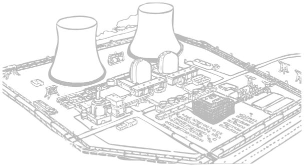

Risk metrics

Eric Marsden

<eric.marsden@risk-engineering.org>

"You can't manage what you don't measure"

Terminology

- A *measure* is an operation for assigning a number to something
- A *metric* is our interpretation of the assigned number
- There may be several measures (measurement methods) for one metric
- Example risk metrics:
  - deaths per passenger kilometre (transportation)
  - probability of failure on demand (systems reliability)
  - value at risk (finance)
- In these slides we focus on **metrics for safety**, rather than for financial risks

"When you cannot measure, your knowledge is meager and unsatisfactory." — Lord Kelvin

RISK ENGINEERING

Context

- ▷ Measurement is a key step in any management process and forms the basis of continual improvement
- ▷ **Safety performance is difficult to measure**
  - success results in the *absence* of an outcome (injuries, losses, health impacts)
  - "Safety is elusive because it is a dynamic nonevent — what produces the stable outcome is constant change rather than continuous repetition" [K. Weick]
  - low incident rates, even over several years, do not guarantee that major accident hazards are controlled
- ▷ There is no single and easy to measure indicator for safety

RISK ENGINEERING

Note: management's obsession with metrics, and the resulting **biases** they introduce into people's behaviour, can have a negative safety impact

- many facets of safety performance can be managed using good professional judgment, without quantitative measures
- Goodhart's law: "When a measure becomes a target, it ceases to be a good measure"

Metrics and modern management

Source: `dilbert.com`

Illustration: difficulties in measuring safety

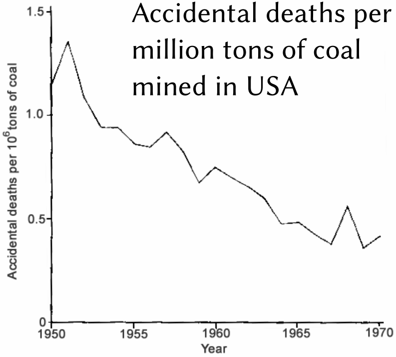

Q: Is coal mining getting safer?

Source: Paul Slovic

Expressing risk to people

- **Individual risk**: risk to any particular individual, either a worker or a member of the public
- **Location-based risk**: risk that a person who is continually present and unprotected at a given location will die as a result of an accident within the site
- **Societal risk**: risk to society as a whole
  - example: *the chance of a large accident causing a defined number of deaths or injuries*
  - product of the total amount of damage caused by a major accident and the probability of this happening during some specified period of time

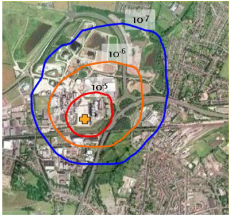

Metrics for individual risk

Individual risk [NORSOK Z-013N]

Probability that a specific individual (for example the most exposed individual in the population) should suffer a fatal accident during the period over which the averaging is carried out (usually a 12-month period).

- Metric: individual risk per annum (IRPA): probability that an individual is killed during one year of exposure
- Measure: $\frac{\text{observed fatality count}}{\text{number of people exposed}}$

RISK ENGINEERING

Metrics for individual risk

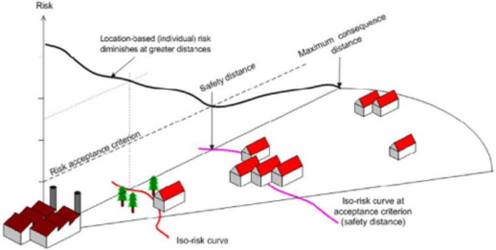

Location-specific individual risk

Annual probability that an unprotected, permanently present individual dies due to an accident at a hazardous site

- is a property of the location, not of the individual
- mostly used in land-use planning
- often represented with iso-risk contours (see figure)

Suggested reading: *Acceptance criteria in Denmark and the EU*, Dutch Environmental Protection Agency

