RiskNZ

WELCOME TO

RiskNZ Lunchtime Seminar: 6 December 2022

RiskNZ is proudly sponsored by:

Premier Sponsors

Sponsors

Empirical Risk Management

Rohan Light

DECISV
GETTING PAST NOW

Rohan’s Risk Empirical

DECISV
GETTING PAST NOW

1. My take on decision risk
2. Some outputs from current fieldwork
3. The Flaw of Averages
4. Hacking Monte Carlo
5. Some causal risk diagrams
6. Some AI risk stuff

DECISV
GETTING PAST NOW

empirical (adj.)
1560s, originally in medicine, "pertaining to or derived from experience or experiments," from Latin empiricus (n.) "a physician guided by experience," from Greek empeirikos "experienced," from empeiria "experience; **practice without knowledge**," especially in medicine, from empeiros "experienced (in a thing), **proven by use**," from assimilated form of en "in" (see en- (2)) + peira "trial, experiment,"

Governance Risk & Compliance Professional (primary risk credential)
Certified Auditor - UK GDPR (AI risk, specialism)
Re-identification Risk Lead (data risk)
Manager Insight and Analytics (client service-delivery decision risk)
Portfolio Risk Specialist, Inland Revenue (enterprise investment risk)
Senior Risk Advisor, Inland Revenue (enterprise compliance risk)
Commercial Analyst, Te Papa Tongarewa

"Why, my man, I am ashamed of you, dodging that way, they couldn’t hit an elephant at this distance"

“General, I dodged a shell once, and if I hadn't, it would have taken my head off. I believe in dodging”

“[Laughing] All right, my man; go to your place”

For a third time the same shrill whistle, closing with a dull, heavy stroke

DECISV
GETTING PAST HOW

Decision risk is the chance and cost of being wrong… where wrong is a better course of action we would have taken… if we could have

Sometimes there are no real options and we have no real agency. Sometimes we don't know what that better course of action was

Sometimes we have few chances and a huge cost regardless. And yet… we still have to make that decision

