outcome of a sequence of events will reveal (in hindsight!) what data was really important, domain experts can make better judgments about the perhaps messy or noisy

context of which these, now critical, data were part and understand why it was reasonable for the person in question to be focusing on other tasks and attentional demands at the time.

- It is probably easier for domain experts to understand the various goals that the person in question was pursuing at the time, and whether these were reasonable given the circumstances, and whether and how these goals may have conflicted with each other (for example, safety versus efficiency, production versus protection). Domain experts can also form a better judgment than outsiders about the reasonability of goal priorities in cases of goal conflicts, especially since the system's preference for one goal over another may have been expressed tacitly, without explicitly stating it. Outsiders would not likely get access to that kind of information.
- For domain experts it is also easier to assess whether any unwritten rules or norms may have played a role in people's behavior. All professions have unwritten rules and national or professional norms, to which members of the profession are supposed to conform. Without conforming to these tacit rules and norms, people often could not even get their work done. The reason, of course, is that written guidance and procedures are always incomplete as a model for practice in context. That means that practitioners need to bridge the gap between the written rule and the actual work-in-practice, which often involves a number of expert judgments and outsiders often have no idea about the existence of these norms, and would perhaps not understand their importance or relevance for getting the work done.

That said, domain experts may have other biases that work against their ability to fairly judge the quality of another expert's performance. There is, for example, the issue of psychological defense: if experts were to affirm that the potential for failure is baked into their activity and not unique to the practitioner who happened to inherit that potential, then this makes them vulnerable too. Sometimes it can be more comforting to think that the errors made by a fellow practitioner would not happen to you; that they really are unique to that other person.

On Question 3

The better safety data is protected from judicial interference, the more likely it is that practitioners could feel free to report. The protection of this safety data is connected, of course, to how the country or profession solves questions 1) and 2). For example, countries or professions that do protect safety data typically have escape clauses, so that the judiciary can gain access "when crimes are committed,"

or in "justified cases when duly warranted," or "for gross negligence and acts sanctioned by the criminal code." It is very important to make clear who gets to decide what counts as a "crime," or "duly warranted" or "gross negligence," because any uncertainty there (or the likelihood of non-domain experts making that judgment) can once again hamper practitioners' confidence in the system and their willingness to report or disclose.

Local Solutions to the Three Questions

Local Solution 1: Do Nothing to Actively Handle the Three Questions

This is a solution that a number of country or professions apply, perhaps because they have not yet been confronted by the consequences of judicial action against practitioners or have themselves seen the difficulty of acting in the wake of failure. This may, of course, be just a matter of time.

1 Who gets to draw the line? In all probability it will be a prosecutor who has become inspired by media reports or other triggers that made her or him look more closely into an occurrence. General risk statutes or other laws can be used to accuse practitioners of, for example, endangering the lives of other people. Access to data to build a criminal case should be relatively easy if the country or profession has not done much or anything to prevent such judicial intrusions into their safety data. The prosecutor draws the line in the first instance, and then the judge (or jury) gets to decide.
2 The role of domain expertise is probably minimal in judging whether a line of acceptability was crossed or not. A prosecutor, for example, has no or limited domain expertise, yet she or he gets to demonstrate whether professional judgments are culpable or not. The judge is not likely to have any domain expertise either.
3 Protection of safety data is not likely to exist, and even if it does, then a country or profession that goes by local solution 1 probably has caveats in its protection so that a prosecutor can open up databases upon suspicion of a crime (and the prosecutor is often the one who decides when that is the case).

Consequences: practitioners may feel uncertain and anxious about whether "they will be next" because the rules of criminalization are left unclear and open to

interpretation. Who gets criminalized for what seems random. A just culture is a long way off, and open and honest reporting could be difficult.

Local Solution 2: The Volatile Safety Database

Some countries or professions who do not actively handle the three questions in legislation or cross-disciplinary arrangements (for example, between their departments of transportation or health on the one hand and justice on the other) spontaneously call for the creation of a next local solution: the destroyable safety database. This means that the safety data that organizations gather themselves are stored in a form that is very easy and quick to destroy. Some safety departments have seriously considered this idea, so as to immunize themselves against prosecution. This is especially the case in countries where the organization's personnel are themselves government employees (such as some hospital workers or air traffic controllers) and can thus be forced, through various statutes and laws, to hand over anything that belongs to the state.

1 Who gets to draw the line? Same as local solution 1.
2 The role of domain expertise: same as local solution 1.
3 Protection of safety data could be guaranteed, as the data will vanish when prosecutorial pressure is applied. The cost, of course, is huge: the disappearance of an organization's safety database (which can in turn violate other statutes).

Consequences: this is not really a practical solution because of the consequences of destroying a database. But that it is being considered in several countries or professions in the first place should serve as an indication of the lack of trust necessary for building a just culture. The relationship between the various stakeholders may be troubled or underdeveloped. The suspicious climate sustained by this solution will not be good for the growth of a just culture.

Local Solution 3: Formally Investigate beyond the Period of Limitation

In almost all countries, prosecutors have a limited number of years to investigate and prosecute crimes. In one country, the investigation of an accident took so long that the so-called period of limitation for any possible charges (seven years in this case) expired. Practitioners sighed in relief. Inspired by such apparently legitimate delaying tactics, stakeholders in other countries and professions have considered deliberately stalling an investigation so that the judiciary could

not get access until the period of limitation has expired. This solution works only, of course, if the judiciary is legally limited in beginning its probe into an occurrence while the formal investigation is still ongoing. In some countries or professions this is indeed the case.

