

To Know Enough to Ask the Questions

A Snapshot of New Zealand News Media Reporting on
AI in 2025

Hannah Clarice Betts | MSc in Science in Society |
Victoria University of Wellington

This presentation:

- Why AI discourse matters now in Aotearoa New Zealand
- Research questions
- Mixed-methods methodology
- Content analysis findings from 246 news articles
- Journalist interview findings
- Discussion, recommendations, and implications

Why This Research Matters Now

- AI adoption in New Zealand is accelerating across government, business, education, and media.
- Regulatory approach is currently “light-touch”



There may be exposure to risks with increasing use

- Public discourse is a medium to negotiate **which risks are acceptable trade-offs** for the benefits of adopting AI.
- Very little is known about the state of public discourse.



State of public discourse is a worthy object of study

- Journalism remains a public forum of discourse:
 - 69% are “highly interested” in news
 - 32% “trust” the news
 - Major news outlets are consistently drawing audiences of 45% of NZers, 35%, 30%, 23% (each).
- Outside of journalism, there is very little public-facing resource supporting NZ public to increase their knowledge about AI.



Two potential roles of news media:

1. A means to see current public discourse
2. A means to increase public AI literacy

Research Questions

Research question: How are risks and benefits of AI presented within New Zealand news media, and why?

I answered this through a mixed-methods explanatory-sequential design.

Phase 1. Quantitative

1. Which risks are referenced, and how often?
2. Which benefits are referenced, and how often?
3. What level of explanatory depth is provided for references to risks?

Phase 2. Qualitative

1. How do journalists approach reporting on high-risk AI applications?
2. Why are we observing any patterns found in Phase 1?

Teaser of findings

Phase 1	Phase 2	Finding

Phase 1: Content analysis

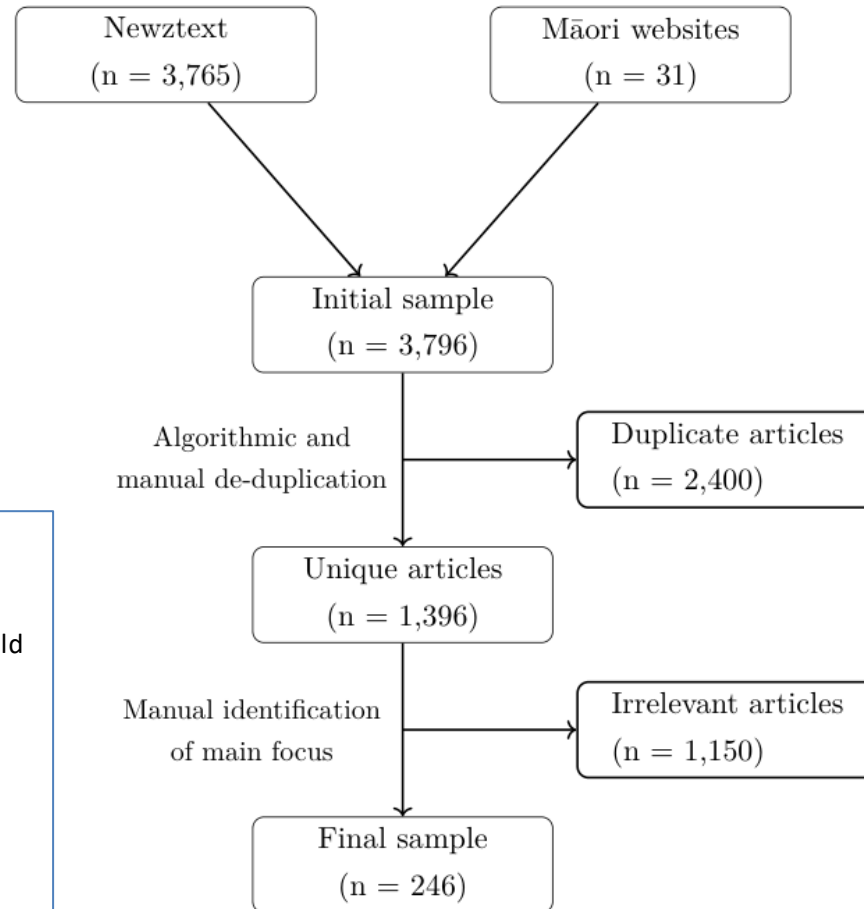


Figure 3.1: Article Selection and Screening Process

Outlets

Newztext database

- NZME (The New Zealand Herald and local newspapers)
- Radio New Zealand
- Scoop
- Stuff (and local newspapers)
- The Spinoff

Manually retrieved

- Te Ao Māori News
- Radio Waatea

I categorised tone of article, and type of AI covered (existing vs. hypothetical).

Risks coded **deductively** using the Domain Taxonomy of AI Risks.

Benefits coded **inductively** because no accepted taxonomy existed.

I looked at whether relevant R and B were mentioned together (an indicator of mention of trade-offs).


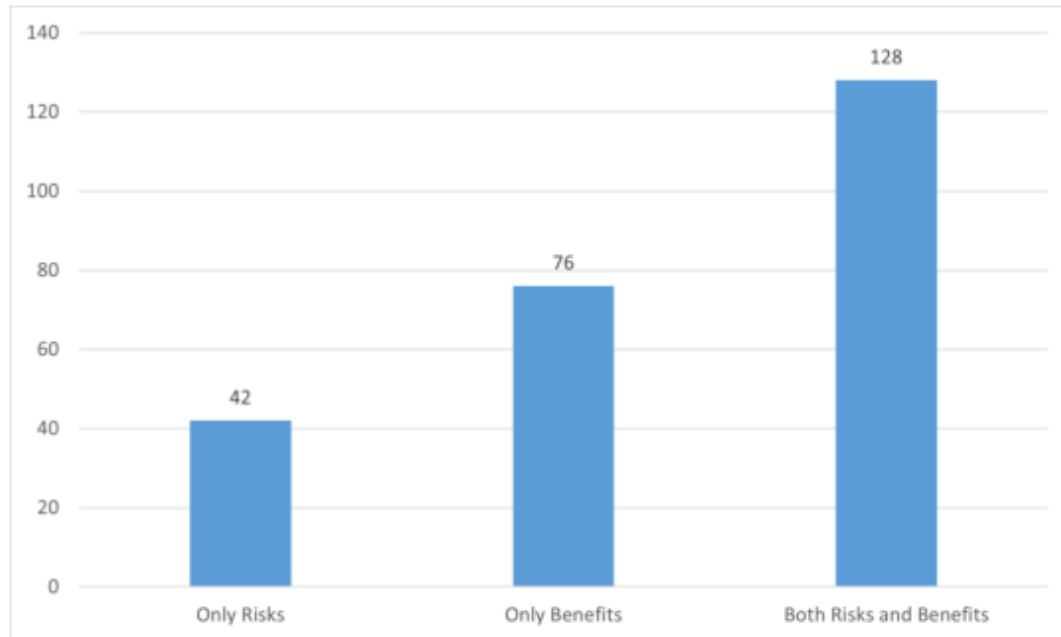
Finally, Risks were analysed for explanatory depth.

