Organisational Amnesia: Is there a cure?

Jonathan Berman and Michael Wright, Greenstreet Berman Ltd, 10 Fitzroy Square, London W1T 5HP

Many major accidents or failures have occurred in long established organisations with well developed safety management systems. Recent examples include Buncefield, and BP Texas City and Deepwater Horizon. Such incidents have underpinned the notions of drift and organisational amnesia. These started to emerge approximately 10-15 years ago, in response to incidents which suggested that high-performing organisations appeared to exhibit gradual degradation in performance that was difficult to self-detect. In some cases the underpinning causes are similar to previous incidents in the same or similar organisations, indicating a 'loss of memory' and failure to learn.

Research into the causes of drift and organisational amnesia identifies two pre-conditions, being: degradation of key aspects of operational standards (such as competencies and safety awareness); failure of oversight processes to detect and reverse the degradation.

There is often a "common cause" effect. The factors that cause degradation in standards also lead to degradation of oversight. It's less clear why that degradation occurs, and why the organisation fails to 'remember' previous events. Is it failure of learning, or failure of memory, or both?

Some of the early research indicated that effective oversight processes need to be independent from, or resistant to, the influences that could degrade oversight itself. The following types of defences were identified: external anchoring to support early identification of drift; a regulator or external authority with the appropriate power and mechanisms to require concerns to be addressed; awareness of, and constant vigilance by senior staff to the potential for drift and a willingness to challenge and be challenged.

One implication is that organisations will struggle to defend against drift without external support – that some form of organisational amnesia is inevitable and requires an independent agency to combat it. However, two new areas of work have emerged, namely Resilience Engineering and 'Safety II'. These hold out hope that organisations can be more self-reliant.

Resilience Engineering advocates being able to react, monitor, anticipate and learn rather than attempting to optimise the process for pre-defined scenarios. Resilience is therefore about managing uncertainty, and improving the learning process. 'Safety II' is oriented towards acknowledging system complexity, pursuit of success (rather than prevention of failure) and increasing options. 'Safety II' refers to the focus on agility and on understanding and supporting success.

This potential shift from coping with uncertainty to managing uncertainty demands different ways of organising and controlling performance that permit greater flexibility. These approaches may provide insights into controlling flexibility and defending against drift – rather than trying to keep the organisation static. The concepts may help develop management defences against 'loss of memory' and hence drift. The challenge for most organisations is how to 'operationalise' resilience concepts.

The paper concludes by drawing together a model of drift and organisational memory, and potential defences, with examples highlighting the lessons learnt for operators of complex major hazard facilities.

Keywords: Organisational Learning; Organisational Memory; Drift; Resilience; Safety II.

Introduction

Amnesia – typically defined as a deficit in memory caused by some form of trauma or disease as a result of which memory is partially or wholly lost (retrograde) or memory traces are unable to be created (anterograde). Whilst the concept applies to people, and the causes might be physical or psychological, it often seems as if organisations are vulnerable to the same apparent effects. Many organisations appear to be condemned to repeat their mistakes and failures, either by failing to recall the causes of past failures and hence not changing their behaviours (retrograde amnesia), or by failing to identify and understand the learning points that should be taken from incidents and failures (anterograde amnesia).

Unfortunately it is relatively easy to find examples of apparent amnesia. The failures at Buncefield bore resemblance to those in previous incidents going back as far as Flixborough. The Hatfield and Grayrigg train derailments shared some similar causes. BP's experiences with Grangemouth, Texas City and Deepwater Horizon appear to show not only that incidents can continue to arise but also that there might even be an implied periodicity: one could postulate a seven-year cycle of incident, response/improvement, gradual deterioration, incident...

One can also look at events beyond the confines of major accident hazards and see the same dysfunctional processes. The failures at the Mid-Staffordshire Hospital Trust did not arise overnight and yet the catalogue of shortcomings and problems did not appear to create effective 'memories' that allowed the organisation to change and improve. The crisis that affected Hospital A&E Departments across the country at the beginning of 2015 appeared to have been well signposted, but effective action was clearly difficult to identify and implement. The rail network in southern England appeared to be paralysed just after Christmas 2014 due to overrunning engineering works. That such major programmes of work overrun when planned for very narrow timeframes is hardly a surprise, and yet there appeared to be little in the way of effective contingency plans available to respond to the ensuing delays. This was not the first time that such 'engineering overruns' had occurred and hence it is legitimate to question why there appeared to be ineffective learning from those previous events. Was it that the

previous failures had not been fully understood, and hence not adequately defended against subsequently? Or that those previous failures were properly investigated, but the corrective measures not properly implemented? Or that those corrective measures had been implemented but then allowed to degrade, perhaps due to a failure to retain the linkages between the control measures and the factors contributing to the original shortfalls.