Fieldwork 1: there is a hidden trap in this handoff…

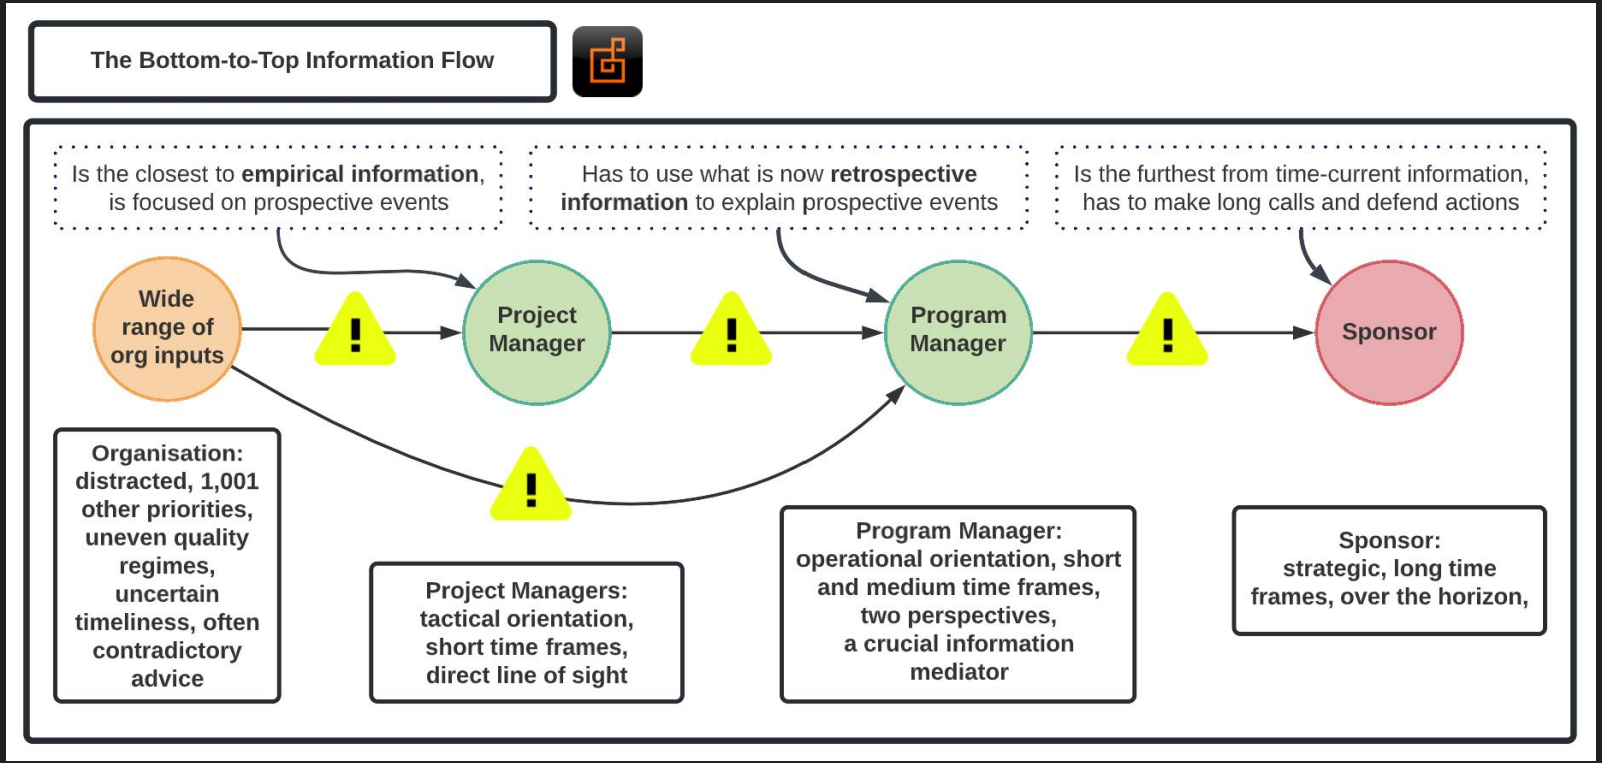

The Flaw of Averages: a plan made using averaged inputs is wrong, on average. Wah?

As a program manager, you look at a set averaged metrics of the component projects. Contextual to the current project state, which can be sampled many times along a timeline, most of those metrics will be wrong

We quickly become interested in how wrong and whether any states of wrong are survivable (or even advantageous)

This is a statistical quirk and I study them in decision making. How to illustrate?

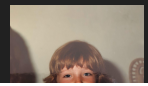

The Mad Hatters Amazing Construction Programme	Funds Approved ($m)
Replica Eiffel Tower on the Waterfront	276.00
12th Floor Giant Spa Pool for Rich Execs	39.00
Underwater Eastbourne Hyperloop	72.00
Personal Sleep Pods on Building Exterior	118.00
Beehive Great Glass Elevator	29.00
	534.00

Note the route the Marshmallow Man took prior to the second truck (the one with the dashcam) crossing his line of travel

The average **position** for his journey is good news: we get **NOT SPLAT** (he sidestepped the first truck)

Yet, sadly, the average **state** of Marshmallow Man in the narrative is **SPLAT**: he's huge, turns slowly and there are just too many trucks to dodge

The inevitability of **SPLAT** makes all the discussion about **NOT SPLAT** kind of cruel

Fieldwork 2: I add probability up front for mini Monte Carlo

Risk Volatility Test

Volatility Test 1: blowing the budget. This test looks at how much more we might have to pay over the lifetime of the project