Metrics for societal risk

- **Fatal accident rate** (FAR): expected number of fatalities per unit of exposure
  - can be expressed per million hours worked, per plane takeoff, per km transported, per hour transported
  - typical formulation: number of company/contractor fatalities per 10^8 hours worked
- **Potential loss of life** (PLL): statistically expected number of fatalities within a specified population during a specified period of time
  - note: when all members of a population are exposed to the same level of risk, $PLL = n \times IRPA$

RISK ENGINEERING

Metrics for societal risk

A **Farmer diagram** or **F-N curve** shows frequency and number of deaths for different accident scenarios

Note:

- ▷ drawn with a logarithmic scale on both axes
- ▷ a lower curve is better

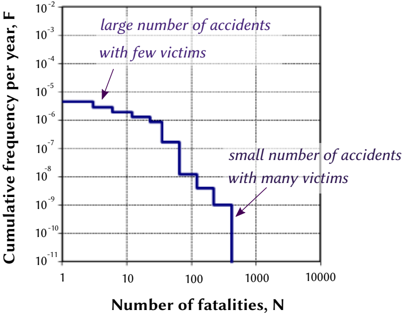

Example F-N diagram

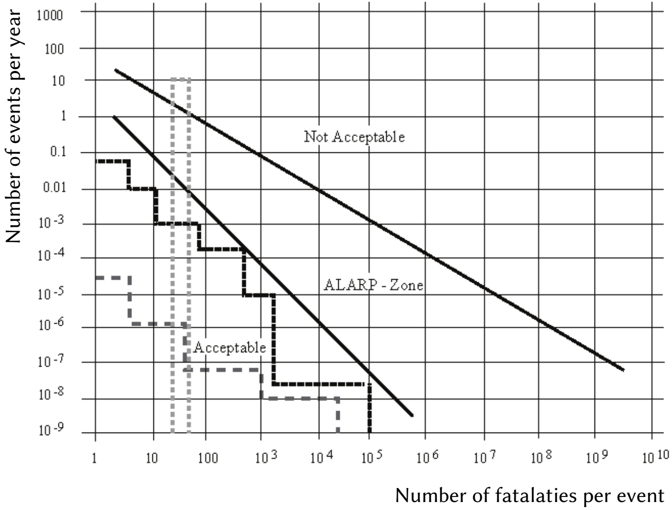

F-N diagram indicating acceptable risks, ALARP zone and non-acceptable risks

F-N diagram for transport accidents

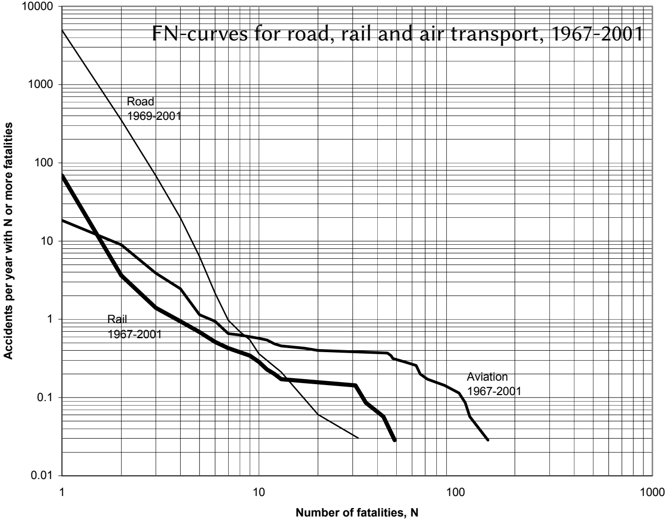

Source: *Transport fatal accidents and FN-curves*, HSE RR073, hse.gov.uk/research/rrpdf/rr073.pdf

