

Special Issue

PROCESS SAFETY INDICATORS

Guest Editors

Andrew Hopkins
A. R. Hill

SRAE 2005

Guest Editor

B. Konic

**safety
science**

Available online at



ScienceDirect

www.sciencedirect.com

CONTENTS

PROCESS SAFETY INDICATORS

Guest Editors

A. Hopkins

A. R. Hale

A. Hale	459	Special Issue on Process Safety Indicators
A. Hopkins	460	Thinking about process safety indicators
L. Allford	466	Process safety indicators: Response to Andrew Hopkins
R. Chaplin	467	Process safety indicators: Response to Andrew Hopkins
S.G. Erikson	468	Performance indicators
M. Hodgkinson	469	Process safety indicators: Response to Andrew Hopkins
B. Ale	470	More thinking about process safety indicators
L.J. Bellamy	472	Process safety indicators: Response to Andrew Hopkins
J. Dyreborg	474	The causal relation between lead and lag indicators
A.I. Glendon	476	Process hazards and analogies: Response to Andrew Hopkins
G. Grote	478	Response to Andrew Hopkins
A. Hale	479	Why safety performance indicators?
L. Harms-Ringdahl	481	Dimensions in safety indicators
P.T.W. Hudson	483	Process indicators: Managing safety by the numbers
U. Kjellén	486	The safety measurement problem revisited
J.-c. Le Coze	490	A taxonomy issue
K. Mearns	491	From reactive to proactive – Can LPIs deliver?
J. Wreathall	493	Leading? Lagging? Whatever!
G.I.J.M. Zwetsloot	495	Prospects and limitations of process safety performance indicators
D.D. Woods	498	Escaping failures of foresight
P. Webb	502	Process safety performance indicators: A contribution to the debate
A. Hopkins	508	Reply to comments

SRAE 2006

Guest Editor

B. Kontic

B. Kontic	511	Innovation and technical progress: Benefit without risk?
M. Padlog	513	Introductory story: Building trust among scientists, authorities and the public
R. Bubbico and E. Salzano	515	Acoustic analysis of blast waves produced by rapid phase transition of LNG released on water
C.D. Gamper and C. Turcanu	522	Can public participation help managing risks from natural hazards?
S. Eiksund	529	A geographical perspective on driving attitudes and behaviour among young adults in urban and rural Norway

Scope of the journal

Safety Science, like its predecessor, the *Journal of Occupational Accidents*, serves as an international medium for research in the science and technology of human safety. It extends from safety of people at work to other spheres, such as transport, leisure and home, as well as every other held of man's hazardous activities.

Safety Science is multidisciplinary. Its contributors and its audience range from psychologists to chemical engineers. The journal covers the physics and engineering of safety; its social, policy and organisational aspects; the management of hazards; the effectiveness of control techniques for safety; standardization, legislation, inspection, insurance, costing aspects and the like.

Safety Science will enable academic researchers, engineers and decision makers in companies, government agencies and international bodies, to augment their information level on the latest trends in the field, from policymakers and management scientists to transport engineers.

The journal focuses primarily on original research papers across its whole scope, but also welcomes state-of-the-art review papers and first hand case histories on accidents and disasters of special significance. The emphasis is on safety hazards, as distinct from health hazards, but may include both.

Most papers will be original publications, but the editors will consider translation into English of important papers previously published elsewhere in another language. In exceptional cases the journal will publish a paper in French or German.

© 2009 Elsevier Ltd. All rights reserved.

This journal and the individual contributions contained in it are protected under copyright by Elsevier Ltd., and the following terms and conditions apply to their use:

Photocopying

Single photocopies of single articles may be made for personal use as allowed by national copyright laws. Permission of the Publisher and payment of a fee is required for all other photocopying, including multiple or systematic copying, copying for advertising or promotional purposes, resale, and all forms of document delivery. Special rates are available for educational institutions that wish to make photocopies for non-profit educational classroom use.

For information on how to seek permission visit www.elsevier.com/permissions or call: (+44) 1865 843830 (UK) / (+1) 215 239 3804 (USA).

Derivative Works

Subscribers may reproduce tables of contents or prepare lists of articles including abstracts for internal circulation within their institutions. Permission of the Publisher is required for resale or distribution outside the institution. Permission of the Publisher is required for all other derivative works, including compilations and translations (please consult www.elsevier.com/permissions).

Electronic Storage or Usage

Permission of the Publisher is required to store or use electronically any material contained in this journal, including any article or part of an article (please consult www.elsevier.com/permissions).

Except as outlined above, no part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without prior written permission of the Publisher.

Notice

No responsibility is assumed by the Publisher for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions or ideas contained in the material herein. Because of rapid advances in the medical sciences, in particular, independent verification of diagnoses and drug dosages should be made.

Although all advertising material is expected to conform to ethical (medical) standards, inclusion in this publication does not constitute a guarantee or endorsement of the quality or value of such product or of the claims made of it by its manufacturer.

Funding Body Agreements and Policies

Elsevier has established agreements and developed policies to allow authors whose articles appear in journals published by Elsevier, to comply with potential manuscript archiving requirements as specified as conditions of their grant awards. To learn more about existing agreements and policies please visit <http://www.elsevier.com/fundingbodies>.

© The paper used in this publication meets the requirements of ANSI/NISO Z39.48-1992 (Permanence of Paper).

For the full Guide for Authors for Safety Science, please visit the journal's homepage at: www.elsevier.com/locate/ssci

Safety Science has no page charges

- L. Pezzullo and R. De Filippo 537 Perceptions of industrial risk and emergency management procedures in Hazmat Logistics: A qualitative mental model approach
- L. Sjöberg 542 Precautionary attitudes and the acceptance of a local nuclear waste repository
- I.O. Lund and T. Rundmo 547 Cross-cultural comparisons of traffic safety, risk perception, attitudes and behaviour
- J. Koehler, L. Dreijerink and R. van Poll 554 Exploring the risk information gap. Research into information supply and information demand of different parties concerned
- M. Gerbec and B. Kontić 561 Implementation of the Seveso II Directive in Slovenia: Survey of implementation and opinions of operators regarding its safety benefits



ELSEVIER

ISSN 0925-7535
47 (4) 459-568 (2009)

cedures in Hazmat
waste repository
cedures and behaviour
ion supply and in-
implementation and

Process Safety Indicators

Guest Editors

A. Hopkins

A. R. Hale

ISSN 0925-7535
(4) 459-568 (2009)



Thinking about process safety indicators

There is nowadays much discussion of the need for indicators that measure how well safety is being managed in major hazard facilities. Two dimensions of safety indicators can be distinguished: personal safety versus process safety indicators, and lead versus lag indicators. The distinction between personal and process safety indicators is relatively clear, but the distinction between lead and lag indicators, while frequently referred to, is rather more problematic. Two influential discussions of the lead/lag indicator dimension have recently appeared – the Baker inquiry following the BP Texas City refinery explosion (Baker, 2007), and a UK HSE publication, *Developing process safety indicators* (HSE, 2006). This paper is stimulated by these two publications and seeks to examine critically the lead/lag distinction. The Baker report also recommends the incorporation of safety indicators into performance pay and the paper goes on to discuss this issue.

Personal versus process safety indicators

Let us be clearer, first, about the distinction between personal and process safety. It is really a distinction between different types of hazards. Process safety hazards are those arising from the processing activity in which a plant may be engaged. Typical process safety incidents involve the escape of toxic substances and the release flammable material which may or may not result in fires or explosions. Many process safety incidents either damage the plant or have the potential to damage the plant. Moreover, they have the potential to generate multiple fatalities. The term process safety originates in the US; in some other parts of the world process safety is referred to as asset integrity or technical integrity.