1 Who gets to draw the line? While prosecutors and judges would still be left to draw the line eventually, other parties can withhold from them both the data and the opportunity to draw a line.
2 The role of domain expertise is interesting in this solution, as those with more expertise of the domain (safety investigators) make a judgment of the potential culpability of the acts they are investigating. If they judge these acts to be potentially (but unjustifiably and counterproductively) culpable, they may stall an investigation until the period of limitation has expired. In this sense, investigators introduce domain expertise into the judgment of whether something is acceptable or not, but they apply this expertise in advance—anticipating how the judiciary would respond to the data they have. Investigators may, of course, lack the domain expertise in the legal area to really make an accurate *ex ante* judgment, but previous experiences or the general climate in the country or profession may give them a good basis for their conjecture.
3 Protection of safety data is pretty strong, but, of course, hinges on the strength of the laws and statutes that prevent the judiciary from getting access to investigation data before the period of limitation has expired. Any legal opportunities that allow the judiciary to get into the formal investigation will directly undermine this solution.

Consequences: a climate of distrust and competition between stakeholders remains strong with this solution. Rather than resolving issues on merit, stakeholders may engage in legal gaming to try to get access (or retain privileged access) to safety data for their own purposes. The climate is not encouraging for the emergence of a just culture.

Local Solution 4: Rely on Lobbying, Prosecutorial and Media Self-restraint

This is different from the previous solutions in that it relies almost entirely on trust between stakeholders. It has been achieved in a few countries (often after intense lobbying of lawmakers and other government officials by industry stakeholders). It has succeeded particularly in countries with strong freedom

of information legislation that leaves their safety data exposed to both media and judiciary.

This local solution depends entirely on the extent of the trust developed and maintained, not on legal protection for any of the stakeholders. Thus, these countries typically have no protection in place for either reporters or safety data, and the judiciary has unfettered access to investigations—in principle. In practice, no prosecutor has dared to be the first to breach the trust built up. Interestingly, this solution seems to work in smaller countries that are culturally inclined towards homogeneity, trust, coherence and social responsibility.

1 Who gets to draw the line? Prosecutors would in principle get to draw the line, but they have so far not dared to draw anything. The proscription against them doing so is not a legal one, but rather cultural or political: going in and upsetting the delicate trust developed between parties is "not done" or politically not wise. But that does not mean it cannot be done. In fact, rules in countries with this solution still make exceptions for the kinds of "crimes" or "gross negligence" that prosecutors should still prosecute. The problem is, of course, chicken-and-egg: how is a prosecutor to find out whether a line was crossed without drawing one?
2 The role of domain expertise has been considerable in building the necessary trust between stakeholders, particularly in convincing other stakeholders (the media, the judiciary) of the value of their self-restraint, so that the entire society can benefit from safer professional systems.
3 Protection of safety data is not legally guaranteed but achieved by cultural convention and/or political pressure.

Consequences: at first sight, this solution comes across as a fraud, and as extraordinarily brittle. After all, there is nothing "on paper:" the entire contract between stakeholders not to interfere with each others' business is left to consensual agreements and trust. Practitioners may feel free to report because historically there is no threat (but can history be a guarantee for the future in this case?). On deeper inspection, though, this solution is as robust as the culture in which it is founded. And cultures can be very robust and resistant to change. This, at the same time, creates a high threshold for entry into such an arrangement: without the right cultural prerequisites, this solution may be difficult to achieve.

Local Solution 5: Judge of Instruction

In the wake of prosecutions of practitioners which were widely seen as counterproductive to safety, some countries have moved ahead with installing a so-called judge of instruction. Such a judge 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.

1 Who gets to draw the line? Initially (and most importantly) it is 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). Other considerations can make it into the drawing of the line too (for instance, the interests of other stakeholders).
2 The role of domain expertise is supposed to be considerable in this solution. In one country's solution, the judge of instruction is supported by a team from the aviation industry to help determine which cases should go ahead and which not. The make-up of this team and their interaction with the judge of instruction are crucial of course. For example, if unions or professional associations are not sufficiently represented, industry representatives may decide that it is in their interest to recommend to the judge to go ahead with prosecution, as it may help protect some of their concerns.
3 Protection of safety data is managed through the judge of instruction. If prosecutors want access to safety data, they will have to go via the judge of instruction, but there are (as usual) exemptions for serious incidents and accidents.

Consequences: a judge of instruction could function as a reasonable gatekeeper—weighing the various interests before a case can even be investigated by a prosecutor. It means, though, that such a judge needs a fair representation of all stakes, and not be susceptible to asymmetric lobbying by certain parties or interests over others. Since it is a rather new solution to the criminalization of human error, there is not a lot of data yet to see whether it works well or not.

Local Solution 6: Prosecutor is Part of the Regulator

A solution that takes domain expertise right up to prosecutor level is one in which the prosecutor him- or herself has a history in or affiliation with the domain, and the office of prosecutor for that particular domain is inside of the regulator.

1 Who gets to draw the line? The prosecutor gets to draw the line (to be confirmed or rejected by a judge or jury), and the prosecutor is a person from the domain and employed by a major stakeholder in the domain.
2 The role of domain expertise is considerable, since the prosecutor comes from the domain and is employed by one of its large safety stakeholders. 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.
3 Protection of safety data is managed as an effect of this arrangement. The regulator has interests in protecting the free flow of safety information (not only as data for its oversight, but particularly for the self-regulation of the industry it monitors).

Consequences: the integration of prosecutor and regulator can prevent unfair or inappropriate prosecution, not only because of the tight integration of domain expertise, but also because of the greater relevance of the laws or regulations that will likely be applied (as the prosecutor works for a body that makes and applies the laws for that particular domain). The risk in this solution, of course, is that the regulator itself can have played a role (for example, 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.