F-N diagram used in a safety case

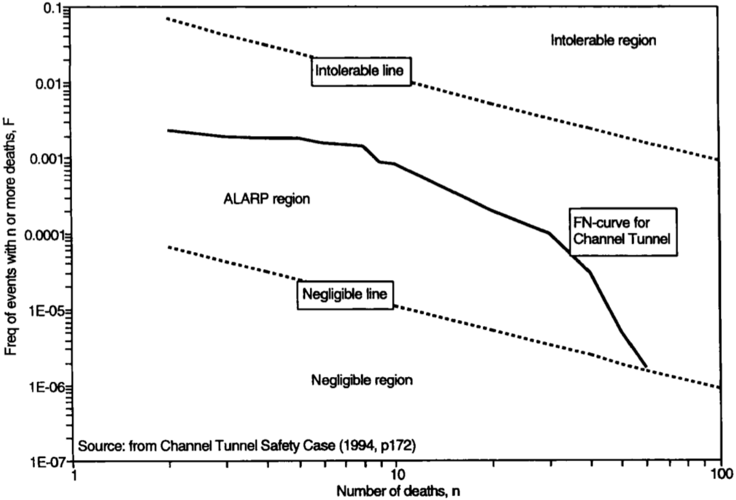

Source: Channel Tunnel Safety Case (1994)

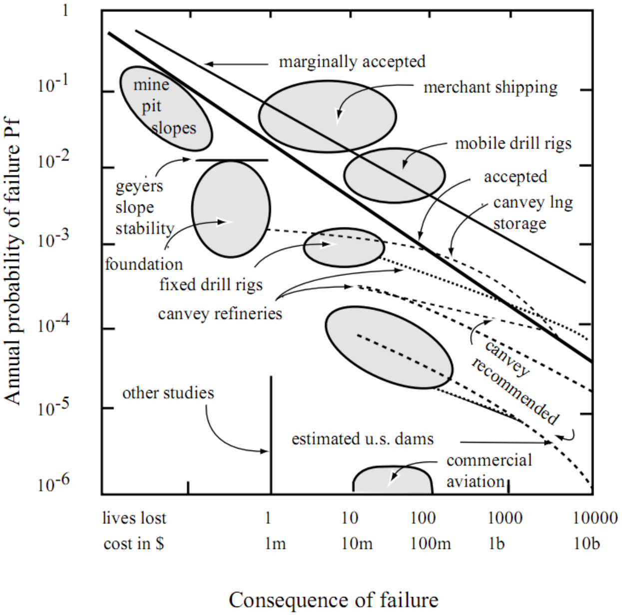

F-N diagram for different socially accepted activities

Source: *Risk and Safety in Engineering*, course notes by M. Faber, ETHZ

Typical occupational safety metrics

- LTRIR: Lost Time Reportable Incident Rate
  - number of hours off work per 200 000 employee working hours (including work-related illness)
- LTIF: Lost Time Injury Frequency
  - number of lost time injuries (fatalities + lost work day cases) per 1 000 000 work hours
- Also used in shareholder reporting on "industrial risk" (as the sole indicator…)
  - this is unfortunate since process safety metrics are equally (or more!) important for indicating the level of safety
  - occupational safety metrics are not correlated with process safety metrics (though there is a widely held view that they are)

Company safety indicators: example

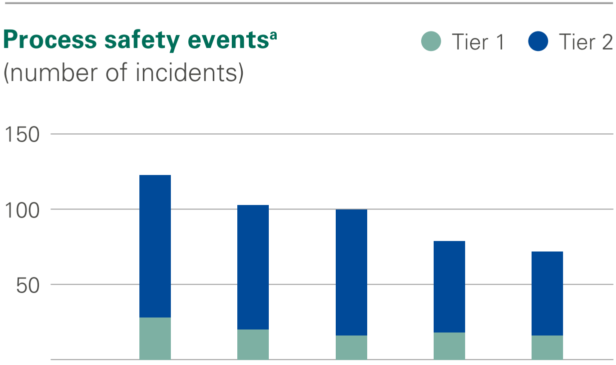

	2014	2015	2016	2017	2018
Tier 1	28	20	16	18	16
Tier 2	95	83	84	61	56

Source: BP's Sustainability Review, 2018, from bp.com

Company safety indicators: example

Recordable injury frequency[b c]