Personal safety hazards, on the other hand, affect individuals but may have little to do with the processing activity of the plant. Typically they give rise to incidents such as falls, trips, crushings, electrocutions and vehicle accidents.

It turns out that most injuries and fatalities are a result of personal safety hazards rather than process hazards and, as a result, injury and fatality statistics tend to reflect how well an organisation is managing personal safety hazards rather than process safety hazards. Any organisation that seeks to assess how well it is managing process safety hazards cannot therefore rely on injury and fatality data; it must develop indicators that relate specifically to process hazards. That is one of the primary conclusions of recent major accident reports (CSB, 2007). Academic writers have also stressed this point. Hopkins has shown how the Esso gas plant at Longford in Australia had an impeccable lost time injury rate and yet was managing its major hazards quite poorly (Hopkins, 2000). He notes that the distinction is well understood in the airline industry, where no one would make the mistake of thinking that an airline's lost time injury rate provided an indication of how well it was managing air safety. Hale argues that the tendency to conflate these two types of safety, that is, safety in relation to the major and minor accidents,

goes back to a mis-interpretation of the Heinrich triangle (Hale, 2001). Many commentators, he says, have wrongly read this triangle as implying that minor accidents are precursors to major ones. This distinction between major and minor accidents, between personal and process safety, is crucial for major hazard facilities, but it is not the main concern of this paper.

The question of interest here is the distinction between lead and lag indicators. Moreover, the focus of the paper is on lead and lag indicators of process safety. In the case of personal safety the distinction between lead and lag indicators is somewhat less problematic. In this context, the term lag indicator generally refers only to injury and fatality rates, while lead indicators are those that directly measure aspects of the safety management system, such as frequency or timeliness of audits. There is of course much more that could be said about lead and lag indicators of personal safety, but in the interests manageability, the present discussion will focus on process safety alone.

Visualising the two dimensions

It is worth reiterating that the lead/lag dimension is quite distinct from the personal versus process safety dimension. This is best seen by arraying the two dimensions in a two-by-two table. Table 1 provides a visual demonstration that it is in principle possible to have lead and lag indicators of personal safety as well as lead and lag indicators of process safety.

I labour this point because at times even the Baker report seems to confuse the two dimensions. One of the main findings of that report was that companies like BP should not assume that indicators of personal safety are at the same time measures of process safety; they should develop dedicated process safety indicators. But the report also contains the following passage.

"The Panel believes that relying exclusively or predominantly on lagging indicators to assess process safety performance is ill-advised. ... BP's reliance on lagging, after-the-fact indicators of process safety performance rather than leading, predictive measures ... impaired BP's ability to measure, monitor and detect deteriorating or degraded process safety conditions and performance... This failure to use a set of effective performance metrics that includes leading indicators increased the likelihood that the organisation would identify the need for improvements or additional controls only after something had gone wrong" (Baker, 2007, p. 194).

The above passage suggests that BP's problem was that it was relying on lag indicators, rather than lead indicators. It is true that BP was relying on indicators such as days away from work injuries. But the problem with such an indicator is not that it is a lag indicator, but that it is a personal safety indicator. In fact one of the Baker panel's most prominent recommendations was that BP should develop

Table 1
The 2-Dimensional indicator space.

	Lead	Lag
Personal		
Process		

a lag indicator based on the number of process-related incidents – fires, explosions and hazardous releases. Contrary to the statement quoted above, in practice, the Baker panel was perfectly happy with lag indicators, provided they were the right kind of lag indicators.

The lead/lag distinction in the Baker report

Consider, now, the meaning of lead and lag. The following discussion will identify several ways in which the Baker report uses these terms.

Harm

Safety management systems ultimately aim to prevent harm, both harm to people and harm to property. One intuitively satisfying idea, then, is that lag indicators are direct measures of this harm, while lead indicators measure things that are in some ways precursors to harm. This is certainly how the terms are used in the case of personal safety, where to talk of lag indicators is almost invariably to talk of injury rates. As for process safety, the Baker report says at one point that lag indicators refer to process-related injuries, incidents and property damage (Baker, 2007, p.184). The emphasis here is on harm, although the word incident is ambiguous in this respect.

However the Baker report does not in practice adhere to this usage. It goes on to say that examples of lagging indicators include the number of unexpected loss of containment incidents, and failures of safety critical instrumentation/alarms (Baker, p.184). Here, the concept of lag indicator has moved well away from harm and includes events that are precursors to harm.

Failure

The above-mentioned examples from the Baker report all have in common the fact that they are *failures* of one sort or another. Perhaps the point is that lag indicators measure failures, regardless of whether they result in harm on any particular occasions. Lead indicators are then a residual category. The problem with this proposal is that there are many kinds of failures that are not normally regarded as lag indicators. For instance, the failure of a system to stay within safety critical parameters is not normally treated as a lag indicator, nor is the failure to comply with procedures. It seems that those who think in these terms draw a line somewhere: certain failures can be treated as lag indicators and others not, but where this line is to be drawn is never made clear.

Evaluating safety management activity

We have not yet considered the literal meaning of the word lag. What do lag indicators lag behind? The Baker report calls them measures of after-the-fact performance. What is the fact they come after? The usual answer is that they lag behind an organisation's safety management activity and cannot therefore give us information about the current state of a safety management system. Hence the need for lead indicators. As the Baker report puts it, lagging indicators of process safety performance suffer the disadvantage that they only suggest corrective actions after an accident (Baker, p.184). By contrast, lead indicators provide feedback on performance before an accident or incident occurs (Baker, p.185).

The fact is, however, that there are many situations in which lag indicators provide a good indication of how well a safety manage-

ment system is performing. The Baker panel itself suggested that the single most important indicator that BP needed to adopt was a composite lag indicator consisting of numbers of fires, explosions, loss-of-containment incidents and process-related injuries (Baker, p.253). It noted that Texas City had had numerous fires and several hundred losses of containment in the previous year (Baker, p.187) and that the refinery needed to focus on these failures and drive the numbers downwards. Any improvement in this composite index would be indicative of improvements in BP's safety management system.

This example serves to clarify the matter. The relevant issue is not whether the indicator is current or after the fact. The issue is whether, in the relevant time period, there are sufficient instances of the events being counted to be able to talk meaningfully about a rate. If there are, then charting this indicator over time will indeed provide evidence of whether the safety management system is improving. On the other hand, if years can go by without the occurrence of a single countable event, it is not possible to compute a meaningful annual rate, nor is it possible to conclude from one occurrence that safety is deteriorating.

Moreover, there is what might be called a zoom effect at work here. We can sensibly compute a fatality rate for an industry, and possibly even for a very large company. Regulators can therefore treat the industry fatality rate as an indicator of how well safety is being managed in that industry and seek to drive it down by encouraging various industry-wide safety initiatives. The same may be true for a very large company. However, once we zoom in on a particular company site we may find that it goes for many years without a fatality. Under these circumstances it will not make sense to treat the fatality rate at the site as an indicator of safety. Management at such a site will need to identify more frequently occurring events, such as injuries, in order to be able to judge how well they are managing safety over time. Similarly, catastrophic accidents occur so infrequently at any one major hazard site that they cannot be used as an indicator of how well process safety is being managed at that site, whereas loss of containment events may be frequent enough to serve the purpose. In short, where harmful events are occurring frequently enough to be able to talk about a rate, this rate can be used to measure safety; where harmful events are rare, we must look to more frequently occurring precursor events in order to be able to measure safety. Putting this another way, there are circumstances in which so-called lag indicators are perfectly good indicators of how well safety is being managed, and circumstances in which they are not. It depends, in part, on the level of zoom. The result is that, if one is interested in how well safety is being managed, the distinction between lead and lag indicators becomes largely irrelevant.

