The Impact of Styles of Thinking and Cognitive Bias on How People Assess Risk and Make Real-World Decisions in Oil and Gas Operations

Ronald W. McLeod, Ron McLeod Limited

Summary
Awareness of the psychological realities of different styles of thinking can provide deep understanding of the choices people make and the actions they take when they are faced with assessing risk and making decisions in real time under operational conditions. At a time when the industry is striving to achieve more with fewer staff and resources, there is a compelling need to understand better how these psychological processes actually influence real-world operations, and to develop practical approaches to mitigating the associated risks.

Introduction
There has been growing awareness over recent years of the risk that "cognitive bias" can present to operational safety. Biases such as normalization of deviance and group think are now widely recognized. The investigation into the 2010 Deepwater Horizon incident led to widespread awareness, at least among the offshore drilling community, of the impact cognitive bias can have on front-line thinking and decision making. Volume 3 of the US Chemical Safety and Hazard Investigation Board's (CSB) investigation into the Deepwater Horizon incident discusses the role that confirmation bias, "a one-sided case-buildng process of unconscious selectivity in gathering and using evidence that supports one's beliefs" (CSB 2016, page 58), had in the events leading up to that incident. In 2012, the International Association of Oil and Gas Producers (IOGP) published their report *Cognitive Issues Associated With Process Safety and Environmental Incidents* (IOGP 2012). The purpose was to improve awareness of how important these cognitive issues can be to safety. Papers given at SPE conferences have dealt with the impact cognitive bias can have on management decision making in oil and gas companies (Sykes et al. 2011) and its role in incidents (Thorogood et al. 2014; Crichton and Thorogood 2015). Shell drew on some of this knowledge in seeking to encourage a culture of "chronic unease" across its global operations (McLeod and Beckett 2012). And Koen (2015) discusses how recent findings in neuroscience are starting to identify the neurological mechanisms underlying these biases. As welcome as these developments are, they represent only the tip of the iceberg of the knowledge, understanding, and opportunities to improve the management of risk that are potentially available.

The purpose of this paper is to illustrate how such knowledge can be operationalized and used to gain deeper understanding of the nature of human error in real-world oil and gas operations. The paper is written from a psychological perspective, though it tries to illustrate the argument with examples relevant to oil and gas operations. While there have been previous attempts to apply this area of knowledge to the analysis of real-world incidents (see, for example, Thorogood et al. 2014; Crichton and Thorogood 2015) and to develop operational interventions (McLeod and Beckett 2012), such attempts have been limited to date, and have lacked the necessary research evidence. There is a compelling need to understand better how these psychological processes actually influence real-world operations, and to develop practical approaches to mitigating the associated risks.

The Psychology of Risk Assessment and Decision Making
Risk assessment and decision making are at the heart of the oil and gas industry, from the operational front line to the boardroom. It can, however, be surprising how rarely the scientific knowledge about these deeply psychological processes is applied to improve awareness and assessment of risk and the decision making, choices, and actions that flow from that assessment at the individual level. The report to the President by the National Commission on the BP Deepwater Horizon Oil Spill and Offshore Drilling (2011) uses the term "decision" 81 times and the phrase "decision making" 18 times. It refers to "risk assessment" in some way approximately eight times. Nearly all of these are references to risk assessment and decision making that take place away from the front line. They are performed by people expected to follow well-structured and approved processes, armed with all the relevant information, and following explicit guidance on how to make rational judgements. These are back-office risk assessments. None of them are concerned with the awareness or assessment of risk or decision making by operators at the front line, or with how those individuals assess the relative priority of the risks they are facing. They are references to risk assessment and decision making made using what psychologists refer to as "slow" or "System 2" thinking. Very few of them deal with the reality of risk assessment and decision making on the front line that will often be carried out by use of "fast" or "System 1" thinking. (These two styles of thinking will be summarized later in the paper.)

Tversky and Kahneman's famous paper "Judgement Under Uncertainty: Heuristics and Biases" (Tversky and Kahneman 1974) was the first systematic, scientific exploration of the ways in which intuition and cognitive bias can influence thinking and judgement. Since that first paper, a large number of scientists and researchers from many disciplines have contributed to what is now an overwhelming and powerful body of knowledge about how people make decisions and judgements in the real world. This work has been influential in many walks of life, perhaps most prominently in economics, the subject for which Kahneman was awarded his Nobel Prize in 2002. It is having a profound influence on how

This paper (SPE 179197) was accepted for presentation at the SPE International Conference and Exhibition on Health, Safety, Security, Environment, and Social Responsibility, Stavanger, 11–13 April 2016, and revised for publication. Original manuscript received for review 27 November 2015. Revised manuscript received for review 4 May 2016. Paper peer approved 14 June 2016.

governments make strategic policy decisions [see, for example, *MINDSPACE: Influencing Behaviour Through Public Policy,* produced by the UK's Institute for Government (Dolan et al. 2010)], and, increasingly, how they think about, and indeed measure, happiness. It is now recognized in its own right as the discipline of behavioral economics.

The knowledge that has flowed from those initial Tversky and Kahneman seeds also has an important role to play in understanding and managing human reliability in safety critical industries. Indeed, understanding this psychological knowledge base, learning from it, and developing practical interventions on the basis of it could be among the most important steps that can be taken to improve safety, environmental control, and human reliability in the oil and gas industry over the coming decades.

Behavioral Economics

The field of behavioral economics grew out of the overlapping interests of two general areas of scientific research:

- The interests of psychologists in the processes and mechanisms of human thought and decision making
- The interests of economists in how real people make economic decisions in the real world

The psychologist Daniel Kahneman and the economist Richard Thaler—leading thinkers from the two traditions—have recently published books giving their personal accounts of the evolution of the field (Kahneman 2011; Thaler 2015). Each book combines the story of the author's career with the emergence of key research findings about how people make decisions in the real world, and the psychological and behavioral theories developed to explain those findings.

One of the great strengths of behavioral economics has been that its research methods have drawn on simple yet highly compelling mental problems that everyone can experience for themselves of the "simplifying shortcuts of intuitive thinking" (Kahneman 2011, page 8). The most famous example is undoubtedly the bat-and-ball problem (see http://uk.businessinsider.com/question-that-harvard-students-get-wrong-2012-12?r=US&IR=T), though there are many others, such as: If it takes five machines 5 minutes to make five widgets, how long does it take 100 machines to make 100 widgets? [The answer is 5 minutes (Frederick 2005).] Thaler offers the example of his students complaining vigorously that his exams were too difficult when he set an exam with marks out of 100, which led to an average score of approximately 70%, but being quite content with an exam having a total of 137 marks available that produced an average score of 96 (also 70%): "To an economist, no one should be happier about a score of 96 out of 137 (70%) than 72 out of 100, but my students were" (Thaler 2015, page 4).