The Mad Hatters Amazing Construction Programme	Funds Approved ($m)	Range of Outcome Likelihoods*				Financial Risk Implications		
		On Budget	Best Case Variance	Middle Range	Worst Case Variance	Best Case Variance $m	Middle Range $m	Worst Case Variance $m
Weights and Ranges		*1-15 %*	*16-31 %*	*32-85 %*	*86-100 %*	*0.23*	*0.53*	*0.93*
Waterfront Eiffel Tower Replica	276.00	**1****	**2**	***4***	***3***	***339.48***	*422.28*	*532.68*
12th Floor Giant Spa Pool for Rich Execs	39.00	**2**	**0**	***7***	**1**	*47.97*	*59.67*	***75.27***
Underwater Eastbourne Hyperloop	72.00	**2**	**1**	***4***	**3**	*88.56*	*110.16*	***138.96***
Personal Sleep Pods on Building Exterior	118.00	**3**	**0**	***6***	**1**	*145.14*	*180.54*	***227.74***
Beehive Great Glass Elevator	29.00	**2**	**1**	***3***	**3**	*35.67*	***44.37***	*55.97*
	534.00	**9**	**4**	**24**	**11**	**656.82**	**817.02**	**1,030.62**

Blue bolded numbers show a for arguments sake possible forward risk portfolio

Red bolded numbers shows the results of 10 simulations, for this portfolio, in current conditions

* Range calculated from a 2017 meta-study, with additional elements via McKinsey, both based on real world construction industry actuals

Volatility Test 2: running late. This test looks at when a project actually delivers and closes

(Note: post-covid median delay is 200 days)

Range of Delayed Delivery Date Likelihoods***

On time	Best Case: up to 60 days late	Steep Spike: up to 250 days late	Worst Case: more than 1 year late	The Deep Down Side: more than 2 years late
1-15%	*16-60%*	*61-83%*	*84-97%*	*98-100%*
1	***6***	***1***	2	0
2	***6***	1	***1***	0
0	***7***	1	***1***	1
4	3	***2***	1	0
1	***4***	3	***2***	0
8	**26**	**8**	**7**	**1**

*** Range calculated from a 2022 machine learning meta-study

Different courses of action?

You take a test to see if you have icthyscrathly disease: it comes back positive. How likely is it that you have icthyscrathly?

The forward probability is the probability of a positive test, given that you have icthyscrathly: P1 (test|disease)

This is the same for most humans because it depends only on the technical capability of the device to perform tests

The inverse probability is the one you're actually interested in - what is the probability that **you** have icthyscrathly, given the positive test: P2 (disease|test)

P1 (test|disease) doesn't behave the same way as P2 (disease|test) does. Many/most organisations answer inverse probability questions with forward probability answers

Epistemic dilemma?

When someone says they want to know **'what to do'**, they are appealing to **instrumental rationality**

When they ask about **'what is true'**, they are appealing to **epistemic rationality**

It is common to see groups who are instrumentally rational… as well as epistemically irrational

Epistemic rationality is tied to empirical and experiential learning. Consider two advisors, one epistemically irrational

Their boss asks a 'what to do' question of the two advisors. Both sources present the same, but one is batsh*t crazy: the advice differs. What to do indeed

You can't be 'in control' w/o controls

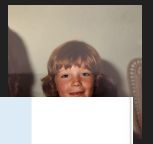

Mad Hatters Amazing Construction Programme

#	These controls consist of activities, systems, policies, devices, practices, processes, techniques or tools that meaningfully reduce the risk to the PoW delivering on its mission	CONTROL SOURCE	CONTROL OWNER	Preventative	Detective	Corrective	Effective	Good	Partial	Ineffective
C1 HRC	Clear decision making processes			☑	☐	☐	☑	☐	☐	☐
C2 HRC	An integrated work plan, status report and risk/issue/decision register			☐	☐	☑	☐	☐	☑	☐
C3 HRC	Relevant, timely communications to staff, and providers			☑	☐	☐	☐	☐	☑	☐
C4 HRC	Daily SitRep Report			☐	☑	☐	☑	☐	☐	☐
C5 HRC	An approved 2DHB Response Framework			☐	☐	☑	☑	☐	☐	☐
C6 HRC	Approved 2DHB minimum service levels			☑	☐	☐	☐	☐	☑	☐
C7 HRC	Service level endemic plans			☑	☐	☐	☐	☑	☐	☐
C8 HRC	Integrated staff shortages management plan			☐	☐	☑	☐	☐	☑	☐

