

Morally Repugnant AI?

An Alternative Lens for Predicting The Future of Work

AI & Society Seminar

22 May 2026

**WILL
ROBOTS
TAKE
MY JOB?**





Drawing the Line: How Professionals Construct Ethical Boundaries in GenAI Use

Dr Rebecca Downes (presenting),

Dr Jocelyn Cranefield, &

Dr Mian Wu

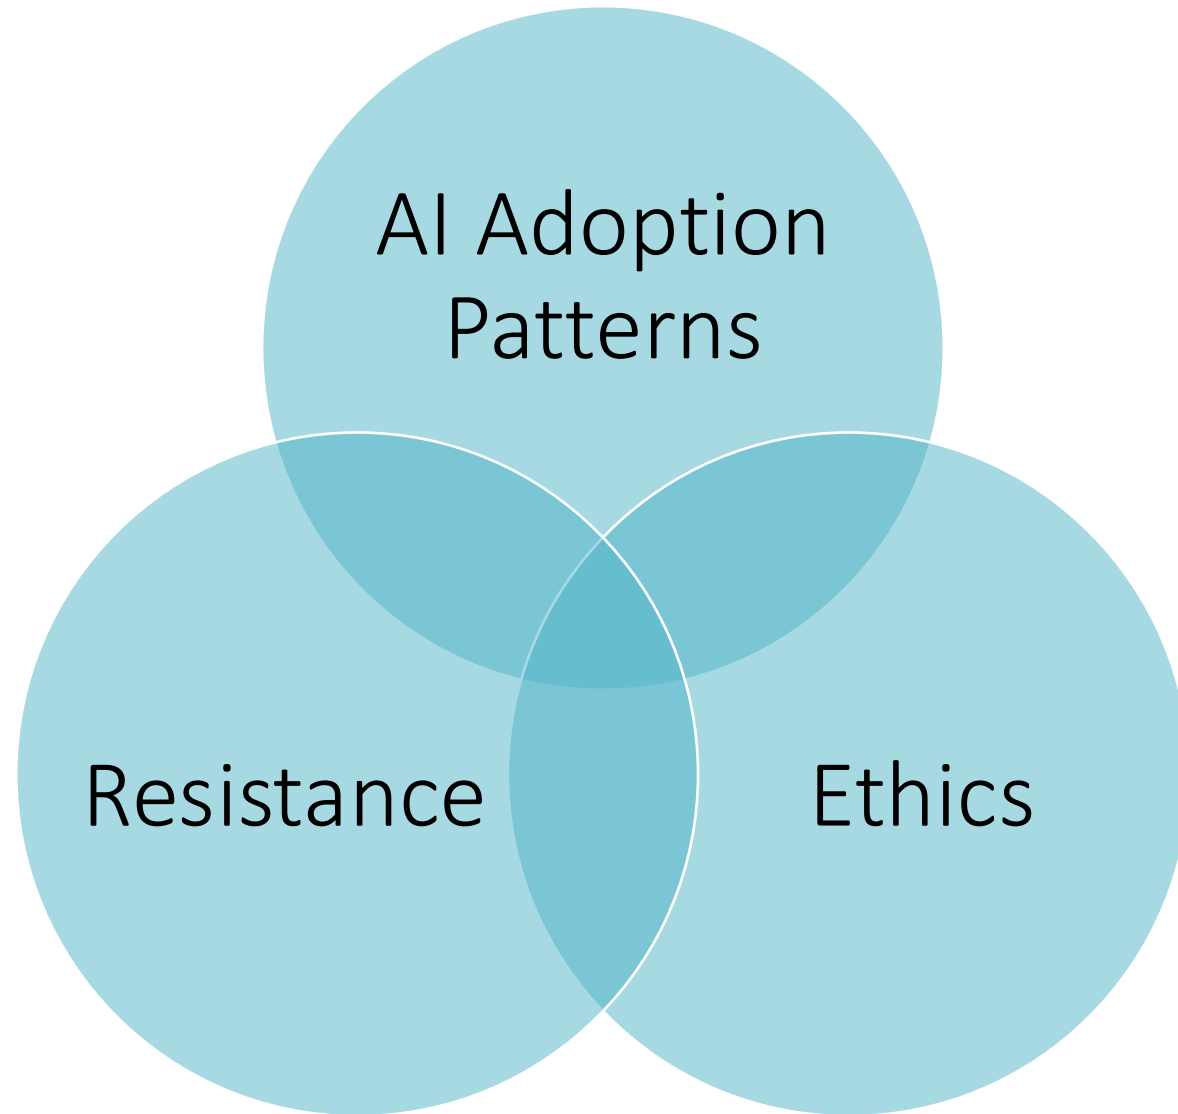


Working Paper 26-017

Performance or Principle: Resistance to Artificial Intelligence in the U.S. Labor Market

