FINDING HUMAN PERFORMANCE PATTERNS IN SAFETY DATA

José A. Blanco
Laurentian University
Centre for Research in Human Development
MSc., Ph.D.

John H. Lewko
Laurentian University
Director of the Centre for Research in Human Development
M.A., Ph.D.

Abstract

This paper was born from the observation that preventable injuries become invisible to managers and workers. It describes a drum model of failure prevention to show how preventable failures become invisible, and what has to be done to see them again. Two industrial cases provided the context to explore these issues from the perspective of managers and workers.

The model couples a drum, representing the technology, with a sorting system representing the management system. The drum contains all of the failures and successes (represented by red and green balls respectively) that the particular technology can produce, and releases them one by one. The sorting system is supposed to accept all the green balls (successes) back into to the drum and reject the red (failures), but it is capable of ignoring some of the red balls and let them back into the drum. These are preventable failures that are ignored.

The model suggests that ignored preventable failures linger: long term failure rates are tied to present failure rates by a conventional decay curve. Decay curves are congruent with available long-term data for a number of industries. The model could be subjected to mathematical analysis, but the focus of this discussion is managerial and not mathematical.

Preventable failures that are ignored become invisible and recur, but they leave traces and patterns. These patterns can lead back to the cause of the failures and thus provide a platform for management to engage those who have to make the preventive changes.

The model links failure rates separately to the prevention capability of the technology and to the ability of the management systems to operate the technology close to its prevention capability. This leads to the conclusion that management systems need to be upgraded before technology; a sharp focus on prevention brings safety and efficiency together.

Introduction

This work deals with finding opportunities for prevention in the patterns that lie hidden within safety data, and using them to slow down the rate of injury and loss. Along the way, we will propose a simple prevention model and borrow from safety practitioners, philosophers and writers to describe the power of the model as a tool to explore the prevention space.

The Value of Models

A model is generally useful if it provides a frame of reference to help us interpret and measure what we observe in the real world. A prevention model that interacts well with real world evidence offers new possibilities for exploring the prevention space, and could eventually lead us towards prediction [a]. Models can be of many types.

One type of model we use often is physical. Models such as crash test dummies or test planes have provided insights into safety gaps that were not otherwise apparent. Another type of model is verbal, and this conference is a good example of how we use verbal models to explore topics and describe results. Yet another type is mathematical. We use mathematical models for many purposes; they appear to be more complex than physical models but are in the end far simpler, as for example flight simulators versus airliners, or CFD versus wind tunnels.

A mathematical model that approximately describes safety data over long periods of time could give prevention a boost. It could assist us in gathering data and interpreting what we measure; it could suggest opportunities for worthy action. At the very least, it could guide the discussions that are needed to align the verbal models of prevention, and thus eliminate a frequent source of misunderstanding between different fields or even within the same company. Because a mathematical model remains the same regardless of language or circumstance, it would widen the range of users who can benefit from its use. A useful model would make prevention more accessible to practitioners, researchers and regulators.

A Simple Prevention Model: the "Result Drum"

We seem to accept the notion that mining, manufacturing, driving, or public health, are safer today than they were in the past. We seem to believe that, over the years, we have identified the causes of some of our earlier failures and learned to prevent them or to neutralize their effects. Machine guards, shatter-proof glass, padded dashboards, remote handling, treated water, pasteurized milk, vaccination are just a few examples of successful intervention following failure. All along, the evidence suggests that we have been learning from our failures and preventing their recurrence or softening their impact. A hopeful prevention model would have to approximate that record of progress if it is to gain acceptance. Let us look at a simple prevention model and see how its predictions match real world data. (Figure 1).

Figure 1. The "Result Drum": A Simple Prevention Model [b,c]

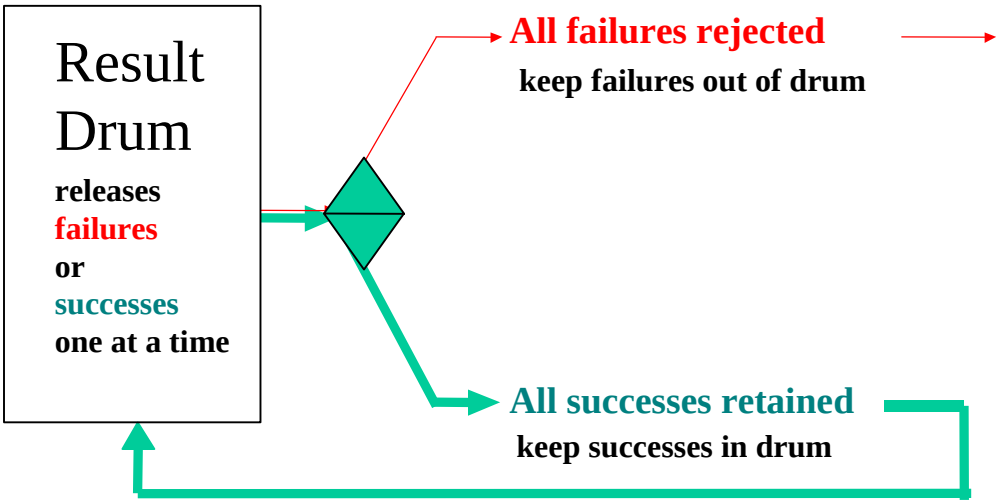

The model we are proposing is a "result drum" that represents a technology and simulates the way the technology can progress from the time it is adopted until it is abandoned. Imagine that the drum contains all of the possible results of that technology, successes as well as failures, and it is not biased in any way. It contains a small number of red balls, representing failures, and a much, much larger number of green ones, representing acceptable results. There should be many more acceptable results (green) than failures (red), or that technology is doomed to fail.

