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

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

Simon Friis
Harvard Business School

James W. Riley
Harvard Business School

Working Paper 26-017

Funding for this research was provided in part by Harvard Business School.

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

Simon Friis*
Laboratory for Innovation Science at Harvard
Harvard Business School

James Riley
blackbox Lab
Harvard Business School

Abstract: From genetically modified foods to autonomous vehicles, society often resists otherwise beneficial technologies. Resistance can arise from performance-based concerns, which fade as technology improves, or from principle-based objections, which persist regardless of capability. Using a large-scale U.S. survey quota-matched to census demographics and assessing 940 occupations (N=23,570 occupation ratings), we disentangle these sources in the context of artificial intelligence (AI). Despite cultural anxiety about artificial intelligence displacing human workers, we find that Americans show surprising willingness to cede most occupations to machines. Given current AI capabilities, the public already supports automating 30% of occupations. When AI is described as outperforming humans at lower cost, support for automation nearly doubles to 58% of occupations. Yet a narrow subset (12%)—including caregiving, therapy, and spiritual leadership—remains categorically off-limits because such automation is seen as morally repugnant. This shift reveals that for most occupations, resistance to AI is rooted in performance concerns that fade as AI capabilities improve, rather than principled objections about what work must remain human. Occupations facing public resistance to the use of AI tend to provide higher wages and disproportionately employ White and female workers. Thus, public resistance to AI risks reinforcing economic and racial inequality even as it partially mitigates gender inequality. 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. By distinguishing performance- from principle-based objections, we provide a framework for anticipating and navigating resistance to technology adoption across domains.

* Corresponding author: sfriis@hbs.edu. The authors contributed equally to this work; author names are listed in alphabetical order.

We thank Bill Kerr and Suzanne Smith for helpful feedback on earlier drafts. We are also grateful to Ezra Zuckerman Sivan, Ray Reagans, Michael Wahlen, Sameer Srivastava, Amir Goldberg, Chengwei Liu, Douglas Guilbeault, Tu Ni, and participants in the MIT Sloan Economic Sociology Working Group (ESWG), Microsoft Research Workshop on AI, Cognition, and the Economy (AICE), and the Nagymaros Conference for valuable comments and discussion. Matthew Nemesure and Samantha Richards provided thoughtful feedback on our public code repository. We are especially grateful to Bernal Cortés and Ashwin Kumar for their excellent research assistance. The authors acknowledge the use of AI tools to assist in the conceptualization of the research and manuscript preparation. This research was approved by the Harvard University-Area Committee on the Use of Human Subjects under protocol numbers IRB24-0098 and IRB24-0756.

INTRODUCTION

Will human labor remain relevant in a future where advanced artificial intelligence (AI) outperforms humans and does so at lower cost? This question has gained urgency as AI systems increasingly outperform humans across cognitive, creative, and emotional tasks (Boussioux et al. 2024; Tu et al. 2024; Manning et al. 2024; Dell'Acqua et al. 2023), once thought to be the exclusive domain of humans. Recent estimates suggest that current AI systems are capable of doubling the efficiency of 47% of all tasks in the U.S. economy (Eloundou et al. 2024), with this proportion expected to grow as AI capabilities expand. The economic logic seems inexorable: as AI becomes more capable and cost-effective than human labor, market forces should drive widespread automation and augmentation of work.

Yet history tells a different story. The adoption of new technologies has never been dictated by technical feasibility and economic incentives alone—social acceptance has imposed surprisingly enduring constraints on otherwise beneficial technologies. Public resistance to genetically modified crops (Frewer et al. 2013), nuclear power (Bauer 1995), and embryonic stem cell research (*National Bioethics Advisory Commission* 1999) persists in many contexts, despite their potential benefits. If certain domains remain off-limits to technological encroachment even when innovation offers clear advantages, could AI face similar resistance in the labor market—not because it is incapable, but because society deems its use fundamentally inappropriate in certain roles?