Simon Friis
James W. Riley

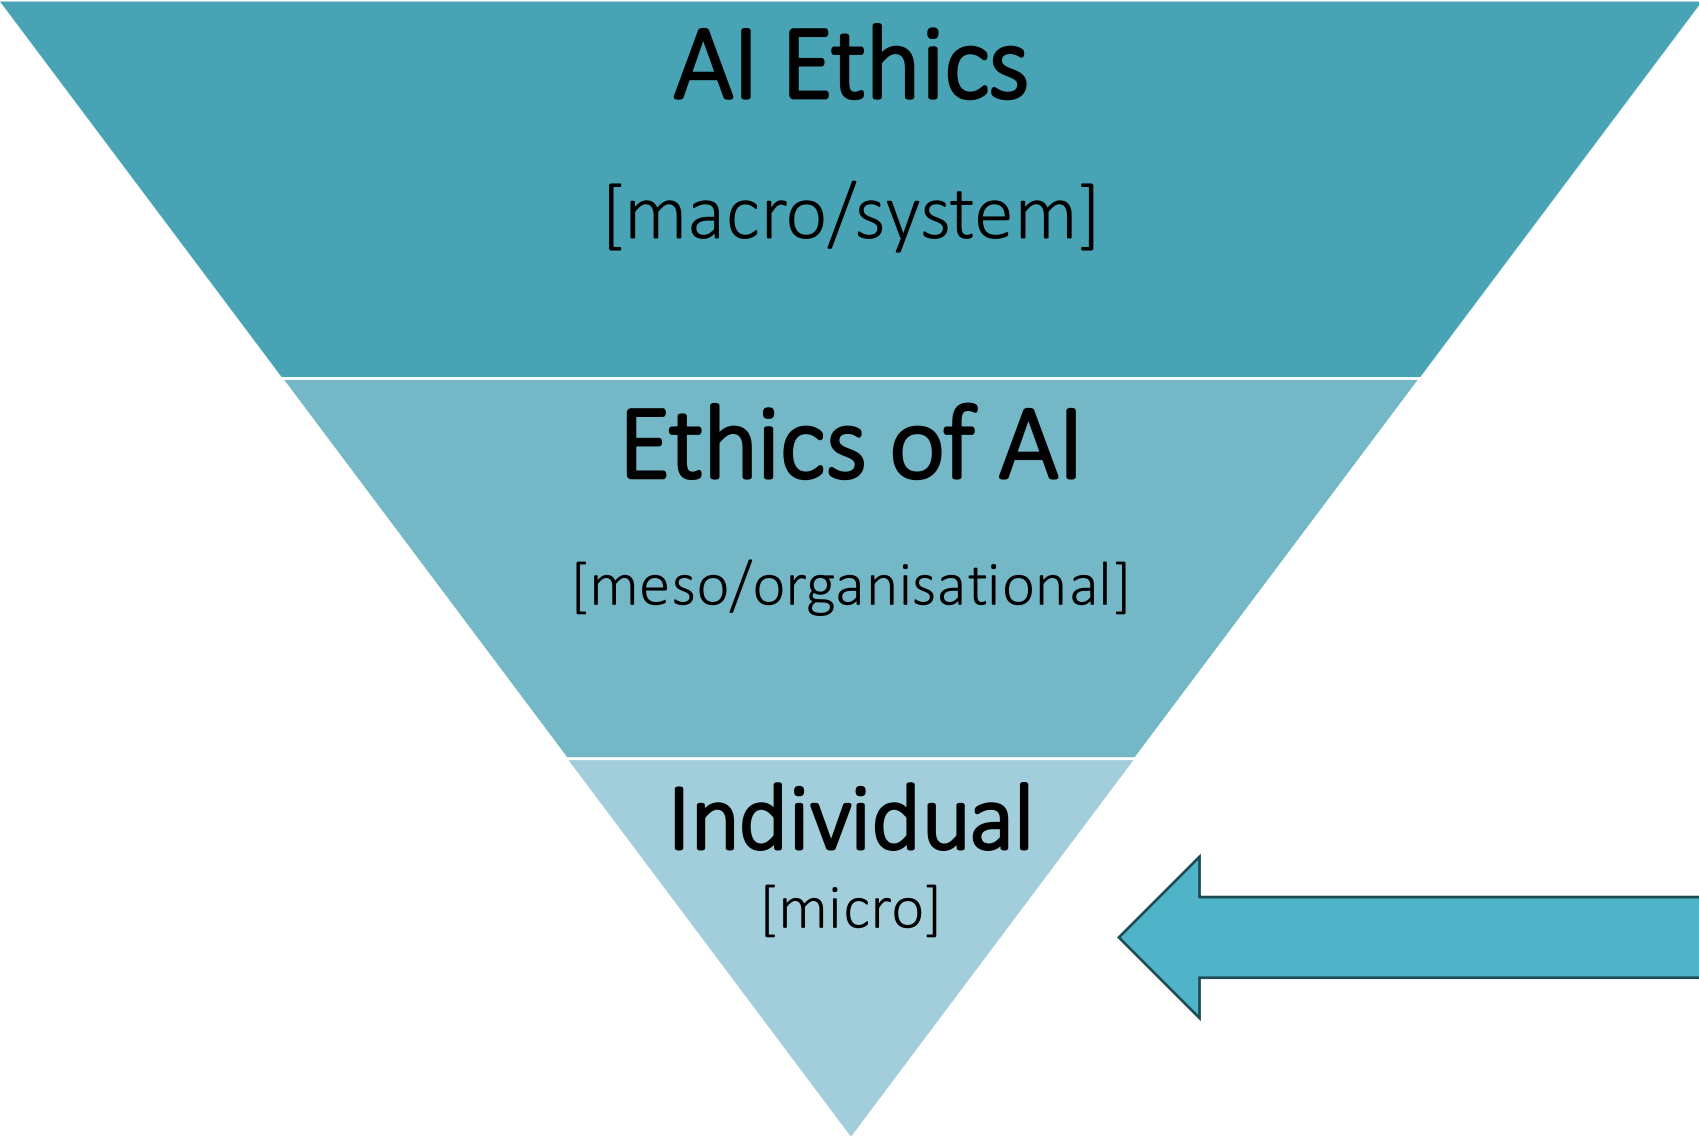
Harvard Business School: <https://www.hbs.edu/faculty/Pages/item.aspx?num=68008>