At time zero, as the technology got under way, we started to draw balls from the "result drum", one by one, at random. They were red or green according to their relative concentrations in the drum. Every time we drew a green ball, we accepted the result and returned the ball to the drum to be drawn another time: we wanted to get that result again. Ideally, every time we got a red ball we rejected it: we wanted to prevent that result from happening again. So, we continued drawing balls, returning the green and rejecting the red. Obviously, as time went on, the number of red balls in the drum dropped, and the chance of drawing a red ball decreased. The concentration of red balls in the drum at any time would be described by a classical decay curve. In prevention language, the drum model suggests that the rate of failures would decrease with time following a classical decay curve such as is shown on Figure 2. The rate of decay would be fastest as we install the technology, slow down with time as we learn to prevent failures, and, in the long term, reach an asymptotic value, the long-term residual failure rate. The model says that this value could be as close to zero as we know how to make it.

Figure 2. A Classical Decay Curve

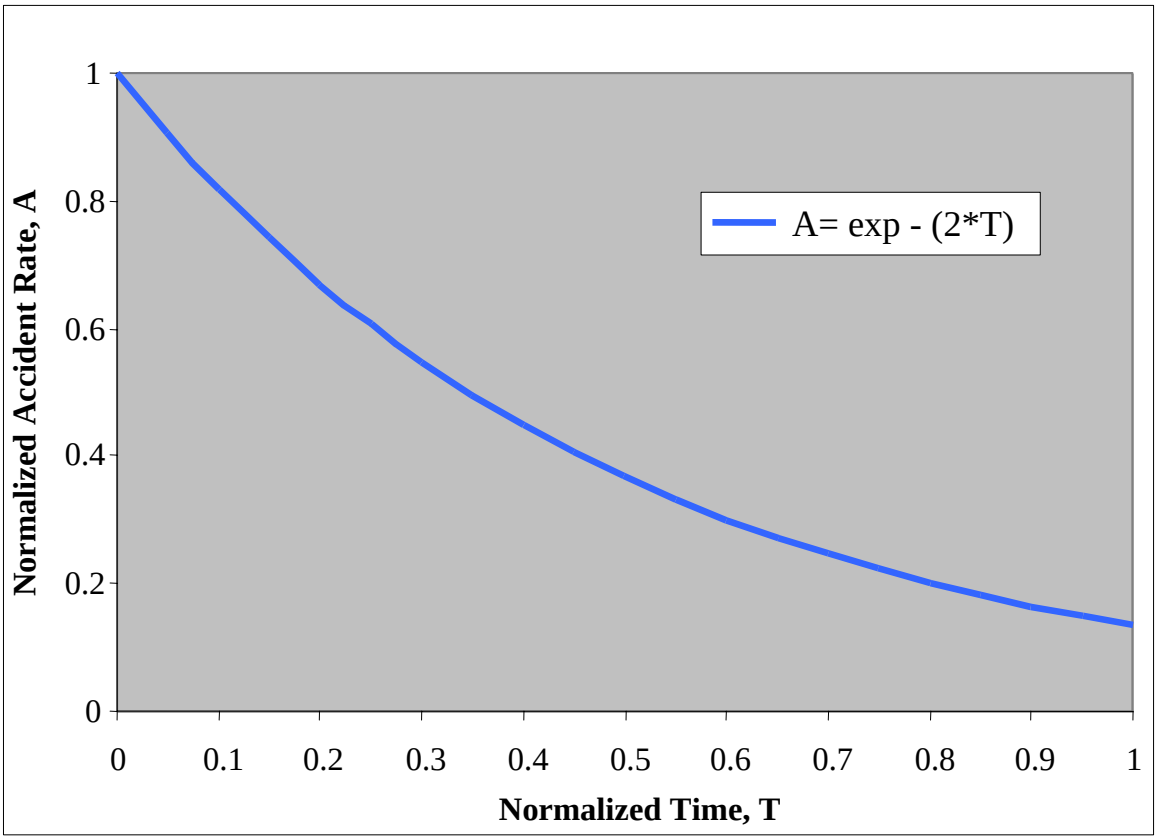

Decay curves can be used to picture the prevention record of a company and compare it with others within the industry. They can also be used to picture an industry as a constellation of companies, each with its own history of prevention, its own "drum" with different concentrations of red and green balls, each with its own "K" constant, that is, its own rates and asymptotes. The aggregate decay curve for the industry would be a composite of the individual curves for companies in that industry. For simplicity, companies within an industry may be grouped according to similar prevention performance. Figure 3 shows an example of composite decay curve that lumps three quarters of an industry with a high "K", and the remainder with a lower "K". The aggregate for the industry would approximate a decay curve, with its equivalent "K", its own rates and asymptote that could be compared to those in other industries or could provide a measure of the success of prevention within that industry.