– workforce (per 200,000 hours worked)

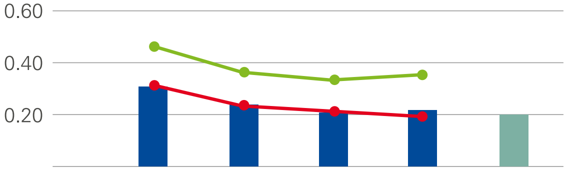

	2014	2015	2016	2017	2018
Workforce	0.31	0.24	0.21	0.22	0.20
Employees	0.27	0.20	0.19	0.20	0.15
Contractors	0.34	0.28	0.22	0.23	0.23

American Petroleum Institute US benchmark

International Association of Oil & Gas Producers benchmark

Source: BP's Sustainability Review, 2018, from bp.com

Illustration in civil aviation

- Typical metrics:
  - number of accidents per million flights
  - number of fatal accidents per million flights
  - number of people killed per year
  - number of hull losses per million flights
- Most widely published metrics concern public air transport operations in scheduled operations, using Western-built aircraft
- Accident rates tend to be higher for:
  - private or military flights
  - cargo operations, test flights
  - non-scheduled operations
  - aircraft built in former Eastern-block countries

IATA definition of an accident

- IATA (trade association for the major airlines) defines an accident as an event where *all* of the following criteria are satisfied:
  - Person(s) have boarded the aircraft with the intention of flight (either flight crew or passengers)
  - The intention of the flight is limited to normal commercial aviation activities, specifically scheduled/charter passenger or cargo service. Executive jet operations, training, maintenance/test flights are all excluded.
  - The aircraft is turbine powered and has a certificated Maximum Take-Off Weight (MTOW) > 5700 kg
  - The aircraft has sustained major structural damage exceeding 1 million USD or 10% of the aircraft's hull reserve value, whichever is lower, or has been declared a hull loss.
- Destruction using military weapons (*e.g.* MH 17 over Ukraine in 2014) not counted as an accident

Source: ICAO Annex 13, `icao.int`

IATA safety indicators for civil aviation

‘Fatal Accidents’ refer to accidents with at least one person on board the aircraft perishing as a result of the crash.

Jet & Turboprop Aircraft

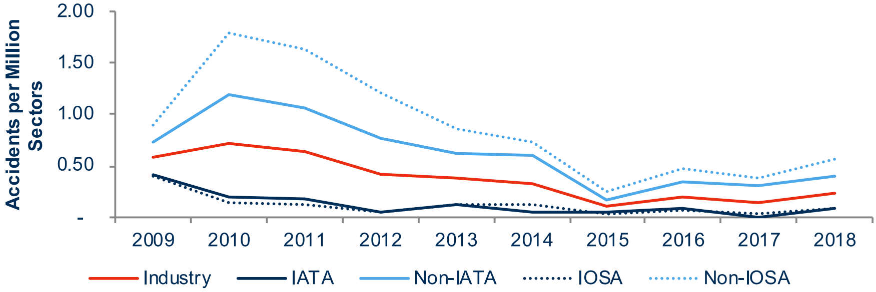

Source: IATA safety report for 2018, `iata.org`

RISK ENGINEERING

IATA safety indicators for civil aviation

All Accident Rate - Industry vs. IATA

This rate includes accidents for all aircraft: it includes Substantial Damage and Hull Loss accidents for jets and turboprops. The All Accident rate is calculated as the number of accidents per million sectors. This is the most comprehensive of the accident rates calculated by IATA.