Tone of coverage

More articles mentioned benefits than mentioned risks.

Most articles mentioned **both risks and benefits** of AI.

BUT if you look at the *tone* rather than *count* of references to benefits/risks, **more articles are positive towards AI**, even if they mention a risk.




Terminology: “affirmed” vs. “negated”

- Some passages explicitly referenced the risk or benefit *not* occurring, not being expected to occur, or not being important even if it did occur.
- I termed this as being “negated”
- References that describe a risk or benefit occurring, being expected to occur, or being important if they do occur were termed as being “affirmed”.

Benefit/Risk	Quote	Decision
RISK 1.3 – Unequal performance across groups	‘A decade ago, facial recognition sparked controversy worldwide over misidentifying black and brown faces much more than white faces, and especially for women [...] But now some systems are down to 0.02 percent inaccuracy, benchmark testing by the US National Institute of Standards and Technology (NIST) shows.’	Reference to Risk 1.3 – affirmed and Reference to Risk 1.3 – negated. States and describes an AI system performing less accurately when presented with different groups, and also asserts that this doesn’t remain a current concern.

Which AI Risks Were Discussed Most? (Domains)



Which AI Risks Were Discussed Most? (Subdomains)






Figure 4.5: Number of New Zealand News Articles that Affirmed or Negated a Subdomain of AI Risk.

Which AI Benefits Were Discussed Most? (Domains)



Which AI Benefits Were Discussed Most? (Subdomains)



Did articles discuss the trade-offs between Rs and Bs?

	BENEFITS							Task automation				Individual agency				Organisational optimisation							Wellbeing					Technological progress					Total
	0	1	1.1	1.2	1.3	1.4	1.5	2	2.1	2.2	2.3	3	3.1	3.2	3.3	3.4	3.5	3.6	4	4.1	4.2	4.3	4.4	4.5	5	5.1	5.2	5.3	5.4				
RISKS	0	8	13	8	5	2	8	4	5	6	15	28	3	18	7	7	13	17	2	0	6	2	0	3	6	10	9	0	5	7	29		
Discriminath & toxicity	1	1	6	1	2	1	1	0	1	0	3	5	2	8	2	2	2	5	1	0	1	0	0	2	0	1	0	1	0	13			
	1.1	0	3	0	1	0	0	0	0	0	1	3	1	4	0	1	2	2	0	0	0	0	0	1	0	0	0	0	0	5			
	1.2	0	1	0	0	0	0	0	2	0	1	3	0	1	0	0	0	1	0	0	1	0	0	0	1	0	0	0	0	10			
	1.3	0	1	0	1	0	0	0	1	0	1	1	2	0	1	2	0	0	1	0	0	0	2	0	0	0	0	0	0	0	6		
Privacy & security	2	4	3	2	1	0	3	1	2	2	4	9	3	6	3	1	5	6	1	0	1	1	0	1	2	3	2	0	1	2	16		
	2.1	0	2	1	0	1	1	0	0	6	4	9	1	5	0	4	3	3	0	0	3	0	0	0	1	2	2	0	1	1	14		
	2.2	0	3	0	1	1	1	0	1	1	1	3	0	3	1	2	0	2	2	0	1	0	0	0	1	0	1	0	1	0	9		
Misinform ation	3	0	1	1	0	0	0	0	0	0	1	1	0	1	1	0	1	0	0	0	0	0	0	1	0	0	0	1	0	2			
	3.1	0	4	3	1	0	1	1	1	8	10	13	1	5	2	4	4	7	0	0	6	0	1	0	1	6	3	0	0	1	20		
	3.2	0	1	3	1	0	0	0	1	1	2	2	0	1	0	2	1	1	0	0	0	0	0	0	0	1	1	1	1	0	13		
Malicious actors & misuse	4	1	0	0	0	1	0	0	0	0	3	3	0	2	0	0	0	0	0	0	1	0	0	1	0	0	0	0	1	7			
	4.1	1	2	2	0	1	0	1	0	0	3	1	2	2	3	1	5	2	4	0	0	0	0	0	1	1	2	1	6	0	18		
	4.2	0	0	2	1	1	1	0	0	0	1	1	2	1	2	0	2	2	3	0	1	1	0	1	0	1	0	0	1	0	6		
	4.3	3	4	3	0	0	0	0	1	6	17	10	5	7	1	3	3	6	0	0	3	1	0	0	2	5	0	0	1	0	40		
Human-computer interactn	5	0	1	0	0	0	0	0	0	0	1	1	1	1	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	1			
	5.1	0	3	1	1	0	1	0	1	1	2	3	0	3	0	1	3	3	2	0	2	0	0	0	3	2	0	0	0	1	9		
	5.2	1	5	4	1	1	3	0	2	6	8	9	0	5	2	6	2	6	2	0	2	0	0	0	4	4	2	0	1	0	23		
Socioeconomic & environmental harms	6	0	1	1	0	0	1	0	1	0	1	1	0	0	0	0	0	1	1	0	1	0	0	1	0	1	1	0	0	0	4		
	6.1	1	3	2	2	0	2	1	2	2	6	4	2	3	1	4	0	2	0	0	1	2	0	0	0	2	2	0	0	0	17		
	6.2	2	6	6	1	0	2	0	2	4	13	12	3	11	1	7	6	11	1	0	2	1	0	2	4	9	5	0	2	4	28		
	6.3	3	5	8	1	0	1	2	1	5	7	10	1	7	1	3	4	7	1	0	2	2	1	3	2	6	4	0	1	4	25		
	6.4	0	2	4	0	2	1	0	1	0	2	2	1	1	3	0	3	2	2	0	0	0	0	0	1	2	0	0	2	1	8		
	6.5	4	12	11	3	3	4	2	3	7	13	18	5	16	8	5	13	14	9	0	5	0	0	1	3	7	6	0	8	5	60		
6.6	0	1	1	1	0	1	0	0	0	1	3	0	1	0	0	0	0	1	0	0	1	0	1	0	1	2	0	1	0	6			
AI system safety, failures & limitations	7	1	2	2	0	0	0	1	1	0	2	2	0	1	1	0	1	1	0	0	0	0	0	1	1	0	0	2	2	3			
	7.1	1	1	1	0	0	0	0	0	0	2	2	1	2	1	0	2	2	0	0	1	0	0	0	1	2	1	0	1	0	4		
	7.2	0	1	4	1	2	1	0	0	0	1	1	2	2	4	0	4	3	5	0	1	0	0	0	1	1	0	0	2	1	9		
	7.3	0	14	9	7	1	9	1	3	7	10	16	3	17	6	5	11	16	3	0	5	0	0	0	6	5	5	1	4	5	39		
	7.4	1	4	2	1	0	0	1	1	2	4	4	1	5	1	3	2	3	0	0	3	0	0	0	3	5	2	0	1	0	9		
	7.5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
	7.6	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
Newly identified Risks	8	0	1	0	0	0	0	0	0	1	2	5	0	1	0	3	2	1	0	0	0	0	0	1	1	1	1	0	0	1	9		
	8.1	1	1	1	0	0	0	0	0	2	0	2	1	2	1	2	2	1	3	0	1	0	0	1	0	0	3	0	2	0	9		
	8.2	3	1	1	0	0	0	0	0	1	1	2	2	2	0	1	1	1	1	0	0	5	0	2	1	1	1	0	2	1	9		
	8.3	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	8		
	8.4	2	2	3	1	0	2	0	1	2	2	5	0	4	0	4	0	5	0	0	1	0	0	1	2	4	1	0	1	2	14		
Total	14	45	38	21	12	27	14	8	34	51	72	20	56	24	27	57	55	26	0	27	17	11	11	13	26	45	11	33	25				