AI Adoption
Patterns

Resistance

Ethics



Drawing the Line: How Professionals Construct Ethical Boundaries in GenAI Use

Dr Rebecca Downes (presenting),

Dr Jocelyn Cranefield, &

Dr Mian Wu



Who did we ask?

- Academics and communications professional (e.g. marketing, comms, etc).
- What tasks they do/do not use GenAI for, and why

Initially noticed two broad categories

INSTRUMENTAL (CAN/CAN'T)

E.G.

Inaccurate outputs

Outdated training data

Lack of local context knowledge

Speed

Efficiency

Improved writing

NORMATIVE (SHOULD/SHOULDN'T)

E.G.

Inauthenticity

Dishonesty

Unfair

Reciprocity and respect

Integrity

Embrace innovation (positive)

Task-type Heuristics: ‘*A but not B*’

Task Type A (Acceptable Use)	Task Type B (Unacceptable use)	Participant Pseudonym (field)	Ethical principles invoked in rationalisations
Getting key points for a lecture	Writing a whole lecture	Leanne (academic)	“That wouldn’t be my contribution” (authenticity)
Creating infographics and graphical art	Generating pictures of products for sale	Isaac (marketing)	“I think trust is the kind of most important thing in in business and in transactions...” (trust, authenticity)
Using GenAI to develop skills in software	Using GenAI for writing	Jason (academic)	“If I ...send it to others, it just feels written by AI. It’s not like honest or real person” (honesty)
Extending an existing promotional image	Creating a whole show poster from just AI	Bella (marketing)	“...that is taking away the integrity of art itself” (integrity)

Overall finding

People evaluate generative AI delegation at a task level and on both ethical and practical dimensions.