At the heart of behavioral economics is the way the human brain so quickly and easily, without our even noticing, comes to an erroneous answer in the former example, and in the latter, leads us into behaviors, attitudes, and emotions that are at conflict with how an economically rational person would be expected to think and behave. The understanding that has been generated through the study of such seemingly innocuous problems provides the opportunity to develop a deeper and psychologically richer understanding of both the causes and the nature of human error when it occurs in industrial settings than has usually been the case to date.

Few who have spent their career in engineering and the oil and gas industry will have had any reason or interest to explore and draw on this rather strange and nonintuitive area of science in their work or thinking. It has certainly (to date at least) had little impact on how industry goes about assessing risk or making decisions about reducing the risk of major accident hazards. However, those two recent books make accessible a great deal of knowledge about how real people, as opposed to what Thaler describes as Econs,* make decisions and assess risks. Together, though most prominently through Kahneman's work (Kahneman 2011), they provide an ideal starting point for anyone with an interest to begin to draw on this knowledge base to improve safety and reliability in industrial operations.

*The term Econs is meant to refer to the kind of person assumed by much of conventional economic theory. Someone who, when faced with economic choices, makes rational decisions in a way that produces economically optimal outcomes.

A Brief Introduction to System 1 and System 2 Thinking

From a psychological perspective, thinking and decision making can be thought of as comprising two distinct systems, or styles, of mental activity—what are referred to as System 1 (fast) and System 2 (slow) thinking. Psychologists still debate the characteristics, and some use different terminology to refer to the two styles of thinking; alternative terms include "automatic" and "reflective" for System 1 and 2, respectively. For the purpose of this article, Kahneman's description of System 1 and System 2 is sufficient [see Kahneman (2011) for a full introduction to the two systems].

System 1 is fast, intuitive, and efficient. It is always "on." It works automatically, requiring no effort or conscious control; you cannot turn it off. System 1 provides the basis for what are referred to as skill-based behaviors. Most of the time, System 1 works perfectly well. It is because of its speed and efficiency that we are able to think, act, and perform in ways that would simply not be possible if the brain had slowly and carefully to seek, process, and think through all of the information and options available to it. But, from the point of view of safety and the levels of human reliability that are now demanded in safety critical industries, System 1 thinking has characteristics that are far from desirable. It does not recognize ambiguity, does not see doubt, and does not question or check. It works through the association of ideas, a near-instantaneous mental network of association in which ideas or feelings trigger other ideas or feelings. If the mental network can quickly produce an interpretation of the world or an answer to a problem that feels comfortable, it will take it. System 1 is "a system for jumping to conclusions" (Kahneman 2011, page 79).

System 2, by contrast, is slow, lazy, and inefficient, but it is careful and rational. It takes conscious effort to turn on and it demands continuous attention; it is disrupted if attention is withdrawn. System 2 works by conscious reasoning. It looks for evidence, reasons with it, takes the time to check, and questions assumptions. System 2 is aware of doubt, and sees ambiguity where there is more than one possible interpretation. Switching between System 1 and System 2 takes effort, especially if we are under time pressure.

From the economics perspective, Thaler describes the two systems in the context of how self-control influences behavior and decision making by use of the metaphor of a "planner" (equivalent to System 2) and a "doer" (equivalent to System 1). Thaler says "... at any point in time an individual consists of two selves. There is a forward-looking 'planner' who has good intentions and cares about the future, and a devil-may-care 'doer' who lives for the present" (Thaler 2015, page 104).

The scientific evidence behind the two styles of thinking is extensive and highly credible. However much it may not be consistent with "the-way-we-think-that-we-think," the evidence is overwhelming: Everyone, irrespective of age, sex, culture, or the organizational controls surrounding how they are expected to work, must be assumed to be susceptible to them at some time. They do not, of course, apply all of the time (that is inherent in the two systems), and different individuls, personality types, and even professions appear to be influenced by them in slightly different ways. But when we are using a System 1 style of thinking, we should anticipate that our thinking will be influenced by cognitive bias and other sources of irrationality.

As far as human reliability is concerned, it is the extent to which System 1 is prone to bias and irrationality, its tendency to jump to conclusions, not to have doubt, and not to see ambiguity that can represent such significant risk. The pace and demands of operational performance and production will often rely on the speed and efficiency of System 1 thinking, and, most of the time, it will deliver

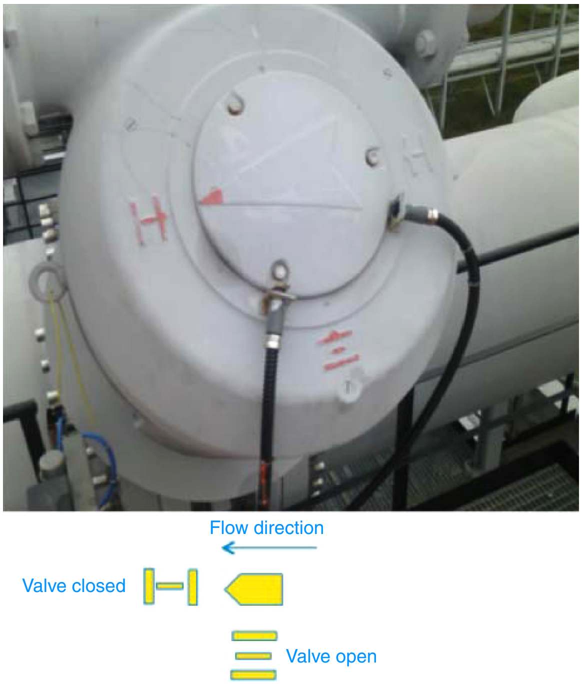

Fig. 1—Position indicator of a large automated isolation valve.

what is needed, safely and efficiently. Safety and environmental protection, by contrast, rely on System 2 thinking—they need care, attention to detail, and checking. They need awareness and the willingness to apply the mental effort necessary to engage System 2.