DECISV
GETTING PAST HOW

Levels of Learning	Activities and Information	Types of Question
	Imagining counterfactuals	"What if I had done..."
	Doing interventions	"What if I do..."
 	Seeing associations	"What if I see..."

Note:

In causal work, the **estimand** and the **query** do not coincide
Here, the estimand is a recipe; **data** are the ingredients applied to the recipe; and the **estimate** is what comes out of the oven

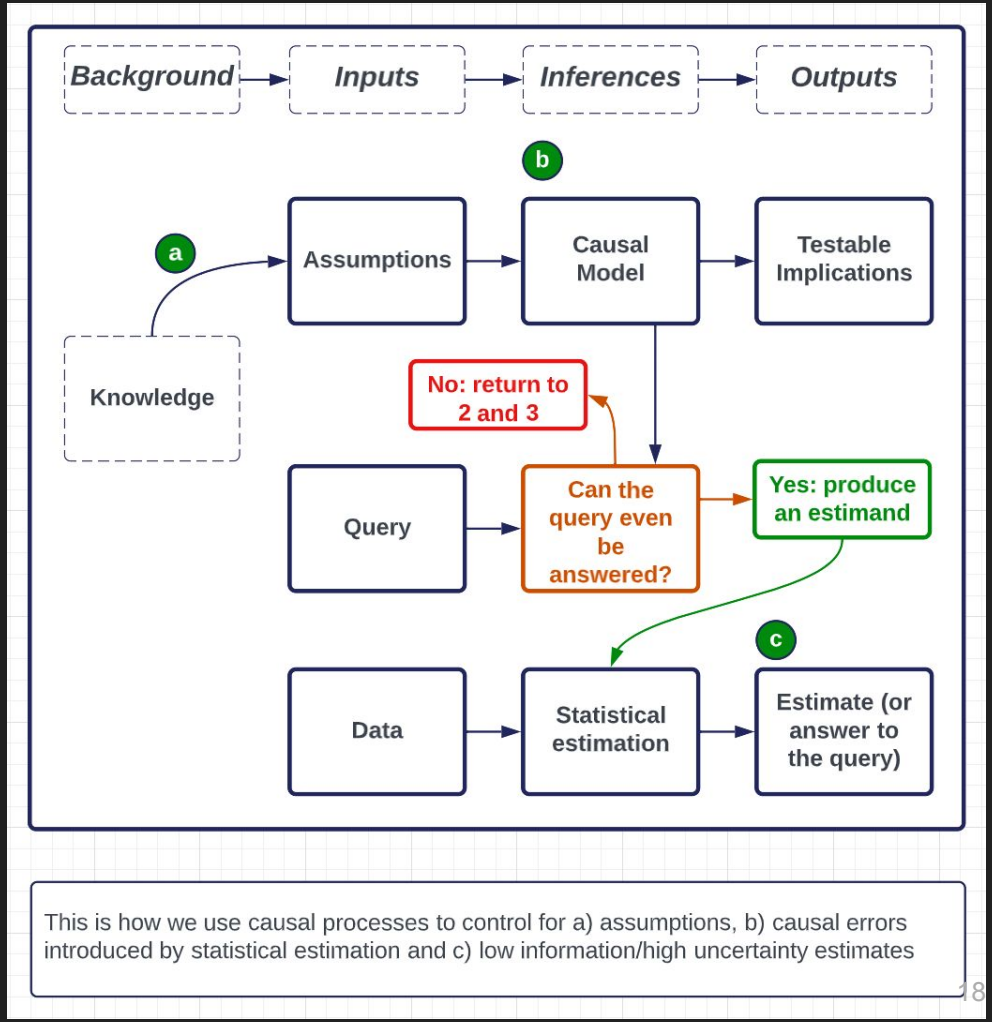

This is how we use causal processes to control for a) assumptions, b) causal errors introduced by statistical estimation and c) low information/high uncertainty estimates

