**Article:**

Rae, Andrew, Provan, David, Aboelssaad, Hossam et al. (1 more author) (2020) A manifesto for Reality-based Safety Science. Safety science. 104654. ISSN: 0925-7535

https://doi.org/10.1016/j.ssci.2020.104654

# About the Manifesto

## Safety science is stagnating

In safety science, we have stopped competing empirically. A key symptom of this problem is the lack of high-quality intervention research. In a healthy field, we should expect to see experimentalists testing theories, feeding back into new theories. For safety, this means researchers and practitioners cooperating to test whether interventions based on the theories work. Instead, systematic reviews consistently show large volumes of publications on any safety theory, but very small volumes of (methodologically poor) intervention evaluations [10]–[12]. We research and teach about variations on many activities - hazard identification, risk assessment, safety cases, safety climate, safety leadership – but we lack an empirical basis on which to prefer any way of managing safety over any other.

It would be pointless to ask, “How did we get into this situation?” All fields of research go through periods of relative stagnation, where activity is dominated by those who are content to research or practice within existing theories. Such tension between present and future consensus is what marks the punctuated equilibrium of scientific advancement in every field [8]. The important question is, “How did we get stuck here?”

In a healthy research field there is empirical growth. Each research program competes to outmatch the others through making novel predictions, and having those predictions confirmed [9]. Safety Science has suffered from a gradual separation between its theoretical, empirical, and practical components. The theorists bicker whilst the empiricists, who should be adjudicating these arguments, are instead trapped within closed theoretical frameworks.

## The cause of the stagnation is a dysfunctional relationship between giants and dwarfs

Isaac Newton once said, “If I have seen further it is by standing on the shoulders of giants.” This was not an original contribution by Sir Isaac, but a refinement of a much older metaphor in which dwarfs could see further than giants by standing on their shoulders. Modern thinking about safety is dominated by “giants” - broad theorists with persuasive ideas – Heinrich, Hollnagel, Hudson, Leveson, Perrow, Reason, Turner, Weick, Zohar and others. Whatever the giants lack in research rigor they make up for with compelling metaphors. They have the power of naming things, and the names, once bestowed, hold and perpetuate the power of the giants.

Around the feet of the giants scurry industrious dwarfs. Each dwarf usually lives within the shadow of a single giant, applying and refining one big idea. Some dwarfs are not even aware of which giant towers over them. Such dwarfs may sincerely believe that they know most of what there is to know about safety, whilst actually only understanding a small slice of extant theory [1].

Is this a fair picture of the safety academic community? No metaphor will ever capture the diversity of the safety researcher population, or the complexity of any individual research career. Still, everyone who regularly reads work within the safety literature is familiar with the gulfs between the work of theorists, the work of empiricists, and the work of practitioners. Each group has its distinctive problems.

The theorists struggle with the foundational questions of the field. What is safety? What does it mean to treat safety as a science? Why do accidents happen? What are the appropriate objects of safety research and practice? They pose broad answers to these questions, but seldom refine their elaborate and eloquent theories into genuinely testable models or practices [2].

The empiricists often appear unconcerned with, or even completely ignorant of, the foundational questions. They align themselves uncritically with the work of a single giant, applying or refining the theories without questioning the foundations. They test whether one variable in a safety climate model mediates two other variables, without considering whether any of the variables have ecological validity, e.g. Gao et al. [3]. They report the results of a single application of an unvalidated method for assessing risk, e.g. Ilbahar et al. [4]. They test whether an intervention changes self-reported behaviour, ignoring the fact that self-reporting is itself a variable behaviour, e.g. Lusk et al. [5].

The practitioners conduct activities that are disconnected from research, guided more by standards and legislation than by either theory or evidence [6], [7]. They accumulate individual experience that speaks directly to the accuracy, relevance and usefulness of safety theory, but seldom collect or record this information in a way that can be credibly and reliably used to update the theory. They thus make very limited contributions to public knowledge.

## Our manifesto is a proposed solution for the stagnation

The title of our paper includes the word “manifesto”. It is a personal statement of policy and commitment from us, with an invitation to others to make the same commitments. Our manifesto is written in first person to indicate that it is intended as self-reflective critique rather than personal attack. If readers see similarities to their own work within the problems we discuss, we hope that they will be provoked and challenged rather than offended.

We introduce the term “Reality-based Safety Science”, intended to evoke similar sentiments to the evidence-based medicine movement, with one important difference. Evidence-based medicine was a reaction to physicians’ over-reliance on intuition, personal experience and theoretical rationale [13]. Often the evidence existed but was not being used by practitioners. The problem we are facing in safety science stems from a lack of evidence production. We are conducting empirical work, but it is a naïve empiricism that is insufficiently informed by a critical understanding of existing theory. We are creating theories, but they are untestable edifices that tower well beyond the supporting evidence. There is systemic pressure placed on researchers to ground their work in untested models, reductionist categories, and proxy measurements, rather than on direct observation and sophisticated analysis of real people doing real work in real organisations.

Hence, our manifesto calls for Reality-based Safety Science - where theory is grounded in rigorous observations of existing practice, and where practice is based on established theory. Reality-based Safety Science is based on the following commitments:

1. We will investigate work as our core object of interest
2. We will describe current work before we prescribe changes
3. We will investigate and theorise before we start measuring
4. We will directly observe the practices that we investigate
5. We will position each piece of research in an appropriate disciplinary context, informed by research practices and recent advances in that discipline;
6. When researching safety methods, we will prioritise real-world case studies over worked examples.
7. We will treat practitioners as respected partners

It is inevitable that any manifesto for the future is a critique of the past. For each "we will" in the preceding list there is an implicit "instead of the way it is usually done". This paper is for publication in a special issue of the journal Safety Science titled The Future of Safety Science. It contains personal opinions and reasoned arguments. Whilst we criticise previous research, we do not criticise those who produced it. They are colleagues we respect, and in some cases, we consider them friends. They are giants, not demons, and we seek to stand on their shoulders rather than cut them down. We do not apologise for the spikes on our boots.

## Why the manifesto is needed

### Our basis for evaluating progress in safety science

It would be arrogant for us, as safety researchers, to pass judgments about the general quality of safety science research as if we were impartial observers. There is a problem, and we are part of it.

Our analysis of this problem is based on Imre Lakatos's method of evaluating scientific progress [14]. For us, as for Lakatos, the appropriate unit of analysis is not an individual paper, or even a single theory, but a "research program" - an evolving sequence of theories. The problem is not "safety research is being executed badly" but "the safety research program is not making sufficient progress".

"Research program" is a scalable term. It can apply to an overall discipline such as "safety science", a narrower group of ideas, such as "safety climate", or a single evolving chain of theories within that group, such as "safety-specific trust".

All research programs have a theoretical "hard core" which cannot be successfully challenged without over-throwing the program, along with a contestable set of auxiliary hypotheses. The auxiliary hypotheses link the hard core with the observed world by providing ways to measure, test and apply the hard core. The auxiliary hypotheses form a protective belt around the hard core – any empirical anomalies are accommodated by adjusting the auxiliary hypotheses rather than by rejecting the hard core.

For example, consider Newtonian physics as an explanation for the movement of planets. Newton's three laws formed the hard core, and knowledge about the planets in our solar system formed the auxiliary hypotheses. Any movement of the planets in violation of the laws could be accommodated by hypothesising new planets.

A research program is “progressive” under two conditions:

1) Each new theory must have greater empirical content than its predecessors – it must describe, explain and predict novel phenomena; and
2) At least some of this novel content must turn out to be true.