Operationalizing Fast and Slow Thinking

This section considers some examples of situations in which the characteristics, biases, and irrationalities associated with System 1 can lead to poor decision making and awareness of risk. It also illustrates how an understanding of the differences between the two systems can allow deeper insight into the nature of the loss of reliable human performance that lies behind many incidents. McLeod (2015) has explored in depth the contribution System 1 thinking may have made in two major incidents: (1) the fire and explosion at the Formosa Chemicals plant in the US in 2005 that led to five fatalities; and (2) the fire and explosion at the Buncefield fuel-storage site in the UK, also in 2005, that (while there were fortunately no fatalities) led to the largest peacetime explosion in the UK, as well as significant damage and disruption.

An Example of a System 1 Error in Operations. Fig. 1 shows the position indicator on a large, automated isolation valve. This valve was involved in an incident in a flare system. Fortunately, there were no injuries, which was a result of both operator action at the time and effective planning of the "line-of-fire" at the worksite, even though the incident was assessed as having the potential for up to three fatalities.

Among a number of factors that contributed to the incident (all human and organizational in nature) was an operator misreading the status of the valve-position indicator shown in Fig. 1. The operator believed the valve was closed when it was, in fact, open. As at most process facilities, the general rule was that a valve is open when the position indicator is in line with the process flow and closed when the indicator is across the flow. The opposite applied in this case: The pictograms showing the valve status were interchangeable based on the valve gearing and were not installed consistently across the plant. The large arrow on the face of the valve, which in the photograph is pointing to the left and is in line with the piping (and therefore in line with the direction of flow), does not indicate that the valve is open. Rather, the purpose of the arrow is to draw the operator's attention to the pictogram on the left-hand edge, showing that the valve is actually closed when it is in that orientation. When the valve is fully open (as it was when the incident occurred), the arrow in the photograph would point straight up and down (i.e., across the direction of flow).

In a conventional human-error analysis, this error would have been classified as a knowledge-based mistake (though, of course, the valve should have been installed in the correct orientation in the first place). It would be concluded that the operator, who was trained and possessed the knowledge needed to read and understand the position of the valve from the pictogram provided, had made an error in applying that knowledge, which, of course, would be correct. However, while it may be correct, it does little more than describe the nature of the error, and provides no real insight into why a trained and competent operator may have made such a potentially critical mistake.

By contrast, recognizing how the characteristics of System 1 thinking would be likely to play out in this situation can provide much deeper insight into why the incident occurred and what needs to be done to guard against similar incidents in the future. An operator busily engaged in a series of activities under time pressure will almost certainly have been carrying out this routine and inherently simple visual task using System 1 (fast) thinking. Looking at the valve indicator in Fig. 1, it is not surprising that someone adopting a System 1 style of thinking would quickly come to a decision on the basis of the large and visually dominant arrow. It would need effortful System 2 thinking to realize that the arrow is simply pointing to the pictogram, and not indicating the state of the valve. And, critically, having quickly decided that the valve was closed, the operator would have no doubt about the decision because System 1 does not question, does not see doubt or ambiguity. There is nothing about the design of this valve indicator that would lead the operator to double check or to question the first conclusion offered by System 1 thinking.

The fact that the incorrect valve position was not picked up by checking with the control room or when the isolations were later verified by a separate individual also has important implications. There have been many examples of incidents during which supervision and cross-checking were not as effective as barriers against incidents as they were expected to be (see McLeod 2016). This was recognized as far back as 1983 in the classic work by Swain and Guttman (1983), which has since provided the basis for most approaches to quantifying human reliability. Unfortunately, those and other similar warnings are frequently overlooked when estimates of the likelihood of human error are made.

System 1 Thinking in the Use of Risk-Assessment Matrices (RAMs). RAMs are probably the single-most widely used method for estimating and prioritizing risk across the global oil and gas industry. Their use is set out in international standard *ISO 31000:2009* (2009), and there are industry guides for the use of RAMs for a range of purposes ranging from contractor management (IOGP 2010) to risk-based inspection (*API RP 581* 2008). There are also innumerable sources that set out how RAMs are to be used to assess risk within individual companies.

Fig. 2 illustrates one type of RAM. Because of their ubiquity and the fact that implementations vary slightly across the industry, it is not necessary or practical to provide more than the following very cursory summary of the key concepts behind RAMs to support the discussion that follows:

- The risk associated with any event is assessed on the basis of a combination of estimates of the likelihood of the event occurring and, if it does occur, the expected severity of the consequences.
- Consequences are usually assessed separately for the potential impact on people, impact on the environment, damage to the asset, and damage to the company's reputation.
- Consequences and likelihood can be assessed either qualitatively or quantitatively. Qualitative assessments usually use

0–5 = Low Risk
6–10 = Moderate Risk
11–25 = High Risk
16–25 = Extremely high unacceptable Risk

Likelihood of the hazard happening	Severity of the potential injury/damage: Insignificant damage to Property, Equipment or Minor Injury 1	Non-Reportable Injury, minor loss of Process or slight damage to Property 2	Reportable Injury moderate loss of Process or limited damage to property 3	Major Injury, Single Fatality critical loss of Process/damage to Property 4	Multiple Fatalities Catastrophics Loss of Business 5
Almost Certain 5	5	10	15	20	25
Will probably occur 4	4	8	12	16	20
Possible occur 3	3	6	9	12	15
Remove Possibility 2	2	4	6	8	10
Extremely Unlikely 1	1	2	3	4	5

Fig. 2—Example of an RAM using qualitative ratings of likelihood and severity of potential consequences.

broad categories: A consequence of 3, for example, might mean "major injury," while a likelihood of C might mean "possible—incident has occurred in our company." Quantitative estimates, by contrast, try to specify either the probability of the event occurring (1 to 5%, for example) or the potential cost of the financial damage that could accrue (more than USD 20 million, for example).

The popularity and widespread use of RAMs stem from a combination of their intuitive nature and the ease of communication of the assessed risks. Formally, RAMs might be used during the course of a hazard and operability study, process-hazard analysis, or other safety analysis activity. They are frequently used in indicating the potential that could have been associated with an incident. RAMs are used to prioritize maintenance activities and are even being used to support competence-management systems (Kedzierski 2016). Less formally, they might be used as part of everyday work activities or meetings called to review design options or to find means of reducing cost or complexity of a project or operation.