Interestingly, a BP internal report into the Texas City fire, the Mogford report, came to a similar conclusion.

"By definition, catastrophic and major process incidents are rare events, and performance measures need to be preferably focused on leading indicators, or at least lagging indicators of relevant, more frequent incidents" (Mogford, 2005, p.144).

Mogford is making clear in this passage that what is important is the frequency of the events to be counted and that whether they are described as leading or lagging indicators is essentially irrelevant.

At this point some further clarification is needed. I have argued above that there are circumstances in which there is no need to make use of precursor events in order to evaluate the effectiveness of a safety management system. This is not to say that companies can or should ignore precursor events. On the contrary, precursor events are warning signs, and companies need to seek out and act on these warnings. Furthermore, they need to develop reporting systems to capture such information. Research shows that virtually every major accident is preceded by warning signs and that the accident would

not have occurred had the organization concerned responded to these warnings (Turner, 1978). These precursor events may be quite unusual and quite ill-defined and therefore not suitable for use as performance indicators. For example, prior to the Longford gas plant explosion there was an unusual cold temperature incident that should have been treated as a warning, but was not (Hopkins, 2000, p.59). The point is that while there is very good reason to seek out and respond to precursor events, this is not necessarily an argument for treating them as performance indicators.

Reactive versus proactive monitoring

So far we have made little progress in distinguishing usefully between leading and lagging indicators. There is, however, one other suggestion in the Baker report about how this distinction might be made. Lagging indicators, it says, are generated by a process of *reactive* monitoring, while leading indicators are the outcome of *active* monitoring. Reactive monitoring amounts to keeping track of undesired events, such as gas releases; active monitoring involves routine, planned testing and inspection. As Baker puts it,

"Reactive monitoring allows an organisation to identify and correct deficiencies in response to specific incidents or trends. . . Active monitoring evaluates the present state of a facility through the routine and systematic inspection and testing of work systems, premises, plant, and equipment. . ." (Baker, p.184)

According to the Baker report, active monitoring gives rise to the following kind of indicator: the percentage of equipment that is past due for inspection. Notice that this is a measure of safety-relevant *activity* and the extent to which that activity has occurred. It is not necessarily a measure of safety. An organisation might score poorly on this indicator if it has fallen behind in its schedule of testing, even though all equipment might be functioning perfectly. There is a converse proposition. An organisation might be scrupulous in its testing schedule yet find that many of the items tested in fact failed the test.


What this suggests is that when considering active monitoring of this type there are really two sets of indicators of interest. First, there are indicators of whether monitoring is being carried out in a timely fashion (percent of tests conducted by due date). Second, there are indicators of the results of the monitoring (percent of cases in which equipment fails the test). The first is a measure of monitoring activity, while the second is a measure of equipment adequacy. The first can be termed an input measure and the second, an output measure. Both, surely, are relevant, yet the Baker report makes no mention of the second. All its illustrations of active monitoring indicators are of the input type, for example, completion of major accident risk assessments and closure of major accident group recommendations.

Summary

None of this is intended as an overall evaluation of the Baker report. It was in fact a ground-breaking document in ways that will not be canvassed in this article. The only point here is that the Baker report does not provide us with a satisfactory account of the distinction between lead and lag indicators. Indeed at times it confuses the lead/lag distinction with the personal/process safety distinction. Nevertheless it notes that effective measuring and evaluation systems use both leading and lagging indicators (Baker, p.23), and it refers to the UK HSE publication on process safety indicators in support of this proposition. I turn therefore to the HSE document.

The lead/lag distinction in the HSE guide

The HSE guide, *Developing Process Safety Indicators*, is an innovative document. Its starting point is the well-known Swiss cheese



model of defence in depth. What is innovative about the guide is that it defines leading and lagging indicators for each of the controls in the risk control system. (see diagram. RCS stands for risk control system). Lead and lag are thus relative terms, relative to the performance of a particular control.

In order to explore precisely how the HSE guide defines lead and lag indicators, let us note first that the process safety indicators in the HSE guide are of three types.

A. Measures of routine, safety-related activity.

e.g. proportion of safety critical instrument/alarm tests done on schedule

B. Measures of failures discovered during routine safety activity.

e.g. proportion of safety critical instruments/alarms that fail during testing

C. Measures of failures revealed by an unexpected incident.

e.g. numbers of safety critical instruments/alarms that fail in use.

The three types are listed down the left side of Table 2 and the following discussion may be more easily understood by referring to this table.

Keeping in mind the three types (A, B and C), let us consider a pair of definitions provided in the HSE guide.

The leading indicator identifies failings or 'holes' in vital aspects of the risk control system discovered *during routine checks* on the operation of a critical activity within the risk control system.

The lagging indicator reveals failings or 'holes' in that barrier discovered *following an incident or adverse event*. The incident does not necessarily have to result in injury or environmental damage and can be a near miss, precursor event or undesired outcome attributable to a failing in that risk control system. (HSE, p.3, emphasis in original)

It appears from these statements that leading indicators are type B and lagging indicators are type C above. Type A measures, concerned with safety-related activity, are not covered by these statements. (Hence the blank cell in Table 2.) This is clearly less than satisfactory.

Interestingly, however, the illustrations of leading indicators provided in the guide systematically include both type A and type B indicators. For instance, suggested leading indicators for the risk control, inspection and maintenance, are as follows:

Percentage of maintenance actions identified which are completed to specified time scale – (type A indicator).

Table 2
Description of indicator types According to different systems of classification in the HSE guide.

Indicator Type	Classification of indicator according to ...		
	Definition	Examples	Input/outcome
A. Measures of safety activity		Lead	Lead
B. Failures revealed by safety activity	Lead	Lead	Lag
C. Failures in use	Lag	Lag	Lag

Percentage of safety critical plant/equipment that performs to specification when inspected or tested - (type B indicator)

In short, although leading indicators are defined as type B indicators, the illustrations given cover both A and B.

This suggests that the definition the guide is using *in practice* is as follows:

Leading indicators relate directly to safety-management activities and are either measures of those activities themselves (type A) or of the results of those activities (type B).

Lagging indicators are measures of unexpected failures occurring in normal operations (type C).

This formulation groups A with B and contrasts these with C (see column 2 in Table 2). It seems like a reasonable way of making the distinction. However the guide does not always adhere to this formulation. It provides the following example of a lagging indicator:

Number of safety critical instruments/alarms that fail to operate as designed, *either in use or during testing* (emphasis added).

Precisely because it covers failure both in use and during testing, it covers both B and C above (see column 3 in Table 2). In so doing it ignores the formal definition and it is inconsistent with the definition that the guide works with in practice for the most part. There is, however, a certain logic to this usage. The guide assumes that lead indicators are measures of input, while lag indicators measure outcomes:

Leading indicators ...can be considered as measures of *process or inputs* essential to deliver the desired safety outcomes. ... Lagging indicators show when a desired safety *outcome* has failed or has not been achieved (HSE, p2, emphasis added)

From this point of view it is of little relevance whether the instrument/alarm fails in use or in testing. Both types of failure demonstrate the presence of a hole in the protective barrier. Both types demonstrate that the plant has been operating in a degraded state, regardless of whether a harmful incident has occurred. Both are outcome measures and, *in that sense*, both are lag indicators.