Figure 3. Two Decay Curves and One Composite

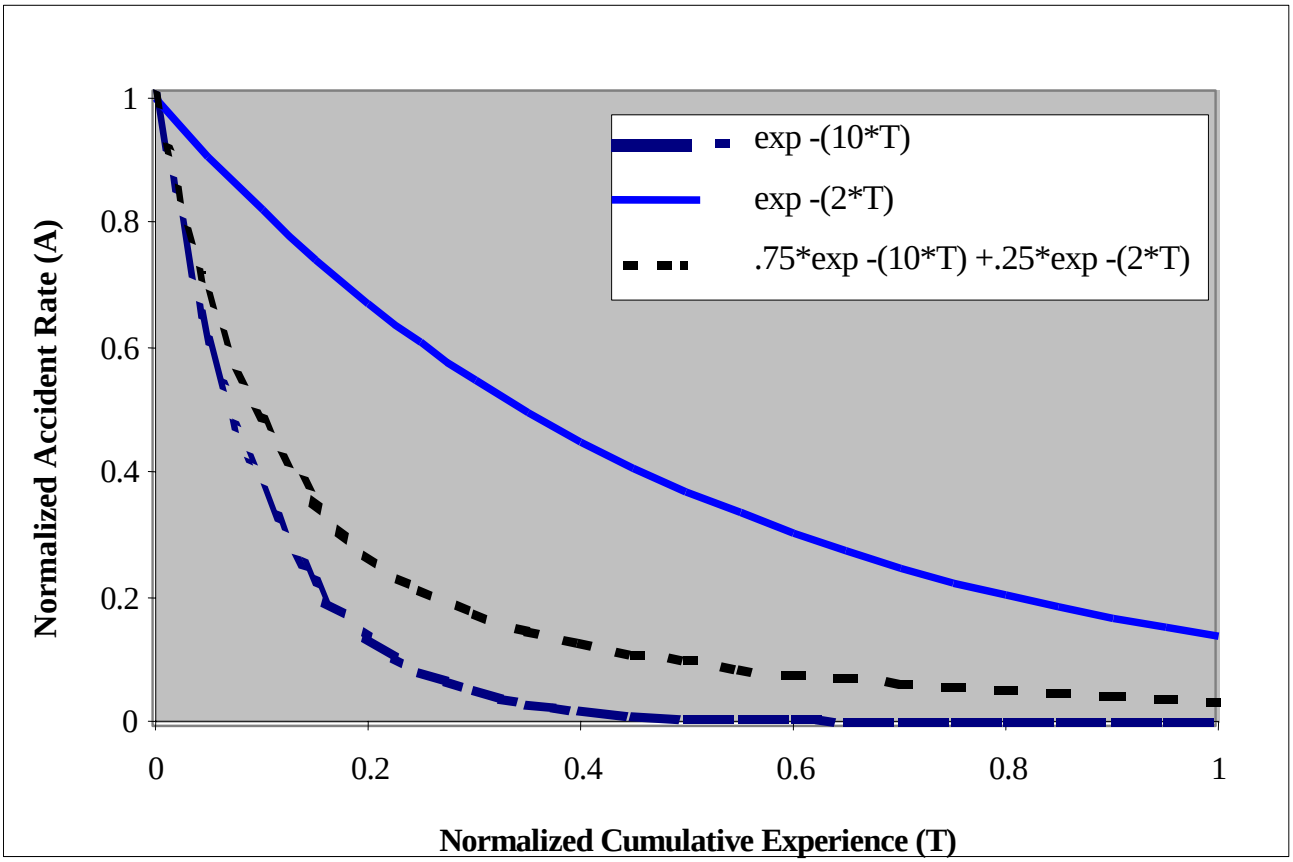

A Test of the "Results Drum" Model [b,c]

Before we continue using this model, we should confirm that it is consistent with the safety information of the last century. Let us see how it fares.
Duffey and Saull[d,e] use failure data collected from many sources to observe how the failure rate changes with time. They observed that the failure rate has been decreasing all along. Figures 4, 5 and 6 are from their published work and are shown with their permission:

Figure 4. 25 Years of Shipping Losses [d,e]

Figure 5. 12 Years of North American Aviation [d,e]

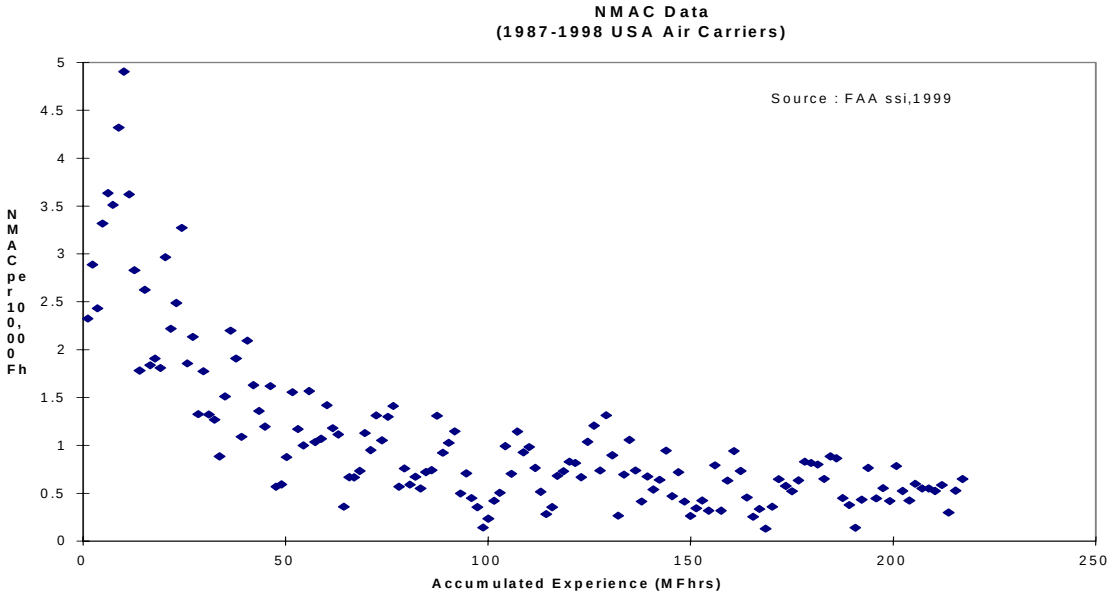

Figure 6. 48 Years of USA Mining [d,e]