DECISV
GETTING PAST HOW

To falsify:

To submit a claim to a known prove-it-wrong test: if that test fails, then the claim is true, to the extent of the failed test

An estimand:

A test mechanism that produces stable estimates of known evidential value

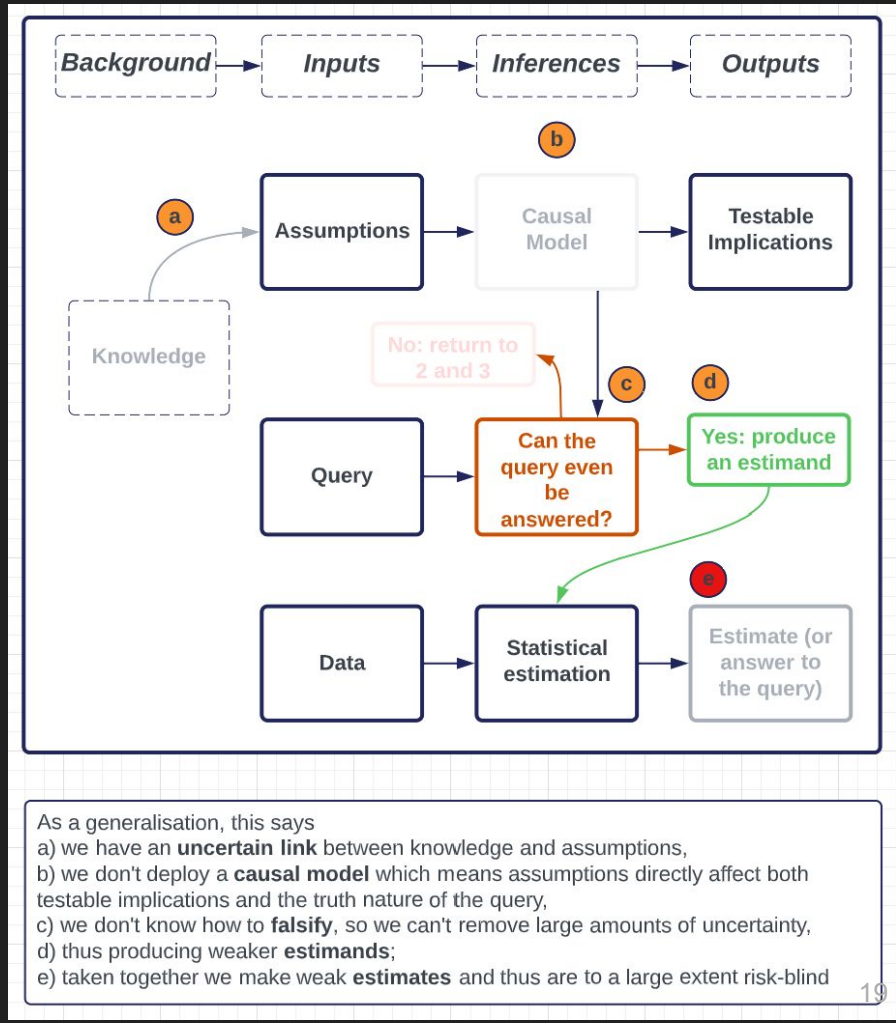

As a generalisation, this says
a) we have an **uncertain link** between knowledge and assumptions,
b) we don't deploy a **causal model** which means assumptions directly affect both testable implications and the truth nature of the query,
c) we don't know how to **falsify**, so we can't remove large amounts of uncertainty,
d) thus producing weaker **estimands**;
e) taken together we make weak **estimates** and thus are to a large extent risk-blind