	2009	2010	2011	2012	2013	2014	Trend	2009-2013
Industry	2.71	2.77	2.63	2.11	2.24	1.92		2.48
IATA Member Airlines	1.78	1.49	1.87	0.74	1.60	0.94		1.49

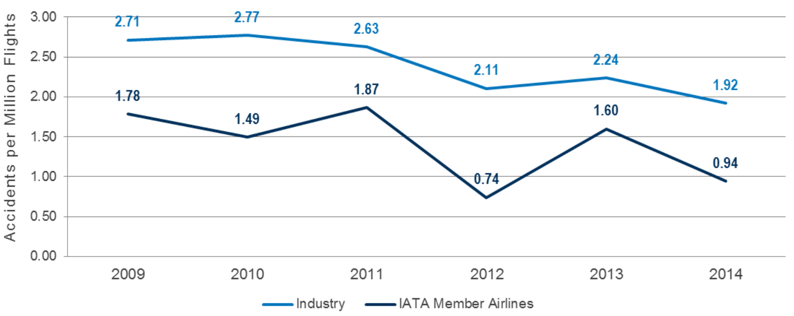

Source: IATA safety fact sheet for 2015, `iata.org`

Illustration: safety performance metrics in oil & gas industry

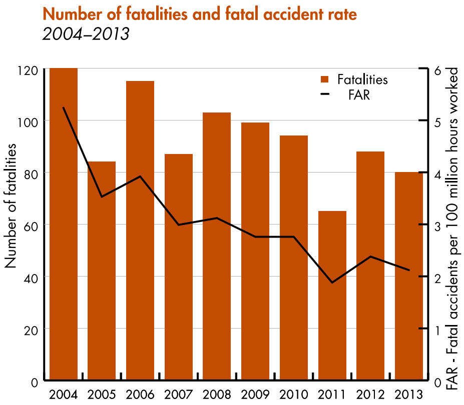

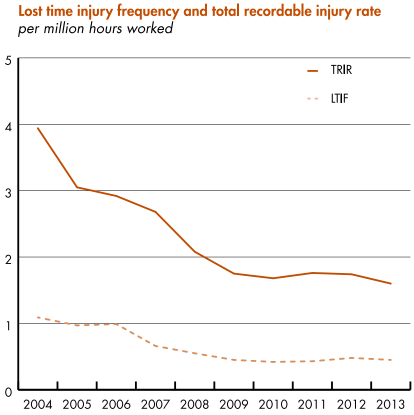

Source: OGP report on *Safety performance indicators – 2013 data*, `iogp.org/bookstore/product/safety-performance-indicators-2013-data/`

Illustration: safety performance metrics in railway transport

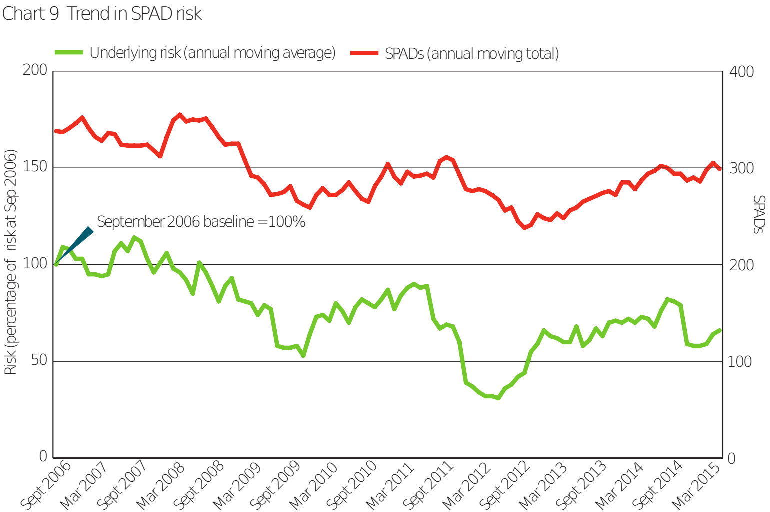

SPAD: Signal Passed At Danger ("burned red light")