A number of important and related issues arise from even a cursory examination of these and other similar events. One is the apparent difficulty that organisations have in learning effectively from past experience and failure. Another is the manner in which organisations seem to find it difficult to detect gradual erosion in performance. Neither of these is new. The challenge of learning from experience has been widely discussed and debated.

Two recent papers at Hazards 24 address this issue. Evans and Martin (2014) discuss the challenge of information management within an organisation, whilst Wilkinson and Rycraft (2014) approach the issue in part from the perspective of learning theory. In both papers there is a clear recognition both of the complexity of the issue, and of its importance. In the Department of Health report 'Organisation with a Memory' (2000), the NHS had a key document that demonstrated that it understood the importance of such processes. The concept of Organisational Drift has been recognised for many years (e.g. Berman and Ackroyd, 2006). Dekker (2011) has discussed the causes and consequences of drift at length. And yet despite the recent work, at times it seems as if we are no nearer to finding solutions that will prevent or restrain amnesia and drift.

One conclusion from the apparently inexorable march of drift and failure might be that we need to find alternative ways of addressing these issues. Over recent years two related themes have emerged that might offer some insight into solutions to organisational amnesia, being Resilience, and 'Safety II'. This paper suggests that drift and 'amnesia' are inter-related, and considers how we might more effectively combat them. A further observation from this paper is that, although its focus may appear to be on High Reliability Organisations, in practice the issue are relevant to all organisations and across all sectors.

Organisational Drift

In their paper on Organisational Drift, Berman and Ackroyd (2006) noted a number of characteristics that defined the challenge. The degradation in performance occurred in high-performing organisations. These were often ones that had been held up as exemplars of good safety performance. Following the 'trigger incident' it became apparent that performance had been gradually degrading over a prolonged period. That gradual degradation in performance had remained undetected by the organisation despite the presence of systems intended to monitor and reveal such changes. The defining characteristic seemed to be the gradual, incremental degradation, which meant that in the short term the changes remained below a threshold of detectability. The implication, in part, is that the organisations did not have effective processes for maintaining an overview of performance that spanned longer periods of time.

A number of possible defences against drift were postulated, including improvements in learning processes with a greater focus on trends and generic implications of incidents, considering more carefully the implications of all change opportunities, considering how standards are monitored (both adherence to standards and the standards themselves), and looking carefully at all aspects of oversight processes including benchmarking and external anchors. External anchors and benchmarking potentially has the benefit of being capable of going beyond validation of the system as delivering what is intended, towards validation of the system as delivering what is needed.

There is a debate that can be had about whether drift is inevitable, and also about whether merely having the capability to recognise drift is a sufficient defence. These are for another paper. What can be taken from the earlier discussions of drift, however, is that there are strong links between the potential defences against drift or its effects and the organisation's 'memory' and its learning processes.

As suggested above, one of the defining characteristics of drift is that the slow change does not trigger concerns about changes in performance – which can be rephrased as a suggestion that the organisation appears unable to recall what it originally set out to achieve or why it imposed the success criteria that it did, and hence that failure of memory prevents it from seeing the gradual change. It's only when we come to redecorate a room at home that we suddenly recognise how shabby it has become – the chipped woodwork and faded wallpaper has been right in front of us for years, but we no longer have sufficient memory of what the room first looked like to trigger an appropriate level of concern. We adapt to the new environment due to the slow pace of change without even recognising that such adaptation is taking place.

What becomes interesting is the question of what makes the learning process and organisational memory less than effective? Most organisations have processes in place to support learning. But these processes are often not coherent or well integrated. Typically they might include the following:

- Event investigation – root cause analysis; action identification; action tracking
- Occurrence reporting – near miss reporting; trend analysis; good practice identification
- Business Improvement Teams – proactive review; process improvement
- Learning and Development – training; staff competence management
- Knowledge Management – records; documentation

Clearly these address inter-related issues, and yet often they are distributed across the organisation and may not always have effective communications processes. But these are likely to be the processes that are best placed to detect gradual change, and to respond to it.

Defences

As noted above, whereas the types of defences that might mitigate the onset or effects of drift can be identified, it is more difficult to identify how best to ensure those defences are effective. A challenge that arises from drift is that the very processes that might be capable of detecting drift (oversight, monitoring, etc) are themselves vulnerable to the factors that encourage drift (complacency arising from good performance, gradual assimilation of change, adaptation to change, etc).