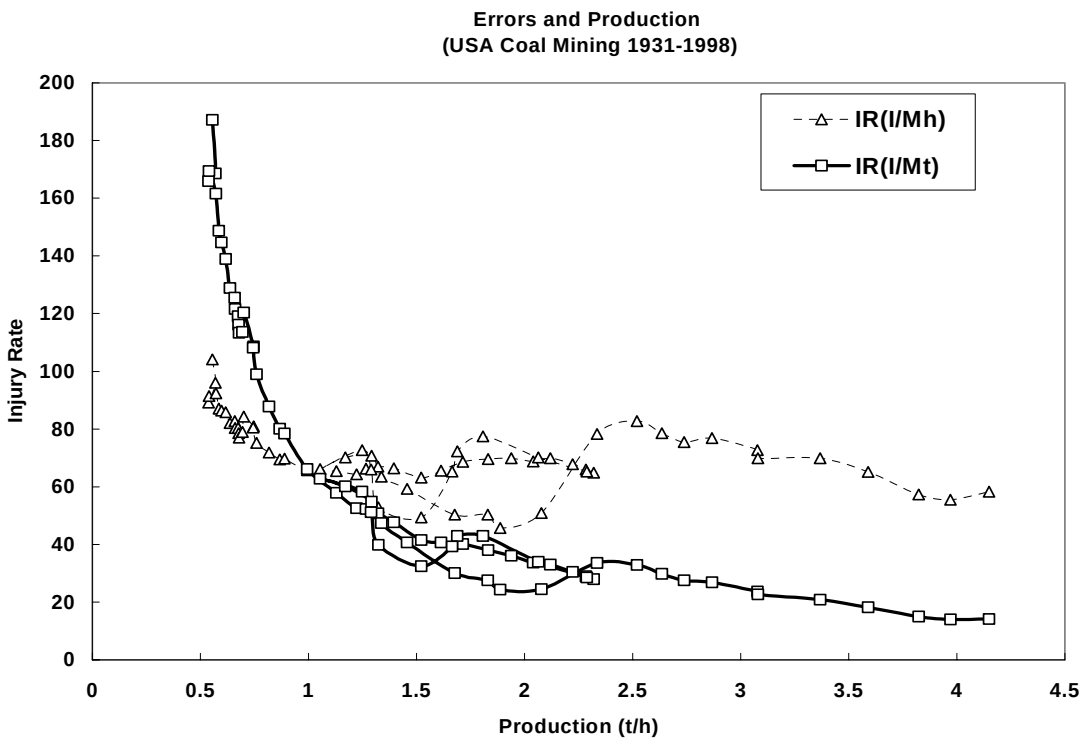

These are decay curves, although you may have noted that the authors do not use "time" as the independent variable. More precisely, Duffey and Saull[d,e] state that, in the industries whose data they have reviewed, the number of failures drops in proportion to the instantaneous failure as the cumulative experience increases. This leads them to propose a classical decay curve using cumulative experience instead of calendar time as the independent variable. "Cumulative experience" allows for variation in the ways that different companies use their technology and approach prevention, and is not hindered by changes in the industry's context. Calendar time, on the other hand, ignores that exposure to injury per production unit is changing all the time as a result of changes in efficiency. For example, car making time per unit is way down, and airliner passenger loads are way up. The "drum" model and the concept of cumulative experience suggested by Duffey and Saull[d,e] both avoid being trapped by the limitations of calendar time. They are not caught either by the need to distinguish between safety and efficiency as drivers of failure prevention.

Duffey and Saull[d,e] observed that the historical failure data follow a time-related decay curve, and the "result drum" model meets that test. Decay curves and their mathematics have been studied in detail and provide useful models in fields such as epidemy prevention, equipment maintenance, reaction kinetics, and radiation. They open the field of industrial failure prevention to new ideas, and should be very useful in the prevention of injury.

Adding Human Factors to the "Results Drum" Model: Voltaire's[f] Caution

Of course, things are not as simple as the model of Figure 1 would suggest: our working model needs to be able to account for human factors. That model assumed a technology within a system of simple prevention decisions: we can tell green from red, and green is in, red is out. But sometimes we ignore failures, either because we do not recognize them, or because we prefer to live with the hazard. Safety daltonism, real or fake. Voltaire[f] warned us about these human complications 250 years ago when he stated: "ce que nous appelons hasard n'est, ne peut être, que l'effet connu d'une cause qu'on ignore". The known effect of a cause that we ignore.

Figure 7 shows a "result drum" that incorporates the option of "ignoring" some of the failures that could be prevented.

Figure 7. A Model for Prevention with Two Decision Levels

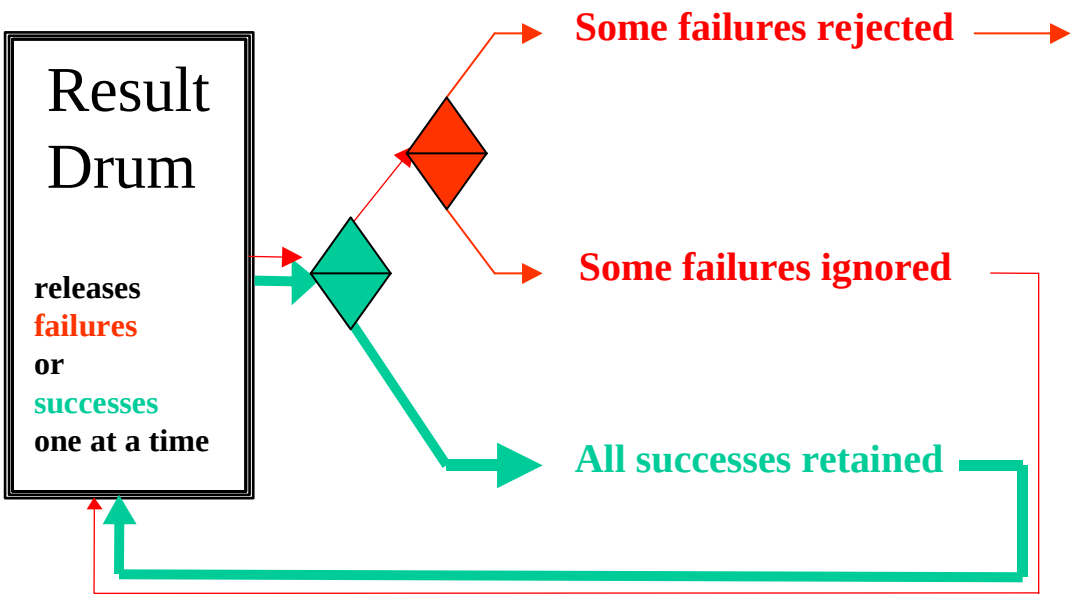

The concentration of red balls in the model of Figure 7 would still follow a decay curve, so the model remains adequate. Both the asymptote (the long-term residual failure rate) and the short-term failure rate would be greater than in the base case because some of the red balls that ought to have been rejected were instead accepted and returned to the drum to be drawn another day. In prevention language, this partial acceptance of red balls to the drum means not using the technology to its full prevention potential. Individual workers, crews, companies or industries that are operating a technology below its potential are causing preventable failures that are due to the system and not the technology. Moreover, the model forecasts that, if they do not change, they will entrench the failures and create a long-term future with more failures.