Local Solution 7: Disciplinary Rules within the Profession

A large number of professional groups (everything from accountants to physicians to hunters to professional sports players) have their own elaborate system of disciplinary rules that are first and foremost meant 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.

1 Who gets to draw the line? The professional's peers get to draw the line between acceptable and unacceptable. There may be pressures, of course, that go outside the actual situation considered, so as to guarantee society's (and the judiciary's!) continued trust in the system (for example, the ATM system) and its ability to manage and rectify itself. This may make it necessary to sometimes lay down the line more strictly so that a message of "we are doing something about our problems" clearly gets communicated to the outside—to the detriment of justice done to an individual professional. Who gets to draw the line for criminally culpable actions is an even larger problem: internal rules are not equipped to handle those, so somewhere there needs to be a possibility for judging whether outside legal action is necessary. This can be the prosecutor's initiative (but then he or she needs enough data to trigger action) or the disciplinary board (but they likely lack the legal expertise to make that judgment).
2 The role of domain expertise is total. Domain expertise is the basis for making the judgment about the right or wrong of somebody's actions, not some externally dictated law or statute. Domain expertise is also used to consider whether to forward a case to the formal judiciary (as there will always be an escape hatch for cases of "gross negligence" and so forth). But it is at least largely domain expertise that gets to draw that line here too.
3 Protection of safety data is likely to be independent of professional disciplinary rules and would need additional legislation for formal protection. However, with a functioning (and trustworthy) internal professional disciplinary system in place, people inside a profession may feel freer to report incidents and concerns.

Consequences: the total integration of domain expertise in the administration of justice makes a solution based on professional disciplinary rules attractive. Not only do domain experts judge whether something is acceptable or unacceptable, they also draw largely from the domain the "rules," written or unwritten, on which basis they make their judgment.

But there is a possible paradox in the justness of professional disciplinary rules. Because disciplinary rules exist for the maintenance of integrity of an entire profession, individual practitioners may still get "sacrificed" for that larger aim (especially to keep the system free from outside interference or undue scrutiny). To remain trustworthy in the eyes of other stakeholders, then, the disciplinary rules may have to wreak an occasional internal "injustice" so as to outwardly show that they can be trusted. This does not necessarily enhance the basis for just culture, as practitioners could still feel threatened and anxious about possible career consequences.

I came across an interesting, spontaneous variant of solving incidents and safety matters internally. An air traffic control center had essentially agreed with some of the airlines frequenting its airspace that they would send reports on near misses and other problems directly to them. This was of course a very short-loop way to learn: a problem was seen and reported directly to the air traffic control center that could do something about it. It also prevented outsiders from making their own judgments about the performance of those involved and meting out any consequences (legal, regulatory). But those other parties (for example, the regulator) felt that accountability was being shortcut—reporting routines established earlier had ensured that reports would go through them. Lessons learned could also be relevant to the wider industry, something that was missed by this local, internal solution.

This page intentionally left blank

12 Not Individuals *or* Systems, but Individuals *in* Systems

In an earlier book,[1] I laid out a choice between the old view and the new view of human error:

- **The old view** sees human error as a cause of incidents. To do something about incidents, then, we need to do something about the particular human involved: suspend, retrain, admonish, charge him or her. Or we have to do something about humans in general: marginalize them by putting in more automation, or rigidify their work by creating more rules and procedures.
- **The new, or systems view**, sees human error as a symptom, not a cause. Human error is an effect of trouble deeper inside the system. To do something about a human error problem, then, we must turn to the system in which people work: the design of equipment, the usefulness of procedures, the existence of goal conflicts and production pressure.

This choice, between old and new view, is founded in decades of research on safety and risk in complex domains. The two alternatives can serve as useful bookends in debates about the causes of mishaps, and what countermeasures you should deploy.

But it leaves an important question unattended: can people simply blame the system when things go wrong? To many, this logical extension of the new view seems like a cop-out, like an excuse to get defective or responsible practitioners off the hook. The new view would seem almost incompatible with holding people accountable.

Indeed, says Pellegrino, systems are not enough.[2] Of course we should look at the system in which people work, and improve it to the best of our ability. But 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. Rather than individuals *versus* systems, we should begin to understand the relationships and roles of individuals *in* systems.[3]

The Discretionary Space for Personal Accountability

A system creates all kinds of opportunities for action. And it also constrains people in many ways. Beyond these opportunities and constraints, we could argue that there remains a discretionary space, a space that can be filled only by an individual care-giving or technology-operating human. This is a final space in which a system really does leave people freedom of choice (to launch or not, to go to open surgery or not, to fire or not, to continue an approach or not). It is a space filled with ambiguity, uncertainty and moral choices.