Despite their popularity and ubiquity, the use of RAMs to assess risk has often been found to be flawed for reasons ranging from their inherent design and the logic behind their scoring systems to lack of consistency in their application. Thomas et al. (2014) reviewed more than 520 papers published by SPE in the last 15 years on the use of risk matrices in oil and gas activities (subsequently reduced to a sample of 30 papers that met the needs of their analysis). Their paper discusses a number of reasons to have serious doubts about the value of risk assessments that are produced by use of "best-practice" RAMs. For example, the relative ranking of risk can depend on whether increasing risk is associated with ascending or descending scores. Thomas et al. (2014) also identified the effect of a cognitive bias that is well-known in psychology as the centering bias: People asked to make judgements using rating scales tend to avoid choosing values at the extremes of the range of possible scores. From the papers they were able to analyze, Thomas et al. demonstrated that up to 90% of the RAM ratings were "centred" (i.e. the users had shown bias against selecting extreme scores irrespective of the actual risks involved). Duijm (2015) has recently reviewed the literature investigating issues with the use of RAMs and proposed an alternative approach that tries to capture the uncertainty inherent in risk assessments.

Fig. 3—View of poor access to cabinet at the rear of pump.

To provide a real-world context for this discussion, consider this example. A human-factors-engineering (HFE) review was carried out as part of the input to a capital project that would be expanding a plant's capacity. The company wanted to ensure any significant HFE deficiencies in the existing plant would not be repeated in the new design. One of the issues identified involved the daily task of checking the oil levels in a large pump that had been procured as a vendor-packaged unit. Performing the check involved entering a nitrogen-filled cabinet, which first had to be vented by opening its front and rear doors. While there was good access to the front doors, access to the rear doors was poor (see **Fig. 3**). Operators were forced to climb over tightly spaced equipment and piping to access the doors to make the cabinet safe to enter.

The challenge that faced the HFE specialist conducting the review, and the team tasked with reviewing and approving the report before it was issued to the project team, was to estimate how big a risk this HFE design issue represented. Specifically, was the risk large enough to justify a design change (which would be expensive and require a vendor to modify a standard design)? How would the individuals involved assess this risk using the company RAM? How much confidence or trust could the organization place in the resulting RAM score? In addition to the centering bias demonstrated by Thomas et al. (2014), there are (at least) three important ways in which the biases and characteristics of System 1 thinking could affect the accuracy of the assessed RAM score—base-rate errors, availability, and substitution.

Base-Rate Errors. Human beings have an extremely poor grasp of statistical reasoning. When we are asked to make judgements about probability, or the relative likelihood of different events, we are poor at taking account of the true underlying frequency of the events of interest in the world. This is true even when we are told what those base rates are, which is known as the base-rate fallacy. We are influenced by often-irrelevant information specific to a particular case. Our System 1 thinking quickly characterizes a situation in terms of previously created stereotypes, and we make judgements of likelihood that are based on those stereotypes—judgements and stereotypes that can bear little relation to the actual frequencies of the events in the world we are trying to estimate.

Even when we are given data describing the actual base rate of an event, if we are also given data that relate only to a specific instance, we are more likely to base our judgement on the specific instance than on the true underlying base rate.

In the case of the HFE design issue illustrated in Fig. 3, there is a recurring rate of injuries, up to and including fatalities, associated with people slipping and falling when using piping or other items of equipment as steps or as a means of accessing equipment. Similarly, there has been a history of fatalities associated with people entering spaces filled with inert gases. There are base rates associated with both of those types of events, but even if those base rates were known and were made available to the individuals tasked with coming up with an RAM assessment of the risks associated with the poorly designed pump unit, judgements of the likelihood of an incident would likely be biased (either higher or lower) by the stereotypes our System 1 thinking has drawn on in thinking about the situation.

Substitution. The substitution bias boils down to System 1 thinking answering an easier question than the one that has been posed: "If a satisfactory answer to a hard question is not found quickly, System 1 will find a related question that is easier and will answer it" (Kahneman 2011, page 97). So how can a valid assessment be made that would support the case either for making a change to the design of the pump unit shown in Fig. 3 or for accepting that the existing risk was already as low as reasonably possible, and that no change was justified? It is a difficult question, and, as with all difficult questions, the chances are that System 1 thinking will simply substitute for it with an easier question unless System 2 intervenes.

More research is required to establish precisely what nature those "easier questions" might take in this or any other situation, though a little speculation might suggest something like whether the individuals making the assessment have ever heard of a serious incident in a similar circumstance or whether they can imagine such an incident. However, these are not the questions that need to be answered. "Requisite imagination" is one of the core psychological attributes for achieving a state of "chronic unease" that is associated with vigilance, a lack of complacency, and awareness of and respect for operational risk (Fruhen et al. 2013). Just as it is difficult to develop and maintain a state of chronic unease, it can be exceedingly difficult to imagine the ways in which events might come together to create the kind of situations that in reality lead people to make decisions and behave in ways that can seem inexplicable. Answering the easier question of whether you can imagine an incident can bear little relationship to the real risk.

Availability. According to the availability bias, people estimate the frequency or likelihood of an event "...by the ease with which instances or associations come to mind" (Tversky and Kahneman 1974, page 208). In terms of the judgements that underpin RAM assessments, people will be less inclined to move toward higher-risk scores (more likely or higher consequence) if they find it difficult to think of examples to justify those scores, irrespective of the nature of the events that are recalled or the real base rate of such events. Availability actually works on the basis of substitution: "The availability heuristic, like other heuristics of judgement, substitutes one question for another: you wish to estimate … the frequency of an event, but you report an impression of the ease with which instances come to mind" (Kahneman 2011, page 132).

For most nonpsychologists, that is not a hugely surprising finding; it is clearly easier to think of examples of things that happen frequently than those that are rare. What is important about availability is that it is the subjective experience of how easily relevant instances are brought to mind that is important, rather than the actual content of the memory itself. "What renders this heuristic error prone is that the experienced ease of retrieval may reflect the impact of variables other than frequency, such as the event's salience or vividness" (Schwarz et al. 1991, page 201). So, rare events that make an emotional impact or have other features that make them highly memorable will be easier to recall (more subject to availability) than more-frequent events that are less salient.