The model also says that tuning up the technology and operating it closer to its prevention potential can reverse the trend. That would produce intrinsically inexpensive prevention, but it is not easy. First we need Voltaire's[f] words to flush it out into the open.

Some Implications of "Ignoring" Causes: Insurance, Prevention as a Budget Variance

When Voltaire[f] wrote "cause qu'on ignore", he could have meant "causes that we do not know" at this time but that we may know later, or "causes that we could know but will continue to ignore". Voltaire[f] seems to suggest that we will see as "hasards" all the failures whose causes we "ignore" regardless of whether the ignorance is real or not.

The implications are profound. We see "hasards" as random, unconnected events that we can not prevent, and thus we seek ways to cope with their effects and put them out of mind. To cope with the effects, we would hedge the risks, for example by buying coverage such as WCB for those failures that are relatively frequent but not too large, maybe liability insurance for the ones that are infrequent but can be very large. Eventually, these failures become invisible as failures and are seen instead as a cost of doing business. Hazard. Fate. We put them out of mind by budgeting a historical level of failures, and, ironically, we entrench failures and expose prevention as a budget variance. You may wish to read Forest L. Reinhardt[g] for a detailed analysis of risk sharing and the surrounding management issues in the context of environmental failure prevention.

But all is not lost. Our drum model says that when we ignore the causes of failures that could have been prevented, we are treating them as if they were inevitable; they will blend into the background and become invisible. But the model also says that the preventable failures that were not prevented will recur, and when they do, they will leave some trace; repeated traces will form patterns.

These traces and patterns are at first invisible to the incumbents, but a careful search can reveal them. This means that the historical failure record of any operation is likely to contain traces from preventable failures that were ignored. This could be as simple as a higher frequency of injury among baghandlers at one airport than another, or more fatal rollovers with a particular combination of tires and suspension than another. Some of these traces may form patterns, and these patterns may reveal opportunities for prevention, provided they are made visible to the incumbents. If the incumbents can not see them, prevention could be delayed until economic or political conditions jolted the incumbents into action. But it is possible to help the incumbents see the hidden patterns in their failure data.

A Case History

I will share with you an example that I know well. I was the manager of a large plant in Canada employing about 1500 people. We had recovered from two 9-month-long work stoppages, major fuel cost increases and a sharp fall in selling prices in the previous six years. We had made major gains in efficiency, productivity, costs and safety. We were feeling tired, but good. A visitor stopped by my office to chat about the improvements we had achieved. We chatted for a while, then he suggested that the rate of injury could not drop any lower because I was satisfied with things the way they were. I asked if he meant that I was holding progress back, and he said, more or less: "some injuries could be prevented but you do not allow that; you seem satisfied".

Discovering Invisible Patterns

I was stunned; we had reduced injury rates by almost a third, and I felt sure he was wrong. But I could not prove it. To prove him wrong, we assembled all the injury reports of the previous 14 years that involved employees still working at the time of the search; we thought 14 years was long enough to show trends. There had been 656 injuries involving 372 employees, and 436 employees had remained injury free.

The data were startling. They were skewed[h.] Clearly the injuries were not distributed evenly across the population. Mechanics got injured more often than operators; electricians more than mechanics. Thirty employees had received 5 lost time injuries each during that period, almost twelve times the average. When seen in batches, many of these injuries seemed preventable, but we had not noticed that before, and were still urging the employees to be careful.

I was the manager, so I can tell you that I did not know; the superintendents and supervisors did not know. You may ask how that is possible, but I can assure you that we did not know, even though we often complained that our WCB costs were excessive. But once we knew, we could not continue "ignoring" preventable failures.

On the Way to Recovery: a Robust "Manager's Prevention Intent"

I still had to contend with the organizational mix of power and accountability. My visitor had pointed out that I was ignoring prevention opportunities, and I had to face the likelihood of the same happening at each level in the organization. Even the workers who had received multiple injuries did not know that multiple injuries should be an exception. All of us had been accepting these injuries as inevitable, as if they were in the nature of the job.

What the evidence showed and discussions confirmed was that management's prevention intent had not been clear and firm enough, and that I had ignored its dilution at every step in the organization that I was supposed to lead and manage.

Penalties, incentives, training did not improve prevention; obviously, having direct power was not sufficient. The invisible remained invisible to the actors, and the data patterns continued, except that now I was aware of them. Management prevention intent had to be made clear, firm and functioning.

Sharing the Prevention Intent: Aligning the Verbal Models of Prevention

The evidence about failures and their skewed distribution was solid, but management knowing about the skewed injury patterns was not enough. Some process, procedure or policy would have to change somewhere in the organization before the preventable failures could be prevented or softened. That meant that the people who had to make these changes had to interpret the evidence of preventable failure the same way management did. They also had to have enough confidence and support to make the changes. If all they could see were the costs of making changes but not the benefits, they would ignore the evidence. We know this from hospitals that continue delivering inadequate drug combinations after the evidence of failure has been circulated. We watch how long it takes for tire and car-makers to acknowledge a fatal mix of loads, suspensions and tires. Obviously, this trait is not industry specific; we often wait till the evidence forces us to act.