Systems cannot substitute the responsibility borne by individuals within that space. Individuals who work in those systems would not even want their responsibility to be taken away by the system entirely. The freedom (and the concomitant responsibility) that is left for them is what makes them and their work human, meaningful, a source of pride. But systems can do two things:

- One is to be very clear about where that discretionary space begins and ends. Not giving practitioners sufficient authority to decide on courses of action (such as in many managed care systems), but demanding that they be held accountable for the consequences anyway, creates impossible and unfair double binds. Such double binds effectively shrink the discretionary space before action, but open it wide after any bad consequences of action become apparent (then it was suddenly the physician's responsibility after all). The same goes when asking dispensation for an unqualified crewmember to proceed with an instrument approach. The system is clear in its routine expectation that a commander will ask such dispensation. And if all goes well, no questions will be raised. But if problems occur on the approach, the request for dispensation suddenly becomes the commander's full responsibility. Such vagueness of where the borders of the discretionary space lie is typical, but it is unfair and unreasonable.
- The other is to decide how it will motivate people to carry out their responsibilities conscientiously inside of that discretionary space. Is the source for that motivation going to be fear or empowerment? Anxiety or involvement? "There has to be some fear that not doing one's job correctly could lead to prosecution," said an influential commentator in 2000 . Indeed, prosecution presumes that the conscientious discharge of

personal responsibility comes from fear of the consequences of not doing so. But neither civil litigation nor criminal prosecution work as a deterrent against human error.[4] 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.[5] The anxiety and stress generated by such accountability adds attentional burdens and distracts from conscientious discharge of the main safety-critical task.[6]

Rather than making people afraid, systems should make people participants in change and improvement. 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.[7]

Haavi Morreim reports a case in which an anesthesiologist, during surgery, reached into a drawer that contained two vials, sitting side by side.[8] *Both vials had yellow labels and yellow caps. One, however, had a paralytic agent, and the other a reversal agent to be used later, when paralysis was no longer needed. At the beginning of the procedure, the anesthesiologist administered the paralyzing agent, as per intention. 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, it turned out that this had happened to them too, and that they were all quite aware of the enormous potential for confusion. All knew about the hazard, but none had spoken out about it.*

The question is of course why no anesthesiologist had flagged this problem before. Anxiety about the consequences of talking about possible failures cannot be excluded: it has squelched safety information before.

Even more intriguing is the possibility that there is no climate in which practitioners feel they can meaningfully contribute to the context in which they work. Those who work on the safety-critical sharp end every day, in other words, did not feel they had a channel through which to push their ideas for improvement. I was reminded of one worker who told me that she was really happy with her hospital management's open-door policy. But whenever she went through that open door, the boss was never there.

The example does raise the choice again. Do you really think you can prevent anesthesiologists from grabbing a wrong vial by making them afraid of the consequences if they do? Or do you want to prevent them from grabbing a wrong vial by inviting them to come forward with information about that vulnerability, and giving you the opportunity to help do something about the problem?

This example also confirms that 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. This also means that 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.

Blame-free is not Accountability-free

Equating blame-free systems with an absence of personal accountability, as some do,[9] is wrong. Blame-free means blame-free, not accountability-free. The question is not whether we want practitioners to skirt personal accountability. Few practitioners do. The question is whether we want to fool ourselves that we can meaningfully wring such accountability out of practitioners by blaming them, suing them or putting them on trial. No single piece of evidence so far seems to demonstrate that we can.

We should convince ourselves that 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. Most practitioners will relish such responsibility. Just as most practitioners often despair at the lack of opportunity to really influence their workplace and its preconditions for the better.

Forward-looking Accountability

"He or she has taken responsibility, and resigned."

We often say this in the same sentence. We may have come to believe that quitting and taking responsibility are the same thing. Sure, they can be. But they don't have to be. In fact, holding people accountable may be exactly what we are *not* doing when we allow them to step down and leave a mess behind.

Accountability is often only backward-looking. This is the kind of accountability in trials or lawsuits, in dismissals, demotions, or suspensions. Such accountability tries to find a bad apple, somebody to blame for the mess. It is the kind of accountability that feeds a press (or politicians, or perhaps even a company's board), who may eagerly be awaiting signs that "you are doing something about the problem." But for you and your organization, such backward-looking accountability is pretty useless beyond getting somebody's hot breath out of your neck.

Instead, you could see accountability as looking ahead. Stories of failure that *both* respond to calls for accountability and allow people and organizations to learn and move forward, are essentially about looking ahead. In those stories, accountability is something that brings information about needed improvements to people or groups that can do something about it. There, accountability is something that allows people and their organization to invest resources in improvements that have a safety dividend, rather than deflecting resources into legal protection and limiting liability. This is captured in what Virginia Sharpe calls "forward-looking accountability." Accountability should lay out the opportunities (and responsibilities!) for making changes so that the probability of such harm happening again goes down.

An explosion occurred at a Texas oil refinery in March 2005, as an octane-boosting unit overflowed when it was being restarted. Gasoline vapors seeped into an inadequate vent system and ignited in a blast that was felt five miles away. The explosion killed 15 people. An internal company study into the accident found that four of the company's US executives should be fired for failing to prevent the explosion, and that even the company's global refinery chief had failed to heed serious warning signals. The company's "management was ultimately responsible for assuring the appropriate priorities were in place, adequate resources were provided, and clear accountabilities were established for the safe operation of the refinery," said the lead company investigator.

Corporate budget cuts had compromised worker safety at the plant, an earlier report had found, and the refinery had had to pay a record fine for worker safety violations at its site. A safety culture that "seemed to ignore risk, tolerated non-compliance and accepted incompetence" was determined as a root cause of the accident. The global refinery chief should have faced and communicated "the brutal facts that fundamentally, the refinery was unsafe and it was a major risk to continue operating it as such."[10]

Calls for accountability are important. And responding adequately to them is too. Sending responsible people away is, of course, one response. But remember from the first chapter that calls for accountability are in essence about trust. About people, regulators, the public, employees, trusting that professionals will take problems inside their practice or organization seriously and do something about them.

This means that just getting rid of a few people (even if they are in positions of greater responsibility) may not be seen as an adequate response. Nor is it necessarily the most fruitful way for an organization to incorporate lessons about failure into what it knows about itself, into how it should deal with such vulnerabilities in the future.