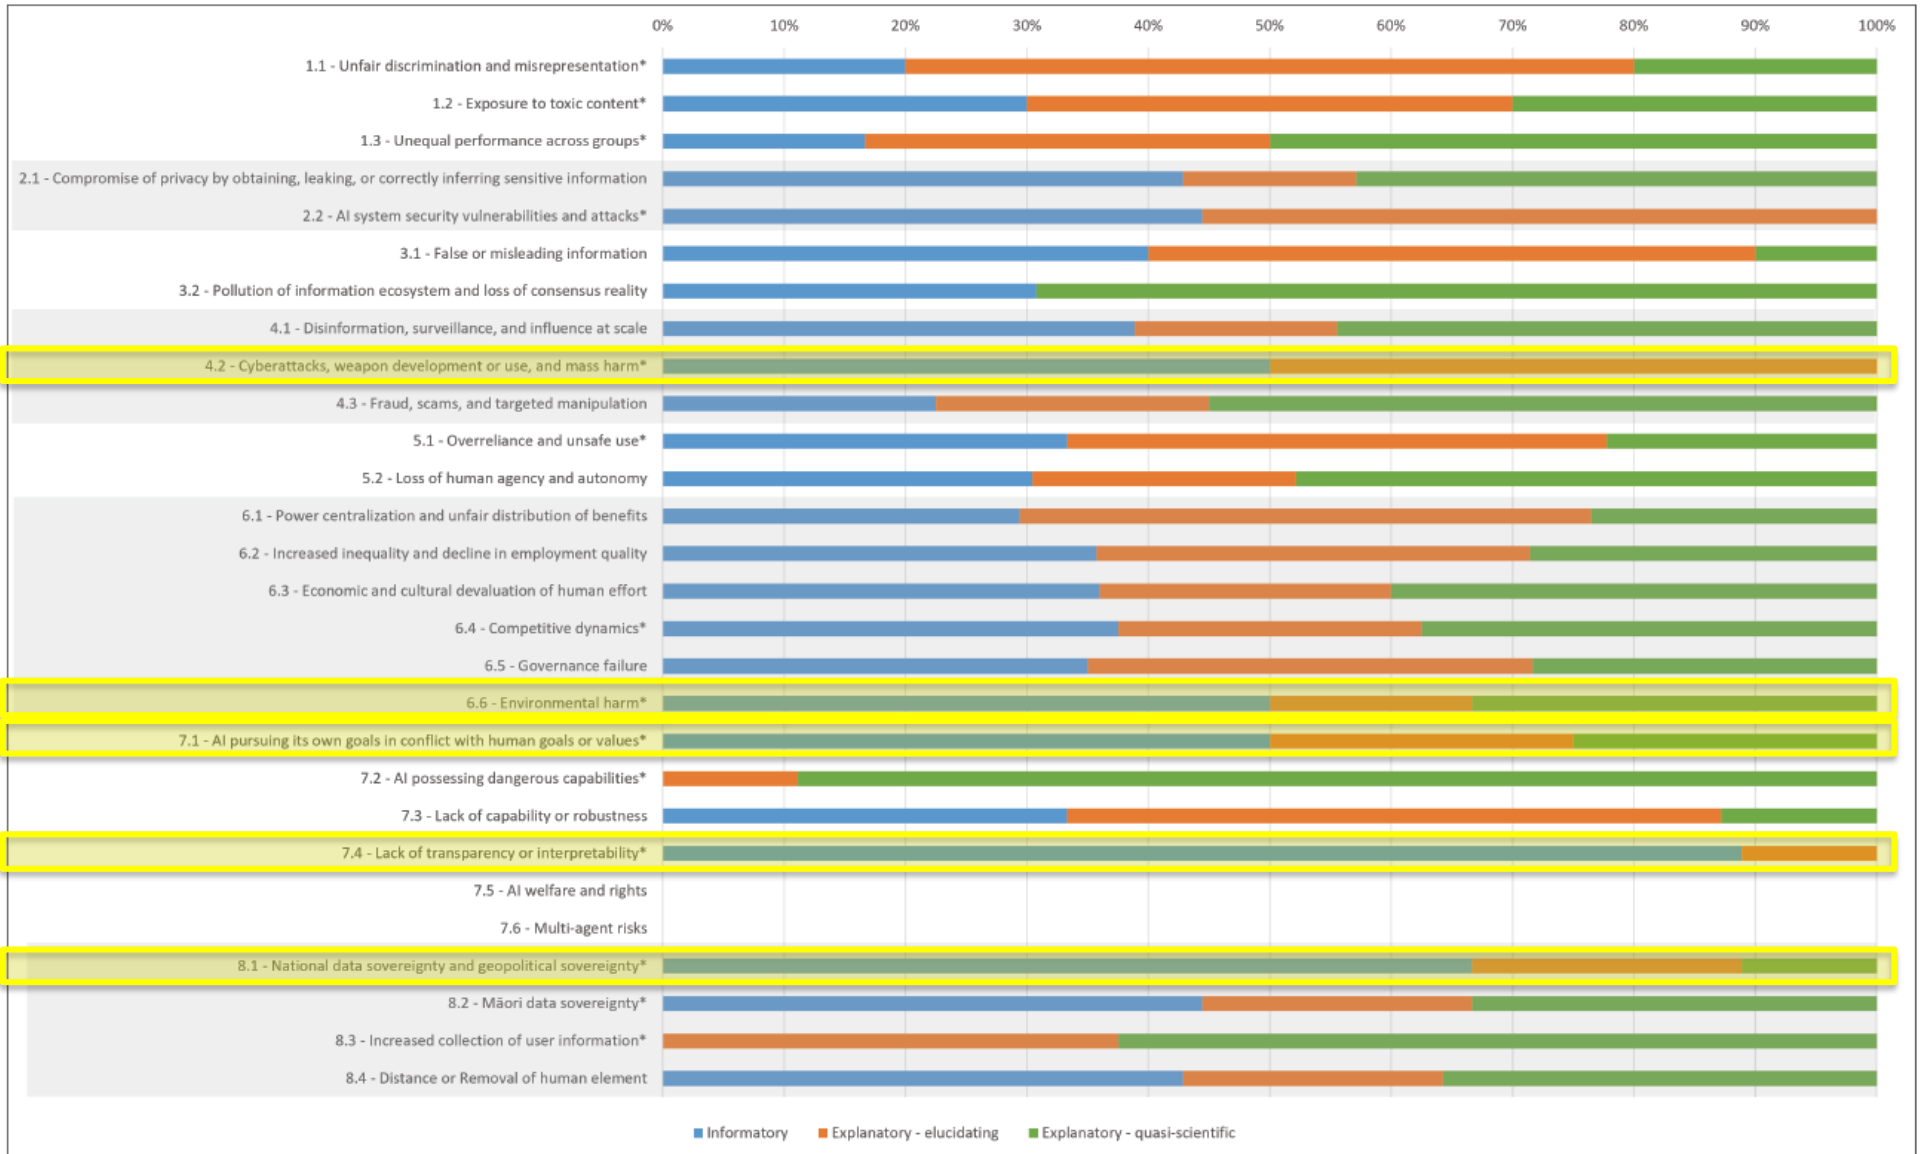
Trade-offs Were Unevenly Discussed

- Some benefits and risks appeared together frequently
- ‘Empowering users’ was often discussed alongside misuse
- Other pairings were largely absent from reporting
- For example: personalisation was rarely linked to manipulation risks
- This suggests missing public discussion around AI trade-offs