Memos and warnings, even threats, don't do the job. What is needed is a model of prevention that is broadly understood and accepted.

Jorge Luis Borges,[i] a modern Argentine writer, tells us why: "las palabras son simbolos que postulan una memoria compartida".

Without a shared memory, he suggests, words are disconnected symbols and there may be no communication, let alone persuasion. To the manager whose memos, pleas and warnings have not worked, Quine[j] points out the perils of assuming that assent or consent without objective confirmation could possibly mean a shared understanding.

The prevention of failure has to start with management setting up appropriate expectations and making sure they do not get watered down at every step. It has to involve those who have to do things differently. It has to involve those exposed to injury because they are the ones who need to learn the causes of failures and develop and adopt the preventive practices. In my case, it started with me being clear and firm about my prevention intent, explaining and adapting it to fit the field examples until I got others to understand the meaning of the intent and to adopt it as if it were theirs.

Persuading the Actors

As I remember, we started to make definite progress when we understood that prevention was not an issue to be "managed". Management had the data but we were not the actors. We offered the data to the groups who had the injuries and we persuaded them that all we could do was to help, only they could solve. And solve they did after they accepted the challenge and saw that management and supervision would support the process, not drive it. Electricians and instrument technicians, after much emotional discussion, discovered the causes of their frequent back injuries. They changed procedures, bought carts with rollers for moving heavy equipment and tools, set up a retraining program and introduced a fast-response, fast-recovery protocol for back injuries. And the heavy-duty mechanics discovered that a six-foot wrench in close quarters was the cause of repeat back and extremity injuries, and they replaced it with an electric torque wrench and retrained themselves. And crews that used to suffer bruises, sprains and the odd fracture, discovered from the cumulative data that most of their injuries arose from slips around the night shift following sunny days on sub-zero weather: roof thawing, walkway freezing. They rearranged their schedules to take charge of sanding the walkways to suit their safety needs. Management ensured the supply of sand and shovels. And we watched the disabling injury rate and the WCB costs drop by half, and the total lost days drop from about 400 to about 100 days per 100 employees per year.

It is worth repeating. We eventually learned that those who were at risk of injury were the ones who had to be able to read the signals in the workplace. Once that the people at risk accepted that there was a problem, they offered practical solutions. Management and supervision had to learn to use the examples of field successes and failures to refine and share our intent about safety and our support for prevention. Step by step we came to a shared understanding. The results are shown in Figure 8. Note that efficiency continued to improve (annual tonnage remained essentially constant) as safety improved year after year.

Figure 8. Normalized cumulative data, 1980-2000

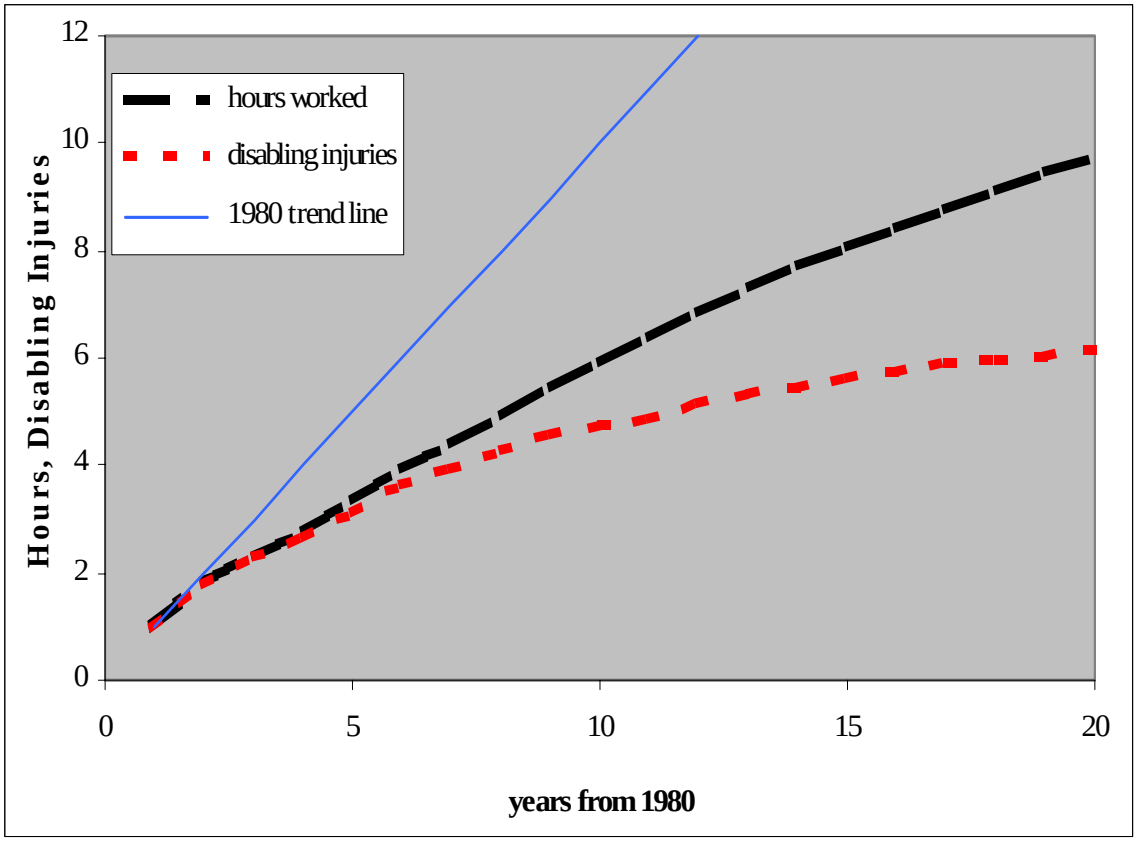

A Shortcut to Results: the "Drum" Model, Voltaire[f], Borges[i], and Quine[j]

But the point of my talk is that it does not have to be done the hard way. What does the drum model say about the red balls that we ignore? It says that by returning them to the drum we ensure that they will be drawn again: failures that have not been prevented will eventually reoccur. Voltaire[f] warns us that when we ignore preventable failures they will seem to be hazards, independent random failures that we can not cope with. But we know that failures leave marks.