This tension between what a technology *can* do and what people believe it *should* be used for poses a central challenge to understanding the long-term impact of AI. On one side, recent studies identify which occupations are "exposed" to AI by analyzing tasks that AI can potentially automate or augment (Eloundou et al. 2024; Brynjolfsson et al. 2018; Frey and Osborne 2017;

Felten et al. 2021). While these efforts establish where AI is technically capable of supplanting human labor, they provide little guidance on the question of whether it is socially feasible to do so.

On the other side, research on algorithmic aversion finds that people often reject AI even when it outperforms human decision-makers (Glikson and Woolley 2020; Freitas et al. 2023; Dietvorst et al. 2015; Longoni et al. 2019). However, many of the mechanisms proposed to explain algorithmic aversion—such as distrust in AI's accuracy, concerns about bias, or unease with opaque "black box" models—will diminish as AI technologies improve over time. Thus, findings that AI is resisted often conflate performance-based objections (e.g., skepticism about AI's capabilities) with deeper, more principled objections that could persist even under "perfect AI." Moreover, while supporting the idea that social acceptance shapes technology adoption, these studies also tend to focus on only one or a few specialized contexts, such as financial or medical advice (Longoni et al. 2019; Bansak and Paulson 2024; Castelo et al. 2019), leaving open the broader question of how prevalent such resistance might be across the hundreds of jobs comprising the labor market.

Decomposing Resistance to AI

To address these gaps, we propose and empirically test a performance–principle decomposition of technological resistance (Fig. 1). Our approach decomposes resistance to AI into two components. The first, performance-based resistance, arises from perceived shortcomings of AI—encompassing concerns about accuracy, reliability, cost, speed, and other technical limitations. We treat these dimensions holistically because they represent interconnected barriers that erode together through technological progress: as AI systems improve, they typically become both more capable and more cost-effective through scale, competition, and innovation.

This form of resistance should recede as technology advances. The second, principle-based resistance, persists regardless of technological progress. It reflects fundamental convictions that certain roles or tasks ought to remain human, be they moral, cultural, or identity-driven. While prior research has hinted that some individuals find the very notion of automated decision-making morally objectionable, no study has systematically disentangled the performance-responsive and enduring principle-based aspects of AI resistance, let alone across a broad swath of occupations.

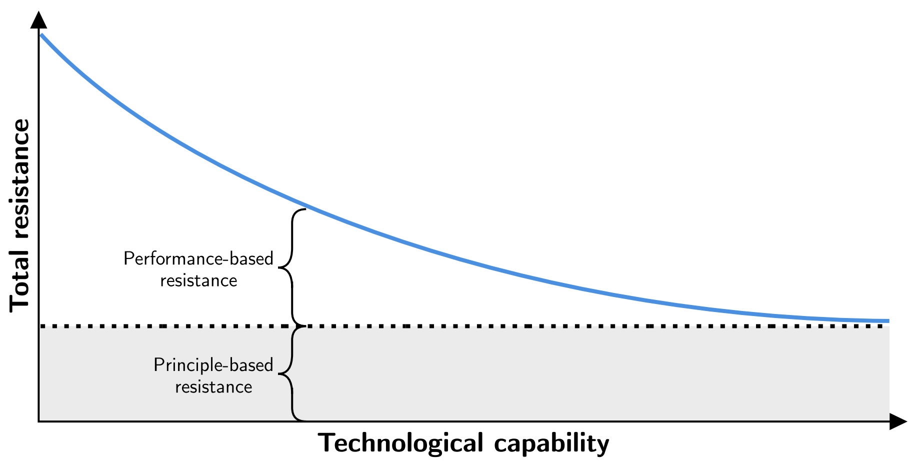

Figure 1. Performance-Principle Decomposition of Technological Resistance. Performance-based resistance decreases asymptotically with increasing technological capability, while principle-based resistance remains constant. The model illustrates how improved technological capabilities and lower costs can 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.