A number of potential defences against drift can be considered, and these can be grouped into three areas:

- Defences that focus on revealing and understanding changes in performance
- Defences that focus on recognising and understanding drivers for change
- Defences that focus on retaining understanding of the safe operating envelope and the reasons for those boundaries

Performance change

In one sense this is the bread and butter of safety management. All organisations identify and monitor key performance indicators. There is much effort devoted to defining both leading and lagging indicators of performance. Lagging indicators can receive a bad press due to their inherent retrospective focus. They are necessary, however, as a form of robust indicator of performance. The parallel importance of leading indicators is incontrovertible as a means of predicting future performance and hence identifying improvement opportunities and potential shortfalls before they are revealed in deteriorating safety. Typically, leading indicators focus on aspects of performance such as behaviours and activities – the performance of the safety management system.

In principle these leading indicators should be an important part of the defences against drift – the KPIs should be chosen such that they can reveal gradual changes in standards and the manner in which they are being attained. In practice they may not always achieve this. It is a truism that you get what you measure – that a focus on a particular KPI elevates that KPI into a goal, rather than a measure of a performance in support of a different goal. A KPI might be identified associated with, say, maintenance revisits – the number of times that planned maintenance has to be repeated due to a defect not being corrected, or becoming apparent after the planned maintenance is completed. This can be considered to be a good indicator of the performance of the maintenance process. Another related KPI might be missed scheduled maintenance, which can be an indicator of tension between production and safety. In both instances it is possible that the behaviours of operations and maintenance become focused on attaining positive measures on both of those KPIs rather than focusing on ensuring effective maintenance. Maintenance tasks are allowed to remain 'open' in order that it is recorded only as a single visit; maintenance priorities are changed to address KPIs rather than actual plant defects; etc.

What becomes eroded is the understanding of why the criteria against which the KPIs have been defined have been selected and, importantly, why a given maintenance schedule has its defined periodicity. It therefore becomes less likely that the significance of gradual changes will be fully recognised. Thus, what is important is that there are effective mechanisms both for identifying changes in performance and, importantly, for those changes to trigger a response. It is not sufficient to have processes that will monitor KPIs unless they also are capable of prompting a response from the organisation when those KPIs are not met.

It is almost inevitable that, following any significant incident or accident, it becomes apparent that some of the indicators of failing performance had identified deterioration but that no organisational response had ensued. The Challenger shuttle disaster revealed that the problem with the insulation tiles was recognised, but the process did not seem capable of marshalling a suitable response. In 2013, an Airbus A320 aircraft had to make an emergency landing when it lost an engine cowling shortly after take-off due to a latch being inadequately fastened after routine maintenance. It then became apparent that a number of previous incidents had occurred with this aircraft type but none had been considered a 'flight safety issue'.

The potential weaknesses in performance monitoring are therefore oriented around the gradual loss of sensitivity of those processes. The inevitable focus on measurable performance (in whatever form) requires also an effective means of retaining sensitivity to the small deviations from defined acceptability. The very nature of performance variability militates against simple 'pass/fail' criteria, but this then leads to 'normalising' deviation, with further pressure to do so coming from commercial constraints (e.g. avoiding an unplanned plant outage that would otherwise be required in order to respond to degraded performance). Methods that avoid normalising deviations become important.

A further potential weakness is the converse of normalising. Variation in performance can be 'anomalised' and hence discounted as a single deviation that is not typical of the system. This then gives 'permission' to ignore the deviation. Thus, both normalising and anomalising reinforce the importance of effective performance review, trending, and hence an adequate organisational memory. Indeed, at the risk of extending the human memory analogy further, one could consider trend and other similar analyses as being representative of the organisation's 'working memory' – and without it there is great difficulty in understanding observed events.

An additional challenge arises from a natural desire to 'make the system work'. One of the strengths of having people as part of the system is their ability to adapt, and yet it is these very adaptations or 'workarounds' that can mask gradual

performance degradation. In almost any organisation there will be various informal measures and arrangements that ensure that the system performs effectively. These are not 'reckless shortcuts', nor violations intended merely make work easier or quicker, both of which are rightly defended against. Rather, these are the small accommodations that ensure that the system functions smoothly. The additional undocumented pre-use check that people have learned to perform. The additional check that an item of equipment is electrically isolated even though it is tagged. The informal communication that work is in progress even though there are appropriate work authorisations in place. In each case these undocumented practices will have arisen through experience – maybe near misses or maybe actual minor incidents. But the very fact that these small accommodations are in place has two effects. One is to demonstrate that some form of drift or change is taking place whilst the other, counter-intuitively, is to mask the precise nature of that drift. Unless there is a focus on revealing the presence of these accommodations, drift will be hidden and yet measures of the frequency or extent of those accommodations might be a powerful indicator of drift.