The opportunity for judgements on both the likelihood and the consequence dimensions of an RAM to be influenced by the availability heuristic seem obvious. Typically, a likelihood scale might have options ranging from the event being "Never heard of in the industry" through to "Has happened more than once per year at the location." How does an individual or a team realistically make such a judgement? It is not usual—it may, indeed, be rare—for any individual or team making an RAM assessment to actually have valid data on whether the kind of events being considered have ever happened "in the company" (unless it was a very small company or had a very short history) or "in the industry." Following the Texas City explosion in 2005, BP was famously criticized for failing to learn from an earlier event at their Grangemouth refinery (Baker et al. 2007). Similarly, following the investigation into Deepwater Horizon, the drilling services provider Transocean was criticized for failing to learn from a similar incident they had experienced only a few months previously in the North Sea (National Commission on the BP Deepwater Horizon Oil Spill and Offshore Drilling 2011). Yet, the RAM model that is considered industry best practice continues to rely on awareness of events elsewhere in the company or the industry.

The reality is that the great majority of people involved in making judgements on where to position a risk on the likelihood and consequence dimensions of an RAM are unlikely to have the facts at hand when the judgements are made. They are likely to be under the influence of availability (and therefore substitution) to a greater or lesser extent.

Central to availability is the fact that the degree of confidence we have in a judgement depends on how quickly and easily examples can be brought to mind. Paradoxically, the more instances that are needed to confirm a judgement, and the more difficult it becomes to think of those instances, the less confident we become in the judgement, despite the additional evidence. Someone performing an RAM assessment might be able to quickly come up with an example of an occasion during which an event led to "major injury or health effect." They would be confident that the risk rating should have a severity of 3. Though, if they were challenged to come up with additional evidence to support the rating, thinking of the additional instances would become increasingly more difficult, reducing their subjective confidence in the assessed risk. The risk would not be lower (it may actually be higher given that additional instances are identified), though the RAM score might well be.

The availability bias can be overcome in a number of ways, not the least of which being when RAM ratings are discussed and reviewed by a group (although even then the dangers of biases such as anchoring and group think are always present). Though, review by a more-senior person is no guarantee that availability will not influence the assessment: Availability is more likely to influence judgements, both when they are made by "knowledgeable novices" and when the individuals making the judgements are (or are made to feel) powerful. So, a senior engineer or manager who is not a domain specialist, but has the authority to change or approve an RAM assessment may actually be more susceptible to the availability bias than the domain specialist who made the original assessment.

The strongest defense against RAM ratings being unduly affected by the availability bias relies on having the awareness and vigilance to be aware of the effects of the bias and to engage System 2 thinking when making or approving RAM ratings. Unfortunately, though, that takes mental effort, and "... many people… apparently find cognitive effort at least mildly unpleasant and avoid it as much as possible" (Kahneman 2011, page 45).

The challenge is to establish how good people and the systems and processes developed for assessing operational risk really are at overcoming base-rate errors, substitution, availability, and other

(a) (b)

Fig. 4—The engines shortly before takeoff (AAIB 2015), (a) showing the gap between the engine cowling and the doors, and (b) showing the door latches in their unlatched state projecting vertically below the cowling.

potential biases in making the judgements relied upon by the RAM model and other risk-assessment processes.

Confirmation Bias and Commitment

Confirmation and commitment are among the more widely recognized cognitive biases. They may even be the most important in leading people to make decisions and act in ways that are unexpected and contrary to what is needed for the safe management of operations.

Confirmation bias refers to the tendency actively to seek out information that supports or confirms what we believe to be the state of the world or to ignore or find ways of rationalizing away information or data that conflict with what we think or expect. Commitment refers to situations in which there are (at least to an objective outsider and perhaps with the benefit of hindsight) clear indications that the world or the state of an operation is not as we think or want it to be, but we nevertheless carry on with the same course of action.

There have been many well-documented instances of both confirmation and commitment bias. They have most frequently been reported in connection with aviation accidents, though they have also been reported in incidents in the oil and gas industry (see, for example, IOGP 2012; Thorogood et al. 2014; McLeod 2015; CSB 2016). Perhaps one of the most-dramatic examples of the ways these biases can create the potential for catastrophic accidents occurred at Heathrow airport in London over the night of 23–24 May 2013. This example is especially powerful and relevant to the oil and gas industry when considered in the light of the relative priority given to all aspects of safety throughout the global aviation industry. That includes everyone with a role in ensuring flight safety being required to undergo a form of crew-resource-management (CRM) training to be aware of the risks from human error and the wide range of factors that can cause or contribute to it—a situation that some parts at least of the oil and gas industry are working toward (see, for example, IOGP 2014a, 2014b; Energy Institute 2014a), but one that remains, as yet, more of an aspiration than a reality for most companies.

Shortly after 0700 on the morning of 24 May 2013, British Airways Airbus A319-131 took off from London Heathrow airport on a routine flight bound for Oslo. During the takeoff and transition to flight, doors covering both engines broke loose, causing damage to the airframe and a fuel pipeline. During the emergency return to Heathrow, the right-hand engine caught fire and was shut down. Fortunately, the aircraft landed safely with no injuries [Aircraft Accidents Investigation Branch (AAIB) 2015].

Over the previous night, two maintenance technicians had been assigned to complete a routine weekly check of the aircraft. Many issues contributed to the incident, nearly all of them to do with human factors, and all having similarities with the way safety critical activities are planned and carried out in oil and gas operations:

1. The technicians put the aircraft into an unflightworthy state without implementing the controls that were intended to protect against the possibility of that unsafe state going undetected before the aircraft's next flight.
2. They interrupted a maintenance task and left the aircraft in the unflightworthy state.
3. They returned to the wrong aircraft to complete the maintenance task.
4. They did not realize they had returned to the wrong aircraft and signed off on the aircraft as being serviceable.
5. A number of opportunities that could have captured and corrected the mistake were missed, involving cross-checking each other's work and intervention by colleagues.
6. Preflight safety inspections did not spot the error, and the aircraft took off in an unsafe state (**Fig. 4**).

For the purposes of this paper, the discussion will concentrate only on the fifth of these issues. The full incident investigation, including a report by a human-factors specialist, is contained in the report by the AAIB (2015).

In brief, the two technicians had, independently, carried out a check of an oil level in each of the aircraft's two engines. To do so, they had to open the engines' fan cowl doors. The checks showed that both engines required topping up of the oil, an event that usually only occurs in less than 1 in 20 inspections. Because they did not have the necessary equipment with them at the time, they delayed topping up the oil until later in the shift. Both technicians left the engine doors unsecured and in a condition contrary to written procedures.