Notes

1 Dekker, S.W.A. (2006). *The field guide to understanding human error*. Aldershot, UK: Ashgate Publishing Co.
2 Pellegrino, E.D. (2004). Prevention of medical error: Where professional and organizational ethics meet. In: Sharpe, V.A. (ed.) *Accountability: Patient safety and policy reform*. Washington, DC: Georgetown University Press, 83–98.
3 See also Berlinger, N. (2005). *After harm: Medical error and the ethics of forgiveness*. Baltimore, MD: The Johns Hopkins University Press.
4 http://www.womens-health.org.nz/publications/WHW/whwjan97.htm#litigation.
5 Dauer, E.A. (2004). Ethical misfits: Mediation and medical malpractice litigation. In: Sharpe, op. cit., 185–202.
6 Lerner, J.S. and Tetlock, P.E. (1999). Accounting for the effects of accountability. *Psychological Bulletin*, 125, 255–75.
7 Dekker S.W.A. and Laursen, T. (2007). From punitive action to confidential reporting: A longitudinal study of organizational learning. *Patient Safety and Quality Healthcare*, 4(5), 50–56.
8 Morreim, E.H. (2004). Medical errors: Pinning the blame versus blaming the system. In: Sharpe, op. cit., 213–32.
9 Pellegrino, op. cit.
10 Texas executives faulted in BP explosion. *International Herald Tribune*, May 4, 2007, 10.

13 A Staggered Approach to Building your Just Culture

Where do you go from here? Building a just culture starts at home, in your own organization. The approach I would like to suggest is a staggered one (see Figure 1). It allows you to match your organization's ambitions to the profession's possibilities and constraints, the culture of your country and its legal traditions and imperatives. Each step in the staggered approach is already a contribution to the creation of a just culture. Each subsequent step is probably more difficult, as it draws in more parties with different backgrounds and persuasions, and larger stakeholder groups and their perspectives and interests. But each step already goes a little bit of the way to reconcile the agendas of different stakeholders. Each step may contribute a little bit to the building of trust between them.

A Staggered Approach: the Steps

Step 1: Start at Home, in Your Own Organization

To lay the basis for the emergence of a just culture in your profession or country, nothing is as important as starting at home, in your own organization. This will allow you to begin building relationships and trust between the first parties that matter: practitioners and their managers. Trust in management is not necessarily widespread among practitioners in all industries, which may have a number of reasons. It might be that managers sometimes come from backgrounds other than the profession, but even if they are ex-professionals, managers can be seen as "outsiders" (or worse: turncoats). Trust that was lost in management because of their positions on industrial or social issues (for example, the application of work time regulations, vacation time) can also spill over into safety issues. So even if management has not acted negatively in relation to an incident before, its

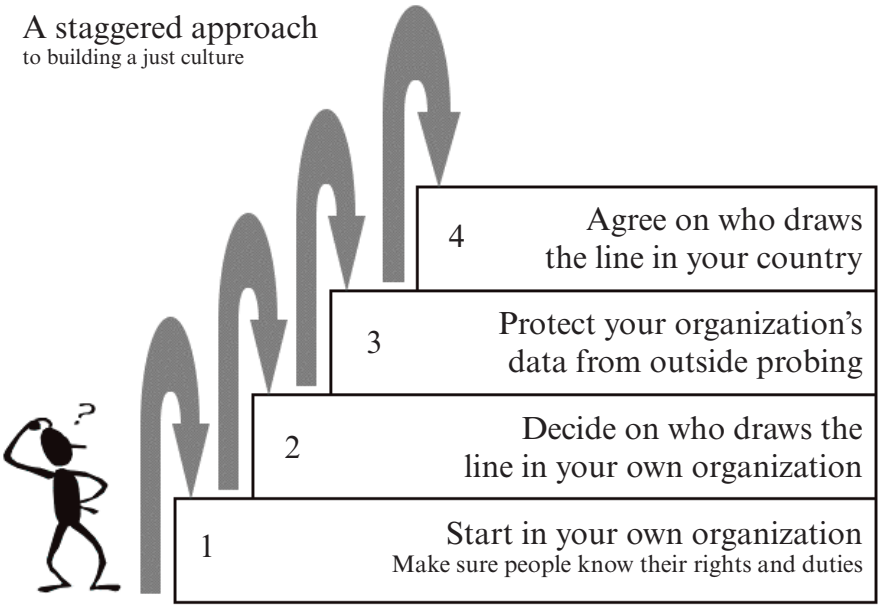

Figure 1 **A staggered approach to building a just culture. Each subsequent step gets more difficult, but each step is already progress in the direction of a just culture**

behavior elsewhere (or perception thereof) can affect the trust that practitioners will have in management handling of safety matters.

You can do a number of things immediately (if you have not already done so today):

- An incident must not be seen as a failure or a crisis, either by management, or by colleagues. An incident is a free lesson, a great opportunity to focus attention and to learn collectively.
- Abolish all financial and professional penalties in the wake of an occurrence. Suspending practitioners after an incident should be avoided at all cost. These measures serve absolutely no purpose other than making incidents into something shameful, something to be kept hidden. If your organization

has these kinds of rules in place, you can count on losing out on a lot of valuable safety information.