Drivers for Change

Whilst it is both necessary and important to be able to recognise change and understand its significance (e.g. monitoring safety incidents), it is also important to be able proactively to identify those factors that might create or aggravate changes in safety performance, in order to be able either to reduce their effects, or to manage the changes if there are reasons for implementing them that are beneficial to the organisation.

Some change is both easy to identify and easy to understand its potential safety significance. Organisational change that arises from a wish to reduce headcount in the organisation is clearly one that needs to be managed effectively. At the same time, more careful consideration might be needed to identify the impact of such changes on emergency response capability.

However, other changes within an organisation can be more subtle and incremental, and these will be more difficult both to identify and to evaluate. The workforce demographics is likely to be changing. If an organisation is perceived as being static or declining, then it might become more difficult to recruit young engineers and hence the overall workforce might be aging. A significant proportion of the UK nuclear industry has entered the decommissioning phase and hence may not appear as attractive in terms of career opportunities. The German nuclear industry is being closed down entirely, albeit over a long period of time. Plant will be aging and investment programmes might alter. New local employers might affect the skill sets available to an organisation when recruiting. Changes in technology might alter maintenance requirements that de-skill relevant staff.

All of these small incremental changes will affect behaviours and performance. Some of the defences against that ensuing drift in performance would be those associated with performance monitoring, but others can be more proactive in terms of identifying the potential for drift and reviewing the changes to prevent it, perhaps by reinforcing relevant standards or ensuring that the changes do not compromise other aspects of safety management. The common theme is to be capable of understanding the factors that give rise to incremental change, and hence better managing the change process. The challenge for drift is that each of these incremental changes, on its own, is not judged as being significant, but that collectively they exert significant pressure for uncontrolled change.

Understanding the Safe Operating Envelope

One of the attributes of drift is a gradual change in performance standards. The conventional leading indicators that monitor adherence to standards and compliance with process will tend to continue to provide positive reports, whilst the actual performance may have changed. It is quite typical that conservative assumptions are built into safety justifications for new plant or equipment, with these being subsequently reviewed in the light of experience. Thus, the defined Safe Operating Envelope (SOE) might initially impose demands to limit inventory to a particular amount. Over time, small deviations from that limit might arise which appear non-consequential. Although the organisation might continue to impose strict adherence to the SOE, the perceived significance of that element of it becomes eroded. For production reasons the request might later be made to relax that element of the SOE – because experience has suggested that it is excessively conservative.

It is at this stage that two defences become significant. One is the imposition of a strict change management process that properly evaluates the impact of change. The other is having the knowledge and expertise properly to evaluate the proposed change. Whilst the change management process may be well documented and comprehensive, it may be excessively reliant on the knowledge and experience of the participants in the process to understand the original reasons for the imposed SOE and hence to be competent to evaluate the impact of change.

It becomes important that there is an appropriate level of corporate knowledge and 'memory' concerning the reasons for the SOE. A water treatment works providing drinking water to a significant number of customers was a critical item of civil infrastructure – loss of the output from the treatment works would create a significant potential health risk. Diverse power supplies came onto the site to reduce the risk of loss of output due to loss of off-site power. Additionally, the site had a set of diesel generators to support the works if there was no grid connection. The arrangements had been in place for many years, when an unrevealed common cause failure off site resulted in loss of grid. The diesel generators failed to auto-start and works output was lost (and the period of loss was extensive because the diesel generators then would not synchronise to a de-energised board). It was subsequently revealed that there had never been a provision for the diesel generators to auto-start. They were not 'emergency generators' – they were 'standby generators' intended to cover for planned interruptions to supply by being manually started before the grid connection was lost (e.g. to cover planned maintenance). Over time, this aspect of the design became overlooked and forgotten – although after the event a few long-serving members of staff recalled the original design intent, the reality was that decisions had been taken over many years based on misunderstanding the role and function of the diesels. The loss of memory of the design intent led to a challenge to the SOE.

Maintaining Defences

From the foregoing, it can be seen that a suite of defences against drift is likely to be required, spanning methods for monitoring performance and detecting change, methods for understanding and managing the drivers for change, and methods for understanding the implications of change.

However, each of these 'tools' will themselves be vulnerable to the effects of drift. In particular, the tools that might be most capable of detecting drift are those that focus on the performance standards against which the organisation is assessing itself and which provide assurance that the standards remain valid and appropriate. To be capable of maintaining suitable scrutiny and oversight of those standards there is a need both for an effective process for learning and 'recalling' previous challenges to the validity and effectiveness of those standards, and for benchmarking the standards and associated performance against external organisations.