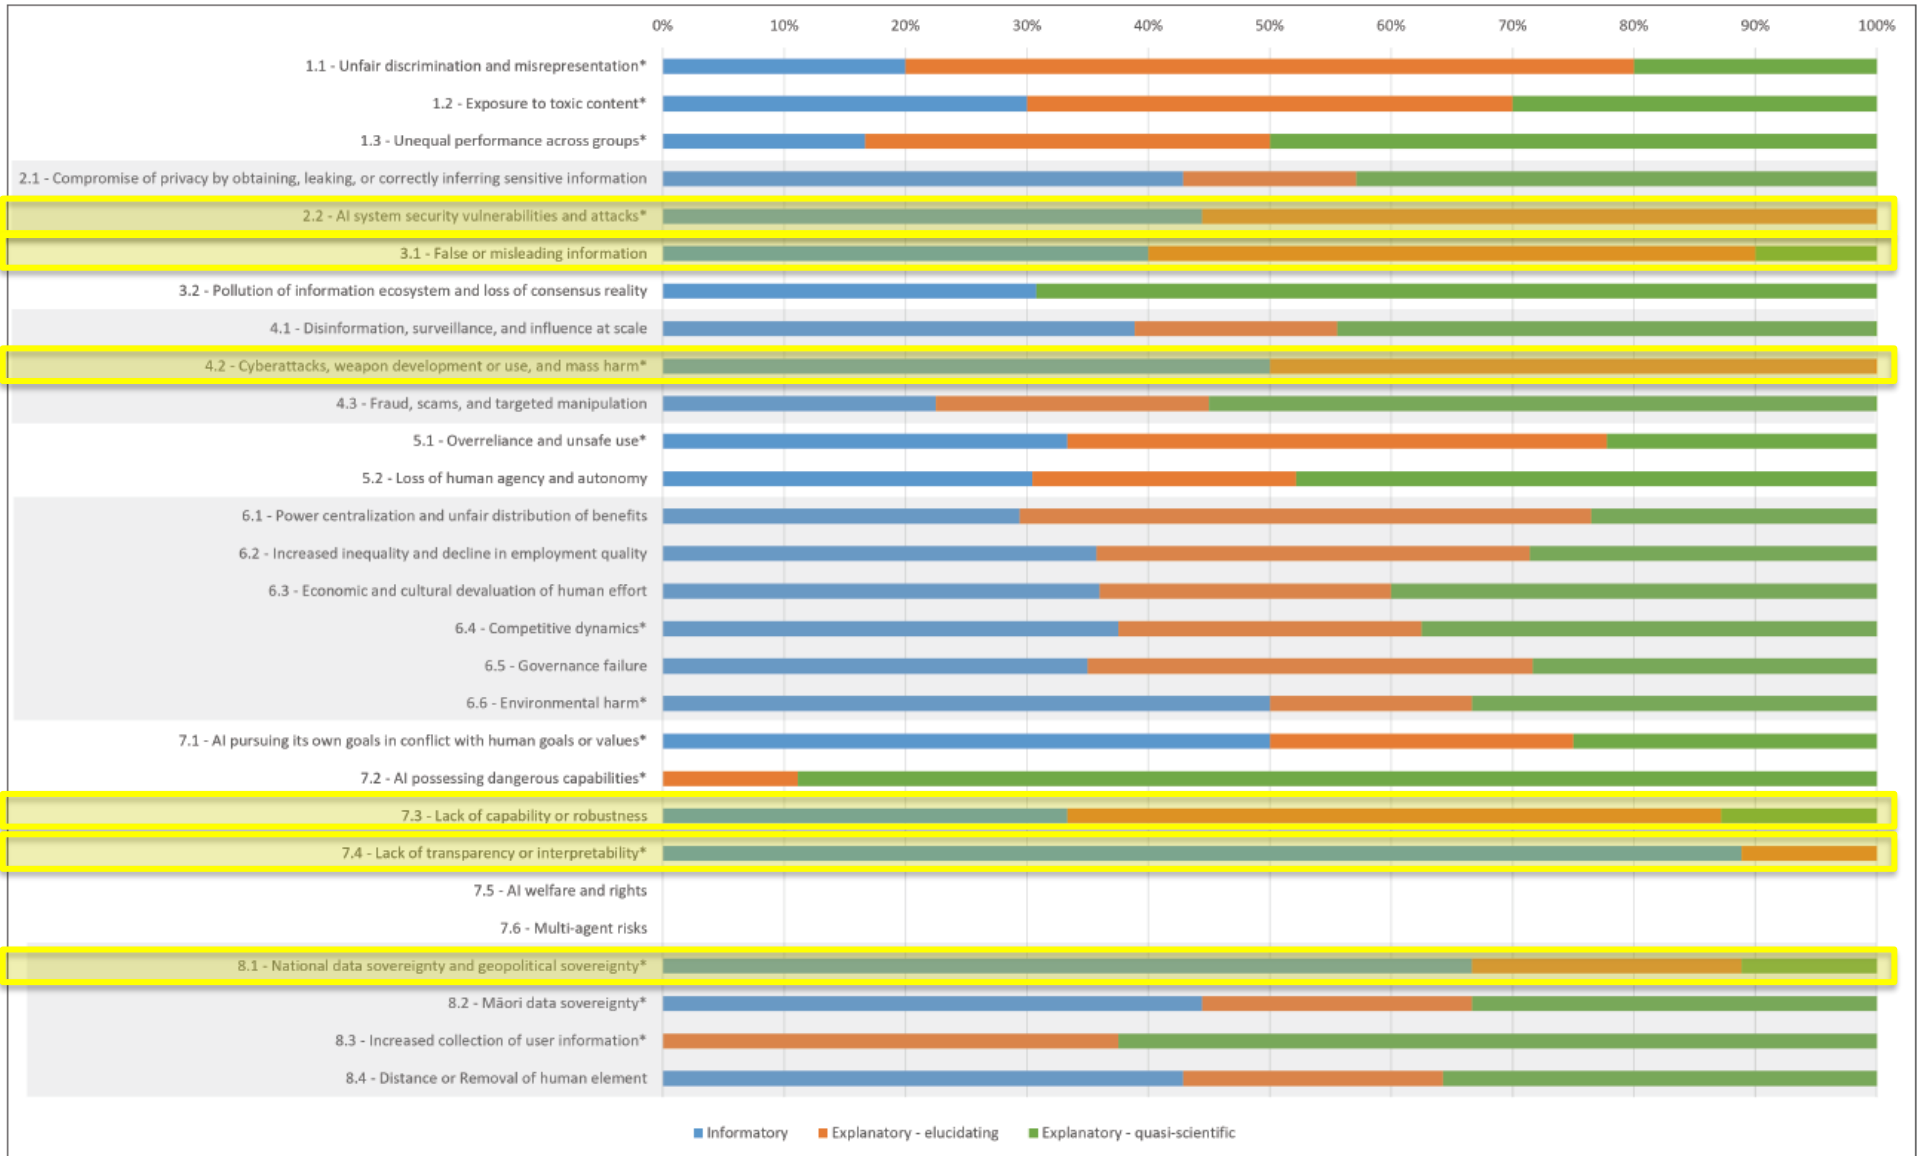
Terminology: Explanatory Depth

	Informatory	Explanatory—elucidating	Explanatory—quasi-scientific
Operationalisation in the present research	References a risk in a manner that presupposes audience understanding of the risk, without definition or explanation.	Introduces the risk by way of a definition or description of the risk.	Develops a reader's understanding of the risk by explaining the risk's causes (social or technical) and the harm caused if the risk eventuates.
Example: Risk 3.1—False or misleading information	'Users were trained and had to check whatever AI came up with.'	'"It's helped our payroll professionals to answer really complex payroll queries and it's accurate to the point it can pass the payroll exam to 93% accuracy." What about the other 7%? Well, this is the thing. As with any large language model, it doesn't necessarily complete a task at all times, but it will get you close enough," Compagnone said.'	'But LLMs are imperfect. While appearing to have knowledge, they can invent facts and details, as they are designed to create answers that are grammatically correct compilations of text, based on predicting what word should come next in a sentence.'