Within the category of principle-based resistance, we focus specifically on moral repugnance toward AI, which we define as the perception that using AI in a particular context is *inherently wrong* irrespective of potential benefits. 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

(Roth 2007; Zelizer 2005; Fiske and Tetlock 1997; Jackson 2023). We extend this line of thinking to technology adoption, positing that certain occupations may trigger a moral line that AI is not permitted to cross, even if it can perform the tasks flawlessly at lower cost.

Using a large-scale survey spanning 940 U.S. occupations, we provide the first systematic evidence of public resistance to AI across the labor market—and the first to decompose that resistance into performance-based and principle-based components. In doing so, we offer a generalizable framework for identifying when resistance may recede with technological improvement and when it reflects deeper moral constraints. This approach clarifies a longstanding conflation across disciplines, ranging from economics to psychology to science and technology studies, about the origins and structure of resistance to automation.

The data provide clarity where few systematic priors exist. Public support for AI-driven automation nearly doubles—from 30% to 58% of occupations—when AI is described as clearly outperforming human workers, suggesting that most resistance is contingent on perceived capability. Yet even under idealized performance conditions, the use of AI is deemed morally repugnant in 12% of occupations, and public support for automation is absent. These roles—centered on care, emotional labor, and moral or spiritual authority—form a bounded domain of principled resistance, where AI is seen as categorically inappropriate. Rather than supporting a generalized moral rejection of automation or a purely performance-driven model of adoption, our findings reveal a sharply delimited moral frontier, where a small subset of sacrosanct occupations remains off-limits, within an otherwise permissive labor market increasingly open to AI as performance improves.

METHODS

Survey Design and Sample

We conducted a survey of 2,357 U.S. adults designed to measure public support for AI automation and augmentation across a comprehensive set of occupations. Participants were recruited via Prolific, a nonprobability internet panel. We applied demographic quotas to match the U.S. adult population on age, gender, ethnicity, and political affiliation. Full details regarding the sampling procedure, inclusion criteria, and participant demographics are provided in the Supplementary Materials.

We analyze resistance at the occupation level because moral boundaries form around socially recognized roles, not abstract tasks. Our approach deliberately deviates from the task-based model of occupational change prominent in labor economics (Autor et al. 2003; Acemoglu and Autor 2011), which sees occupations as bundles of tasks that are differentially automated or augmented. While this framework correctly recognizes that technologies affect occupations through specific tasks, moral acceptance operates differently: the same abstract task, such as providing advice, may be acceptable when performed by AI financial advisors but unacceptable from AI spiritual counselors.

Occupations were drawn from O*NET, a comprehensive database of occupations in the U.S. labor market (*National Center for O*NET Development*). While O*NET provides nearly 20,000 occupation-specific task statements that could capture both task content and occupational context, analyzing public attitudes at this granular level would be intractable, requiring either extensive respondent batteries or prohibitive sample sizes. Given our primary goal of mapping resistance across the entire U.S. labor market, we focus on O*NET occupation taxonomy. This approach achieves comprehensive labor market coverage while maintaining survey tractability.

Future work could examine occupation-specific tasks within targeted subsets of morally contested occupations, providing granular validation of the broader patterns we identify.

We included all O*NET occupations except 76 residual categories (e.g., "Managers, All Other") that were too broad for meaningful evaluation, yielding 940 distinct occupations for evaluation. Each participant evaluated a random subset of 10 occupations, resulting in a final analytical sample consists of 23,570 participant-occupation ratings.

Support for Automation and Augmentation: Current vs. Advanced AI

For each occupation, participants saw the official O*NET title and standardized description. We employed a within-subject design where participants rated each occupation twice. First, participants rated their support for AI automation and augmentation under current AI capabilities. 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. These questions began with "If AI outperforms humans..." and asked whether AI should be used to significantly assist workers (augmentation) or fully automate core tasks (automation). To provide participants with a consistent frame of reference, we defined core tasks for them as "the essential responsibilities that define a job, even if they do not take up the most time in a typical workday."