Source: UK RSSB annual safety report, 2015

Illustration: safety performance metrics in railway transport

Chart 7 Potentially higher-risk train accidents

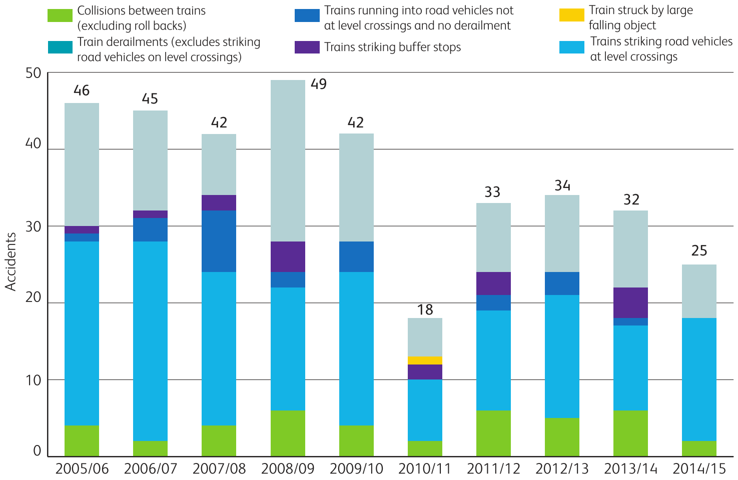

Source: UK RSSB annual safety report, 2015

Illustration: typical criteria used for nuclear power

- Typical society-level criterion: *"The use of nuclear energy must be safe; it shall not cause…"* (Finland)
- Typical technical targets, expressed probabilistically:
  - average core damage frequency (CDF) should be $< 10^{-4}$ per reactor year
  - large early release frequency (LERF) should be $< 10^{-5}$ per reactor year (accidents leading to significant release to atmosphere prior to evacuation of surrounding population)
- Note: actual (observed) CDF is around 10^{-3} per year worldwide!
  - 11 nuclear reactors out of 582 have suffered serious core damage over 14 400 reactor years
  - rate of 1 in every 1309 reactor years

Source: *Probabilistic Risk Criteria and Safety Goals*, Nuclear Energy Agency, 2009, `oecd-nea.org/nsd/docs/2009/csni-r2009-16.pdf`

Illustration: typical criteria for dam failure

> *The individual risk should be considered in terms of the "maximally exposed individual" that is permanently resident downstream of the dam. Typically the maximally exposed individual is exposed to the hazard significantly more than 50% of the time. The maximum level of individual risk should generally be less than 10^{-4}/year.*

— Canadian Dam Association guidelines

Interpreting and using metrics

Selecting relevant metrics

Questions to help you select safety metrics / KPIs that support safety management while minimizing unwanted side effects:

- What data do we need to really **understand safety**, not just as an *absence* of undesired events, but as a *presence* of something?
- Could some of our safety metrics **encourage under-reporting** of certain events?
  - watch out for the risk of developing a **target culture** (where meeting the numerical target becomes more important than operating safely and providing quality)...
- Is the **scope** of the measured undesirable events defined in a precise way?

RISK ENGINEERING

Watch out for "**watermelon safety metrics**"

Green on the outside, but red when you dig a little under the surface...

Risks of misuse of safety metrics

- Watch out for situations where safety management becomes a **bureaucratic exercise**, where risk metrics are misused to justify the *status quo* rather than identifying sources of progress

- Quoting safety researcher Sidney Dekker:

> *In a world where safety is increasingly a bureaucratic accountability that safety professionals need to show up, and to a variety of stakeholders who are located far away from the sharp end, it makes sense that safety gets organized around reportable numbers. Numbers are clean and easy to report, and easy to incentivize around. They are gratefully inhaled by greedy, if stunted and underinformed consumers: insurers, boards of directors, regulators, media, clients.*

Source: `safetydifferently.com/the-failed-state-of-safety/`

Source: `dilbert.com/strip/1998-11-22`

A disconnect

What most companies measure in terms of risk

What is most important for safety

The disconnect between these two has to be reconciled by safety professionals and other workers

RISK ENGINEERING

Interpretation and use of quantitative risk targets

Some issues to consider in the use of risk targets:

- Are all initiating events considered?
  - terrorism, loaded jet airplane striking facility…
- What are the consequences of not achieving the target?
  - immediate shutdown, obligation to revise safety case, warning from regulator…
- Are risk targets **revised periodically** to account for society's desire for continual improvement of safety performance?
- Are risk targets the same for new facilities (state-of-the-art design) and old facilities?