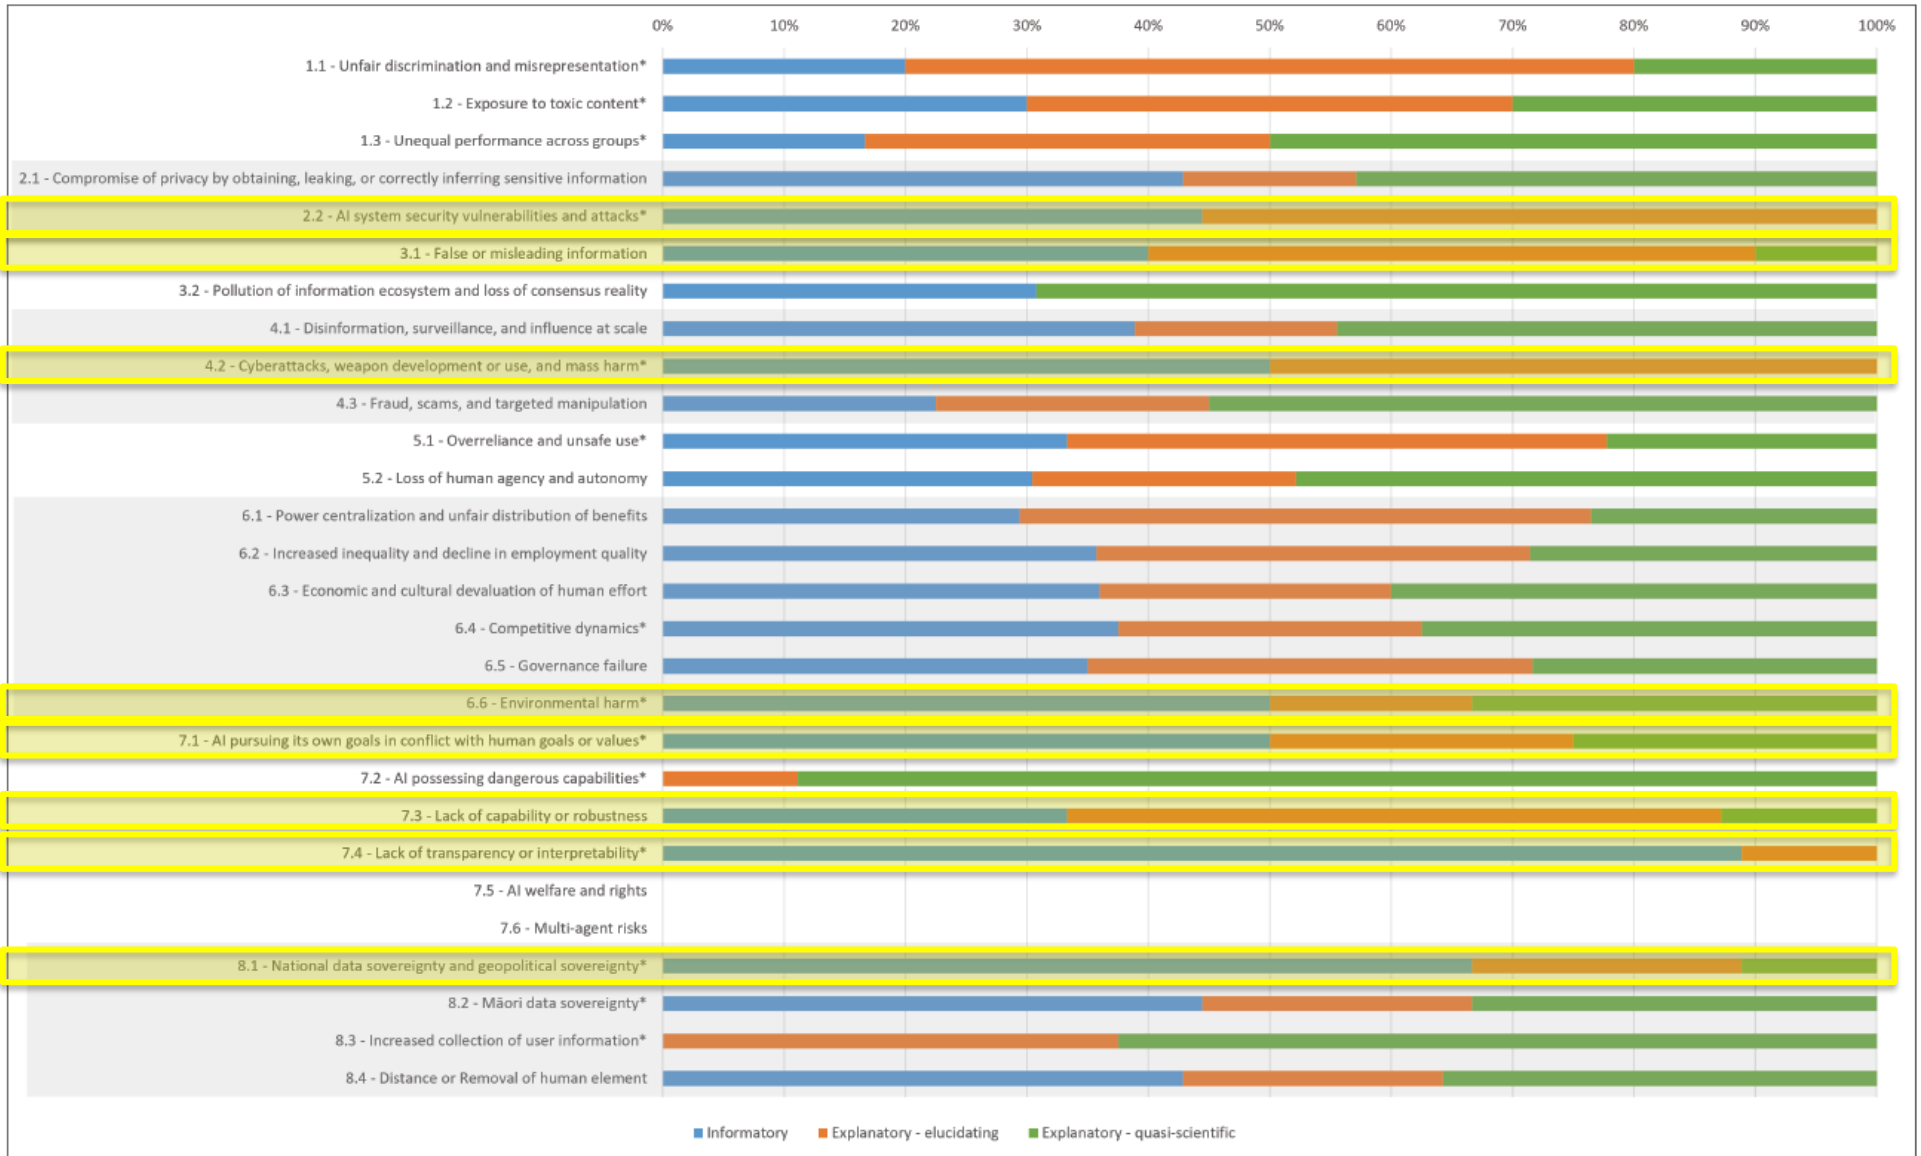
Depth of Risk Explanations: High proportion of 'informatory'.



Depth of Risk Explanations: low proportion of 'quasi-scientific'



Readers get very little explanation about these risks:



Summary of potential blind spots in coverage:

code	Not reported often	Not reported in detail
1 Discrimination & toxicity		
1.1 Unfair discrimination and misrepresentation.		
1.2 Exposure to toxic content.		
1.3 Unequal performance across groups.		
2 Privacy & security		
2.1 Compromise of privacy by obtaining, leaking, or correctly inferring sensitive information.		
2.2 AI system security vulnerabilities and attacks.		
3 Misinformation		
3.1 False or misleading information.		
3.2 Pollution of information ecosystems and loss of consensus reality.		
4 Malicious actors & misuse		
4.1 Disinformation, surveillance, and influence at scale.		
4.2 Cyberattacks, weapon development or use, and mass harm.		
4.3 Fraud, scams, and targeted manipulation.		
5 Human-computer interaction		
5.1 Overreliance and unsafe use.		
5.2 Loss of human agency and autonomy.		

code	Not reported often	Not reported in detail
6 Socioeconomic & environmental harms		
6.1 Power centralization and unfair distribution of benefits.		
6.2 Increased inequality and decline in employment quality.		
6.3 Economic and cultural devaluation of human effort.		
6.4 Competitive dynamics.		
6.5 Governance failure.		