This means that the failure record of any operation probably contains marks from preventable failures that were ignored, and the marks are likely to form patterns, and the patterns can reveal opportunities for prevention. Borges[i] tell us that we have to make sure that those who can do something to prevent the failures see them and their causes, because if they don't, they can not prevent them. Quine[j] want us to use objective examples to confirm that we share our understanding. And this requires the manager to make sure that the prevention intent fits the internal evidence and the trends within the industry, and that it is shared at every level in the organization.

Examples of Pattern Analysis and Prevention from Other Fields

You may wish to read about D. Kim Rossmo's success in Vancouver using geographic patterns of crimes[k] or about how Jack Maple[l] used crime pattern analysis to guide crime prevention in New York City and in New Orleans. The principles of prevention are the same: failure leaves traces and repeat failures leave patterns. That is why we need accurate, timely intelligence about failures and relentless assessment and follow-up, whether the failures are crime or injury.

Little and Mount[m] describe how geographic patterns guided Dr. John Snow to act on a cholera outbreak in Victorian London, and how higher death rates in maternity wards run by men compared to those run by women guided Dr. Ignaz Semmelweiss in 1850 Vienna to introduce basic hygiene and save many lives. Paul Ewald[n] provides a modern and very powerful take on issues of models, patterns, and prevention. But just the same, Dr Semmelweiss dismissal and Professor Ewald's writings show that even a good model with a proven record of prevention is not enough: the incumbents must be persuaded of its validity and benefits. And power helps: Dr. Snow got away with removing the handle in the suspect well to save lives and prove his point. Prevention regulators, researchers and implementers need to keep these examples and cautions in mind.

A Second Case History: Finding Opportunities for Prevention in Safety Data

The subject of this case is a mine with about 1000 people, also in Canada. This is a highly competent outfit with a leading safety performance within their industry. The safety data covers 54 months up to date, and the purpose of the work is to find failure patterns and use them to help the supervisors to focus on injury prevention. The fundamental difference between the two cases is that in the first one the manager, who was the driver, had the power to remove the pump handle and make it stick, whereas in this case the researchers have even less than Semmelweiss.

The Safety Data

We had planned to use the injury database of the official safety organization. It was clear that those who prepared the incident reports and those who worked with the database did not share a common understanding of prevention. They also used different codes and different measures of time worked. The incident report is a menu-driven form with many questions, and the resulting database is huge and opaque. We condensed the information using conventional tools. Time charts to track trends, Mindmapping to simplify a formidable amount of data into useful Pareto pictures. Bar charts to identify workgroups or issues with unusual results. Poisson statistics to validate data.

We found evidence of human factors. Menu-driven forms luring supervisors into unusual coding. The links between assessed risk and recommended prevention weak. Most of the preventive options listed in the menu were not used. The form seemed to overshadow prevention.

We found little evidence of feedback: each injury seemed to stand alone. As the drum model predicted, many failures had become invisible, and there were traces and patterns that could have lead to prevention but were invisible even though the aggregate data showed incongruous groupings. Almost a replica of Case 1 except for better rates.

Bringing Our Partners Along

Once we were reasonably certain that the patterns we found were statistically significant, we asked our partners for help in interpreting them. Our questions resonated with the managers and union safety leaders. We took care to keep the discussions open and non-threatening. They have been willing to imagine reasons why employees might not respond to hazards they ought to know about, or why they might propose accident prevention steps that seem unfocused or ineffective, or why they might not see the preventive actions they recommend through to completion. They may even consider redesigning the investigation form so that it fits their failure experience and stays focused on prevention as it provides the necessary formal records.

Moving New Possibilities of Prevention Into the Workplace

Case 2 differs from Case 1 in that in Case 1 the manager was involved in the research and could influence the crews whereas in Case 2 the researchers have no formal influence over the manager or the crews.

The absence of direct influence is a handicap because improving prevention means changing some practices, procedures or policies somewhere in the organization. The manager has to believe in the logic of any prevention initiative (all benefits exceed all costs) and has to show support, or the initiative will wither. Researchers (and, to some extent, regulators) need to consider ”benefits” and acknowledge “costs” from the perspective of the managers.

Moving Along the Manager’s Line of Authority

Changing practices is difficult at any time. Imagine then what would happen to a suggestion to change a practice if the supervisors, the union, or the workers thought that they would have to make more effort to get less benefit. The supervisors know how hard it is to erase accepted practices and implementing new ones, especially because these come without guarantee. They would have to persuade the workers that the proposed practices would bring to the crew more benefits than headaches. That could only be done by involving the supervisors and their crews in interpreting the problems and defining the solutions from the start. This in turn requires the manager to make a personal investment. Rational managers and supervisors who do not see patterns in the failures, and benefits in their prevention, can not be blamed for thinking it is best to ignore the researchers.

Moving Across Other Lines of Authority

It should be obvious that implementing new processes or practices could get even testier if the costs show up in one manager’s area and the benefits show up in another’s. Implementing such changes could require personal investments from a higher level of management.

Things get even more complicated if implementing the new process or practice would conflict with an existing policy. Then, the corporate managers have to believe that it is in their best interest to let the line managers in their watch tinker with policies or change processes or procedures. These difficulties are real. Huls[0] uses data from the mining industry to highlight the difficulties of adopting innovation even when strong economic benefits are at stake.