Later in the shift, and having in the meantime completed a number of both daily and weekly checks on similar aircraft, they set out to complete topping up the oil; however, they went to the wrong aircraft, an event that is not unknown, but that is rarely reported to management (it is known in the industry as an "aircraft swap error"). It is, however, three events that occurred subsequently to arriving at the wrong aircraft that are of interest to this discussion of the effects of confirmation bias and commitment on human reliability in safety critical activities:

1. On arriving at the (wrong) aircraft, they found that the engines, which they believed they had previously left open and unsecured, were now closed and secured. According to the AAIB report "They thought this was strange, but they reasoned that a third party must have closed the fan cowls during their absence of approximately three hours" (AAIB 2015, page 34).
2. Having reopened the engine covers to top up the oil level, they found the oil no longer needed topping up. They apparently rationalized this by concluding that the engines must have cooled

down in the four hours since the engines had stopped, and therefore that residual oil would have drained back into the sump.
3. Despite subsequently mentioning both of these unexpected findings to colleagues in the crew room, no one questioned whether they had been at the right aircraft.

The crucial facts (a) that that they were easily able to create explanations about why both doors they had left open were now closed and why the engines no longer needed topping up with oil, and (b) that they appeared to experience no doubt or uncertainty about those explanations, are characteristic of System 1 thinking.

To rationalize away one surprising finding is understandable and entirely human. For two separate technicians, both trained and competent (including having completed the aviation industry's mandatory CRM training), to rationalize away two surprising findings in the space of a short period of time when performing an activity that they must have known was flight critical, gets into territory that is profoundly unsettling. Though, neither of these surprises was sufficient to make either technician stop and question what they were doing.

The incident must also raise questions about how effective CRM-type training in nontechnical skills (something on which the drilling community is placing great emphasis) can really be in overcoming cognitive bias (IOGP 2014a, 2014b; Energy Institute 2014a). Concerns about the usefulness of CRM training for aircrew have been raised in previous aviation incidents [see, for example, the report into the loss of Air France flight AF447 over the South Atlantic in 2009 (BEA 2012)]. Thorogood (2015) has also raised concerns over expectations about the value that CRM-type training can deliver in well control.

Apart from the implication of the effects of cognitive bias on safety critical decision making and task performance, one feature that is especially notable about the Heathrow incident is how ineffective was the opportunity for an operator to challenge and correct the decisions and actions of their colleague. This was also noted earlier in this paper in the discussion of the incident with the isolation valve. There is a real question for the industry in general terms about to what extent such cross-checking can ever be relied on when people work closely in teams.

System 1 Thinking and Fatigue. Is there perhaps something about the people who work in oil and gas or the way oil and gas activities are organized and managed that makes their operations robust or resilient against the effects of these cognitive biases? It is true that the subjects studied in many behavioral-economics experiments will predominantly have been students (such is the nature of much psychological research). But that does not mean the evidence does not generalize to real-world industrial operations. Even to operations in which there is strong safety leadership, good safety culture, and comprehensive and effective systems, including procedural controls over how work is organized and performed.

Shift work, often involving 12-hour shifts (sometimes more), and rotating between day and night work, are common across the industry. Many people work rotations of 14 or 28 consecutive 12-hour shifts, often with at least one change between day and night working in the middle of the rotation. Many people travel for long periods, sometimes over multiple time zones, before arriving for the start of a rotation at an offshore or remote work site. For these and other reasons, fatigue is increasingly recognized as a significant risk (IPIECA/IOGP 2007; IOGP 2014b; IOGP-IPIECA 2014; IPIECA/IOGP 2015; Energy Institute 2014b; *API RP 755* 2010).

A fundamental difference between System 1 and System 2 styles of thinking is that it takes effort to engage and apply System 2 thinking. One of the most-important effects of fatigue is a general reduction in motivation and energy and in the willingness or ability to apply effort. You do not need to be a psychologist to speculate that one of the effects of fatigue is, therefore, likely to be to make people less likely to go to the effort to apply System 2 thinking in critical (indeed, in any) situations.

People who are fatigued should be expected to be more likely to be subject to cognitive bias and to interpret the world, make judgements, assess risks, and make decisions in ways that are not a rational reflection of the actual state of the world or the evidence or information available to them. They should be expected to be more likely to jump to conclusions, to rationalize away unexpected findings, not to doubt their decisions or to consider alternative explanations, and to be less likely to apply the mental effort needed to stop and check what they are doing. These are all characteristic of System 1 thinking, so it is possible, perhaps likely, that the oil and gas industry may actually be more exposed to biased and irrational thinking than many other areas of life.

In the case of the incident at Heathrow airport discussed previously, the AAIB's human-factors specialist conducted an analysis of the hours worked by the two technicians in the previous 3 weeks. They concluded that the technicians were likely to have been experiencing a relatively high level of fatigue (though the estimated fatigue was lower than will be found in many oil and gas operations, especially immediately following the transition from day to night working). It is perhaps not surprising that the technicians involved should have been prone to System-1-thinking-based cognitive bias, that they would easily find explanations for the surprising finding that the fan cowl doors they knew they had left open, were now closed, and that the engines that had previously been found with low oil levels, now no longer needed topping up.

Conclusion

This paper has sought to illustrate the operational significance of what is a very substantial and high-quality body of scientific knowledge about the nature of human thought processes and different styles of thinking. Additionally, it has tried to illustrate the potential implications of "fast" or System 1 thinking in terms of the awareness and assessment of operational risk and decision making. The paper has touched on only a small portion of the psychological knowledge that is available.

Much of the oil and gas industry and other safety critical industries appear to hold an implicit belief that people who are competent, in a fit state to work, working in a strong safety culture, and who are provided with work instructions, standard operating procedures, and other forms of work controls, will in some way be immune to the biases and shortcuts associated with System 1 thinking, that management systems and other controls over work will in some way ensure or enforce System 2 thinking in critical situations. There is no evidence that such a belief is true. Indeed, the scientific knowledge gathered over the previous 4 decades strongly suggests such a belief will be false in many (perhaps the majority of) real-time work situations.