We have, then, three different approaches in the guide to determining whether indicators are lead or lag: a formal definition, the system of classification revealed in most of the examples, and the input/outcome distinction. These different approaches generate the different entries in the cells in Table 2.

The table is read as follows:

Where indicators are classified according to the formal definition, type A indicators are unclassified, type B are lead, and type C are lag.

Where indicators are classified as are most of the examples in the guide, type A indicators are lead, type B are lead, and type C are lag.

Where indicators are classified according to the input/outcome distinction,

type A indicators are lead, type B are lag and type C are lag.

I conclude, therefore, that the HSE document does not provide us with a single, consistent account of the lead/lag distinction.

It is important to note that this analysis of the meaning given to lead and lag indicators in the HSE guide is not intended as an overall evaluation of that document. The value of the guide is that it draws attention to the role that performance indicators can play. Given that safety management systems nowadays are about putting in place a system of risk controls, the real purpose of performance indicators must be to evaluate the effectiveness of those controls. This can be done by monitoring failures of the controls in use, as well as monitoring the functioning of those elements of the safety management system designed to ensure that these controls are working effectively. The HSE guide demonstrates admirably just how this can be done. The guide clearly identifies the three types of indicators discussed above. It is only when it comes to grouping the three types into two that inconsistencies appear.

What emerges from this discussion, I think, is that, in the quest for process safety, the important thing is to identify measures of how well the process safety controls are functioning. Whether we call them lead or lag indicators is a secondary matter. Companies I have studied that are actively seeking to identify indicators of process safety do not make use of the lead/lag distinction in any systematic way. They use indicators of failure in use, when these are available, as well as indicators arising out their own safety management activities, where appropriate, without thought as to whether they be lead or lag. These indicators include:

- Tardiness in completing maintenance, inspection or corrective action tasks;
- measures of plant or process failure in use, such as plant trips and alarm rates;
- failure rates following testing;
- plant over-rides (by-passing of safety devices);
- leaks, and so on.

Improving performance in relation to these indicators must enhance process safety.

A final comment here concerns the relative priority to be given to the three types of indicators described above. Which of the three provides the best indicator of safety? The answer would seem to be type C. The aim of the system of controls is to prevent a major incident. Where a particular control fails in use, we are correspondingly closer to that incident. To draw on the Swiss cheese model, these are cases in which a potential incident has penetrated some of the barriers, but not all. They are, to varying degrees, near misses. Such failures are therefore undesirable events, to be avoided if at all possible. Where such failures occur with sufficient frequency to be able to talk about a rate, the top priority must be to drive these failure rates downwards. Where failures are so infrequent that it is not possible to talk of a rate, the priority shifts to identifying measurable aspects of the safety activities designed to ensure that the controls remain in place. The strategy is to move upstream in the chain of precursor events only so far as is necessary to identify countable occurrences.

The idea that, where possible, safety indicators should be based on the undesired events themselves, rather than their precursors, is applicable more generally. We know, for instance, that Chinese coal mines kill well over 10,000 miners each year. Suppose the industry was to respond by increasing the number of safety audits carried out at each mine. This of course is a step in the right direc-

tion, and in the absence of any other indicators we would be inclined to say that safety had improved. But the acid test in this case is the fatality rate. If this rate remained unchanged we would certainly conclude that the management of safety had not improved, no matter what the lead indicators were telling us.

Driving improvement

Safety indicators are only worth developing if they are used to drive improvement, that is, if the organisation focuses attention on them in such a way as to make a difference. This seems like a truism, but it needs to be stated. BP at Texas City measured loss of containment incidents (eg spills), and the figure got steadily worse from 2002 to 2004, increasing by 52% in the two year period, from 399 to 607 (Baker, p.187; CSB, p.168). This was a critical process safety indicator, but it was not a measure that mattered to BP. If such indicators are to drive improvement, ways need to be found to hold senior managers accountable for them. One obvious way is to build them into their remuneration packages.

The inquiries following the Texas City fire provide some useful information on bonus pay systems at that site. Most of the workforce received a bonus payment based on the overall performance of the refinery. There was no way that individual performance could influence the payment. According to the US chemical safety board (CSB) report, 50% of the bonus was determined by cost leadership, (that is, cost cutting), while safety determined only 10% (CSB, p.149). Furthermore, safety was measured as the number of OSHA-recordable injuries, clearly a measure of personal safety.

For senior managers, the situation was different. They had individually constructed personal performance contracts with their immediate superiors which served as the basis of bonus payments. According to the CSB,

"The contracts consisted of weighted metrics for categories such as financial performance, plant reliability and safety. The largest percentage of the weighting was in financial outcome and cost reduction. The safety metrics included fatalities, days away from work case rate, recordable injuries, and vehicle accidents; process safety metrics were not included. HSE (health safety and environment) metrics typically accounted for less than 20% of the total weighting in performance contracts". (CSB, p.149. See also Baker, pp.28-29)

Process safety, then, was completely missing from the incentive system at Texas City. But quite apart from this issue, what is striking about these bonus schemes is the low weight given to safety in comparison with cost reduction. One wonders in these circumstances whether the bonus system directed much attention to safety at all. The answer appears to be that at Texas City it did. Prior to the accident in March 2005, the total recordable injury rate at the site was one third the industry average (CSB p.197) and in 2004 the injury rate had improved to such an extent that BP made an additional payment to the whole Texas City workforce.¹ The bonus system did seem to have driven the reported accident rate to very low levels.

This apparently anomalous outcome points to something important about the way bonus systems can work. The literature on management incentives systems suggests that managers are influenced not just by economic incentives but also by praise and criticism (see, for instance, Nalbantian, 1987; Armstrong and Murlis, 1988). These more subjective influences can be very powerful. A bonus is not just a monetary incentive; it is a symbolic statement that the recipient has done well. This is especially true if the award of the bonus follows a performance review in which the person concerned is con-

gratulated for performing well according to the specified criteria. In these circumstances the bonus becomes a psychological reward.

For the most senior people in a corporation it is hard to see how safety bonuses can provide any significant financial motivation. The financial benefits they receive go beyond those specified in the performance agreements of most of their subordinates and include payments based on share holdings and share price movements. Financial returns from these sources can far outweigh other sources of income. In these circumstances, the impact of safety bonuses must be symbolic, affecting reputation and pride, rather than exercising any real financial leverage.

Assuming that incentive schemes do have an effect, either in material or symbolic terms, one of the clearest lessons arising from the Texas City accident is the need to include measures of process safety into remuneration systems. The Baker panel made the following recommendations (Baker, p.251):

A significant proportion of total compensation of refining line managers and supervisors (should be) contingent on satisfactory meeting process safety performance indicators and goals... A significant proportion of the variable pay plan for non-managerial workers ... (should be) contingent on satisfactory meeting process safety objectives.

The need for caution

A warning must be sounded at this point. In deciding on performance indicators for inclusion in such pay schemes, it is important to recognize that the moment there are consequences attached to performance with respect to an indicator there is an incentive to manage the indicator itself rather than the phenomenon of which it is supposed to provide an indication (see, for instance, Goddard et al, 2002).