6.6 Environmental harm.		
7 AI system safety, failures & limitations		
7.1 AI pursuing its own goals in conflict with human goals or values.		
7.2 AI possessing dangerous capabilities.		
7.3 Lack of capability or robustness.		
7.4 Lack of transparency or interpretability.		
7.5 AI welfare and rights.		
7.6 Multi-agent risks		
8 Risks not mentioned in MIT Domain Taxonomy		
8.1 National Data Sovereignty and geopolitical sov.		
8.2 Māori Data sovereignty		
8.3 Increased collection of user information		
8.4 Distance or removal of the human element		

Findings

Phase 1

More positive than
negative

Governance and job
loss mentioned often

Certain risks aren't
mentioned often/at all

Benefits often focus on
individual or orgs

Some trade-offs
referenced, others not

Explanation provided
about risks varies

Phase 2

Finding

Phase II: Journalist Interviews

- 8 journalists who reported on high-risk AI applications were interviewed
 - (defined by EU AI Act “high risk” or applications that had implications for national sovereignty or Te Tiriti o Waitangi).
- Semi-structured interviews explored reporting practices and barriers.
- Journalists mostly described AI reporting as mostly incidental to existing beats.
- Most did not consider AI reporting fundamentally different from their other reporting (i.e. they didn’t follow any distinct practices).
- Three major themes emerged from thematic analysis.