Ethical evaluations are based on heuristics, which are better than rational models of ethical decision-making for GenAI delegation decisions.

This kind of decision-making is largely absent in dominant models of AI delegation and use.

Why Heuristics Makes Sense

GenAI is flexible – lots of potential decisions!

Many models of technology adoption/use assume rational ethical decision-making processes – heuristics may be a better fit

Heuristics are developed at a task level, crafting (and simplifying) a complicated relationship for how GenAI fits into work

So

- Some evidence that normative (should/shouldn't) concerns impact use decisions
- Evaluations are heuristic and task based
- Policies or social norms might override individual heuristics (or conflict with them)
- Suggests a mechanism by which usage patterns develop

Working Paper 26-017

Performance or Principle: Resistance to Artificial Intelligence in the U.S. Labor Market

Simon Friis
James W. Riley

Harvard Business School: <https://www.hbs.edu/faculty/Pages/item.aspx?num=68008>

Overview

- Resistance can arise from performance-based concerns, which fade as technology improves, or from principle-based objections, which persist regardless of capability.
- For most occupations, resistance is based on performance concerns
- Yet a narrow subset remains categorically off-limits because such automation is seen as morally repugnant.
- Authors argue that: These findings clarify the “moral economy of work” in which society shields certain roles not due to technical limits but to enduring beliefs about dignity, care, and meaning

(pg 1-2)

The concept of moral repugnance has a rich tradition in economics and sociology, helping to explain why some market transactions, such as organ sales, are widely considered taboo and thus remain off-limits despite economic incentives (pg 4)

Data:

- N = 2,357 U.S. adults
- Each participant evaluated a random subset of 10 occupations, resulting in a final analytical sample consists of 23,570 participant-occupation ratings.
- For each occupation, participants saw the official O*NET title and standardized description

- Questions began with "*Given current capabilities...*" and asked whether AI should be used to assist workers with core tasks (augmentation) or fully automate the core tasks done by workers (automation).
- They were then instructed to "*Imagine that AI has advanced to the point where it can outperform humans in core tasks in this job and does so at a much lower cost*" and rated their support under these advanced conditions.

(pg7)

Performance-based concerns dominate:

30.0% (95% ClopperPearson exact CI: [27.1%–33.0%]) of occupations receive majority support for automation under current AI capabilities.

This increases to **58.3%** [55.1%–61.5%] receive majority support in the advanced scenario. We can call these “morally unprotected” roles

But some roles (~12%) remain categorically off-limits regardless of AI’s capabilities (morally protected).

(pg 9)