Beware the McNamara fallacy

> *The first step is to measure whatever can be easily measured. This is okay as far as it goes. The second step is to disregard that which can't be measured or give it an arbitrary quantitative value. This is artificial and misleading. The third step is to presume that what can't be measured easily really isn't very important. This is blindness. The fourth step is to say that what can't be easily measured really doesn't exist. This is suicide.*

– [Smith 1972]

More: article J. Kingston (2017), *The McNamara fallacy blocks foresight for safety*, in proceedings of ESReDA seminar *Enhancing safety: the challenge of foresight.*

Criteria for evaluating risk metrics

- Validity: reflects an important aspect of risk
- Reliability: can be clearly defined and repetitively calculated across analyses
- Transparency: possible to evaluate with respect to informative and normative content
- Unambiguity: precise analytic boundaries
- Contextuality: captures relevant decision factors
- Communicability: adaptable to communication
- Consistency: provides unambiguous advice
- Comparability: applicable across different systems
- Specificity: relevant to the particular system
- Rationality: logically sound
- Acceptability: politically acceptable

Source of this list: *Risk Metrics: Interpretation and Choice*, I. L. Johansen & M. Rausand, `frigg.ivt.ntnu.no/ross/reports/risk-metric.pdf`

RISK ENGINEERING

Misuse of safety metrics: illustration

- The UK health service NHS uses metrics to measure the performance of hospitals
  - and sets associated quantified targets that healthcare centers must meet
- One target: anyone admitted to an emergency room must be treated within 4 hours
- Some managers accused in 2016 of requiring patients to be left in ambulances during busy periods rather than admitting them
  - strategy called “stacking” which can improve performance according to this metric (delays the “clock starts ticking” moment)
  - clearly very bad for safety of patients
- Recall Goodhart’s law: “When a measure becomes a target, it ceases to be a good measure”

Source: `telegraph.co.uk/news/9637865/`

RISK ENGINEERING

Another illustration of misuse of safety metrics

After the 2010 Macondo/Deepwater Horizon catastrophe, executives from Transocean, the company operating the drilling rig, granted themselves a record bonus, based in part on the company's "safety performance". A statement justifying executive pay states "Notwithstanding the tragic loss of life in the Gulf of Mexico, we achieved an exemplary statistical safety record". The bonuses were based on TRIR, an occupational safety metric, little affected by the Macondo disaster, and clearly a poor proxy for management performance on safety in that year.

Goodhart's law

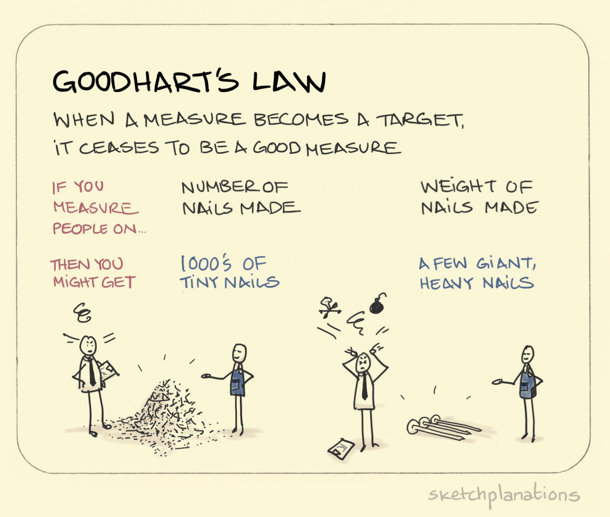

Once upon a time, there was a factory in the Soviet Union that made nails. Unfortunately, Moscow set quotas on their nail production, and they began working to meet the quotas as described, rather than doing anything useful. When they set quotas by quantity, they churned out hundreds of thousands of tiny, useless nails. When Moscow realized this was not useful and set a quota by weight instead, they started building big, heavy railroad spike-type nails that weighed a pound each.