The idea of “novelty” is relative to both other theories and to “common sense” expectations. If a theory makes predictions that no one else makes or expects, and those predictions turn out to be true, the program has made strong progress. If only the most obvious predictions of a theory turn out to be true – particularly if rival theories make the same predictions – the theory cannot claim to have made a novel contribution.

There is room for a new theory to make wrong predictions, particularly if these can be explained away by adjusting the auxiliary hypotheses. However, once a program bogs down in constant adjustment of auxiliary hypotheses to explain away wrong predictions, at the expense of novel true content, the program has become degenerate.

At first, Newtonian physics predicted new planets, and these were eventually discovered. One thing that couldn’t be explained was the changing orbit of Mercury. Various hypotheses were put forward, but none of these panned out. Eventually the degenerate Newtonian physics was replaced by General Relativity, with a new hard core that could satisfactorily explain Mercury’s orbit. General Relativity did more than explain things that Newtonian physics could not – it made new predictions that also turned out to be true.

## Our evaluation of safety science

To the extent that safety science makes progress, it does so by adopting and customising progressive research programs from related fields. Once those programs become part of safety science, they usually cease making progress. Key examples of this include the adoption of organisational culture as “safety culture”, behavioural psychology as “behavioural safety”, and theory X and theory Y as “Safety I and Safety II”. There are isolated examples of progressive research programs within safety science, but these programs are usually related to specific technologies or biological processes. In terms of novel and confirmed empirical content, safety science is usually where research programs come to die.

This indictment of the field is not necessarily a judgement on any individual publishable unit of research. Exquisite rigour can be found within a degenerate research program. This can include novel ideas that turn out not to be true, or empirical investigations that challenge the current auxiliary hypotheses. An overall research program, though, must be judged on its forward progress.

We ask our readers to reflect on the following research programs:

“Safety climate” was introduced in 1980 by Dov Zohar as a specific type of organisational climate [15]. The concept built upon existing organisational science research outside of the field of safety and made new empirical predictions specific to the causation and prevention of accidents. Forty years later, how much progress has been made, according to Lakatos's criteria? There are many variations on the theme of safety climate. There are many instruments for measuring safety climate, and many commercial programs for improving safety climate. But where is the sequence of theories making progressively novel and empirically confirmed predictions? It took the safety community almost three decades before a few authors [16], [17] started to seriously question the conceptual and methodological limitations inherent in culture and climate surveys/questionnaires, or the fact that organisational climate scholarship has not been convincing in resolving its own critical problems. These problems include confusion with constructs like job-satisfaction and leadership, and the theoretical inconsistency of aggregating individual/psychological climate perceptions to represent organisational wide climate. Such criticisms were active in organisational climate scholarship long before they were acknowledged by safety climate researchers [18]–[20].

"System-Theoretic Accident Model and Processes" (STAMP) was introduced by Nancy Leveson in 2004 as a new accident model [21]. STAMP applied General Systems Theory to the specific problem of accident causation. The STAMP research program had similarities with other "systems thinking" approaches to safety, but included novel empirical content. STAMP has been widely adopted, and is the subject of many "case study" papers - e.g. Ouyang, Hung, Yu & Fei [22] among many others - but has made no theoretical progress since its introduction. STAMP is almost exactly the same theory as it was in 2004. Models such as Functional Resonance Analysis Method (FRAM) [23] present alternate applications of General Systems Theory to safety, with a similar proliferation of derivative papers, but safety science is yet to find a case where the models make different, testable predictions, so that one model can be empirically preferred over the others [24]. Meanwhile, outside of safety, the pursuit of a universal General Systems Theory has been replaced by domain-specific and application-specific approaches to modelling complexity [25].

"Safety Cases" were first mentioned in the journal Safety Science in 1996, as a reference to an already wide-spread industry practice [26]. For a short time, researchers enriched the existing practice by adapting Toulmin's models of argumentation structure to understand how safety arguments were constructed, undermined and adapted. This theoretical work was very quickly transformed into a simplified notation for representing safety cases, along with a set of prescriptive practices for applying and reviewing the notation [27]. There is little to no work that investigates whether (and if so, how) safety cases lead to safer systems. In 2016, a cross-industry review published in Safety Science and co-authored by some of the original researchers, suggested, as an article highlight, that "Research about effectiveness of safety cases is required" [28]. Twenty years of academic discussion of a pre-existing industrial practice has delivered a wealth of elaboration, and a near-total absence of evaluation. Meanwhile, safety cases (or their equivalent) are used by practitioners and regulators in almost every international safety critical industry (e.g. aviation, maritime, rail, oils and gas, nuclear) as the central instrument of confirmation that a system or technology is safe.

Safety climate, STAMP, and safety cases were all empirically interesting theories. They offered new ways to interpret and explain existing data, and they made novel predictions. All three approaches were quickly and widely adopted by industry, creating a potential wealth of data to test the predictions and progress the theories. Instead, researchers devoted their efforts to elaborating and applying the theories in ways that did not add empirical content through prediction or confirmation. There are more tools, more methods, more guidance, and more case studies of applications, but very little more evidence about what works or doesn't work.

### The pattern of program stagnation

To avoid giving specific offence, the three examples above are representative, and hold no special status. The same charges can be laid against behavioural safety, risk assessment, normal accidents, high reliability organisations, resilience engineering and Safety II. There is no shortage of examples that fit the following pattern:

1. An existing research program from a field that aligns with or overlaps with safety science;
2. a translation of the program into the safety domain, usually in a way that adds novel empirical content through the specific application to safety;
3. widespread industry adoption of the early ideas and models presented by the program, before they have been empirically tested; and
4. a large body of research literature that elaborates and applies the early ideas and models, whilst neither increasing the volume of confirmed empirical content, nor making novel predictions for testing.

Davis [29] makes a compelling case that all interesting social theories attack the prevailing assumptions of their audience. Rigorously constructed theories that fit in with the status quo may be accepted as true, but they do not have real value, because they do not change what we already know. New safety theories are often interesting, but they seldom advance knowledge beyond this initial splash of attention. When they are not advancing whole new theories, most publications in safety aren't even interesting in a Davis sense.

Reality-based Safety Science commits itself to advancing the hard core of safety. We will seek to identify and challenge the assumptions that shape safety practice. We will transform dogmatic assertions into testable predictions. We will test those predictions.

## Commitment 1 - We will investigate work, rather than accidents, as our core object of interest

### Issue giving rise to this commitment

Safety science studies two main categories of things – accidents, and work.

Of these two categories, accidents are the most interesting and least useful objects of investigation. They can almost never be studied except through secondary data that has already been filtered and interpreted by investigators. As rare and complex events, each accident affords too many interpretations to challenge and update existing theories. Accidents are good communication tools [30], but they are dangerous distractions for safety researchers. Accidents do not read, understand and abide by our models of accident causation.

The study of accidents consistently leads safety researchers into the trap of drawing conclusions about how work is, and how work should be, based on single instances of work that, by definition, have unusual outcomes. As Hollnagel [31], Dekker [32], and Amalberti [33] have discussed eloquently and at length, it is not possible even to determine what makes an accident unusual by studying that accident.

## Foundation of this commitment

Setting aside accidents, then, this leaves "work" as the core object of interest for safety science. Work is a catholic concept that includes engineering and design, management, regulation, analysis, social interaction, education, and many other activities. Safety researchers and practitioners often differentiate themselves based on the domain of work they examine. The safety of design work is often treated as a special sub-discipline under labels such as "system safety" or "safety engineering". The safety of medical work is also often treated as a special subdiscipline under the label "patient safety".

For the purpose of this commitment, we do not differentiate between different domains of work. We recognise that there are bespoke challenges to studying design and medical work, but we do not accept that there are fundamental epistemological differences between researching work programming aviation software and researching work building fences. Each type of work is performed by humans, and can be influenced by the psychological state of the worker and the social meaning that the worker gives to the work. Each type of work can be represented, with imperfect fidelity, through standardised models and procedures. Each type of work varies in its performance, and may lead to an accident.