In recent years, the oil and gas industry has been asking what more it needs to do to identify and manage the risks that can be associated with major incidents. On the basis of the arguments set out in this paper, inescapable conclusions include that the industry needs

- To recognize that sensitivity to and awareness of risk, and the real-time decision making that flows from them, involve deeply psychological processes.
- To recognize that those psychological processes are complex; indeed, far more complex than those involved in disciplines such as well or process engineering. While they usually support the levels of skilled and efficient human performance the industry relies on, they are also subject to many sources of irrationality and bias that can lead to poor risk awareness and poor decision making.
- To identify the steps that can be taken, both in back-office risk assessments and those made in real time under operational pressures, to avoid the biases associated with System 1 thinking that can lead to ineffective management of operational risks. For example, it should be possible to develop an operationalized variant of the cognitive reflection test developed by Frederick (2005). Such a test could offer the potential to identify an individual's

susceptibility to System 1 thinking errors and, perhaps, screen extreme-scoring individuals out from safety critical activities that have a particularly high dependence on real-time decision making. Use of such a test in awareness training would in itself help to raise awareness of the risks of System 1 thinking.
- To identify practical steps that can be taken, whether in the design of work systems and the work environment or in working practices and procedures, that can be effective in disrupting System 1 thinking and engaging System 2 thinking in critical situations.

Addressing the needs in the third and fourth bullets in the preceding list requires investment in applied research, drawing on the combined skills and experience of psychologists and those with operational experience.

There are many cognitive biases associated with System 1 thinking. Kahneman himself at one point reviewed more than 40 of them, and, crucially, he concluded that they all favor “hawks” over “doves” (i.e., they are all inclined to lead people to make riskier decisions than they otherwise would have made had they engaged System 2 thinking).

To conclude, here are two final quotes from *Thinking, Fast and Slow* (Kahneman 2011) that should give anyone with any concern or responsibility for human reliability in safety critical operations cause to stop and question the real nature of the thinking and decision making that are so essential to safety critical operations:
- “When people believe a conclusion is true, they are also very likely to believe arguments that appear to support it, even when these arguments are unsound. If System 1 is involved, the conclusion comes first, and the arguments follow” (Kahneman 2011, page 45).
- “...neither the quantity nor the quality of the evidence counts for much in subjective confidence. The confidence that individuals have in their belief depends mostly on the quality of the story they can tell about what they see, even if they see little. We often fail to allow for the possibility that evidence that should be critical to our judgement is missing—what you see is all there is” (Kahneman 2011, page 87).

Acknowledgments
I am grateful to Neville Moray, John Thorogood, and Johan Hendrikse for their comments on a draft version of this paper. I am also grateful to SPE’s technical editors for their comments and suggestions on the submitted manuscript.

References

Aircraft Accidents Investigation Branch (AAIB). 2015. Aircraft Accident Report 1/2015–Airbus A319-131, G-EUOE, 24 May 2013. Formal Report, Air Accidents Investigation Branch, Hampshire, UK (14 July 2015). https://www.gov.uk/aaib-reports/aircraft-accident-report-1-2015-airbus-a319-131-g-euoe-24-may-2013.

API RP 581, Risk-Based Inspection Technology, second edition. 2008. Washington, DC: American Petroleum Institute.

API RP 755, Fatigue Prevention Guidelines for the Refining and Petrochemical Industries, first edition. 2010. Washington, DC: American Petroleum Institute.

Baker, J. A., Bowman, F. L., Erwin, G. et al. 2007. The Report of the BP U.S. Refineries Independent Safety Review Panel. Final Report, BP U.S. Refineries Independent Safety Review Panel (January 2007).

Bureau d’Enquêtes et d’Analyses pour la sécurité de l’aviation civile (BEA). 2012. Final Report on the Accident on 1st June 2009 to the Airbus A330-203 Registered F-GZCP Operated by Air France Flight AF 447 Rio de Janeiro–Paris. Final Report, BEA, Le Bourget, Franch (July 2012), http://www.bea.aero/en/index.php.

Crichton, M. T. and Thorogood, J. L. 2015. Confirmation Bias in a Routine Drilling Operation: A Case Study. *Contemporary Ergonomics and Human Factors 2015, Proceedings of the International Conference on Ergonomics & Human Factors 2015*, Daventry, Northamptonshire, UK, 13–16 April 2015, eds. S. Sharples, S. Shorrock, and P. Waterson, 143–150. Taylor & Francis.

Dolan, P., Hallsworth, M., Halpern, D. et al. 2010. MINDSPACE: Influencing Behaviour Through Public Policy. Report prepared for the Institute for Government (2 March 2010). http://www.instituteforgovernment.org.uk/sites/default/files/publications/MINDSPACE-Practical-guide-final-Web_1.pdf.

Duijm, N. J. 2015. Recommendations on the Use and Design of Risk Matrices. *Safety Science* **76:** 21–31. http://dx.doi.org/10.1016/j.ssci.2015.02.014.

Energy Institute. 2014a. Guidance on Crew Resource Management (CRM) and Non-Technical Skills Training Programmes, first edition. London: Energy Institute.

Energy Institute. 2014b. Managing Fatigue Using a Fatigue Risk Management Plan (FRMP), first edition. London: Energy Institute.

Frederick, S. 2005. Cognitive Reflection and Decision Making. *Journal of Economic Perspectives* **19** (4): 25–42. http://dx.doi.org/10.1257/089533005775196732.

Fruhen, L. S., Flin, R. H., and McLeod, R. W. 2014. Chronic Unease for Safety in Managers: A Conceptualisation. *Journal of Risk Research* **17** (8): 969–979. http://dx.doi.org/10.1080/13669877.2013.822924.

International Association of Oil and Gas Producers (IOGP). 2010. HSE Management – Guidelines for Working Together in a Contract Environment. OGP Report No. 423 (June 2010). http://www.ogp.org.uk/pubs/423.pdf.

International Association of Oil and Gas Producers (IOGP). 2012. Cognitive Issues Associated With Process Safety and Environmental Incidents. OGP Report No. 460 (July 2012). http://www.ogp.org.uk/pubs/460.pdf.

International Association of Oil and Gas Producers (IOGP). 2014a. Crew Resource Management for Well Operations Teams. OGP Report No. 501 (April 2014). http://www.ogp.org.uk/pubs/501.pdf.

International Association of Oil and Gas Producers (IOGP). 2014b. Guidelines for Implementing Well Operations Crew Resource Management Training. OGP Report No. 502 (December 2014). http://www.iogp.org/pubs/502.pdf.

International Association of Oil and Gas Producers/International Petroleum Industry Environmental Conservation Association (IOGP-IPIECA). 2014. Assessing Risks From Operator Fatigue: Guidance Document for the Oil and Gas Industry. OGP Report Number 492 (October 2014). http://www.ipieca.org/news/20141007/ipieca-ogp-launches-new-guidance-assessing-risks-operator-fatigue.