One of the key defences against drift noted in earlier work was the maintenance of robust external oversight, recognising that the same factors that caused drift in performance were likely to cause drift in internal oversight arrangements. External oversight can take a number of different forms, but a common feature is the ability to relate the arrangements and performance in one organisation to those observed in another. The challenge that arises is the ability to understand the differences between organisations and hence to understand the significance of observed similarities and differences and to draw appropriate conclusions.

Over recent years there has been a steady improvement in the ability and willingness of organisations in one industry to seek out lessons from others. The Haddon-Cave report into the RAF Nimrod accident has become almost required reading across most high-reliability organisations, even though few of them have anything to do with aviation. Healthcare has looked to motor-racing to improve teamwork, although, interestingly, more recently there has been some transfer in the opposite direction. Healthcare is learning from civil aviation, and so forth. The two-way exchange between healthcare and motor sport illustrates another factor, which is that different organisations progress and drift at different paces, and hence whereas one sector might be more advanced at one moment in time, drift might then lead to a reversal. No single industry sector should be held up as the permanent exemplar in respect of safety performance.

This underlines the importance of considering memory and learning as an integral part of performance maintenance and the control of drift – there need to be effective mechanisms for ensuring that learning takes place using both internal and external sources, and that there are effective mechanisms for ensuring that the learning is maintained and refreshed. The results of learning need to remain accessible over time.

Current Developments

The picture painted above can be taken to suggest that drift is an inevitable consequence of operating in a dynamic world. Alternatively, it could be interpreted as suggesting that simple awareness of the potential for drift would be sufficient to maintain adequate and robust arrangements that defend against the effects of drift – its slow and incremental nature providing ample time and opportunity for an alert organisation to respond appropriately.

In practice neither of these positions is entirely correct. One characteristic that is repeatedly cited as a key feature of High Reliability Organisations is 'mindfulness', which includes the ability critically to self-reflect on performance. It is likely that a mindful organisation can make significant progress in defending against drift, but concrete arrangements also are necessary if the effects of drift are to be countered effectively.

Over the past few years, two related areas of research concerning high-reliability organisations have emerged, being Resilience Engineering and 'Safety II'. Both of these inform consideration of the control of drift and the importance of organisational memory.

Resilience

Resilience Engineering started to emerge at about the same time as the concept of drift was being considered. Hollnagel et al (2006) provides an example of the resilience concept. It emerged in response to the challenge of responding to the unexpected. Classically the intent within any safety management system is to reduce uncertainty and to be demonstrably defended against adverse events and failures. Resilience is a response to the recognition of complexity in systems, and hence to the challenge of attempting to predict all outcomes despite the inevitability of emergent properties within complex systems. Grote (2012) illustrates the balance that organisations seek to achieve between formalised control of predictable events and flexible response to uncertainty. She notes that stability comes from central planning, high levels of standardisation, high levels of automation and engineered systems, and constrained operations. Conversely, flexibility comes from local planning, reduced standardisation, technology that supports rather than controls, high degrees of freedom and feedback control. The challenge is to find arrangements that support both of these domains – stability and flexibility. She also notes that Resilience is about a shift from seeking to reduce uncertainty, to seeking to manage uncertainty.

Stability is 'comfortable' but almost inevitably there will be sets of circumstances that reveal shortcomings in the standardised arrangements – the causes of Deepwater Horizon, or Buncefield, etc. However, the adoption of flexible arrangements becomes a challenge for many organisations because it reduces the perceived levels of control. Personnel will adapt and respond to emerging situations. It then becomes more important that there is effective learning arrangements and 'memory' in place to enable robust decisions to be taken. There is a need to be able to recall not only what actions were taken previously, but also what were the circumstances under which those actions are judged appropriate.

Resilience is likely to offer significant benefits in terms of improving the response to abnormality, but it may both create increased opportunity for drift and conversely increased opportunity for control of drift. Permitting greater flexibility could increase drift – through normalising deviations from accepted practices and arrangements. This therefore demands alternative approaches to the manner in which deviation is managed. At the same time, if this threat arising from increased flexibility is recognised, then the greater focus on the management of flexibility could act to enhance the ability to detect and correct drift. There is currently little guidance as to how to balance the competing demands of control and flexibility, and this presents one of the key challenges for operationalising Resilience concepts. It would appear that successful approaches to doing so are likely to be highly reliant on the organisation's ability to understand and learn from experience and hence to understand how flexibility can be best applied and controlled.