Safety science is interested in the aspects of work that make it safe or unsafe. As a matter of practical scope, rather than a foundational assumption, safety science is concerned with aspects of work that are generalisable across organisations and domains. For example, research to develop more stable airframes, or to write software that more closely matches its specifications, may certainly be relevant for safety. However, the primary audiences for such work are airframe engineers and software developers. Safety science is not directly interested in how to design airframes or write software. The work of both airframe designers and software developers may be studied to draw broader conclusions about how organisations support design workers in creating safe designs. That is the remit of safety science.

One type of work that is unquestionably generalisable is work done by, or at the direction of, safety practitioners. This "safety work" is not strictly necessary for the accomplishment of business goals, and would often not take place if safety were not a concern in its own right [34]. Despite the wealth of literature telling safety practitioners how to do their jobs, there is surprisingly little empirical investigation of safety work [35].

## Specifics of this commitment

Almost all research questions in Reality-based Safety Science should be questions about work.

- How does work happen?
- How do workers make sense of the work that they do?

- How does work vary in the short term? What factors cause it to vary or stay the same?
- What events that occur during work are meaningful for workers? How do workers interpret these events?
- How does work change over longer periods of time? How are practices shared and improved? How does work respond to external influences?
- Who performs work? How does their identity influence the work?
- Where does work take place? How does the nature of work change according to its environment?
- How is work organised? What is the effect of the organisation of work on the conduct of work?
- What counts as core work, and what is discretionary? What is the relationship between non-core work activities (in particular safety work) and the conduct of core work?

Note that these are not intrinsically questions about safety, but any answers to these questions form the building blocks of new theories of safety. Studying work gives researchers access to data that the existing theories of safety already say is important, and, unlike the study of rare accidents, offers new instances of such data to every safety researcher entering the field.

Reality-based Safety Science commits itself to studying generalisable phenomena relating to work. We will pay particular attention to both safety work, and those aspects of operational work that safety work is intended to influence.

Reality-based Safety Science eschews artificial distinctions between different types of work, particularly where such distinctions are used to divide the safety research community.

Reality-based Safety Science recognises the limits of accident research for anything other than very preliminary theory building.

## Commitment 2 – We will describe current work before we prescribe changes

### Issue giving rise to this commitment

It is very tempting to attempt to influence safety practice by writing prescriptions. Young researchers, in particular, fall prey to the idea that advancing safety requires proposing new tools and techniques. This idea can be reinforced by academic program rules that require a "novel contribution", and implicitly or explicitly recognise novelty of methods more readily than novelty of observation, analysis or evaluation. As a result, there are thousands of Master and PhD theses each presenting a new method for hazard analysis. Collectively, this work forms a methods lottery. In the unlikely event that a method is adopted by industry, a researcher can build an entire career out of applying and refining the method. Much more likely, such research fails to answer any question that a practitioner is interested in the answer to.

Current theories suggest that accidents happen when work varies in uncontrolled ways [21], when it fails to adapt to changed circumstances [36], when it drifts into routinely unsafe practice [37], or when it was organised with insufficient safeguards [38]. These theories, to the extent that they say anything about how safety should be practiced, are cautionary tales about the limited ability of technical solutions to solve socio-technical problems.

Any prescriptive safety analysis method that claims to be informed by socio-technical safety theory is a contradiction in terms. If the safety research community is serious about acknowledging the socio-technical and cultural conceptions of organisations as open systems, then it is time for us to give up on providing "solutions" that treat safety as an output from a closed mechanical system.
The type of research that will best support current work practices is research that is deeply informed by those practices. It is unreasonable to expect to advance the industrial use of safety cases by studying how academics use safety cases, or to improve hazard identification by enhancing a published method that no one currently uses.

Foundation of this commitmentAt the heart of Reality-based Safety Science is a deep interest in the current practice of safety work and operational work. By describing these things back to safety practitioners in new ways, we seek to give them improved understanding and capability to do their jobs.

"Describing" encompasses much more than raw data collection. Descriptive research encompasses:

- Making direct and indirect observations of the thing being studied (which in turn includes a wide range of data collection methods);
- analysing and modelling the thing being studied; and
- assessing and evaluating the thing being studied.

Reality-Based Safety Science is based on a virtuous cycle of studying current practice in order to advance theory, and applying theory to advance current practice.

Given the current state of knowledge and practice in safety, Reality-based Safety Science is likely to be dominated by descriptive research for the near future. Once a better understanding of current practice has been achieved, there will be many opportunities for intervention research. At present, however, safety interventions are very seldom based on theory [39]. This is inevitable, given how divorced the theories are from safety practice, and how a-theoretical most safety practice is.

How should a safety practitioner, encountering Safety II for the first time, change their day-to-day activities? How should a supervisor take recent advances in safety climate theory into account when performing inductions? Where does systems thinking fit into a safety management system? These should be questions with easy and well researched answers. They are not.

## Specifics of this commitment

Reality-based Safety Science commits itself to publishing research with clearly expressed research questions, where the answers to those questions are helpful to practitioners. We will describe real work, in real organisations, in new and interesting ways. We will build our theories of safety from the answers to these questions, so that the theories have immediate relevance.

Reality-based Safety Science is cautious about the creation and refinement of new safety work methods as a goal of safety research. We recognise that methods can be a practical way to both communicate and test theory, but we are concerned that methods are seldom used in this way. The magnification of “swiss cheese” from a simple cartoon into the most widely applied “theory” of accident causation [2] is a cautionary tale, not a success story.

## Commitment 3 - We will investigate and theorise before we start measuring

### Issue giving rise to this commitment

Most quantitative research in safety involves the measurement of attributes of phenomena that are interesting because they are believed to be related to safety. Safety researchers usually adopt the jargon of behavioural psychology by referring to these measurements as “constructs”.

Safety scientists measure constructs relating to, amongst others:

- leadership (e.g. leaders’ self-reported styles);
- climate (e.g. worker perceptions of organisational priorities);
- culture (e.g. organisational response to errors);
- behaviour (e.g. compliance rates with wearable equipment rules); and
- individual perception of safety (e.g. self-reported pro-social behaviour).

Unfortunately, we started quantifying these constructs long before the phenomena were qualitatively investigated or theorised (at least within the safety domain). As a result, to the extent that there are theories of safety leadership, safety climate, safety culture, safety behaviour or safety perception, these theories are assemblages of relationships between constructs that lack ecological validity. We do not have adequate descriptions or explanations of the phenomena, so we do not know what real-world referents the constructs are truly associated with.

In many cases, it is not even clear whether different constructs are different attributes of the same phenomena, different theoretical conceptualisations of the same attributes, or discrete but related phenomena. For example, what is the ontological relationship between an individual’s perception of their leaders’ commitment to safety, and their organisation’s safety climate? Is perception a dimension of safety climate, an alternate conceptualisation of safety climate, or something separate from safety climate? This ontological problem needs to be resolved before it even makes sense to investigate the causal relationship between the perception and the climate.

As Ioannidis [45] points out in his famous paper “Why most published research findings are false”, when a large number of comparisons are made between variables, with low prior plausibility for any particular outcomes, there will be a high rate of published false positives. This applies particularly to significance testing using p-values. To have a high chance of being true, a statistical comparison needs to start with a plausible hypothesis, and a sufficient understanding of the context to control for most sources of variation. Otherwise it is just searching for patterns in noise.