International Petroleum Industry Environmental Conservation Association/International Association of Oil and Gas Producers (IPIECA/IOGP). 2007. Managing Fatigue in the Workplace: A Guide for Oil and Gas Industry Supervisors and Occupational Health Practitioners. OGP Report Number 392, http://www.ogp.org.uk/pubs/392.pdf.

International Petroleum Industry Environmental Conservation Association/International Association of Oil and Gas Producers (IPIECA/IOGP) 2015. Fatigue in Fly-In, Fly-Out Operations: Guidance Document for the Oil and Gas Industry. IOGP Report 536 (June 2015). http://www.ipieca.org/publication/fatigue-fly-fly-out-operations-guidance-document-oil-and-gas-industry.

ISO 31000:2009, Risk Management—Principles and Guidelines. 2009. Geneva, Switzerland: International Organization for Standardization.

Kahneman, D. 2011. *Thinking, Fast and Slow.* New York City: Farrarm Straus and Giroux.

Kedzierski, B. A. 2016. Royal Dutch Shell’s Journey to Strategic Competence Management. *J Pet Technol* **67** (1): 46–51.

Koen, S. 2015. Safety Leadership: Neuroscience and Human Error Reduction. *Safety and Health* (25 October), http://www.safetyandhealthmagazine.com/articles/13159-safety-leadership-neuroscience-and-human-error-reduction.

McLeod, R. W. 2015. Designing for Human Reliability: Human Factors Engineering in the Oil, Gas, and Process Industries. Elsevier.

McLeod, R. W. 2016. Issues in Assuring Human Controls in Layers-of-Defences Strategies. *Chemical Engineering Transactions* **48:** 925–930. http://dx.doi.org/10.3303/CET1648155.

McLeod, R. W. and Beckett, S. 2012. *Chronic Unease: Psychology and Practice.* Presented at the 4th International Conference on Human and Organisational Factors in the Oil, Gas and Chemical Industries, Aberdeen, UK.

National Commission on the BP Deepwater Horizon Oil Spill and Offshore Drilling. 2011. Deep Water: The Gulf Oil Disaster and the Future of Offshore Drilling—Report to the President (January 2011). https://www.gpo.gov/fdsys/pkg/GPO-OILCOMMISSION/pdf/GPO-OILCOMMISSION.pdf.

Schwarz, N., Bless, H., Strack, F. et al. 1991. Ease of Retrieval as Information: Another Look at the Availability Heuristic. *Journal of Personality and Social Psychology* **61** (2): 195–202.

Sykes, M. A., Welsh, M. B., and Begg, S. H. 2011. Don't Drop the Anchor: Recognizing and Mitigating Human Factors When Making Assessment Judgements Under Uncertainty. Presented at the SPE Annual Technical Conference and Exhibition, Denver, 20 October–2 November. SPE-146230-MS. http://dx.doi.org/10.2118/146230-MS.

Swain, A. D. and Guttman, H. E. 1983. Handbook of Human Reliability Analysis With Emphasis on Nuclear Power Plant Applications: Final Report. NUREG/CR-1278, Prepared by Sandia National Laboratories for the Division of Facility Operations, Office of Nuclear Regulatory Research, US Department of Energy, Washington, DC under contract DE-AC04-76DP00789 (August 1983). http://pbadupws.nrc.gov/docs/ML0712/ML071210299.pdf.

Thaler, R. H. 2015. *Misbehaving: The Making of Behavioural Economics*. New York City: W. W. Norton & Company.

Thomas, P., Bratvold, R. B., and Bickel, J. E. 2014. The Risk of Using Risk Matrices. *SPE Econ & Mgmt* **6** (2): 56–66. SPE-166269-PA. http://dx.doi.org/10.2118/166269-PA.

Thorogood, J. L. 2015. The Macondo Inflow Test Decision: Implications for Well Control and Non-Technical Skills Training. Presented at the SPE/IADC Drilling Conference and Exhibition, London, 17–19 March. SPE-173123-MS. http://dx.doi.org/10.2118/173123-MS.

Thorogood, J. L., Crichton, M. T., and Bahamondes, A. 2014. Case Study of Weak Signals and Confirmation Bias in Drilling Operations. *SPE Drill & Compl* **29** (3): 304–310. SPE-168047-PA. http://dx.doi.org/10.2118/168047-PA.

Tversky, A. and Kahneman, D. 1974. Judgement Under Uncertainty: Heuristics and Biases. *Science,* New Series, **185** (4157): 1124–1131. http://www.jstor.org/stable/1738360.

US Chemical Safety and Hazard Investigation Board (CSB). 2016. Investigation Report Volume 3: Drilling Rig Explosion and Fire at the Macondo Well—Human, Organizational, and Safety System Factors of the Macondo Blowout. Report No. 2010-10-I-OS, US Chemical Safety and Hazard Investigation Board, Washington, DC (12 April 2016), http://www.csb.gov/assets/1/7/Macondo_Vol3_Final_staff_report.pdf.

Ronald W. McLeod is an independent consultant and author. He has more than 30 years of applied experience as a human-factors professional in industries ranging from defense and aerospace to oil and gas. Until early 2014, McLeod was Royal Dutch Shell's Global Discipline Lead for Human Factors, responsible for the company's technical standards and competence in human-factors engineering. He was also one of Shell's global subject-matter experts on operator fatigue. McLeod is an honorary professor of engineering psychology at Heriot-Watt University, Edinburgh, and a visiting academic at the Loughborough University Design School. He has authored or contributed to numerous industry guides and best practices on aspects of human factors and fatigue-risk management published by organizations including SPE, IOGP, IPIECA, the Energy Institute, and the UK Process Safety Leadership Group. McLeod's first book, *Designing for Human Reliability: Human Factors Engineering in the Oil, Gas and Process Industries,* was published by Elsevier in 2015. He holds an honours degree in psychology, an MSc degree in ergonomics, and a PhD degree in engineering and applied science. McLeod is a member of SPE; a registered member, and on the council, of the UK Chartered Institute of Ergonomics and Human Factors; and is a member of the US Human Factors and Ergonomics Society. He was a member of the organizing group for SPE's 2012 Human Factors Summit and was a contributing author to the resulting SPE white paper on human factors. McLeod has presented at various SPE technical conferences and is an SPE Distinguished Lecturer for the 2016–17 lecture season.