Source: Jono Hey, `sketchplanations.com/goodharts-law`, CC BY-NC licence

The cobra effect

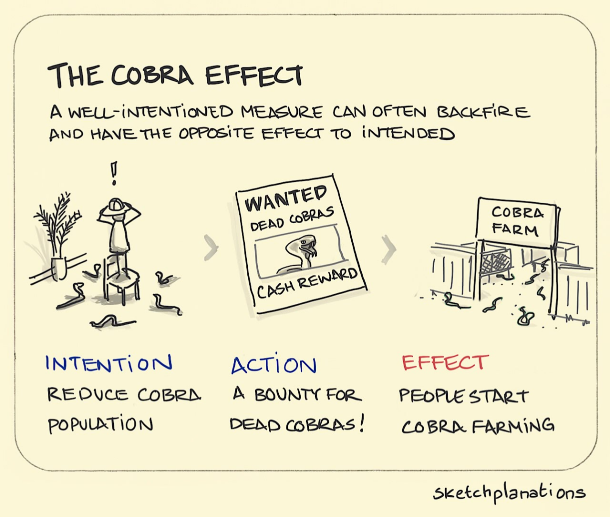

Source: Jono Hey, `sketchplanations.com/the-cobra-effect`, CC BY-NC licence

Complementary information sources

- Risk metrics and KPIs are tools used for an analytic and mechanistic view of risk management
  - risk as a product of event probability and event consequences
  - safety as a system attribute that can unproblematically be measured and monitored
- Social scientists suggest there is another dimension of risk, which is continually interpreted and debated by a community of practice
  - safety seen as a positive capacity for control, rather than as the absence of hazardous events
  - significant discussion between different professional groups may be needed to analyze the safety implications of an incident
  - the organization's ability to cope with these failures is as important as the engineering measure of their severity
- Suggestion: these viewpoints are complementary
  - both contribute to improving safety in complex systems

RISK ENGINEERING

Image credits

- Tape measure on slide 2: vintspiration via `flic.kr/p/bkG7dX`, CC BY-NC-ND licence
- Tape measure (slide 4): antony mayfield via `flic.kr/p/5UDjAw`, CC BY licence
- Glen Canyon dam (slide 27): Ashwin Kumar via `flic.kr/p/XvMxQj`, CC BY-SA licence
- Watermelon (slide 30): sama093 via `flic.kr/p/wHrX4v`, CC BY-NC-ND licence
- Ambulances (slide 35): Greg Clarke via `flic.kr/p/JftB5v`, CC BY licence
- Deepwater Horizon rig (slide 38): US Coast Guard, public domain
- Books (slide 43): FutUndBeidl via `flic.kr/p/cdaEDL`, CC BY licence

Further reading

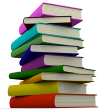

- *A guide to measuring health & safety performance*, UK Health and Safety Executive (2001), `hse.gov.uk/opsunit/perfmeas.pdf`
- CCPS book *Guidelines for Process Safety Metrics*, Wiley, 2009 (ISBN: 978-0470572122)
- Chapter *Risk measurement and metrics* of the free textbook *Enterprise and individual risk management*, available online
- Metrics for financial risk: see the slideset on *Estimating Value at Risk* from `risk-engineering.org/VaR/`

For more free content on risk engineering, visit `risk-engineering.org`

RISK ENGINEERING

Feedback welcome!

Was some of the content unclear? Which parts were most useful to you? Your comments to `feedback@risk-engineering.org` (email) or `@LearnRiskEng` (Twitter) will help us to improve these materials. Thanks!

@LearnRiskEng

fb.me/RiskEngineering

For more free content on risk engineering, visit `risk-engineering.org`