Indicators of safety-related activity are inherently dubious from this point of view. The problem is that safety-related activities can vary in terms of their quality as well as their quantity, and it is often possible to increase quantity by sacrificing quality. For instance, if performance is being assessed by the number of audit corrective actions that have been closed out or completed, the result can be improved not only by putting more resources into closing out these actions (the intended outcome), but also by decreasing the quality of the close outs (an unintended outcome). If performance is measured in terms of the numbers of people who have undergone training, the quantity can be improved by reducing the quality of the training. These are perverse outcomes. They are focused on managing the measure rather than managing safety. By contrast, where a failure rate is being used as a performance indicator, particularly a rate of equipment failure, there is less possibility of perverse outcomes such as described above. Of course a rate of equipment failure can be deliberately understated, but this requires some element of dishonesty, while reducing the length of a training program in order to increase the number of people trained involves no such dishonesty. The fact that managers can in good conscience modify safety related-activity in this way is what makes indicators of safety related-activity especially problematic. The Baker report seems unaware of this issue.²

This is not to say that performance indicators of safety-related activity should be abandoned; merely that extreme care should be taken before including them in bonus pay systems. Where there is poor performance with respect to such an indicator, the first reaction should be to investigate the reasons. It may be that there are not sufficient resources to carry out the activity effectively. Penalizing poor performance in these circumstances is bound to lead to perverse outcomes.

¹ It is ironic that this improvement in the figures occurred despite three fatalities in that year, two of which were process related. CSB p. 161. The OSHA injury rate excludes fatalities. CSB p.149.

² It does refer to other types of limitations at pp.190-191.

While it is more meaningful to count failures than activities, this does not mean that counting failures is straightforward. As noted earlier, one of the principal recommendations of the Baker report was that BP should develop a single summary indicator of process incidents, including fires, explosions, hazardous releases and process related injuries/fatalities. The report recognized that each of these failure types would need to be carefully defined. So, for instance, it defined fire as follows (Baker, pp.258–9):

- a leak that results in a flame,
- a tangible indicator of fire (eg soot on the inside of distillation tower),
- a fire on a scaffold board in a process unit,
- a fire in a vehicle parked by an operating unit,
- a 120 V shorted switch,
- a ground fault on electrical heat tracing,
- a phase to ground or phase to phase short on electrical power distribution,
- a fault in a motor control centre,
- a fault in electrical switch gear.

The definition also included a list of types of fire that would *not* be included in this process safety indicator, such as a fire in a rubbish can in an office building. It is clear that even such an apparently discrete and countable event as a fire needs to be carefully defined if it is to be part of an indicator used to drive performance.

Tailoring performance agreements to individuals

The recommendations of the Baker report distinguish between incentive systems that apply to senior managers and those that apply to lower employees further down the hierarchy. Clearly it is the senior managers who have the greatest influence on safety and their incentive structure is therefore critical. In particular it is critical that their performance agreements focus on things that it is within their power to influence. From this point of view the problem with many performance indicators is that it is not immediately obvious how the activities of any given manager may influence them. The personnel manager, for instance, may feel that he or she has little or no control over the number of gas leaks, fires or other undesired events occurring at a site. If so, an incentive system that highlights such events has little potential to influence behaviour of that manager. Putting it another way, measures of this nature are not sufficiently targeted to motivate the performance such a manager.

On the other hand, it may be that the personnel manager can indeed influence safety by ensuring that safety critical roles are adequately staffed. Such a matter is far too specialized to be included in any incentive scheme applying to other employees, but it may be particularly appropriate in the case of a personnel manager. The point is that senior managers must have performance agreements that identify safety-relevant activities peculiarly within their control and for which they can sensibly be held personally accountable.

Moreover there is no need for these personalized goals to be quantifiable. Some qualitative estimate of the extent to which they have been achieved is all that is necessary in order to be includable in an incentive payment scheme. In one company I have studied, the CEO writes into the performance agreements of his direct reports a series of objectives, some of them safety-related, and then evaluates each person in relation to his or her set of objectives, on the following scale: meets most; achieves all; exceeds most; exceeds all. Pay is linked to this scale in a pre-determined manner.

For operations and engineering managers, some of the site-wide performance indicators discussed in the Baker report and the HSE guide may be appropriate measures for inclusion in personal performance agreements, precisely because they have direct influence

over these matters, but for other managers, such as personnel and finance managers it may require some ingenuity to devise appropriate safety-related goals. It is vital that this be done, however, because matters within their sphere of influence, such as staff overload and cost cutting, are routinely identified as contributory factors in major accident investigations. Corporate HSE managers can be particularly useful here. Ensuring that appropriate HSE objectives are included in the performance agreements of very senior executives is one of the most important functions a corporate HSE manager can perform.

Once these objectives are chosen for the most senior executives they can be cascaded downwards, so that the performance agreements of immediate subordinates encompass matters within their direct control, and so on down the line. In this way the motivational structures of various levels in an organization can be aligned.

None of this is to suggest that site-wide or company-wide performance indicators are can be dispensed with in the incentive schemes devised for senior executives. If nothing else, these measures serve to remind senior people of overall corporate goals and encourage them to align their behaviour with these goals wherever possible.

Conclusion

This paper has been an attempt to think through the meaning of process safety indicators. In particular I have examined the meaning of the terms leading and lagging in two recent influential publications and found that they are not used with any consistency. Nor do I think there is much point in trying to pin down a precise meaning since in different contexts these terms are used to draw attention to different things. The most important point to emerge from the HSE document is that process safety indicators must be chosen so as to measure the effectiveness of the controls upon which the risk control system relies. Whether they be described a lead or lag is ultimately of little consequence.

The paper also reflects on the Baker panel proposal that process safety indicators be included in incentive pay schemes. This would be an important step forward, but care must be exercised to ensure that it does not lead to attempts to manage the measure as opposed to managing safety. Given that it is senior executives who have the most influence on safety, the aim must be to include appropriate indicators of process safety in their performance agreements. There is scope for a great deal of creativity in targeting these agreements to safety-relevant matters that lie within their control.

References

- Armstrong, M., Murlis, H., 1988. *Reward Management*, Kogan Page, London.
- Baker, J., 2007. *The Report of the BP US Refineries Independent Safety Review Panel*.
- CSB, 2007. *Investigation Report, Refinery Explosion and Fire*. US Chemical Safety and Hazard Investigation Board.
- Goddard, M., Davies, H.T.O., Dawson, D., Mannion, R., McInnes, F., 2002. Clinical performance measurement part 2: avoiding pitfalls. *Journal of Royal Society of Medicine*, 95 549–551.
- Hale, A., 2001. Conditions of occurrence of major and minor accidents. *Institution of Occupational Safety and Health Journal* 5(1), 7–21.
- HSE, 2006. *Developing process safety indicators: a step-by-step guide for chemical and major hazard industries*, UK Health and Safety Executive.
- Hopkins, A., 2000. *Lessons from Longford: The Esso Gas Plant Explosion*, Sydney, CCH.
- Mogford, J., 2005. *Fatal Accident Investigation Report*, December 9, p. 144.
- Nalbantian, H. (Ed.), 1987. *Incentives, Cooperation and Risk Sharing: Economic and Psychological Perspectives on Employment Contracts*, Rowman & Littlefield, New York.
- Turner, B., 1978. *Man-Made Disasters*, Wykeham, London.

Andrew Hopkins
Professor of Sociology, Australian National University
E-mail address: andrew.hopkins@anu.edu.au



ELSEVIER

Contents lists available at ScienceDirect

Safety Science

journal homepage: www.elsevier.com/locate/ssci

Process safety indicators: Response to Andrew Hopkins

In his excellent analysis of process safety indicators, Andrew Hopkins references the recently published landmark Baker Panel report in pointing out the current confusion with interpretation of leading and lagging indicators for process safety. The lead and lag distinction has inhabited process safety circles for several years but until now very few have dared to challenge its value. He argues persuasively that the distinction is unimportant compared to that of obtaining a meaningful count which over time can serve to drive process safety improvement. Andrew Hopkins then turns the spotlight on the link between process safety performance and individual reward. Again he is correct to urge caution in the use of reward systems which reflect so called safety activity indicators since these

can be engineered to inflate pay at the expense of quality. There is no doubt that the development of indicators is currently the hot topic within the process safety community and this timely contribution helps to dispel much of the confusion which continues to stall genuine progress in this field.