If the application of Resilience concepts is likely to be reliant on learning processes, then it raises the question of what are likely to be valid and valuable sources of learning. Almost by definition, Resilience is oriented towards response to uncertainty, and hence each response will vary, and will vary for different reasons. This imposes a challenge for learning, particularly given that effective resilience will reduce the frequency of recordable incidents and hence reduce the availability of data. This in part reflects the drivers for Safety II

Safety II

Safety II is a further concept that has emerged in the past few years (e.g. Hollnagel, 2014; Eurocontrol 2013). Safety I refers to the conventional approach that has epitomised safety management for many years, whereby a reductionist approach is applied. Systems are analysed, scenarios are considered, and defences put in place to prevent credible failures. Where failures and shortfalls arise, they are investigated and analysed, and control measures implemented to prevent the failures from recurring. Whilst this approach is understandable, and a natural reaction to system shortfalls, it presents a situation of diminishing returns. High reliability systems are, by definition, robust and well-defended against failure. There are therefore few opportunities to obtain data that reveal shortfalls. If a system has a claimed human failure rate of 1.0E-05 per demand then it might be necessary to wait a long time before it is observed to fail such that corrective actions can be implemented.

This led to the recognition that there is an alternative approach, which has become termed Safety II. This focuses on success, and on improving the understanding of how successful performance is being achieved and maintained. Within high reliability organisations there should be significant opportunities to observe success, and hence to understand those arrangements that are providing the assurance of reliable performance.

This re-orientation towards maintaining success might offer significant benefits in respect of control of drift. It has the potential to provide a means of understanding variability in performance, and to legitimise acceptable variability without normalising deviation from formal arrangements. However, if it is to be effective, arrangements in support of Safety II have to be properly linked in to learning arrangements and organisational memory. There has to be a mechanism for recording and understanding the positive and negative impacts of flexibility and hence to understand how successful performance is being supported. Importantly, it also requires effective methods for disseminating the understanding that emerges.

Thus, whilst Resilience and Safety II are potential routes towards improved control of drift and performance degradation, they are strongly reliant on the effectiveness of organisational learning and memory. Whereas Safety I tends to be focused on stability and keeping the organisation in the same state, Safety II and Resilience are focused on agility and maintaining flexibility. The notion of an organisation in which change is expected and managed further increases the demands on memory in order to recognise and understand those changes. It is the inability to recognise that change is occurring which both defines drift and gives rise to failure.

Model

Incidents and major accidents appear to occur, in part, because of failures of organisational memory. Those incidents demonstrate repeated instances of similar failures, and of erosion of management controls. The gradual erosion of those controls, and the degradation in safety performance appears also to be characteristic of organisational drift. There appear to be strong links between the ability to control drift, and the effectiveness of organisational learning and memory. An indicative model emerges from this, as represented in Figure 1.

The model illustrates both the nature of drift – the interactions between Factors, Change and Drift – and the role of organisational learning and memory in defending against drift. The different elements of learning, comprising Investigation, Knowledge Management, Learning and Development, Occurrence Reporting and Business Improvement/External Learning all play a part in defending against different aspects of drift and performance degradation. For example, Knowledge Management can play a part in detecting erosion of standards and degradation of oversight, and can also support proactive identification of factors affecting drift. It is less capable of directly controlling change. Similarly, Occurrence Reporting might offer benefits in respect of monitoring performance degradation, but may not directly reveal change mechanisms.

The model should not at this stage be over-interpreted – it is intended to be indicative. However, a key message is the manner in which there are different aspects of Organisational Learning and Memory each of which interacts with the others, and needs to be integrated with the others in a coherent manner.

It is this integration of the elements of learning that is considered to underpin the ability of an organisation to be effective in the maintenance of high performance. Investigation alone is insufficient without an effective link to the Learning and Development process (so that the new knowledge becomes embedded in the competence of relevant personnel), and also

linked to Knowledge Management, so that the information is accessible to other staff as the need arises. Occurrence Reporting, particularly with a focus on success, needs also to be integrated with Business Improvement processes and with Knowledge Management. Safety II is predicated on the ability to take lessons from successful performance and to embed them into the organisational arrangements.