In safety science we have developed the habit of tinkering with flying cars. We do not know if the cars are touching the ground, and if not, what miraculous force is keeping them in the air. These are important issues to investigate, but instead we ask how the heater is connected to the radiator. This question matters, we say, because no one has yet answered whether the relationship between cabin temperature and engine temperature is mediated by the capacity of the heat exchanger.

Knowing the relationships between two constructs in a safety model is meaningless unless that relationship has significance for the practice of safety. If neither construct is observable except in a psychometric survey, then neither construct is grounded in the practice or experience of safety. It is important to understand the real-world meaning of our constructs before we worry about the statistical relationships between them.

## Foundation of this commitment

This is not to say that quantitative research in safety is inherently unreliable. Quantitative research is necessary and important - but a rigorous quantitative study starts with a theory. From the theory comes a proposed relationship between variables, and from that relationship comes a hypothesis. Disproving the hypothesis should shake the theory to its core, requiring careful rebuilding to resurrect the theory.

The concept of "trust", particularly as investigated by Stacy Conchie [40]–[42], illustrates the desirable blend of theory building and measurement. Prior to this research program, trust was referred to uncritically as a desirable part of safety culture. Conchie and her colleagues began with qualitative investigations to build a theory of safety-specific trust and distrust. This theory suggested that trust and distrust both had positive and negative connotations for safety and made predictions about how organisations could change trust and distrust in ways that would improve safety. The research program then moved into a quantitative phase that tested how organisations could change worker trust.

Whilst "safety specific trust" and "safety climate" papers often appear superficially similar – using survey results to test the relationships between dependent and independent variables – the key difference is that trust has a strong mechanistic explanation for the relationships being tested. This means that each quantitative result has clear implications for how leaders should seek to manage safety. Often the quantitative results lead on to more qualitative investigation of safety practices. As Conchie et al [43] wrote:

"If management is to achieve safer behaviours from employees, then it is important that we understand precisely what underpins leader-focused strategies, and other strategies, that appear to promote safe behaviour. This may require research examining the role of presumptive trust processes, but it is also likely to require further study of trust formed through the longer-term development of relationship and actual experience. Only then can we identify what it is that needs to be done in an efficient way."

It is tragic that this work is mostly cited to support the unsophisticated claim that "safety leadership" is important for safety, usually in the introduction to a survey-analysis study that eschews the sophisticated and empirical theory of trust that Conchie and her colleagues developed in favour of unvalidated constructs. See Huang et al [44] as a typical example of this. Huang cites Conchie positively, and then uses "supervisory safety communication" as a construct to test hypotheses such as "Supervisory safety communication will moderate the relationship of group-level safety climate with safety performance."

### Specifics of this commitment

Reality-based Safety Science commits itself to describing the real-world phenomena that create, correlate with, emerge from or are otherwise associated with safety or its absence. We will, where possible, identify measurable aspects of the phenomena. When seeking to change these measurable aspects, we will confirm that the phenomena itself is changing, not just the measurement.

Reality-based Safety Science eschews putting "safety" in front of phenomena such as "leadership" and "climate" and assuming that a new phenomenon has thereby been identified.

Reality-based Safety Science recognises the limits of statistical significance testing for any purpose other than testing already plausible relationships between properly theorised constructs.

## Commitment 4 – We will directly observe the practices that we are investigating

### Issue giving rise to this commitment

Precision is not a solution to epistemic uncertainty. A mercury thermometer in Brisbane cannot measure the temperature in Chicago. Replacing the thermometer with a network of precision thermistors will not help. Using expert judgement through an Analytic Hierarchy Process to weight the input of each thermistor to the temperature calculation will still not help.

Data that does not match the research question is bad data.

Here are some common types of inappropriate data use in Safety Science, along with examples:

- self-interested reports used as measures of quality, such as safety personnel reporting on the quality of safety management in their organisation, e.g. Santos et al. [46] and Stolzer et al. [47];
- self-reported behaviour used as measures of behaviour, such as individuals describing their own safety conduct, e.g. Kievik et al. [48];
- guesses at frequencies or risks being represented as actual risk, such as individuals ranking hazards based off their own perception of risk, e.g. Sanni-Anibire et al. [49];

- arbitrary measures of importance used as objective measures, such as individuals ranking the importance of hazards or control measures, e.g. Andrić and Lu [50]; and
- frequency of reported events used as frequencies of events, such as historical reports of accident frequencies used as an outcome variable in studies to determine the relationship between safety climate and injuries, e.g. Young [51].

Inappropriate data stays as inappropriate data no matter how it is processed. Dekker and Nyce [52] describe the problem of "ontological alchemy" in human factors, where subjective judgements are transmuted into apparently objective numbers. Safety Science also experiences this problem, as well as several other forms of alchemy:

- using fuzzy logic, neural networks, or other algorithms that operate by differentially weighting data items in order to combine multiple dubious sources of data, e.g. Liu and Tsai [53];
- following Analytic Hierarchy Process, Delphi, Multi Criteria Decision Analysis or other expert decision-making mechanisms to reach a social consensus on an empirical question, e.g. Janackovic et al. [54];
- using "big data" techniques to identify clusters between variables, without an underlying research question, e.g. Carter et al. [55]; and
- reporting the outputs of unvalidated quantitative risk assessment models as objective measures of risk, e.g. Zhou and Liu [56].

It is no co-incidence that all of the examples in the preceding lists involve quantitative survey research. Most data problems in safety arise from an inability by researchers to directly observe the phenomena that they are interested in measuring. It is irrelevant whether this inability comes from cost, difficulty in gaining access, demand to produce publications, or methodological difficulties – the consequence is a futile attempt to transmute lead into gold.

## Foundation of this commitment

The central activity for Reality-based Safety Science is the examination of work practices. There are many different legitimate ways to find out about work, including interviews, surveys, document analysis and electronic measurements. All of these methods make epistemological sacrifices. The reality of work is different from:

- how the rules say that work is done [57];
- how organisations formally understand and represent their work [58]; and
- how workers describe the work in interviews and documents [59].

Unfortunately for the state of knowledge in safety science, it is often this very gap between work practices and how those activities are understood and represented (sometimes called work-as-done versus work-as-imagined) that is critical to understanding what makes work safe and unsafe [60]. As researchers, if our own data is merely a re-representation of work-as-imagined, we can never really see the problems that we are trying to describe.

Examining work practices means more than simply documenting what is going on at a particular place and time. To understand a practice, we need both chronicles and discourses – the actions and their meanings [61]. A researcher observing work should look at the relationships between practices, how these practices are produced and re-produced, what their underlying assumptions and meanings are, and what this might imply in terms of workplace tensions and power relations. None of this is possible, however, without observing what is going on.

For example, when investigating “safety leadership”, we should be trying to identify and observe the practices that make up leadership. We will, of course, interview safety leaders, but interviews provide information about the interpretation and meaning of the practices, not information about the practices themselves. To understand the practices, we should see what a leader does. We need to see for ourselves what goes on in formal meetings and informal interactions. We should compare those observations to what both the leader and others say about what happened.

### Specifics of this commitment

In the absence of reliable and consistent ways to measure the likelihood of most types of accidents, Reality-based Safety Science must steer away from making broad causal claims. This can be achieved by observing and describing mechanisms, and by measuring the operation of those mechanisms. For example, if a particular leadership behaviour is thought to enhance safety by improving a particular aspect of worker knowledge, it may be possible to find ways to measure both the behaviour and the knowledge. The overall claim about enhancing safety can probably not be directly evaluated.

Reality-based Safety Science commits itself to observing and measuring the phenomena that we seek to describe. We will use proxy measurements only to the extent that the proxy has been demonstrated to be a reliable and consistent indicator of the actual phenomenon of interest. We will use numerical methods to explore and explain data, rather than to obscure problems.