The researchers need to take into account these organizational complications. It is much easier to find ways to work with managers and supervisors if the environment favours prevention. The WCB, regulators, other governmental agencies and insurance companies may be

able through their actions to create and sustain a favourable environment for prevention strategies.

Other insights from the drum model

Shifting to a Higher "K": Getting Closer to the Technology 's Potential

The "drum" model suggests that long-term residual failure rates are linked to the short-term failure rates, which are sensitive to practices, procedures and policies. Thus, it turns out that, within the range that their choice of technology permits, a company's prevention management systems and intent determine present and future failure rates. Not a revelation, but often overlooked.

The drum model suggests that a company that needs to improve its failure rates has two rational options: either improve their "K", or adopt a technology with a lower failure capability, that is, a higher intrinsic "K". Interestingly, learning to improve prevention is common to both options.

A company trying to raise its "K" can start by searching its failure record for traces of any preventable failures it had been ignoring. This requires looking at habitual things in a new light; it requires determination. Learning about failures and their prevention is the first step to sharing the new knowledge and creating the mix of persuasion, incentives and disincentives that encourage adoption of the knowledge. Some examples were discussed in the context of Cases 1 and 2.

Switching From Old to New Technologies

A company that wishes to replace its technology with one that has a lower failure capability also needs to search its record for traces of preventable failures it had been ignoring. In the drum model, a new technology means a new drum with its full complement of red balls and green balls. To get the best chance of a good prevention payback from the new technology, the new drum should not contain any red balls that represent preventable failures that are already recurring within the company or within the industry.

Adopting New Systems from Other Industries

Similar considerations apply to industries. Duffey and Saull [b,c] show that different industries such as airlines and shipping have very different long-term residual failure rates. The drum model suggests that these long-term failure rates are linked to their present failure rates and are therefore affected not just by their choice of technology, but also by the way their management systems use that technology.

So, even if an industry seeking to lower its long-term residual failure rate can not adopt technologies from unrelated industries, they may be able to achieve the lower rates in other ways. They could raise their "K" by searching their failure records for traces and patterns and proceed to improve prevention as we discussed above. Or they may try to adopt management systems that are successful at preventing failure in other industries and get a yet higher "K".

Adapting geographical and other pattern management systems to police work [j] and copying incident investigation and management systems from the airline industry are examples of grafting failure avoidance techniques from systems proven in other industries. In model language, industries that can not import technologies with a higher "K" may still be able to reach a lower long-term residual failure rate by adapting from other industries management systems that can bring them closer to their prevention potential.

Safety and Efficiency

Some people sense that safety and efficiency are separate measures of human performance competing in a zero sum game, one at the expense of the other, whereas others see them as two inseparable measures of human performance. As we have seen, both the drum model and cumulative experience integrate safety and efficiency. I confess my bias: I see safety and efficiency as the two faces of Janus, and I am happy to see the beginnings of a theoretical bridge bringing them together.

Figure 9. Janus [b,c]

Janus:

Roman god of gates and doors, beginnings and endings, is represented with a double-faced head, each looking in opposite directions.

Safety and efficiency, constantly watching in each direction to attain and protect sustainable performance.

REFERENCES

a. Bross, I.D.J., 1953, *Design for Decision*, Macmillan Company, pp.161-182

b. Blanco, J., November, 2000, *Safety Management Systems*, Safety Through Partnership, NAV CANADA/ Transport Canada, Toronto

c. Blanco, J.A., February, 2001, *Safety Management Systems*, The Sixth Northamerican Trilateral Safety and Security Meeting, Ixtapa, Mexico

d. Duffey, R.B., Saull, J.W., September, 2001, *Errors in Technological Systems*, paper for the World Congress Safety of Modern Technical Systems, TUV Saarland Foundation, Saarbruecken, Germany

e. Duffey, R.B., Saull, J.W., November, 1999, *On a Minimum Error Rate in Complex Technological Systems*, paper for the Joint Conference on Enhancing Safety in the 21st Century, Rio deJaneiro, Brazil

f. Voltaire, 1865, *Dictionaire Philosophique: Atomes*, Librairie de L. Hachette et Cie, Boul. Saint Germain 77, Paris, Tome douzieme, p. 361

g. Reinhardt, F.L., 2000, *Down to Earth: Applying Business Principles to Environmental Management*, Harvard Business School Press, Boston, Massachusetts, USA

h. Blanco, J.A., March, 2000, *The Return On Investment Of Safety Management*, paper for the joint Symposium on Human Factors in Aviation Maintenance, Vancouver, Canada

i. Borges, J.L., 1999, *El Congreso*, el Libro de Arena, Biblioteca Borges, Alianza Editorial,S.A., Madrid, Spain, p. 53

j. Quine, W.V., 1987, *Quiddities; An Intermittently Philosophical Dictionary*, The Belknap Press of Harvard University Press, Cambridge, Massachusetts, USA

k. Atkinson, B., August, 1996, *Wheredunit: Cybercops make house calls*, The Globe and Mail, Toronto, Canada, p. D8

l. Maple, J., Mitchell, C., 1999, *The Crime Fighter; How You Can Make Your Community Crime-free*, Broadway Books, Doubleday, 1540 Broadway, New York, N.Y., USA

m. Little, M., Mount, K., 1999, *Prevention and Early Intervention with Children in Need*, , Ashgate Publishing Limited, Gower House, Croft Road, Aldershot, Hants, England, pp. 15-21

n. Ewald, P.W., 2000, *Plague Time: How Stealth Infections cause Cancers, Heart Disease and Other Deadly Ailment*, The Free Press, aA Division of Simon & Shuster, Inc.,1230 Avenue of the Americas, New York, N.Y., USA

o. Huls, B.J., 1994, *When Innovations Occur: Their Effect on a Large Mining Company*, Innovations in Mineral Processing: Proceedings of the Innovations in Mineral Processing Conference, Laurentian University, Sudbury, Canada