Lee Allford

*Manager Operations,
European Process Safety Centre,
Rugby, United Kingdom*


E-mail address: lallford-epsc@icheme.org.uk



Contents lists available at ScienceDirect

Safety Science

Journal homepage: www.elsevier.com/locate/ssci



Reply

Process safety indicators: Response to Andrew Hopkins

These comments are based on some 20 years in the chemical industry in UK and Belgium and then latterly within UK Construction.

Essentially (as most of the commentators have said) the massive and inappropriate emphasis in recent years upon personal safety indicators (such as Lost Time Injuries [LTI]) as opposed to process indicators had a negative effect on industry. From my own experience with managing directors at first hand, it has often led to 'self gratification' by business heads who see this as a simple numeric. This is particularly seen in construction which is a relatively immature work environment. When the data give a positive picture everyone relaxes. This is particularly so in managements where 'simplicity' is the byword. Sadly such managements cannot then see the wood for the trees, especially where entrenched views have been held over a period of time. Operational pressures within the dynamics of business are largely the reason for this position being adopted. It is just easier to see a low LTI rate as success and not to look further.

All of this leads to support for the points made by Hopkins. The only point I would like to add is that this discussion must lead us in

the direction of more emphasis on 'modern safety management', including organisational resilience measures, rather than traditional measures of safety correction. By this I mean audit, review and the neat circle of traditional timely improvement). Such indicators can certainly point out the 'state of health' of an organisation as much as its performance at a point in time.

Within the MCG (UK Major Contractors Group) I have been able to share a simple, so-called RAG (Red Amber Green) process which is now used by many UK contractors, and is added into their traditional audit processes for construction sites. This process is monitoring those lead indicators of performance (or non performance) and allows a net view to be taken of the site status (green for OK, amber for degraded, or red for critical) and thus the level of appropriate/capable management control. This has proved a valuable tool in the safety armoury, when linked to an allied course of intervention.

Robin Chaplin
Balfour Beatty, United Kingdom
E-mail address: robinchaplin@hotmail.co.uk



ELSEVIER

Safety Science

journal homepage: www.elsevier.com/locate/ssci

Reply

Process safety indicators: Response to Andrew Hopkins

I agree totally with Hopkins that the Baker Report missed the distinctions between and mixed its terminology for personal safety, process, leading and lagging indicators. I also like his table 1 with the two dimensions of indicator purpose and indicator instance. I suspect that BP would have had indicators that could have been put in each box of that table 1, but that these indicators were only known and understood at lower levels such as by supervisors. I imagine they would have had a leading personal safety indicator like, Job Safety Analysis (JSA) or similar process. It may have been that there was just no active monitoring of the effectiveness of that process. Likewise, I expect that Loss of Containments (LOC, a process lagging indicator) were being measured by the Environmental Health Safety & Risk (EHSR) Department and I expect that the inspection department would have been able to count the number and age distribution of temporary clamps on lines (a leading process indicator), if there were asked to. However, I suspect that no one was formalizing all this into an active monitoring system that the Baker Investigators could easily understand. Hence, I suspect BP had a lot more leading metrics working than the report discovered, probably because most processes were informal.

As for leading indicators, I look for them to be predictive, whilst lagging are retrospective indicators. The goal is to avoid posthumous indicators! Hence, both can be used: leading as predictive, lagging as "near misses" which we can learn from.

Caltex do report and monitor events where a safety device did not operate in a formal governance system. With respect to our terminology, we call the operators/inspection checking of the equipment, "checks". "Active Monitoring" is the analysis by superintendents of the checks, both for compliance, results and effectiveness. We have "audits" (both internal and external) which we differentiate from monitoring because audits are performed by a third party and not the line management. We use a model "MR ADA" (Measure, Record, Analyse, Decide, and Act) for monitoring,

where the supervisor performs the **MR A** and the manager and supervisor perform the **DA**.

I agree with Hopkins' point on safety activities being a check on compliance. Our company, like many others, has an emphasis on compliance. Companies must be careful this does not create the unintended behaviour of just completing the safety checks and not giving weight to the results of the checks. For example, it is important that companies understand that they can be 100% compliant on checking safety devices like pressure safety valves (PSV) and have managers think that all is well. However, if the checks are finding a 1% failure rate and managers are either not hearing about these results or not reacting to them, the monitoring system is failing.

I agree with Hopkins' concern with the unintended outcomes for metrics. This comes down to leadership (what my boss measures or show interest in, interests me hugely.) and everyone understanding the purpose and intent of the metrics. We use a model here of **FBI**, people either **F**udge the data, **B**lame the system, or actually **I**mprove the system, to meet the metrics. We try to avoid the **F** and **B** by documenting every metric's purpose and how it is measured and make sure our leaders provide the context for the metric.


I think we need to recognise that a good monitoring system does require a mature organisational culture. There can be issues with people working in a performance based culture where efforts are made to "not" report issues that might look like "they" are not performing. In some countries, losing face is worse than almost everything!

M. Hodgkinson

Operational Excellence and Risk Manager,

Caltex Refining, Lytton Refinery,

Queensland, Australia



Letter to the Editor

Response to Andrew Hopkins

In his paper, Andrew Hopkins does an excellent job of teasing out the relevant categories for differentiated types of safety indicators and of showing how difficult it is to use these different categories in a consistent manner. In the discussion of safety indicators, he stays on a purely descriptive level, however. From that perspective he argues and rightly so, that the distinction between lead and lag indicators is irrelevant, but what is relevant is whether there are sufficient instances of the indicator to look at trends within companies and to compare between companies. He also makes the very interesting and important point that depending on the level of analysis, different indicators may be useful due to differences in their occurrence. Fatality rate, for instance, (hopefully) is not a good indicator at company level, but it might very well be for certain industrial sectors.

However, the article misses out on one issue central to the debate about safety indicators, namely that sound knowledge about cause-and-effect relationships is needed to predict safety performance from whatever set of indicators. Hopkins does mention that lead indicators are important as warning signals and should be paid attention to in order to avoid negative safety outcomes. But at the same time he argues that if lead indicators relate to rather singular events they should not be used as performance indicators; this goes back to his point that indicators need to be based on sufficiently frequent events. What he leaves out completely from the discussion is the fact that one important reason why precursor events are not acted upon is because their causal link to worse events happening later on can often only be established after the fact. This is also true for lead indicators that in his argumentation would be good ones, because they are sufficiently frequent, such as certain technical systems not functioning properly during inspection. Whether this is a good lead indicator does not only depend on its frequency, but much more on whether a causal link can be established to negative safety outcomes. The fact that an indicator has more frequent occurrences definitely does help to establish such a link empirically, but the frequency as such is irrelevant from the perspective of trying to promote good safety performance. Instead, Hopkins' focus on frequency may lead people to focus on indicators purely because they are frequent, but which happen to be completely irrelevant for increasing safety.

Hopkins makes the statement, that irrespective of them being lead or lag indicators, "improving performance in relation to these indicators must enhance process safety". However, lead indicators

are crucial exactly in this respect, as they (given that there is a causal link to safety performance) allow companies to prevent negative safety outcomes much more proactively and, therefore, more effectively than lag indicators. From this perspective, I strongly disagree with his statement that "where possible, safety indicators should be based on the undesired events themselves, rather than their precursors". Only if companies and auditing and inspecting bodies focus on precursors, can the actual accidents and incidents be prevented. The very difficult issue is to establish good precursors that have sufficient predictive validity. Also, what accident and incident analyses teach us, is that it is usually not single precursor events, but particular patterns of events that lead to negative safety outcomes. Gaining more knowledge about these patterns and building competence within companies and regulatory bodies for recognizing relevant patterns is crucial.