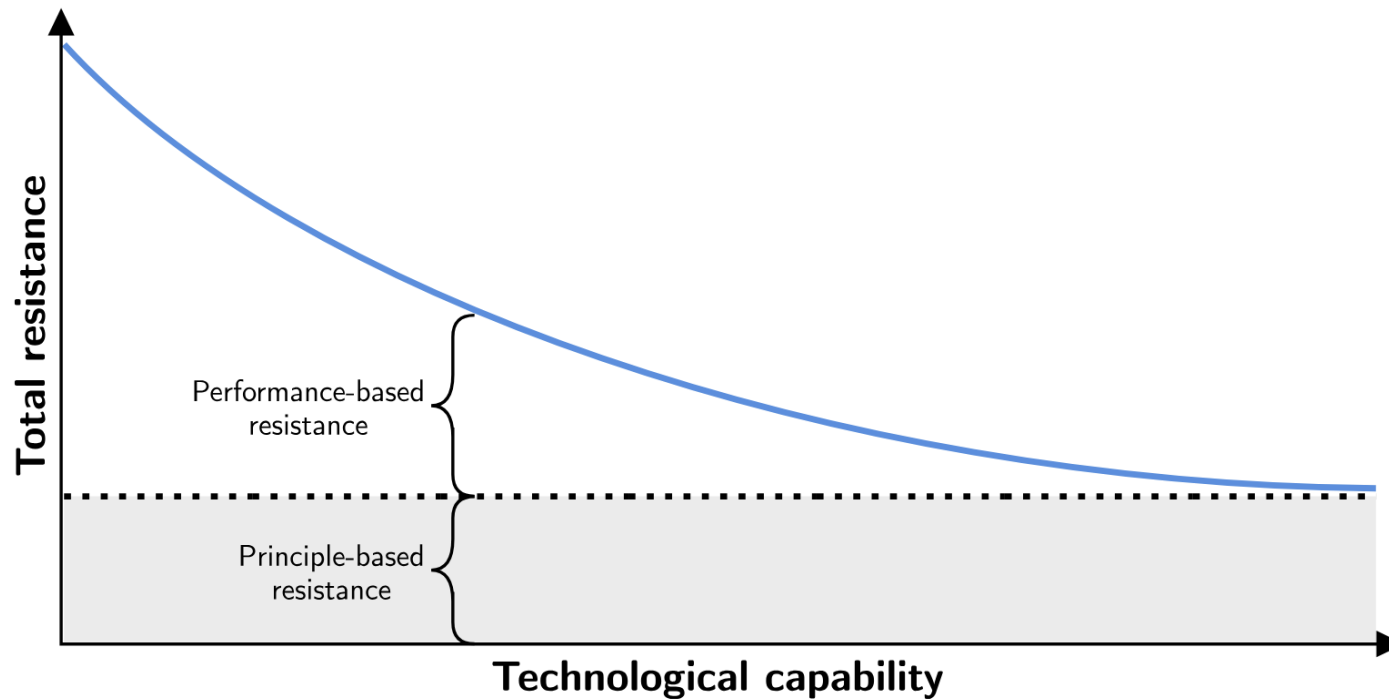


Figure 1. Performance-Principle Decomposition of Technological Resistance. Performance-based resistance decreases asymptotically with increasing technological capability, while principle-based resistance remains constant. The model illustrates how improved technological capabilities and lower costs can overcome initial opposition until reaching an irreducible moral floor that persists regardless of benefits. All else equal, occupations in which the use of AI is considered morally repugnant, the greater the principle-based resistance.

Moral Repugnance:

Participants answered a 7-item scale that tapped into the belief that using AI for a given occupation is inherently wrong, regardless of how well the AI performs.

Sample items included: “No matter how advanced AI becomes, this job should remain off-limits to machines,” and “I would feel betrayed if I found out AI was being used in this job.”