Figure 1: Model of Drift and Organisational Memory

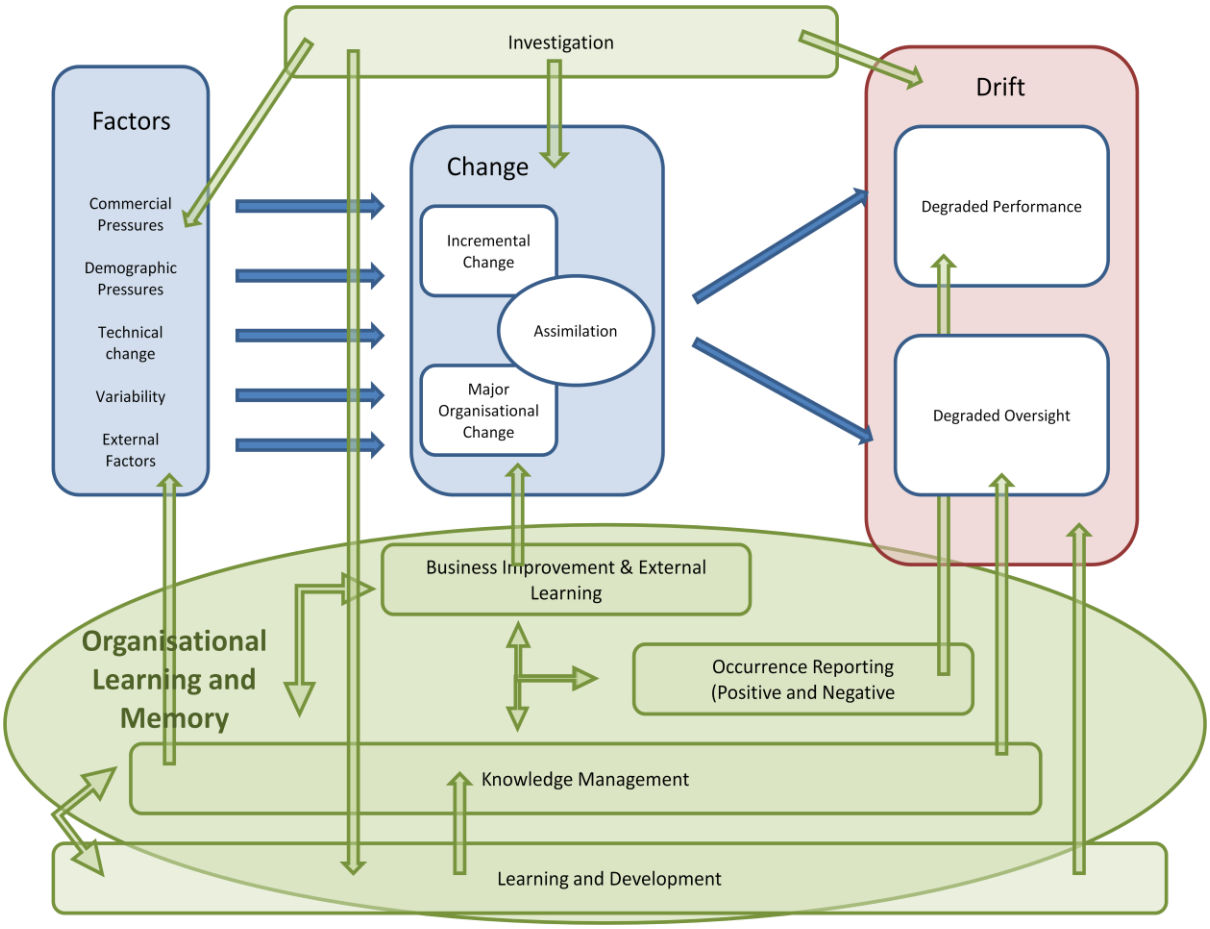

Recommendations

A number of potential recommendations can be derived from this perspective on drift and organisational learning.

Improving learning:

- Better integrate the different activities associated with learning within the organisation
- Ensure an effective focus on learning from success
- Consider how improved knowledge and understanding is cascaded down to the relevant personnel – learning requires change in behaviour
- Consider how subtle changes in performance can be identified
- Consider how changes in the organisation/arrangements can be fully assessed for their unintended consequences
- Consider how to ensure that investigation does not focus only on understanding causes of failure

Improving memory

- Consider the adequacy of knowledge management arrangements
- Consider how information can continue to be made accessible over time

- Ensure that robust audit trails remain to underpin understanding of change – the drivers for change, the rationale for the planned change, the implementation plan etc
- Ensure that all change is captured

Optimising resilience

- Identify where flexibility is required or beneficial
- Identify how flexibility is permissioned and supported
- Identify and monitor criteria for ensuring flexibility is being controlled
- Utilise learning processes to provide feedback control of flexibility

Deflecting Drift

- Integrate the information from all sources of learning
- Maintain sensitivity to weak signals
- Avoid normalising deviation
- Use organisational memory to validate the acceptability/non-acceptability of deviations

Conclusions

A key message from this review of drift and memory should be that organisations need to be flexible if they are to be able to resist drift towards poor performance. Incidents and major accidents appear to occur, in part, because of failures of organisational memory and the learning processes. Organisations appear to struggle both to generate an effective understanding of the causes of incidents and to maintain a proper and accessible memory of those causes. Whilst most organisations appear capable of responding to events and implementing change, the adequacy of new controls erodes over time. This may be, in part, a consequence of an attempt to control and stabilise the organisation, whereas it might be more profitable to recognise the dynamic nature of high-reliability, and hence to seek to ensure that there are robust learning processes that are sensitive to gradual change, and capable to ensuring that only beneficial change is allowed to remain.