Any Programme-of-Work Four Basic Risk Scenarios

CCHV EPMO 1/8/22

We labour under uncertainty
- These are causal risk models: they use the language of **causal science** to model prospective risk, without losing sight of innovation opportunities.
- They encourage **matakite**, foresight and situational awareness for the purposes of risk avoidance and benefit realisation

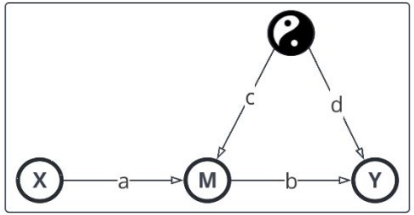

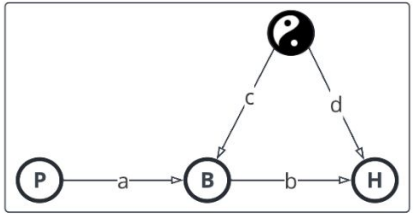

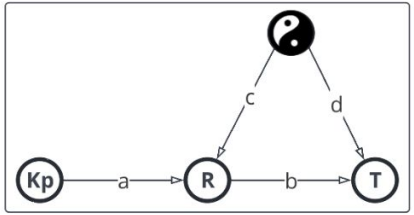

There are two sides to uncertainty (and two sides to each of those two sides)
- Some of these effects are helpful (**opportunities**), while others are not (**threats**)
- We manage opportunities as **innovation** and the threats as **risk**, all the time not fully understanding how the two relate

How to read it
- We are applying **intervention** X for effect (a) on M, which is a 'mediator', some other thing that is intended to produce a secondary effect (b) upon **target** state Y
- However, there are **unknown, unseen, unestimated and/or unconsidered** effects (c and d) on both our mediator and target state

How to read it
- Our context is **change investment**: building and sustaining value for Te Whatu Ora and local communities
- The model says **project** P is an **intervention** on some **health capability** group B to improve h**ealth outcomes** H
- Note: the project is only a temporary entity

How to read it
- We can speak te reo **Māori**!
- The operation of the causal model remains unchanged: it speaks all human languages because humans think causally about things
- The project is the ***kaupapa*** Kp, the business group is the ***rōpū*** R, the health outcomes are those of the people, the ***tangata*** T

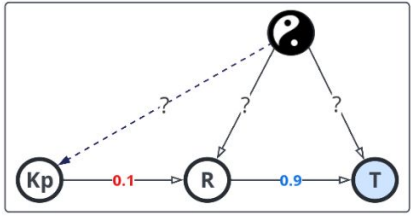

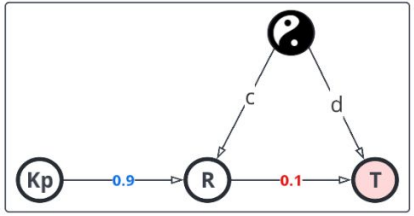

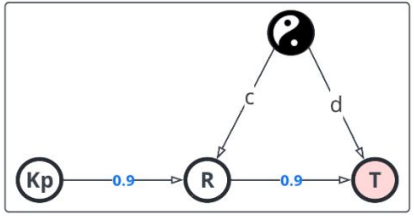

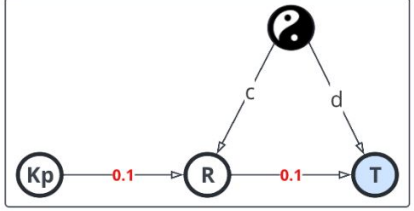

How to read it
- This is the **first** basic change investment risk case: the effect (a) of the kaupapa on the rōpū did not have an impact on a sucessful outcome
- A common situation is where the kaupapa has insufficient independence from other rōpū, who might have a conflict of interest over the kaupapa