In relation to his stance that safety indicators should always have the purpose of improving safety performance, Hopkins ends his article with a discussion of the value of linking performance on safety indicators to individual pay. I agree fully with all the pros and cons he presents, which are the same as those put forth generally in discussions of pay-for-performance systems (e.g. Rynes et al., 2005). What his discussion shows additionally, though, is that safety indicators can serve very different functions. It would have been very useful, if the article had started out by describing and possibly categorizing different functions safety indicators may have. For the most part, Hopkins seems to assume that safety indicators serve a purely descriptive function and especially one of comparing safety performance over time and across companies, sectors etc. Only in the later part of the article does he focus on driving improvement as the main purpose for safety indicators. His arguing for high frequency of occurrence being the core quality of a good indicator may be appropriate for purely descriptive and comparative indicators; it definitely is not when actively improving safety in individual companies is the main purpose.

Reference

- Rynes, S.L., Gerhart, B., Parks, L., 2005. Personnel psychology: performance evaluation and pay for performance. *Annual Review of Psychology* 56, 571–600.

Gudela Grote
ETH Zürich, Switzerland
E-mail address: ggrote@ethz.ch



ELSEVIER

Contents lists available at ScienceDirect

Safety Science

journal homepage: www.elsevier.com/locate/ssci

Reply to comments

Andrew Hopkins

Australian National University, Canberra, Australia

Introduction

First let me thank the editor and all those who took the time to respond to the original article (hereafter, "respondents"). It is rare that we get an opportunity to reflect on the basic assumptions of what we are doing. It is probably only when we are challenged that we become aware of what those assumptions are. This debate has revealed that the concepts that we all use are not nearly as clear as we think, and that our understandings of what they mean are very varied. Andrew Hale wrote an editorial some years ago about the confusion in the way we use the term "culture". I suspect that we are similarly confused about the idea of indicators. This discussion has certainly made me aware that the way I use the term shifts subtly from one context to another. I suspect we are all in the same boat in this respect. The main thing I want to do in my reply is highlight this diversity of understandings.

Can a single event be an indicator?

Most understandings of the word "indicator" assume that we are dealing with sufficiently large numbers to be able to talk sensibly about increases and decreases in a period of time. On this basis, a single event cannot count as an indicator. That is not to say that we are uninterested in single events. On the contrary, HROs are vitally interested in single events, which may constitute weak signals that something is wrong. As I mentioned in the original article, Esso experienced a cold temperature incident in the months before the Longford explosion which was a clear warning of what was to come. It was ignored with disastrous consequences. I have chosen not to treat single events of this nature as performance indicators. One might of course treat the *number* of such weak signals reported as an indicator (the more the better), but not the single event itself. However, several respondents take the view that it is appropriate to describe single warning events as indicators. This difference in terminology is one of the sources of misunderstanding. I need to stress that in this reply I shall continue to assume that indicators are based on a sufficient number of instances to be able to identify change over time.

E-mail address: andrew.hopkins@anu.edu.au

0925-7535/\$ - see front matter © 2008 Elsevier Ltd. All rights reserved.
doi:10.1016/j.ssci.2008.07.020

The distinction between lead and lag indicators

The main point of my article was that the distinction between lead and lag is not clear, and in any case, it may not be important to make this particular distinction. Let me try to summarise what respondents have had to say about this.

As preliminary point, most want to maintain the distinction, but not all. Mearns says, "perhaps we should move away from the concept of leading and lagging indicators and focus instead on key performance indicators for safety". Wreathall titles his piece "Leading? Lagging? Whatever!", which suggests that in his view the distinction may not be important.

The dominant opinion is that the distinction is important, but there is a variety of views about just what the distinction is. These can be grouped into two categories, which I shall call relativist and absolutist.

Relativist

The HSE document takes the relativist position: indicators are lead or lag relative to a particular risk control. I argued that the HSE did not maintain the distinction consistently and I shall not repeat that argument here.

Among the respondents, Erickson also adopts a relativist view.

"A specific indicator could be lagging in relation to one objective, while leading in relation to another. Failure at testing is for example lagging in relation to the performance of the individual barriers, while leading in relation to the performance of the overall process safety management system".

This is fair enough, as far as it goes, but it is not a complete account of the distinction.

Dyreborg argues the relativist position at some length. He says:

"A lag indicator (in relation to one risk control) could in the next instance be included in measurements of precursors to other, maybe more serious effects or incidents, and thus become a measure used as a leading indicator. The same measure can represent both an input indicator (lead) and an output indicator (lag) depending on the causal relationships in focus".

I am afraid I cannot follow this argument. Suppose I have a set of risk controls as follows:

1. procedures: operators are supposed to keep a process away from critical limits. If this fails,

Performance measurement

The preceding discussion is not intended to belittle measurements of how well risk controls are functioning. But to call these lead indicators is conceptually a very different usage of the word indicator. This second usage relies on a causal model of how accidents occur, that is, a model that identifies precursor events or risk factors. Once these are known (or assumed) safety depends on putting in place a variety of risk controls – technical, behavioural and managerial – such as maintenance, PPE and management walk-arounds. The performance of these controls can be measured, using such measures as number of component failures, rates of PPE usage, and frequency of walk-arounds. These are essentially performance measures, that is, measures of how well the various risk controls are functioning. Although these performance measures are frequently called indicators, none by itself is an indicator of how well the safety management system is functioning. Each is a direct measure of how well a particular risk control is performing. Performance measurement of this nature is vital for safety, and improvements in terms of these measures do indeed lead to reductions in personal injury rates, as Hale notes.

If questions are raised about the validity of such measures, the issue is not whether the indicator is a valid indicator of some unmeasured variable, but whether it is predictively valid, that is, whether it correlates with some measure of harm. What is at stake here is not the validity of the indicator, *qua* indicator, but the validity of the causal model that is being assumed.

Several respondents have understood me to be arguing against the use of risk control performance measures. I hope it is clear, now, that is not the case. All I am saying is that it is not helpful to call them lead indicators. Arguably it is not even helpful to call them indicators.

Leading in time

In the case of leading indicators, there is a third usage that can be identified in the discussion: leading indicators are variables that predict the future state of some other variable of interest. Various respondents note that this meaning has been inherited from economics. However, it does not always translate simply into the safety arena.

Mearns provides one example where it does. She argues that safety climate is a leading indicator which predicts injury rates at a future time. But unlike most other respondents, she does not treat safety climate as an indicator that is of direct use in deciding where to intervene. As she says, "the focus is now on understanding the antecedents of positive safety climates". Safety climate thus becomes an intervening variable, even a conceptual variable. What strikes me about her description of safety climate research is that it is basic research aimed at understanding the causes of accidents, rather than an attempt to identify indicators of immediate practical value. To talk of safety climate as an indicator, and more particularly a leading indicator, is rather different from the usages discussed above.

Mearns treats safety climate as an indicator that leads, in time, in relation to personal injuries. When we come to process safety, however, I would argue that so-called lead indicators are not being used to predict future events. Risk control performance measures are measures of the *current* state of those controls. If we treat them

as indicators of the state of the safety management system, they are indicators of the *current* state. The timing of process safety accidents is essentially random, depending on the simultaneous failure of all controls, and it is generally not possible to correlate these indicators with future process safety accidents. In other words, the idea that lead indicators lead in time is not an appropriate usage when we are talking about process safety.