Reality-based Safety Science eschews self-reporting of things that could and should be counted.

Reality-based Safety Science recognises the limits of quantitative surveys for measuring anything other than individual psychological constructs.

## Commitment 5 — We will position each piece of research in an appropriate disciplinary context, informed by research practices and recent advances in that discipline

### Issue giving rise to this commitment

Safety is not a fully independent research discipline. It does not have canon literature. It does not have conventions for the design, execution or publication of research. A researcher who only cites safety literature, or who bases their methods and practices on what they read in the safety literature, is unlikely to be performing work that would be recognised at a high standard outside of safety science.

Some common examples in safety are:

- work situated within safety culture or safety climate that does not take into account advances in organisational theory outside of safety [63];
- work involving the elicitation and processing of expert opinions that does not consider the methodological problems with expert opinion aggregation raised in the economics literature about forecasting [64];
- work involving risk assessment that does not engage with the conceptual difficulties of defining and characterising risk [65]; and
- discussion of behaviour change that uses theories from social psychology that failed when serious attempts were made to replicate the original experimental results [66].

(The citations above are to works which describe rather than exemplify the problems.)

## Foundation of this commitment

Safety science draws on many research disciplines for ideas, perspectives, and methods. As a non-exhaustive list, papers published in the journal Safety Science reference work from behavioural psychology, organisational psychology, engineering, social science, biomedicine, marketing, administrative science, mathematics, law, and human resources. All of these fields study the way work is performed. What makes something part of "safety", rather than any of these other fields, is interest in the way work is associated with the causes and consequences of accidents.

It is legitimate and important to have a field of "safety" research. The study of safety is concerned with things that don't happen - dynamic non-events [62]. These non-events can be overlooked unless specially examined. There is also a lot in common between safety in one field and safety in another. It makes sense to talk about patient safety in the same journal as airline safety, if there are lessons from one that apply to the other. However, the link to accidents is seldom enough to fully differentiate a "safety" topic from its associated disciplines. In this sense, safety science is multidisciplinary rather than interdisciplinary. Safety may be approached using the methods and theories of many disciplines, but it does not stand outside those disciplines. Any advance in the "parent" discipline – in particular advancement, criticism or contradiction of existing theories - is likely to apply also to the safety application of theories from that discipline.

For any given safety topic, there will be many potentially relevant parent disciplines. Where, for example, should a researcher or practitioner look for knowledge about safety in a moving crowd? Crowds are discussed by researchers interested in civil rights, emergency services, urban design, complex systems modelling, psychology, fluid mechanics, computer vision, animal behaviour, and social history. No individual researcher can have a comprehensive understanding of all of the knowledge that may be relevant for a particular safety topic.

## Specifics of this commitment

Reality-based Safety Science commits itself to identifying the parent discipline for every project. We will situate our projects within the current literature of the parent discipline. We will abide by the methodological standards and norms of the parent discipline. We will subject our work to the scrutiny of experts in the parent discipline. We will seek to make findings that advance the parent discipline, rather than merely apply it to safety.

Reality-based Safety Science eschews the overuse of “safety” as a keyword in literature searches, because of the risk this causes of ignoring relevant advances in parent disciplines.

## Commitment 6 – When researching safety methods, we will prioritise real-world case studies over worked examples

### Issue giving rise to this commitment

When making or evaluating claims about safety methods, there are several layers of context for each claim. A method might vary in effectiveness based on many factors, including:

- whether it is applied by researchers, or by practitioners;
- whether it is applied by its creators, or by a third party after training;
- whether it is applied to a specially chosen example, or a typical example;
- the amount of time and resource available to apply the method; and
- the completeness of information available to apply the method.

A worked example considers a method, applied by researchers, to a specially chosen example that can be tailored to suit the method, invariably with scale and complexity smaller than a real industrial application but information and resources more readily available than on a real project. Such a method application can answer some questions about the properties of the method, such as how much effort is required, or whether the method is capable of finding a particular result (for example whether a hazard modelling process can, in principal, describe a known hazard). There are times when answers to such questions are important and interesting. These occasions are much rarer than the use of worked examples in the safety literature.

Safety literature is currently dominated by inappropriate use of worked examples, mislabelled as “case studies”, at the expense of actual real-world case studies.

Of particular and recent concern is the treatment of uncertainty within risk assessments and safety cases. Several types of problem dominate real-world application of risk assessment [68]:

- Incorrect specification of the scope of the assessment;
- failure to consider major sources of risk;
- inappropriate selection of data;
- incorrect assumptions;
- inappropriate use of models; and
- systematic errors in the conduct of the assessment.

The safety academic community has shown very little concern for these problems. They are recognised in industrial literature – see e.g. Crawford [69] – but the academic focus has been on how to measure and represent uncertainty within assessment methods, as illustrated by Denny, Pai and Habli [70]. This distracts from rather than solves any of the real-world problems with uncertainty.

Arguably, there is a spiral where researchers become increasingly concerned with the technical intricacies of modelling reality as a way to hide from the foundational problem that safety models do not represent reality. Springer, Haas and Porowski [71] describe this as a general problem in applied policy research. Researchers, in order to make problems tractable, ignore details that are necessary for solutions that work in the real world. They write: “Policy makers and implementers cannot escape from the messy and wicked elements of problem reality; they live within them. Policy researchers can and often do retreat into the relative safety of tame problems and technical solutions. This can allow a claim of authority through conformity to method” [71]. We suggest that the broader problem is a narrow conception of “rigour” which sometimes manifests as conformity to method, and sometimes as a pre-occupation with logical correctness at the expense of real-world validity.

### Foundation of this commitment

The term “case study” has been abused so often and thoroughly in safety research that we are reluctant to provide this commitment without some basic definitions.

By case study, we mean “an empirical inquiry that investigates a contemporary phenomenon within its real-life context; when the boundaries between phenomenon and context are not clearly evident; and in which multiple sources of evidence are used” [67].

In contrast, a worked example is an explanatory device that applies a method, step by step, to a well understood problem, in order to illustrate how the method works.
When used to examine safety practices, the key differences between a case study and a worked example are:

1. A case study occurs in a real-life context. The phenomena under investigation would occur whether or not a researcher was present to observe it.
2. In a case study, researchers try to limit their own influence, so that any outcomes can be ascribed to the activities of practitioners, not the activities of researchers.
3. A case study involves research questions. A worked example is used to communicate research that has already been performed.

Case studies are vitally important for advancing the state of knowledge about safety methods. Without knowing how methods are used in practice, and the challenges and problems with the methods in practice, researchers are poorly equipped to suggest improvements.

### Specifics of this commitment

Reality-based Safety Science commits itself to suggesting improvements to methods based on a close examination of how those methods are currently used. We will seek to improve the methods in ways that will make them easier and more effective to apply. We will make specific and testable claims about the improvements. We will test those claims.

Reality-based Safety Science eschews the use of the term “case study” for any safety method application performed by researchers.

Reality-based Safety Science recognises the limits of worked examples for answering most meaningful research questions.

## Commitment 7 – We will treat practitioners as respected partners

### Issue giving rise to this commitment

Safety researchers have a bad habit of treating practitioners like misbehaving children. No sub-field of safety is exempt from this. Software safety researchers blame practitioners for failing to adopt methods that the researchers think are obviously better than current practice [72]. Accident theorists don't understand why investigators persist in applying outdated models [73]. Safety II researchers are ostentatiously sympathetic to the local expertise and situation-specific constraints of almost every worker, except those who have "safety" in their job title [74].