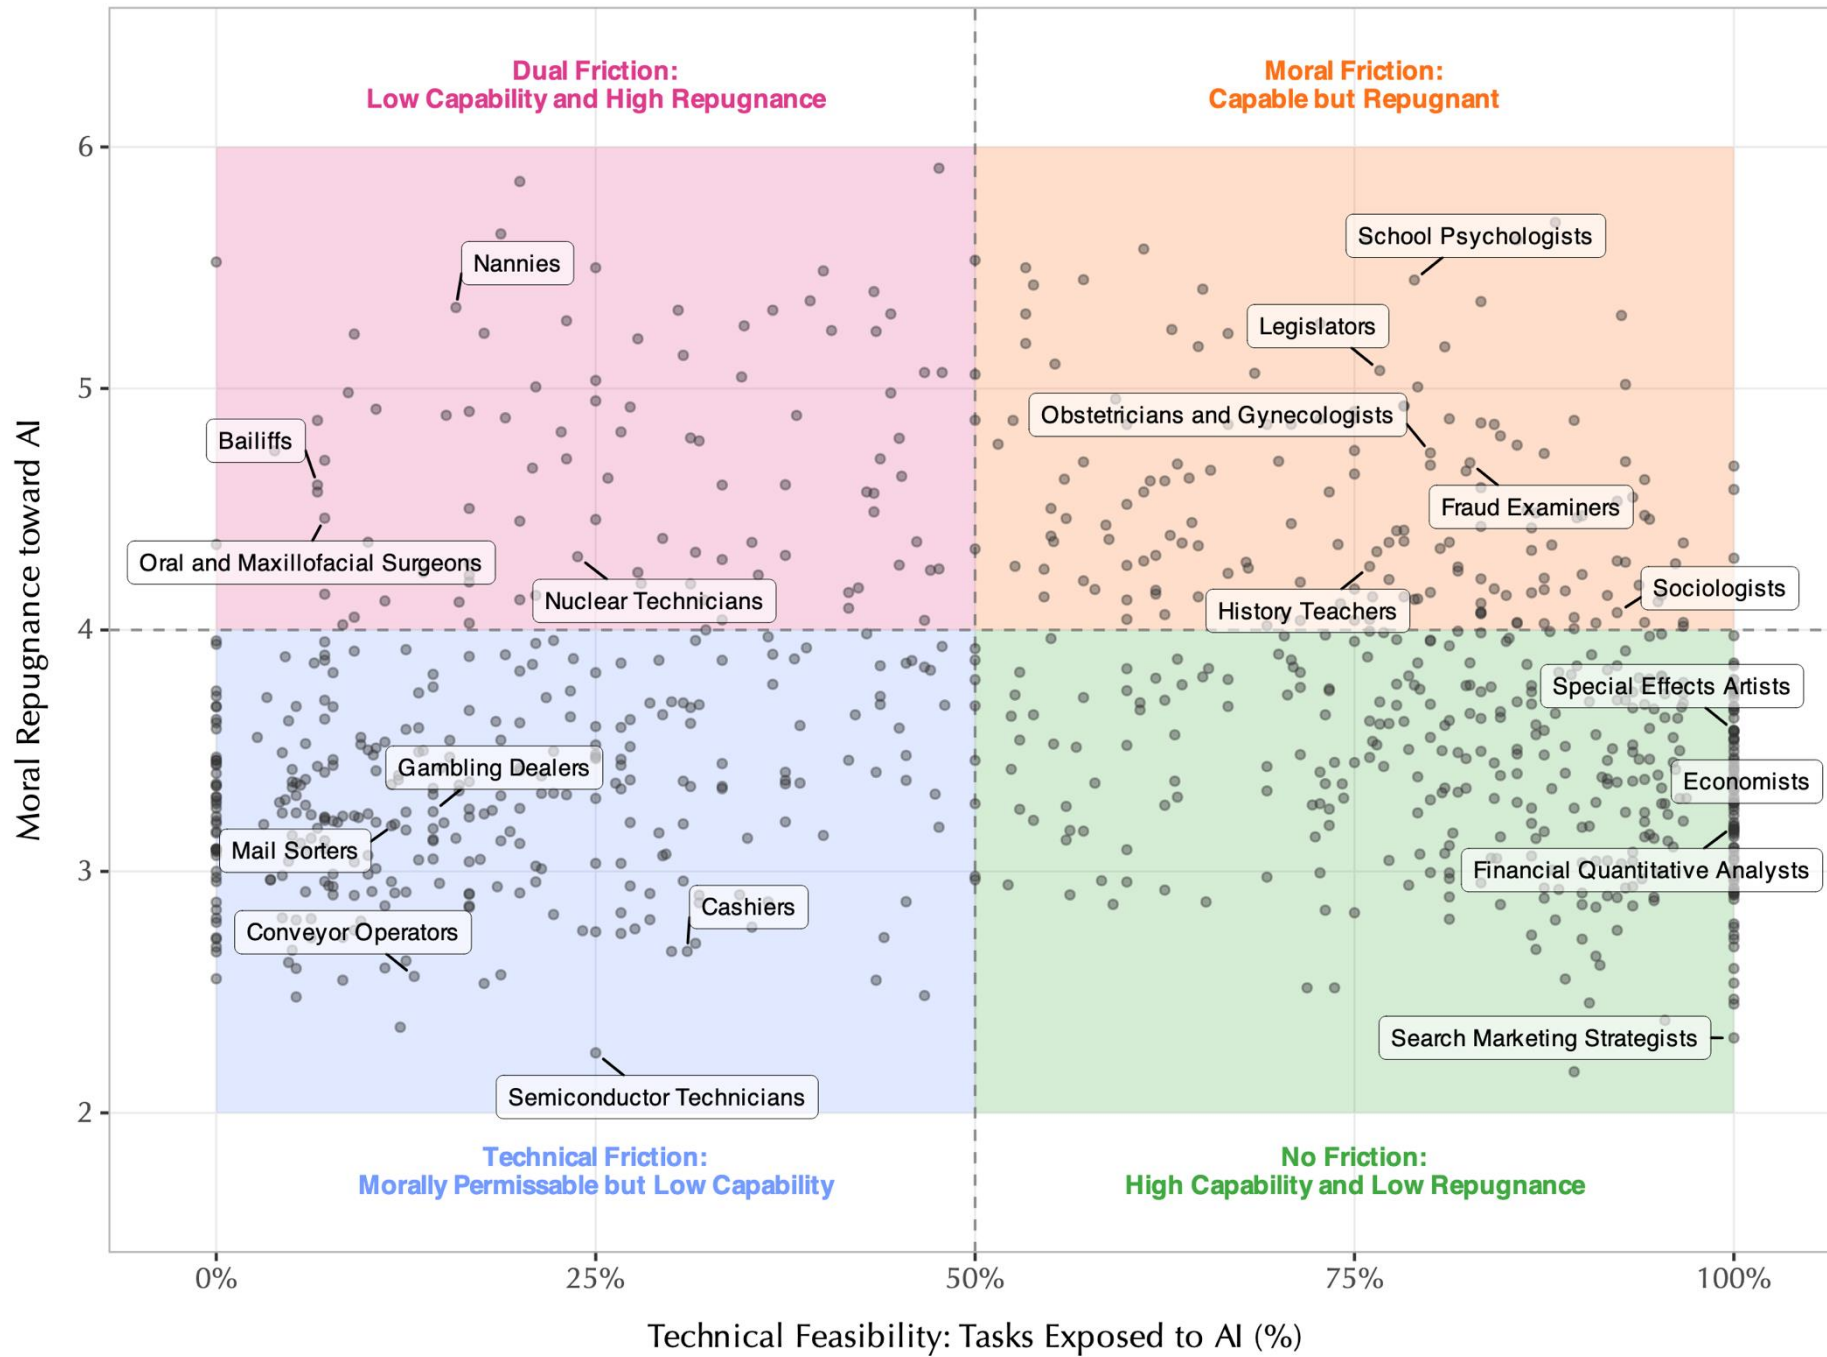
Responses were averaged to produce a single moral repugnance score (1–7).