How to read it
- This is the **second** basic change investment risk case: the effect (a) of the kaupapa on the rōpū was excellent, but the rōpū had no impact on an unsuccessful outcome
- A common situation is where the rōpū has insufficient mana with other rōpū, who might have a conflict of interest over tangata

How to read it
- This is the **third** basic change investment risk case: both kaupapa and rōpū were successful, but the tangata felt no improvement
- This indicates deep seated, strategic or cultural obstacles, which may or may not have been identified by Te Whatu Ora, the kaupapa or the rōpū

How to read it
- This is the **fourth** basic change investment risk case: both kaupapa and rōpū were unsuccessful, but the tangata felt improvement
- This may indicate significant untapped or unknown innovation opportunities; or something greater within tangata is empowering and enabling change

DECISV
GETTING PAST HOW

Epistemic irrationality bad causal model herd effect

DECISV
GETTING PAST HOW

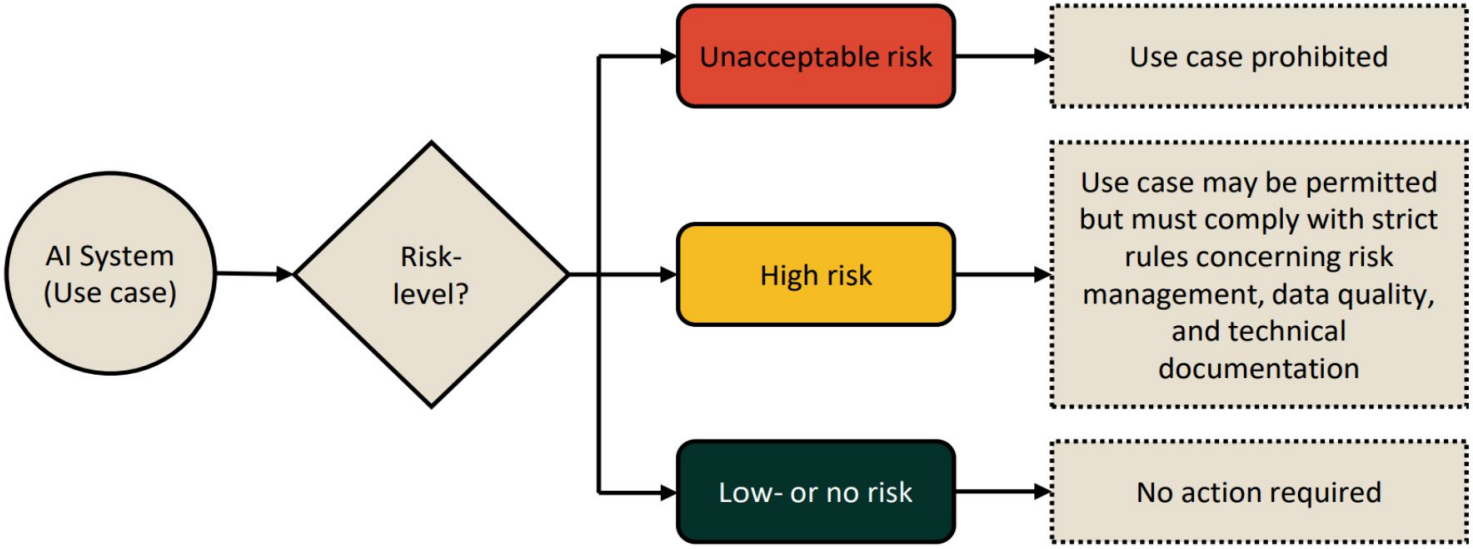

Figure 1: *Risk categories for AI use cases under the AIA [14]*

DECISV
GETTING PAST HOW

Risk-based approach

Digital ethics
(hard ethics)