If the published literature were taken as an accurate representation of the relationship between academia and industry for safety, the lifecycle of a typical innovation in safety would be:

1. A researcher develops and presents a new method
2. Researchers (often the same as or affiliated with the first researcher) apply the method in hindsight to show how it could have prevented a famous accident
3. Industry practitioners apply the method with no formal measurement, evaluation or comparison, and declare the application to be successful
4. Further researchers elaborate the method or show how to apply it in new situations
5. Industry practitioners apply the method in a wider range of situations, with no formal measurement, evaluation or comparison, and declare the application to be successful
6. Researchers publish reviews of the application of the method, and declare it to be successful because of its wide adoption and the large number of papers about it

We could have given illustrative citations in the above list, but we have not. There are simply too many papers in safety that could be accurately summarised as "Yet Another ACRONYM Paper" - readers may insert their own acronym - perhaps selecting from STAMP, GSN, FTA, HAZOP, WBA, ICAM - or join us in condemning the general phenomenon of YAAPing at safety conferences.

### Foundation of this commitment

Safety researchers and safety practitioners should be working together, but this should be in genuine knowledge-producing partnerships, rather than leader-follower relationships. Such partnerships should recognise the expertise that each party can provide. Researchers base their expertise on systematic public knowledge. Their contribution to research comes through their skills as researchers – they are expert in study design, data collection, and analysing and interpreting data. Safety practitioners are expert in managing safety within their organisational context. Safety practitioners have access to detailed, local data that is often hidden to outsiders.

Researchers are well-positioned to advise practitioners on how to test and measure interventions. They may be able to describe broad patterns and theories to help practitioners to select which interventions are most likely to be effective. However, a medical researcher would never tell a doctor what to prescribe to a particular patient. Medical researchers know that only doctors have access to local knowledge about the circumstances, history and preferences of the patient. Similarly, safety researchers should be providing practitioners with options rather than constraints [74].

Following the standard practice in medical trials, safety researchers could develop the habit of making their research programs known to industry, and inviting eligible organisations to volunteer as participants. They could also establish mechanisms for practitioners to regularly report back on what they are seeing in the population that requires further enquiry.

Such a partnership should work both ways. In the absence of specific local knowledge to the contrary, practitioners should be guided by the best available evidence of what generally works. Where there is no such evidence, practitioners should proceed cautiously, and should cooperate with researchers to advance the state of evidence, rather than operating beyond it.

This applies particularly to practitioners working on behalf of regulatory or standards bodies. Formal prescriptions for practice that extend beyond the available evidence unnecessarily constrain both researchers and practitioners. “I would like to improve safety, but then I would be non-compliant” is both a common refrain and a damning indictment on the state of safety regulations. There is some evidence that safety researchers routinely become implicated in this problem by overstepping role boundaries to act as regulatory practitioners [74], [75].

### Specifics of this commitment

Reality-based Safety Science commits itself to treating safety practitioners as respected participants or partners in safety research. We will defer to practitioners on local, practical knowledge, and seek to communicate knowledge that generalises beyond local circumstances and practical applications.

Reality-based Safety Science eschews making recommendations for standards or regulation, except as outputs from research that directly studies work performed by regulators or constrained by regulation.

## Conclusion

Any reader who agrees with our manifesto so far probably has their own pet explanation for how safety science came to be in such as sorry state.

There is a tradition in safety of stakeholders describing each other as constraints preventing high-quality empirical research. The day-to-day work of safety practitioners is heavily constrained by organisational objectives and directives [7]. The safety management systems of organisations are responsive to legislative and regulatory requirements rather than evidence of what works. Regulators are waiting on researchers to tell them how to incorporate new theory into regulatory practice, and in the absence of this guidance have no choice but to base their rules on political priorities. Researchers can only conduct real-world research that organisations are willing to fund and engage with [76]. Each party blames the others for the lack of empirical research and evidence-based practice.

This excuse-making must stop, because even the status-quo provides many opportunities for Reality-based Safety Science research. Organisations are “experimenting” - in the informal sense of innovating - constantly. Even when organisations are merely undertaking common safety activities, these tasks are poorly documented in the academic literature. The current evidence base is so low that novel research data can be generated just by keeping good research records of current safety practice.

No organisation - either business or regulator - needs permission from anyone else to collect data about its own activities. No organisation should be wary of seeking researcher assistance, because there is a direct financial return on this investment. Organisations are currently spending money on safety. It is unknown whether and where this expenditure is effective or needed. Better data collection will lead to either cost saving, or more effective deployment of existing money.

The wariness to engage in novel safety research stems not from inherent structural constraints, but from previous insufficiently grounded safety research. Business and regulators are wary, with good reason, that “safety research” will consist of new tools and practices that are expensive to implement, rather than actionable information about the status quo [6].

One early reviewer of the manuscript suggested that safety researchers could learn from Rasmussen’s safety envelope model [77]. The gradient of economic performance pushes researchers to create publications and attract citations. The gradient of least effort encourages desktop research based on accident reports, surveys, or blue-sky theorising. The only counter-gradient consists of researchers holding themselves and each other to account.

We conclude this paper with a personal message to you, the reader. If you are a safety researcher, we ask you to take a position on the commitments in this manifesto. If you disagree with them, please feel free to debate them, in public. If you agree with the commitments, please say so, and start holding yourself, and us, to them.

If you are a reviewer, please feel free to refer directly to these commitments when rejecting papers. Peer review should be respectful, constructive, and helpful. Allowing a fellow researcher to continue along a path of degenerate research is neither respectful nor helpful. If we hold to these commitments, the volume of published Safety Science research will decrease drastically. The journal publishers will not be happy, but the world will be a safer place. Which do we care about?

If you are a regulator or funder of research, we ask you to prioritise work that aligns with the manifesto. Our intent with this request is not ideological purity. It is important that a wide range of methods, approaches, and conceptualisations are permitted within safety science. However, we believe strongly that work that does not follow this manifesto deserves a significantly higher level of scrutiny to justify its value.

If you are an educator, we invite you to teach current and future practitioners how to identify and apply Reality-based Safety Science. We ask you to teach them how to recognise and engage critically with work that breaks the commitments in this manifesto. We consider it important for practitioners to understand when they are operating within the current evidence base – and should therefore be applying current theory – and when they have moved beyond the current evidence base – and should therefore be following good research practice.

If you are a safety practitioner, we invite you so engage with safety science research. Contribute to our understanding of the real-world problems and practice of safety and work. Support the development or meaningful research questions and research programs. Collect data. Understand the genuine developments in safety science and incorporate these findings in your professional practice.

If you are supervising young researchers, this manifesto is a demand, not an invitation. Do not allow their time and talents to be wasted. If the commitments aren't quite right, we are willing to be corrected on the detail, but not on the overarching message. Safety Science is a degenerate research program. It must grow up in order to have a future.

## Bibliography


[1] A. J. Rae, "Open questions and closed minds: mapping the gaps and divisions in the safety body of knowledge," presented at the Australian Safety Critical Systems Conference, Brisbane, 2015, p. 6.

[2] J. Reason, E. Hollnagel, and J. Paries, "Revisiting the Swiss Cheese model of accidents," *J. Clin. Eng.*, vol. 27, pp. 110–115, 2006.

[3] Y. Gao, Y. Fan, J. Wang, X. Li, and J. Pei, "The mediating role of safety management practices in process safety culture in the Chinese oil industry," *J. Loss Prev. Process Ind.*, vol. 57, pp. 223–230, Jan. 2019.

[4] E. Ilbahar, A. Karaşan, S. Cebi, and C. Kahraman, "A novel approach to risk assessment for occupational health and safety using Pythagorean fuzzy AHP & fuzzy inference system," *Saf. Sci.*, vol. 103, pp. 124–136, Mar. 2018.