Theme 1: An Ideal of Simplicity

Journalists aim to condense and simplify technical content — keeping to big-picture issues over specific details

42.6% of risk references were informatory

Presupposes audience understanding — no definition or explanation given

28.7% elucidating (defined but not explained)

Readers told what a risk is, but not why or how it arises

Only 28.7% quasi-scientific

Explains the mechanism — but reserved only for story-critical content

Theme 2: Access Is Key



Academics: Enablers

Provide angles, context, and seeded new stories

Valued as reliable sources and story catalysts

Knowledge gaps in journalists filled by academic commentary

Some journalists relied on academics to 'surface' important risks



Non-Academic Sources: Barriers

Often less forthcoming — exercise choice about info shared

OIA requests needed to access government AI data

Private sector has NO transparency requirements

High scepticism from journalists — but scepticism didn't always translate to practice

Theme 3: Knowing 'Enough'

AI knowledge is simultaneously 'unnecessary' AND 'crucial'

'You don't need to understand AI in depth'

Journalists are generalists; they rely on expert sources for subject-matter knowledge. It's impossible to be expert in every beat. Accuracy of the published article is what matters.

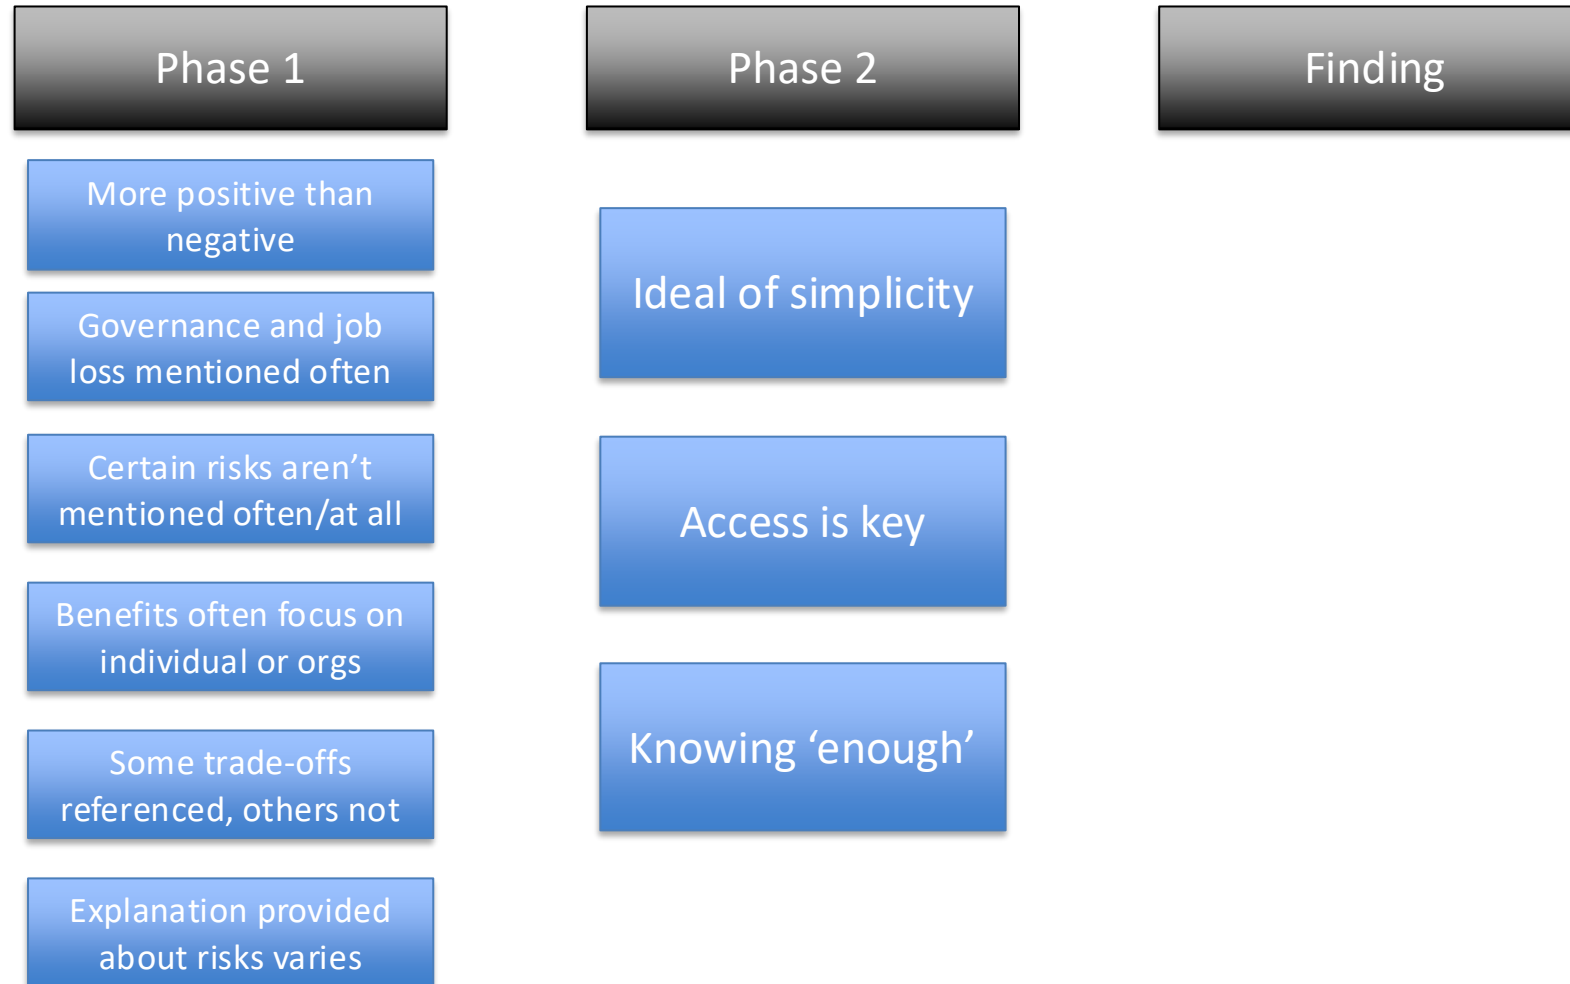
'Deeper knowledge changes the story'

Journalists who knew more about AI recognised more risks — bringing issues of public interest into the article that sources didn't raise. Knowledge enables 'asking the right questions.'

Note: this is how journalism works – not necessarily a flaw. But in specific circumstances, this can create challenges for the profession:

- Knowledge of AI risks is still emerging. Knowledge gaps in journalists and their sources shapes which risks entered public discourse.
- Industry sources can choose not to mention risks – potentially leading to one-sided coverage.