Conclusion

The responses to my original paper demonstrate that the safety science community holds a diverse set of understandings of the distinction between lead and lag indicators. In my view this is not just diversity; most of these understandings themselves turn out to be problematic on close examination. If we intend to go on using these terms we need to think harder about what they mean. At the very least we should define them carefully every time we use them. Even the concept of an indicator turns out to have multiple meanings that can sometimes lead to confusion.

The safety science community is drawn from a wide range of disciplines, each with its own conceptual apparatus and set of shared meanings. Those meanings are not necessarily shared across the whole safety science community, with the result that we sometimes talk past each other. This debate highlights the need for us to work towards an agreed set of meanings.

Addendum

The contribution to this debate by Woods is rather different to the others. It is an interesting account of resilience theory. It stimulates two responses in me that are not strictly on the topic of this debate – hence an "addendum".

First, a resilient organisation, as described by Woods, seems indistinguishable from a high reliability organisation. I will not detail all the similarities here, except to say that the description of the hospital emergency department stretching to cope with high workload is powerfully reminiscent of the description in the HRO literature of how air traffic control centres mobilize additional resources at times of heaviest workload. I hope that resilience theorists will some day explain the difference, if there is any, between these two ideas.

Second, I read the Columbia accident a little differently. Wood sees the problem as NASA lacking any "mechanism to assess the risk that the organisation is operating nearer to safety boundaries than it realizes". Yet as he himself points out, NASA had already been told very clearly in the Stephenson report that it was dangerously close to the boundary. In my view the problem was that NASA's organisational structure prevented it from acting on this knowledge. One of the most important recommendations of the Columbia Accident Investigation Board was that an independent Technical Authority be established within NASA, with the capacity to intervene in the shuttle program and, if necessary, veto launch decisions. This suggests that the issue is better seen as one of inappropriate organisational structure, rather than maladaptive functioning. It may be out of place to raise these matters here, since they introduce a whole new debate. Perhaps *Safety Science* could host such a debate?

2. a critical alarm sounds. If this fails,
3. a cutoff or trip mechanism should operate. If this fails,
4. pressure relief valves release flammable hydrocarbon to a closed system. If this fails,
5. flammable material is released via an emergency flare system. If for some reason this fails and flammable material is released to atmosphere,
6. ignition control policies should minimized the risk of ignition. If these fail, we have a fire or explosion.

I hope this is a realistic sequence. It is inspired by the Texas City explosion, although it is not an exact description of that accident sequence.

Take the indicator, "frequency with which a critical alarm fails to operate in use". This is a lag indicator in relation to a particular risk control, the critical alarm. According to Dyreborg it is also a lead indicator with respect to the next risk control, a cutout mechanism. It is hard to see how this is so. Of course, the failure of the alarm is an opportunity for the next control to either operate as it should, or fail in use, but the only indicator for the cutout mechanism that we can construct from this scenario is something like "number of times the cutout fails in use". In short, the lag indicator of the earlier control is *not* a lead indicator for the next.

Dyreborg's own example of what he is talking about does not clarify the issue. He argues that an injury rate is a lag indicator from a company point of view, but a lead indicator from the perspective of a regulatory authority, because it can be used to decide which industries to target. This is problematic for two reasons. First, the idea that an injury rate could ever be regarded as a lead indicator is so contrary to all usages that it is bound to lead to endless confusion. Second, there seems to be an assumption that injury rates cannot be used by companies to determine intervention strategies. On the contrary, provided injuries are occurring relatively frequently, the data can certainly be analysed at company level to identify the hazards or accident mechanisms involved and thus to determine intervention strategies.

In summary, Dyreborg takes the relativist view of the distinction between lead and lag, but it is not clear just how we are to understand the distinction that he makes.

Absolutist – bowtie

The contrasting view, expressed by several respondents, is that the distinction is an absolute one. This view is expressed most clearly by Hale and by Hudson. It is based on the bowtie model, with a non-arbitrary cutting point, or knot, or top event. All indicators to the left, or upstream of the cutting point are lead, and all to the right, or downstream, are lag indicators. Moreover, as Hale notes, the bowtie model invites us to conceive of lead and lag as two ends of a continuum, rather than a dichotomy. The crucial point for the following discussion is that the knot in the bowtie is defined as the point at which control of the process has been lost.

The response by Webb in this debate provides us with a wonderful opportunity to compare theory and practice in this matter. Here are two so-called "lag" indicators adopted by his company:

1. Safety related protective system called into operation, e.g. emergency depressurization/flaring, operation of pressure relief devices.
2. Safety related unplanned shutdowns ... covers emergency stoppage of an individual plant, either automatically or by operator intervention, to prevent a loss of containment.

The question provoked by these examples is this: do these cases involve a loss of control? In every scenario envisaged by these indi-

cators, what we are seeing is the operation of one of the preventive controls, that is, one of the controls designed to prevent a catastrophic event. So, for incidents involving flaring and unplanned shutdowns to prevent a loss of containment, we are still one or more barriers away from disaster. From this point of view, the process remains under control. Accordingly, from the bowtie perspective these are lead indicators, not lag indicators.

What do we make of this? Either Webb is wrong in wanting to call these things lag indicators, or the bowtie model is not as good a solution to the problem as it at first appears. If we accept that it is appropriate to call Webb's indicators lag indicators, and I suspect that intuitively most of us would want to do that, then the question arises: how far back along the chain of causation do we go before lag becomes lead. For instance, Webb talks about the operation of a pressure relief device as a lagging indicator. What about the activation of a critical alarm that requires operators to throttle back a process? In short while the bowtie model in theory appears to provide a clear and non-arbitrary distinction between lead and lag, when we examine the situation more closely matters are a good deal less clear cut.

There is a second, quite different problem with the bowtie model that is readily grasped with an example. A deluge system is a fire fighting system and, as such, it is part of the mitigation system, designed to operate *after* the loss of control. Consider an indicator like "percentage of times a deluge system fails on testing". Is this a lead or lag indicator? Given that it refers to a barrier on the right side of the bowtie, it might well be considered a lag indicator, but it is also a proactive measure of preparedness and from that point of view, a lead indicator. Kjellen's response provides one answer. In his Section 4.1.1 he talks about performance measures of the status of fire and explosion barriers under the heading of leading indicators. If we accept this resolution of the problem then, as I understand it, the bowtie model no longer distinguishes between lead and lag indicators.

To summarise, the bowtie model does not provide us with a fully satisfactory account of the distinction between lead and lag. Nor do the relativist definitions discussed earlier. With due respect to the respondents, I remain unconvinced about the value of this distinction.

What is an indicator?

Several respondents took exception to my claim that, where possible, safety indicators should be based on harmful outcomes. In thinking about this objection I realized that I was working with a particular meaning of "indicator", what I might call a literal meaning of "indicator". On this view, an indicator is an indicator of something else; it is a surrogate for some other, unmeasurable variable. Wreathall puts it well:

Indicators are proxy measures for items identified as important in the underlying model(s) of safety. As such they are uncertain and often only distantly connected to idealized measures that are rarely available in practice".

If we ask the question: what is the unmeasurable variable of which lead and lag indicators are both indicators, it is surely something like "the state of the safety management system". If we have meaningful data on accident rates, I would argue that this is a more valid (in the technical sense) indicator than indicators based on failure rates of various risk controls. My point is simply that if the risk control measures suggested that the safety management system was improving, but the accident rate was actually getting worse, we would surely want to conclude that the state of the safety management system was deteriorating.