[5] S. L. Lusk, O. S. Hong, D. L. Ronis, B. L. Eakin, M. J. Kerr, and M. R. Early, "Effectiveness of an Intervention to Increase Construction Workers' Use of Hearing Protection," *Hum. Factors J. Hum. Factors Ergon. Soc.*, vol. 41, no. 3, pp. 487–494, Sep. 1999.

[6] D. V. Eerd, S. Cardoso, E. Irvin, R. Saunders, T. King, and S. Macdonald, "Occupational safety and health knowledge users' perspectives about research use," *Policy Pract. Health Saf.*, vol. 16, no. 1, pp. 4–19, Jan. 2018.

[7] D. J. Provan, A. J. Rae, and S. Dekker, "An ethnography of the safety professional's dilemma: Safety work or the safety of work?," *Saf. Sci.*, vol. 117, pp. 276–289, Aug. 2019.

[8] T. S. Kuhn, *The Structure of Scientific Revolutions*, New ed of 3 Revised ed edition. Chicago, IL: University of Chicago Press, 1996.

[9] I. Lakatos, "The role of crucial experiments in science," *Stud. Hist. Philos. Sci.*, vol. 4, no. 4, pp. 344–354, 1974.

[10] M. M. Lehtola *et al.*, “The Effectiveness of Interventions for Preventing Injuries in the Construction Industry: A Systematic Review,” *Am. J. Prev. Med.*, vol. 35, no. 1, pp. 77–85, Jul. 2008.
[11] L. S. Robson *et al.*, “The effectiveness of occupational health and safety management system interventions: A systematic review,” *Saf. Sci.*, vol. 45, no. 3, pp. 329–353, Mar. 2007.
[12] J. Vilela, J. Castro, L. E. G. Martins, and T. Gorschek, “Integration between requirements engineering and safety analysis: A systematic literature review,” *J. Syst. Softw.*, vol. 125, pp. 68–92, Mar. 2017.
[13] L. L. Leape, D. M. Berwick, and D. W. Bates, “What Practices Will Most Improve Safety?: Evidence-Based Medicine Meets Patient Safety,” *JAMA*, vol. 288, no. 4, pp. 501–507, Jul. 2002.
[14] I. Lakatos, *The methodology of scientific research programmes*. Cambridge University Press, 1978.
[15] D. Zohar, “Safety climate in industrial organizations: theoretical and applied implications.,” *J. Appl. Psychol.*, vol. 65, no. 1, p. 96, 1980.
[16] F. W. Guldenmund, “The use of questionnaires in safety culture research – an evaluation,” *Saf. Sci.*, vol. 45, no. 6, pp. 723–743, Jul. 2007.
[17] J. I. Håvold, “Measuring Occupational Safety: From Safety Culture to Safety Orientation?,” *Policy Pract. Health Saf.*, vol. 3, no. 1, pp. 85–105, Jan. 2005.
[18] W. H. Glick, “Conceptualizing and Measuring Organizational and Psychological Climate: Pitfalls in Multilevel Research,” *Acad. Manage. Rev.*, vol. 10, no. 3, pp. 601–616, Jul. 1985.
[19] R. M. Guion, “A note on organizational climate,” *Organ. Behav. Hum. Perform.*, vol. 9, no. 1, pp. 120–125, Feb. 1973.
[20] R. E. Johannesson, “Some problems in the measurement of organizational climate,” *Organ. Behav. Hum. Perform.*, vol. 10, no. 1, pp. 118–144, Aug. 1973.
[21] N. Leveson, “A New Accident Model for Engineering Safer Systems,” *Saf. Sci.*, vol. 42, no. 4, pp. 237–270, 2004.
[22] M. Ouyang, L. Hong, M.-H. Yu, and Q. Fei, “STAMP-based analysis on the railway accident and accident spreading: Taking the China–Jiaoji railway accident for example,” *Saf. Sci.*, vol. 48, no. 5, pp. 544–555, 2010.
[23] E. Hollnagel, *FRAM: the functional resonance analysis method: modelling complex socio-technical systems*. CRC Press, 2017.
[24] E. Grant, P. M. Salmon, N. J. Stevens, N. Goode, and G. J. Read, “Back to the future: What do accident causation models tell us about accident prediction?,” *Saf. Sci.*, vol. 104, pp. 99–109, Apr. 2018.
[25] A. Vespignani, “Modelling dynamical processes in complex socio-technical systems,” *Nat. Phys.*, vol. 8, no. 1, pp. 32–39, Jan. 2012.
[26] N. A. Stanton and P. Marsden, “From fly-by-wire to drive-by-wire: Safety implications of automation in vehicles,” *Saf. Sci.*, vol. 24, no. 1, pp. 35–49, Oct. 1996.
[27] S. P. Wilson, T. P. Kelly, and J. A. McDermid, “Safety Case Development: Current Practice, Future Prospects,” in *Safety and Reliability of Software Based Systems*, 1997, pp. 135–156.
[28] M. A. Sujan, I. Habli, T. P. Kelly, S. Pozzi, and C. W. Johnson, “Should healthcare providers do safety cases? Lessons from a cross-industry review of safety case practices,” *Saf. Sci.*, vol. 84, pp. 181–189, Apr. 2016.

[29] M. S. Davis, “That’s Interesting!: Towards a Phenomenology of Sociology and a Sociology of Phenomenology,” *Philos. Soc. Sci.*, vol. 1, no. 2, pp. 309–344, Jun. 1971.
[30] A. J. Rae, “Tales of disaster: the role of accident storytelling in safety teaching,” *Cogn. Technol. Work*, May 2015.
[31] E. Hollnagel, *Safety-I and Safety-II*, New edition edition. Farnham, Surrey, UK England ; Burlington, VT, USA: Ashgate, 2014.
[32] S. W. A. Dekker, *Ten Questions About Human Error: A New View of Human Factors and System Safety*, 1st ed. CRC Press, 2004.
[33] R. Amalberti, *Navigating Safety: Necessary Compromises and Trade-Offs - Theory and Practice*, 2013 edition. New York: Springer, 2013.
[34] A. J. Rae and D. J. Provan, “Safety work versus the safety of work,” *Saf. Sci.*, Jul. 2018.
[35] D. J. Provan, S. W. A. Dekker, and A. J. Rae, “Bureaucracy, influence and beliefs: A literature review of the factors shaping the role of a safety professional,” *Saf. Sci.*, vol. 98, no. Supplement C, pp. 98–112, Oct. 2017.
[36] D. D. Woods and M. Branlat, “Basic patterns in how adaptive systems fail,” *Resil. Eng. Pract.*, pp. 127–144, 2011.
[37] S. W. A. Dekker, *Drift into Failure*. Farnham, UK: Ashgate, 2011.
[38] J. Reason, “Human error: models and management,” *BMJ*, vol. 320, pp. 768–770, Mar. 2000.
[39] A. R. Hale, F. W. Guldenmund, P. L. C. H. van Loenhout, and J. I. H. Oh, “Evaluating safety management and culture interventions to improve safety: Effective intervention strategies,” *Saf. Sci.*, vol. 48, no. 8, pp. 1026–1035, Oct. 2010.
[40] S. M. Conchie and C. Burns, “Trust and Risk Communication in High-Risk Organizations: A Test of Principles from Social Risk Research,” *Risk Anal.*, vol. 28, no. 1, pp. 141–149, 2008.
[41] S. M. Conchie and I. J. Donald, “The functions and development of safety-specific trust and distrust,” *Saf. Sci.*, vol. 46, no. 1, pp. 92–103, 2008.
[42] S. M. Conchie, I. J. Donald, and P. J. Taylor, “Trust: Missing Piece(s) in the Safety Puzzle,” *Risk Anal.*, vol. 26, no. 5, pp. 1097–1104, 2006.
[43] S. M. Conchie, P. J. Taylor, and I. J. Donald, “Promoting safety voice with safety-specific transformational leadership: the mediating role of two dimensions of trust,” *J. Occup. Health Psychol.*, vol. 17, no. 1, pp. 105–115, Jan. 2012.
[44] Y. Huang, R. R. Sinclair, J. Lee, A. C. McFadden, J. H. Cheung, and L. A. Murphy, “Does talking the talk matter? Effects of supervisor safety communication and safety climate on long-haul truckers’ safety performance,” *Accid. Anal. Prev.*, vol. 117, pp. 357–367, Aug. 2018.
[45] J. P. A. Ioannidis, “Why Most Published Research Findings Are False,” *PLoS Med.*, vol. 2, no. 8, 2005.
[46] G. Santos, S. Barros, F. Mendes, and N. Lopes, “The main benefits associated with health and safety management systems certification in Portuguese small and medium enterprises post quality management system certification,” *Saf. Sci.*, vol. 51, no. 1, pp. 29–36, Jan. 2013.
[47] A. J. Stolzer, M. A. Friend, D. Truong, W. A. Tuccio, and M. Aguiar, “Measuring and evaluating safety management system effectiveness using Data Envelopment Analysis,” *Saf. Sci.*, vol. 104, pp. 55–69, Apr. 2018.