At the occupation level, calculated a mean repugnance score across all respondents who rated that occupation.

- Morally **repugnant** = 95% bootstrapped confidence interval of its mean moral repugnance score lies entirely above the neutral midpoint (4.0);
- Morally **permissible** = interval lies entirely below 4.0;
- Morally **ambivalent** = overlaps 4.0.

Table 1. Selected occupations in the top/bottom 50 ranked by repugnance score. Occupations were selected to illustrate the range of occupations in these categories. Repugnance score indicates mean with 95% confidence intervals, calculated via stratified bootstrap with 10,000 replications.

Rank	Occupation	Repugnance score
<i>Most repugnant</i>		
1	Clergy	5.91 [5.44, 6.31]
2	Childcare Workers	5.86 [5.31, 6.26]
4	Marriage and Family Therapists	5.64 [5.12, 6.04]
6	Administrative Law Judges, Adjudicators, and Hearing Officers	5.62 [4.82, 6.10]
9	Athletes and Sports Competitors	5.52 [4.82, 6.01]
11	Craft Artists	5.50 [5.00, 5.95]
17	Police and Sheriff's Patrol Officers	5.40 [4.89, 5.81]
27	Funeral Attendants	5.28 [4.54, 5.89]
48	Barbers	5.03 [4.41, 5.52]
50	Actors	5.01 [4.43, 5.57]
<i>Least repugnant</i>		
904	Derrick Operators, Oil and Gas	2.72 [2.26, 3.29]
911	Cashiers	2.67 [2.10, 3.42]
913	Janitors and Cleaners, Except Maids and Housekeeping Cleaners	2.67 [2.15, 3.32]
924	Data Entry Keyers	2.55 [2.05, 3.25]
926	Segmental Pavers	2.55 [2.10, 3.05]
927	Biostatisticians	2.54 [2.14, 3.04]
930	Switchboard Operators, Including Answering Service	2.52 [2.02, 3.27]
936	Transportation Planners	2.38 [1.92, 2.93]
938	Search Marketing Strategists	2.31 [1.91, 2.89]
940	File Clerks	2.17 [1.69, 2.84]



So...

If we want to predict how AI might impact the future of work, thinking about morals might generate more useful predictions than thinking about technological capabilities alone.

I would argue that task level evaluations are important, though harder to measure at scale.

Lingering questions...

- Does this seem/feel right?
- Why do you think some roles are morally off-limits? Do you think they should be? Why does AI generate this type of resistance?
- What does the future of work look like, if this is true? Is it better or worse?

Thanks!

Dr Rebecca Downes

Senior Lecturer, School of Management

Te Herenga Waka - Victoria University of Wellington

rebecca.downes@vuw.ac.nz