FORHUMANITY

Relevant Legal Frameworks,

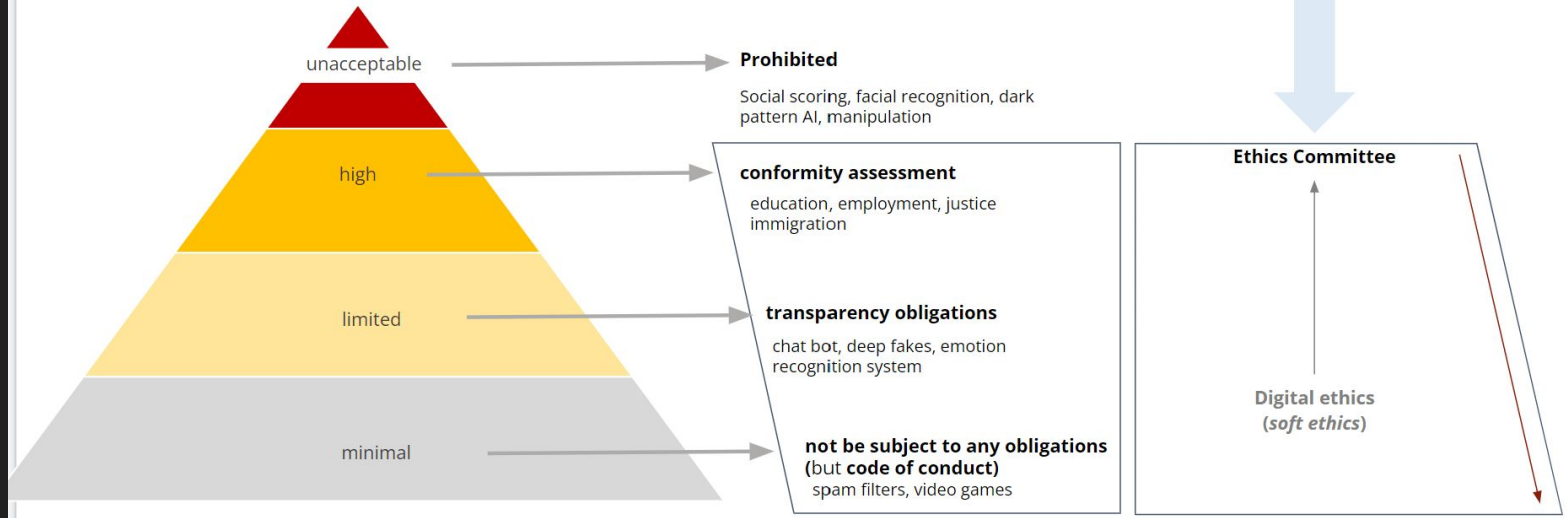

Type of risk

DECISV
GETTING PAST HOW

Documentation

FORHUMANITY

General Oversight and Accountability	The organisation shall conduct an **Ethical Risk Assessment** for all instances of **Ethical Choice** found in the algorithmic lifecycle including documenting:	Public Disclosure Document

Identification of the nature of the decision (e.g. binary or multi-faceted)

a

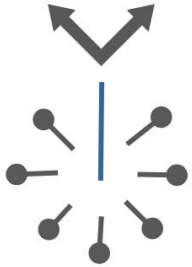

An unbiased display of Pros and cons for each side of the choice

b

c

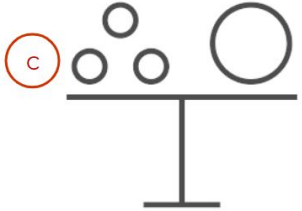

An unbiased display of tensions and Trade-offs for each side of the choice

Decision reached

d

e

Residual risk associated with the choice

Traceability of the deployment of the choice

f

g

Receipt of Traceability when the deployment is satisfied

DECISV
GETTING PAST HOW

Effective observation doesn’t need lots of data

Do you have any questions?

Type them into the Q&A platform

RiskNZ

Thank you for joining us today