[48] M. Kievik, E. F. J. Misana-Ter Huurne, J. M. Gutteling, and E. Giebels, "Making it stick: Exploring the effects of information and behavioral training on self-protectiveness of citizens in a real-life safety setting," *Saf. Sci.*, vol. 101, pp. 1–10, Jan. 2018.
[49] M. O. Sanni-Anibire, A. S. Mahmoud, M. A. Hassanain, and B. A. Salami, "A risk assessment approach for enhancing construction safety performance," *Saf. Sci.*, vol. 121, pp. 15–29, Sep. 2019.
[50] J. M. Andrić and D.-G. Lu, "Risk assessment of bridges under multiple hazards in operation period," *Saf. Sci.*, vol. 83, pp. 80–92, Mar. 2016.
[51] S. Young, "From zero to hero. A case study of industrial injury reduction: New Zealand Aluminium Smelters Limited," *Saf. Sci.*, vol. 64, pp. 99–108, Apr. 2014.
[52] S. W. A. Dekker and J. M. Nyce, "From figments to figures: ontological alchemy in human factors research," *Cogn. Technol. Work*, vol. 17, no. 2, pp. 185–187, Jan. 2015.
[53] H.-T. Liu and Y. Tsai, "A fuzzy risk assessment approach for occupational hazards in the construction industry," *Saf. Sci.*, vol. 50, no. 4, pp. 1067–1078, Apr. 2012.
[54] G. L. Janackovic, S. M. Savic, and M. S. Stankovic, "Selection and ranking of occupational safety indicators based on fuzzy AHP: A case study in road construction companies," *South Afr. J. Ind. Eng.*, vol. 24, no. 3, pp. 175–189, Nov. 2013.
[55] E. A. Carter, B. J. Westerman, A. E. Lincoln, and K. L. Hunting, "Common game injury scenarios in men's and women's lacrosse," *Int. J. Inj. Contr. Saf. Promot.*, vol. 17, no. 2, pp. 111–118, Jun. 2010.
[56] Y. Zhou and M. Liu, "Risk Assessment of Major Hazards and its Application in Urban Planning: A Case Study," *Risk Anal.*, vol. 32, no. 3, pp. 566–577, 2012.
[57] J. Pfeffer and R. I. Sutton, *The Knowing-Doing Gap: How Smart Companies Turn Knowledge into Action*, 1 edition. Boston, Mass: Harvard Business School Press, 2000.
[58] P. Selznick, "Institutionalism 'Old' and 'New,'" *Adm. Sci. Q.*, vol. 41, no. 2, pp. 270–277, 1996.
[59] D. Silverman, "Qualitative research: meanings or practices?," *Inf. Syst. J.*, vol. 8, no. 1, pp. 3–20, 1998.
[60] E. Hollnagel, "Why is Work-as-Imagined Different from Work-as- Done?," *Resilient Health Care, Volume 2*, 02-Mar-2017. [Online]. Available: https://www.taylorfrancis.com/. [Accessed: 22-Sep-2019].
[61] D. Nicolini and P. Monteiro, "The Practice Approach: For a Praxeology of Organisational and Management Studies," in *The SAGE Handbook of Process Organization Studies*, 55 City Road: SAGE Publications Ltd, 2016, pp. 110–126.
[62] K. E. Weick, "Organizing for Transient Reliability: The Production of Dynamic Non-Events," *J. Contingencies Crisis Manag.*, vol. 19, no. 1, pp. 21–27, Mar. 2011.
[63] F. W. Guldenmund, "The nature of safety culture: a review of theory and research," *Saf. Sci.*, vol. 34, no. 1–3, pp. 215–257, Feb. 2000.
[64] A. J. Rae and R. Alexander, "Forecasts or fortune-telling: When are expert judgements of safety risk valid?," *Saf. Sci.*, vol. 99, pp. 156–165, Nov. 2017.
[65] F. Goerlandt, N. khakzad, and G. Reniers, "Validity and validation of safety-related quantitative risk analysis: A review," *Saf. Sci.*, 2016.
[66] B. D. Earp and D. Trafimow, "Replication, falsification, and the crisis of confidence in social psychology," *Front. Psychol.*, vol. 6, 2015.
[67] R. K. Yin, *Case Study Research: Design and Methods: (Applied Social Research Methods, Volume 5): 005*, Third Edition. Sage Publications, Inc, 2003.

[68] A. J. Rae, R. Alexander, and J. McDermid, "Fixing the cracks in the crystal ball: A maturity model for quantitative risk assessment," *Reliab. Eng. Syst. Saf.*, vol. 125, pp. 67–81, May 2014.
[69] J. Crawford, "What's Wrong with the Numbers? A Questioning Look at Probabilistic Risk Assessment," *J. Syst. Saf.*, vol. 37, no. 3, Mar. 2001.
[70] E. Denney, G. Pai, and I. Habli, "Towards Measurement of Confidence in Safety Cases," in *2011 International Symposium on Empirical Software Engineering and Measurement*, 2011, pp. 380–383.
[71] J. F. Springer, P. J. Haas, and A. Porowski, *Applied Policy Research: Concepts and Cases*. Routledge, 2017.
[72] J. Davis *et al.*, "Study on the Barriers to the Industrial Adoption of Formal Methods," presented at the Lecture Notes in Computer Science, 2013, vol. 8187, pp. 63–77.
[73] P. Katsakiori, G. Sakellaropoulos, and E. Manatakis, "Towards an evaluation of accident investigation methods in terms of their alignment with accident causation models," *Saf. Sci.*, vol. 47, no. 7, pp. 1007–1015, Aug. 2009.
[74] P. G. Almklov, R. Rosness, and K. Størkersen, "When safety science meets the practitioners: Does safety science contribute to marginalization of practical knowledge?," *Saf. Sci.*, vol. 67, pp. 25–36, Aug. 2014.
[75] J. A. McDermid and A. J. Rae, "Goal Based Safety Standards: Promises and Pitfalls," presented at the Safety-critical Systems Symposium, 2012.
[76] F. Lamm, "The challenges of researching OHS of vulnerable workers in small businesses," *Small Enterp. Res.*, vol. 21, no. 2, pp. 161–179, Jan. 2014.
[77] J. Rasmussen, "Risk Management in a Dynamic Society: a Modelling Problem," *Saf. Sci.*, vol. 27, no. 2–3, pp. 183–213, 1997.