Findings



Phase 1

More positive than negative

Governance and job loss mentioned often

Certain risks aren't mentioned often/at all

Benefits often focus on individual or orgs

Some trade-offs referenced, others not

Explanation provided about risks varies

Phase 2

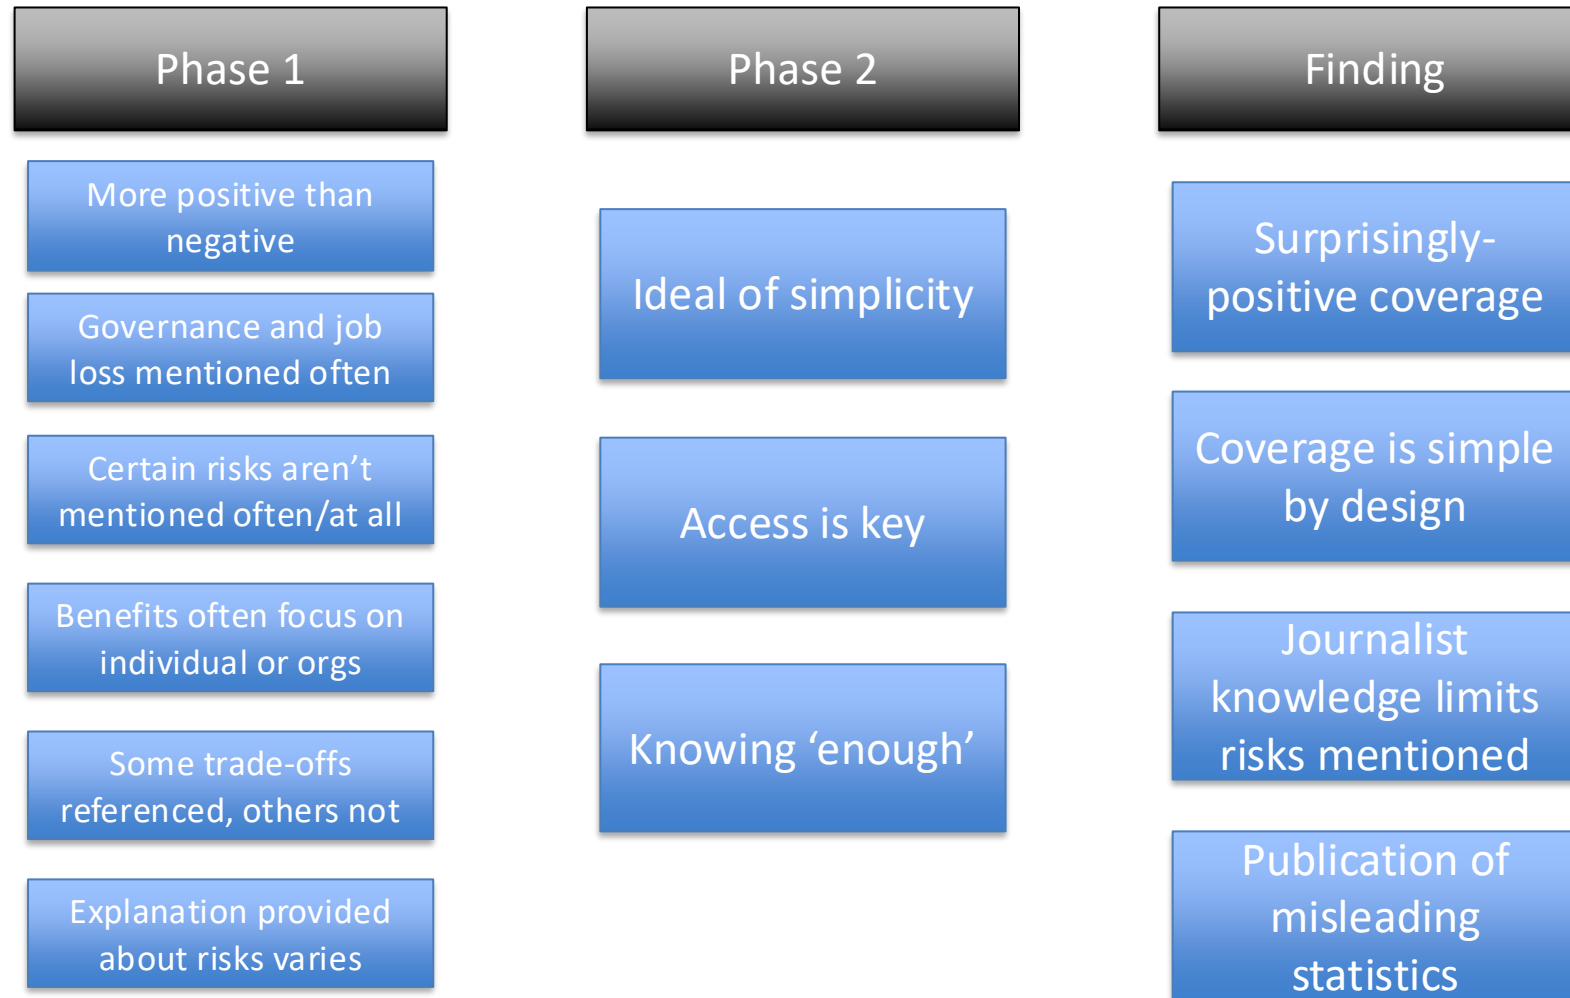
Ideal of simplicity

Access is key

Knowing 'enough'

Finding

Findings



Revisiting the Content: the Publication of Misleading Statistics.

This was, yeah, [I'm] not very proud of the story because this... later on there's another story, I think that's related to this, and that shows you that... Yeah, someone said to me, "Oh, look, that's a really small sample to test it on. So what would that have proved?" [P2]

- During interviews, most participants expressed **high distrust** in non-academic sources
- However, this scepticism **didn't always translate into practice**, potentially due to resource and time constraints
- I looked at 3 articles I discussed with participants:
 - P2 reported on accuracy of a government model, as judged by a third-party evaluator.
 - P4 presented an accuracy statistic that sounds good, but doesn't represent the industry standard for determining accuracy.
 - P7 reported the accuracy of a model. "And the [company said the] data said that it came out at 92%, I had no reason not to believe them."

Key challenges for AI coverage in NZ

CHALLENGE I No 'Go-To' Summary of AI for Journalists

Unlike climate journalism — where the IPCC provides a consensus reference document — AI journalism has no equivalent. The International AI Safety Report exists but wasn't mentioned by any of my 8 interviewees, and covers only GPAI. Without a reliable reference, journalists' own knowledge gaps become public discourse gaps.

CHALLENGE II Reporting Without Transparency in Tech

Without legislated transparency requirements, journalists cannot scrutinise claims from AI developers/deployers. Even sceptical journalists published potentially misleading statistics because there was no independent data to cross-reference. The EU's Digital Services Act demonstrates a workable model.

Conclusion

- AI discourse in New Zealand is active but uneven
- Coverage is often more positive and less technically detailed than expected
- Journalists face major knowledge and transparency barriers
- The quality of AI discourse may shape future democratic decision-making
- Public understanding depends on the ability to 'know enough to ask the questions'

Questions?

Thank you