- 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. This is not only devastating for them, but for every practitioner and, by extension, the organization, as incidents are once again seen as something to be kept concealed, out of view. Reintegrate these practitioners into the operation smoothly and sensitively, being aware of the possibility for stigmatization by their own colleagues.
- Implement, or review the effectiveness of any debriefing programs or critical incident/stress management programs you may have in place to help practitioners after incidents (and if you don't have any in place, start building such programs pronto!). 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 deal with that practitioner in the wake of an incident (other then perhaps relieving him or her temporarily to deal with the stress and aftermath). Aim to decouple an incident from what may look like a performance review of the practitioner involved. Any retraining of the practitioner involved in the incident will quickly be seen as punishment (and its effects are often quite debatable), so this should be done with utmost care and only as a last resort.
- 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 in how many incidents it has, but in how it deals with the incidents that it has its people report.
- Be sure 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. One union had prepared little credit-sized cards on which it had printed the practitioner's rights and duties in the wake of an occurrence (for example, to whom they were obliged to speak (for example, investigators) and to whom not (for example, the media)). Even

> in a climate of anxiety and uncertainty about the judiciary's position on occurrences, such information will give practitioners some anchor, some modicum of certainty about what may happen. At the very least this will prevent them from withholding valuable incident information because of misguided fears or anxieties.

Starting at home, in your own organization, will allow you to lay the basis for a just culture. Without you laying that basis first, don't count on anybody else (like your judiciary or legislature) to do it for you.

Step 2: Decide Who Draws the Line in Your Organization

One important decision for an organization is 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). It is also 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.

Whether a practitioner should undergo retraining, for example, is something that should be discussed not only with the practitioner in question (rather than just handed down from above), but also checked with a group of peers who can consider the wider implications of such a measure in the wake of an incident (for instance, on the reputation of that practitioner, but also on the way incidents will be seen and treated by colleagues as a result).

Empowering and involving the practitioner him- or herself in the aftermath of an incident is the best way to maintain morale, maximize learning, and reinforce the basis for a just culture.

Step 3: Protect Your Organization's Data from Undue Outside Probing

Protecting your organization's data from outside probing should not be left to chance (as in, say, "the prosecutor has not previously shown interest, so why would he or she now?"), and probably not just to cultural convention or political pressure either.

The creation of trust between stakeholders is, of course, very important, and in this case it means that the judiciary will be willing to let an organization handle its own data when it has been given the assurance and confidence that the organization will come to it if a case is really likely to be culpable. This in turn hinges on the decisions made in step 2: who draws the line between acceptable and unacceptable inside the organization? This person or group will

likely be the one that has to give the judiciary the confidence that it knows what it is doing and that it will treat the organization's data with integrity and fairness.

None of this, however, is likely to automatically inspire practitioners to freely report without fear. It is best to try to enshrine the protection of the organization's critical data in law.

As said in Chapter 3, however, you have to think through the consequences of this step. One problem is that it can lock information up even for those who rightfully want access to it, and who have no vindictive intentions. Remember the patient, for example, or the victim of a transportation accident (or the family), whose main aim is to find out something specific about what happened. The protection of reporting can make such disclosure to such parties more difficult.

Step 4: Decide Who Draws the Line in Your Country

Having non-domain experts draw the line between acceptable and unacceptable practitioner performance is fraught with risks and difficulties. This is where the advantages of local solutions that somehow meaningfully integrate domain expertise in the drawing of the line nationally come in. The use of expert witnesses during a trial is not likely to do this meaningfully, as that role is always rather constrained and testimony limited. In fact, this may be too late a stage to wait to bring in domain expertise.

It could be profitable to start a discussion with the prosecuting authority in your 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.

Local solutions that rely on internal professional disciplinary rules (to which the judiciary delegates its legal oversight) achieve a total integration of domain expertise in the determination of where the line is drawn. These solutions are currently not very common. Local solutions that otherwise integrate domain expertise (such as a integrating prosecutor and regulator in one organization, or having a judge of instruction supported by a team of domain experts) have advantages too. But they do not go so far as to really create a "jury of peers" that is able to judge performance in context. Any delegation to a greater degree of domain expertise, of course, does require that the judiciary can be confident

that cases will be handled fairly and without prejudice in favor of colleagues (who may be seen to try to protect one another).

Not Bad Practice, but Bad Relationships

Unjust responses to failure are almost never the result of bad performance. They are the result of bad relationships.[1] [2] You can see this in almost any situation where we want to talk of just culture.

The strongest predictor, for example, of whether a physician will be sued is the extent to which patients feel they are treated with respect, honesty, and personal interest. The nature and gravity of the injury matters much less, if much at all. Injuring a patient during medical care is a severe breach of the fiduciary relationship between caregiver and patient. Patients typically feel betrayed, and angry. Restoring that relationship, or at least managing it wisely, is often the most important ingredient of a successful response.

Managing relationships between patients and doctors, if not restoring them, is one major aim of mediation, a form of alternative dispute resolution (ADR) in medicine. Mediation restores communication between the two parties, often (if not always) with the help of a mediator. What is said is kept confidential by law, thus making mediation a safe place for showing remorse, for introspection and the exploration of corrective actions without it being seen as admitting liability. In contrast to litigation or criminal-legal action, mediation allows apology, expressions of regret, compassion to occur much more naturally. Mediation is also much more flexible in allowing different outcomes. Compensation does not have to be money (indeed, it most often is not in ADR). In addition to agreeing to care for the injured party in whatever way necessary (medical or otherwise), mediation can inspire changes to procedures, augmenting of education, or other changes that respond to a patient's desire to never see this happen again.[3]

Here is another example of the importance of relationships. Whether employees will see management responses to failure as just depends not so much on the response (or on the bad performance that triggers it). Rather, it depends to a great extent on the existing relationship between management and employees.

We did extensive fieldwork among firefighters, to see how they learn from failures that occur during their emergency responses. If firefighters felt that they could come forward with their errors, then it was due largely to the relationship with their supervisors and

their managers. In one station firefighters worked in close concert with their management, allowing an atmosphere where reporting errors and suggesting changes was normal, expected and without jeopardy for any of the parties. Conversely, at a larger urban fire station with a distrustful industrial relationship, there was less bottom-up participation in decisions involving work-context, less firefighter involvement in learning, and much greater suspicion that any reported errors would not be responded to fairly.