This paper has sought to provide an overview of the inter-relationships between drift and organisational learning/memory. It has noted the emerging promise of Resilience Engineering and Safety II, but linked these to the importance of an effective learning and knowledge management process. It has emphasised the need for those various elements of learning to be integrated effectively.

There are various tools emerging that might help to guide some of these activities. However, there is a need for caution in implementing what appear to be 'ready-to-wear' solutions rather than bespoke systems. Some of the principles that might underpin improved performance are set out above. The manner in which they are implemented must take account of the characteristics of the organisation. Rather than focus on 'resilience' or 'Safety II' as a starting point, it is suggested that these should be considered themes that can offer guidance when reviewing learning arrangements.

Some tools are starting to emerge that attempt to support the development of resilience. Some of these are based on the premise that failure arises from an aggregation of everyday variables (and hence is not consistent with the concept of root cause). They emphasise the need to focus on 'work as done' rather than 'work as specified'. Whilst the concept of understanding how the organisation behaves in reality is sensible, and is consistent with the suggestions above that there needs to be a focus on how organisational behaviour differs from what might have been expected (as an indicator of drift), these approaches need also to be consistent with how memory is structured and hence how we track and monitor performance.

Memory (whether human or organisational) is focused on how work was done last time and what went right/wrong, rather than supporting a focus on what is different about today. Whilst resilience starts to focus on flexibility and the need to accommodate uncertainty and variability, there is a need also to ensure that learning considers 'process' and how the system performs rather than being oriented towards the outcomes alone. It is this focus on enhancing both the understanding of performance and the manner in which that understanding is retained in 'memory' which appears to offer real potential for the control of drift and the assurance of success.

It would be inappropriate to offer platitudes as to how past accidents could have been avoided, using the benefit of hindsight. However, this discussion of drift and memory suggests the following lessons for operators of complex major hazard facilities – most of which have already been learned, but many of which might be gradually forgotten:

- Don't assume that compliance with process and standards is sufficient – those standards may no longer be appropriate or may not be applied in the manner expected. There is a need constantly to review the validity of the internal and external standards that are being applied.
- Don't assume that Investigation is equivalent to Learning – investigation highlights what should be changed, but learning is the integration of those changes in a sustainable manner. Actions tend to be oriented towards making changes – learning is about retaining the understanding of why they were needed and what they are achieving.

- Don't focus only on the identification of root causes of failure – ensure that there is also proper identification of how success has been achieved. Formalise the understanding of success and hence the factors that could threaten it.
- Don't assume that the knowledge gained from operational experience (whether failure or success-based) is automatically accessible by all personnel. Learning requires constant and active intervention.
- Don't assume that all changes in performance will be readily apparent or understood. Enable trend analysis to highlight areas of gradual change. The dynamic nature of an agile organisation means that failures can become 'anomalised' and hence discounted.
- Don't assume that the ability to detect and evaluate change will remain robust. Accommodation to small change can mask the effects of those changes.
- Don't underestimate the importance of an integrated approach to learning and memory. None of the elements of learning are sufficient on their own and hence a coherent approach will maximise the benefits
- Don't allow deviations to be normalised. All deviations should be challenged, justified and managed.

References

Berman, J. and Ackroyd, P. (2006) Organisational Drift: a challenge for enduring safety performance. In Hazards XIX. IChemE Symposium Series 151. IChemE Rugby

Dekker, S. (2011) Drift into Failure. Ashgate

Department of Health (2000) An Organisation with a Memory. HMSO

Evans, J. and Martin, G. (2014) Corporate Memory – the Collective Brain and the Struggle against Amnesia. In Hazards XXIV, IChemE Symposium Series 159. IChemE, Rugby

Eurocontrol (2013) From Safety-I to Safety-II: A White Paper

Grote, G. (2012) Safety Management in different high-risk domains: all the same? Safety Science, 50, 1983-1992

Hollnagel, E. (2014) Safety-I and Safety-II. Ashgate

Hollnagel, E. Woods, D and Leveson, N. (2006) Resilience Engineering: Concepts and Precepts. Ashgate

Wilkinson, J. and Rycraft, H. (2014) Improving organisational learning – why don't we learn effectively from incidents and other sources? In Hazards XXIV, IChemE Symposium Series 159. IChemE, Rugby