Here is a final example: the relationship between the judiciary (or even one particular prosecutor) and a profession can be a strong determinant of legal action in the wake of an incident. A prosecutor may suspect that there is deliberate stonewalling, that legally accessible information is being deliberately withheld (remember: "omertà"). This, however, is often a response to earlier action by the prosecutor. Both are evidence of a relationship gone sour, and one way forward is to simply go and talk together. In the wake of several cases mentioned in this book, I have seen professional representatives propose exactly that. Talking, of course, is about the possible restoration of a relationship.

If bad relationships are behind unjust responses to failure, then good relationships should be seen as a major step toward just culture. Good relationships are about openness and honesty, but also about responsibilities for each other, and bracketing (this is yours, this is mine). Good relationships are about communication, about being clear about expectations and duties, and about learning from each other.

Perhaps this can come as somewhat of a relief. "Justice" and "culture" are two huge concepts. They are both essentially contested categories: what either means will forever be open to debate and controversy. They are basically intractable, unmanageable. A relationship, on the other hand, is manageable. At least half of it is in your hands. So if you want to do something about just culture, that is probably where to start.

Notes

1 Morreim, E.H. (2004). Medical errors: Pinning the blame versus blaming the system. In: Sharpe V.A. (ed.) *Accountability: Patient safety and policy reform*. Washington, DC: Georgetown University Press, 213–32.

2 Berlinger, N. (2005). *After harm: Medical error and the ethics of forgiveness*. Baltimore, MD: The Johns Hopkins University Press.

3 Dauer, E.A. (2004). Ethical misfits: Mediation and medical malpractice litigation. In Sharpe, op. cit., 185–202.

This page intentionally left blank

Epilogue

If professionals consider one thing "unjust," it is often this: split-second operational decisions that get evaluated, turned over, examined, picked apart and analyzed for months—by people who were not there when the decision was taken, and whose daily work does not even involve such decisions.

British special operations officer Christopher Sherwood may just have felt that way in the aftermath of a drug raid that left one man dead.[1] It was midnight, January 15, 1998, when Sherwood and 21 other officers were summoned to the briefing room of the Lewes police station in East Sussex. They had body armor, special helmets and raid vests (sleeveless vests with two-way radios built in). They might need to immobilize somebody tonight, the briefing began. A raid was mounted, and Sussex was in need of officers who could shoot. Police intelligence had established that a suspected drugs dealer from Liverpool and his associates were in a block of flats in St Leonards, near Hastings. They were believed to be trying to work their way into the extensive drugs trade on the British south coast, and to have a kilogram of cocaine with them.

One of the men was James Ashley, previously convicted of manslaughter. The other Thomas McCrudden, thought to have stabbed a man before. Both men were briefed as violent and dangerous. And most likely armed. The purpose of the raid was to capture the two men, and to confiscate their contraband.

As usual with intelligence, however, it was incomplete. Where in the block of flats they were going to be was not known. There were no plans over the flats either—all would have to be searched, and as quickly as possible: with speed and surprise.

Equipped with rifles (fitted with flashlights on the barrel) and automatic pistols, and up to 60 rounds of ammunition each, the officers proceeded in convoy to the Hastings buildings. None of them were in any doubt about the threat awaiting them, nor about the uncertainty of the outcome. "You get out on the plot, and you never, never know how it's going to end," one veteran explained later. "Your heart is pounding …."[2]

After quietly unloading, and making their way to the block in the dark, six officers took up positions outside the target building. The other officers were divided into pairs, accompanied by an officer with an "enforcer," capable

of breaking down front doors and other obstacles. Each group was assigned a specific flat to search, where one officer would cover the left side of whatever room they entered, the other the right side.

"You know you are walking into danger," commented another officer later. "You know you may be in a situation where you have to kill or be killed. It's a hell of a responsibility."[3]

Christopher Sherwood was one of the officers who went to Flat 6. He was 30 years old, and for carrying that "hell of a responsibility," he was getting paid £20,000 per year (about US$35,000 US per year at that time).

The door went down under the impact of the enforcer and Sherwood veered into his half of the room. Peering through his gun sight into the dark, he could make out a man running toward him, one arm outstretched. The officer's time was running out quickly. Less than a second to decide what to do—the figure in the dark did not respond, did not stop. Less than two feet to go. This had to be Ashley. Or McCrudden. And armed. Violent, dangerous. Capable of killing. And now probably desperate.

Sherwood fired.

Even if there had been time for the thought (which there almost certainly was not), Sherwood would rather be alive and accountable than dead. Most, if not all, officers would.

The bullet ripped into the gloaming assailant, knocking him backward off his feet. Sherwood immediately bent down, found the wound, tried to staunch it, searched for the weapon. Where was it?

Screaming started. The lights came on. A woman appeared out of the bedroom and found Sherwood bent over a man flat on the ground—Ashley.

Ashley, splayed on his back and bleeding, was stark naked. He was unarmed. And soon dead, very soon. It was determined later that Sherwood's bullet had entered Ashley's body at the shoulder but deflected off the collarbone and gone straight into the heart, exiting through the ribcage. Ashley had died instantly.

Whenever a police officer fired a fatal shot, an investigation is started automatically. It was in this case. Officers from the Kent police force were appointed to investigate. They found systemic failure in the Sussex force, including concocted intelligence, bad planning, misapplication of raid techniques, and a wrong focus on small-time crooks. Kent accused Sussex of a "complete corporate failure" in researching, planning and executing the raid.

Sherwood, devastated that he had killed an unarmed man, was interviewed for four days. He maintained that, given the knowledge available to him at the time, he had acted in self-defense. Not long thereafter, however, Sherwood read that the investigator had prepared reports for the Crown Prosecution Service

and the Director of Public Prosecutions, though it was added that nobody knew at that point whether criminal charges were going to be brought.

They were. A year and a half after his shot in the dark, Sherwood was charged with murder.

"Why should anyone want to risk their career or their liberty," an ex-firearms officer reflected, "if, when something goes wrong, they are treated like a criminal themselves?"[4]

The "complete corporate failure" that had sent Sherwood into the building, together with other officers loaded with 1,200 rounds of ammunition, faced consequences too (such as admonition letters). Nothing as serious as murder charges.

A review of the armed raid by a superintendent from the National Firearms School had said that there had been no need for firearms to be used. Another superintendent, with responsibility for the guidelines for the use of firearms, said that this case had not met the requirements. The tactic for searching the flats, known as "Bermuda," was also extremely risky. "Bermuda" was originally designed for rescuing hostages from imminent execution. Sussex Police claimed that their inspiration for using "Bermuda" for arresting suspects came from the Royal Ulster Constabulary in Northern Ireland. The RUC denied it. Sussex Police's own memos had warned as early as 1992 that "risk factors are high and, as such, it should only be considered as a last resort." Their specialist tactical adviser had been warned by the head of the police's National Firearms School that "Bermuda" was too dangerous for such circumstances.[5]

Meanwhile, the enquiry discovered that there had been meetings quiet between senior officers and some of those involved in the shooting that had been kept. After those discoveries, the Kent enquiry stopped cooperating with the Chief Constable of Sussex, and informed the Police Complaints Authority that they suspected a cover-up. Sussex countered that Kent was bullying and incompetent. The Hampshire police force then took over the investigation instead. Yet in this defensive finger-pointing aftermath, nothing stood out as much as Sherwood's murder charge.

The murder charge may not have been connected only to this case. The Police Complaints Authority was facing a national outcry about the apparent impunity with which officers could get away with shooting people. In the previous ten years, police in England and Wales had shot 41 unarmed people, killing 15 of them. No police officer had ever been convicted of a criminal offence, and most involved were not even prosecuted.[6] In this case, it seemed as if Sherwood was to be a sacrificial lamb, the scapegoat, so that it would be obvious that the police force was finally doing something about the problem.[7]

After a year and a half of mulling over a split-second decision, people who had not been there to share the tense, menacing moments in the dangerous dark, opined that Chris should have made a different decision. Not necessarily because of Sherwood, but because of all the other officers and previous incidents, the public image of the police service, and the pressure this put on its superiors.[8]

It would not have been the first time that a single individual was made to carry the moral and explanatory load of a system failure. It would not have been the first time for charges against that individual to be about protecting much larger interests. Many cases in this book point in similar directions. What it raises may seem troubling. These sacrifices violate Aristotelian principles of justice that have underpinned Western society for millennia. Such justice particularly means refraining from *pleonexia*, that is, from gaining some advantage for oneself by blaming another, by denying another person what is due to her or him, by not fulfilling a duty or a promise, or not showing respect, or by destroying somebody else's freedom, reputation.[9] Holding back from *pleonexia* puts an enormous ethical responsibility on organizational leadership. Nothing can seem more compelling in the wake of a highly public incident or accident than to find a local explanation that can be blamed, suspended, charged, convicted. The problem and the pressure it generates is then simply packed off, loaded onto somebody who can leave, slide out of view, or get locked up—taking the problem along. But it is not ethical, and it is not likely to be productive for the organization and its future safety and justness.

The Injustice in Justice

Pursuing justice in court will always produce truths and lies, losers and winners (and more losers). Even if a scapegoat eventually gets exonerated, interests will have been lined up against each other in a way that makes any kind of reconciliation really difficult. By treating error as a crime, we ensure that there *always* will be losers, whatever the outcome of a trial. Since it divides people into groups of adversaries, we guarantee that there will always be injustice in justice, whether the practitioner gets off the hook or not. Common interests dissipate, trust is violated, shared values are trampled or ignored, relationships become or stay messed up.

On May 2, 2001, Chris Sherwood was cleared of any blame at the Old Bailey in London when the judge, Mrs Justice Rafferty, instructed the jury to find him not guilty. There was no evidence of any intention to kill, she argued, other than that he had fired in self-defense. Justice prevailed, at the same time

that injustice prevailed. What about the cover-up during the aftermath? And what about the victim's family or his girlfriend, the woman who had stumbled upon Ashley's still-warm corpse? No recourse for them, no justice, popular opinion found. Many commentators cried foul.[10] The reaction meant that the pressure for the Police Complaints Authority to show its teeth, and for others to charge and convict, would probably remain.

No just culture—no peace for those who do the work every day.

Notes

1 Seed, G. and Palmer, A. (1999). A shot in the dark. London: *The Sunday Telegraph*, January 24, 23.
2 Ibid.
3 Ibid.
4 Ibid.
5 Bond, P. (2001). *British police acquitted of wrongdoing after shooting unarmed man*, June 1. http://www.wsws.org/articles/2001/jun2001/kil-j01.shtml.
6 Ibid.
7 Seed and Palmer, op. cit.
8 Ibid.
9 Rawls, J. (2003). *A theory of justice*. Cambridge, MA: Harvard University Press.
10 Bond, op. cit.

This page intentionally left blank

Index