This design was constructed to "turn off" performance-based concerns—including both capability and cost limitations that predictably improve with technological progress—to isolate principle-based resistance. By having the same individual evaluate the same occupation under both conditions, we can observe how support changes when performance constraints are removed. The shift in support captures performance-based resistance, while resistance that persists despite stipulated AI superiority indicates principle-based objections, suggesting a moral floor that technological advancement may not overcome. To the extent that participants with strong priors about AI's limitations (e.g., inherent lack of empathy) do not fully update when asked to imagine superior AI, our estimates of performance-based resistance are conservative. Such skepticism would cause us to understate the role of performance concerns and overstate principled objections, making our finding that most resistance is performance-based particularly robust.

All responses used a 7-point Likert scale (1 = Strongly Disagree, 4 = Neither Agree nor Disagree, 7 = Strongly Agree). We classify participants as supporting AI when their rating equals or exceeds the neutral midpoint (≥4). This coding decision reflects our interest in identifying social barriers to adoption: scores below the midpoint indicate active opposition that could generate resistance to AI implementation, while the midpoint and above represent passive acceptance through enthusiastic support. Ambivalent respondents (those selecting 4) may not champion AI adoption, but they are unlikely to organize against it or avoid services that employ it. At the occupation level, we measure support as the percentage of participants meeting this threshold for automation or augmentation in that role.

Measuring Moral Repugnance

To assess whether resistance to AI extended beyond concerns about performance, we developed and validated a measure of moral repugnance (see SM2 for scale construction and validation procedures). While still being instructed to consider a scenario in which AI had advanced to the point where it outperforms humans at core tasks in the job and does so at lower cost, 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, we calculated the mean repugnance score across all respondents who rated that occupation. We categorized the use of AI in an occupation as *morally repugnant* if the 95% bootstrapped confidence interval of its mean moral repugnance score lies entirely above the neutral midpoint (4.0); as *morally permissible* if the interval lies entirely below 4.0; and as *morally ambivalent* if it overlaps 4.0. Details of the bootstrapping and classification procedure are provided in SM3.

RESULTS

As AI Improves, Resistance Fades: More Than Half of Occupations Receive Majority Support for Automation Under Advanced AI

A central question is how public support for AI automation changes under advanced AI. Figure 2 contrasts the distribution of support for automation and augmentation across the 940 occupations under current capabilities versus advanced capabilities. We find that 30.0% (95% Clopper–Pearson 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. This shift suggests that a substantial portion of existing resistance is performance-based: respondents are willing to reconsider their opposition once AI is described as unequivocally superior to human workers in core tasks. However, technological advancement alone is insufficient to garner majority support for automation in the remaining 41.7% of occupations.

Augmentation garners much higher acceptance than full automation. Even under current AI, respondents endorse AI as a tool to assist humans in 94.4% [92.7%–95.7%] of occupations, rising marginally to 95.6% [94.1%–96.9%] under advanced AI. This highlights a gap between approval for "AI helping humans" and "AI completely replacing humans," reinforcing that reservations about AI often center on displacement rather than collaboration.

Morally (Un)Protected Occupations

Public resistance to AI is not uniform across occupations. While some roles see widespread support for automation, others remain categorically off-limits, regardless of AI's capabilities. To better understand this variation, we measured moral repugnance toward AI across all 940 occupations, finding substantial variation in moral repugnance across occupations (Fig. 3). To systematically analyze these patterns, we classified occupations into three categories reflecting the average stance toward AI use: morally permissible (36.8%, n = 346), morally ambivalent (51.3%, n = 482), and morally repugnant (11.9%, n = 112), based on whether their moral repugnance scores fell significantly above, below, or around the neutral midpoint of 4.0 (see Table 1 for examples of occupations in which the use of AI is considered morally permissible versus morally repugnant). These classifications allow us to investigate how moral repugnance moderates both baseline support for AI and responsiveness to technological improvements.

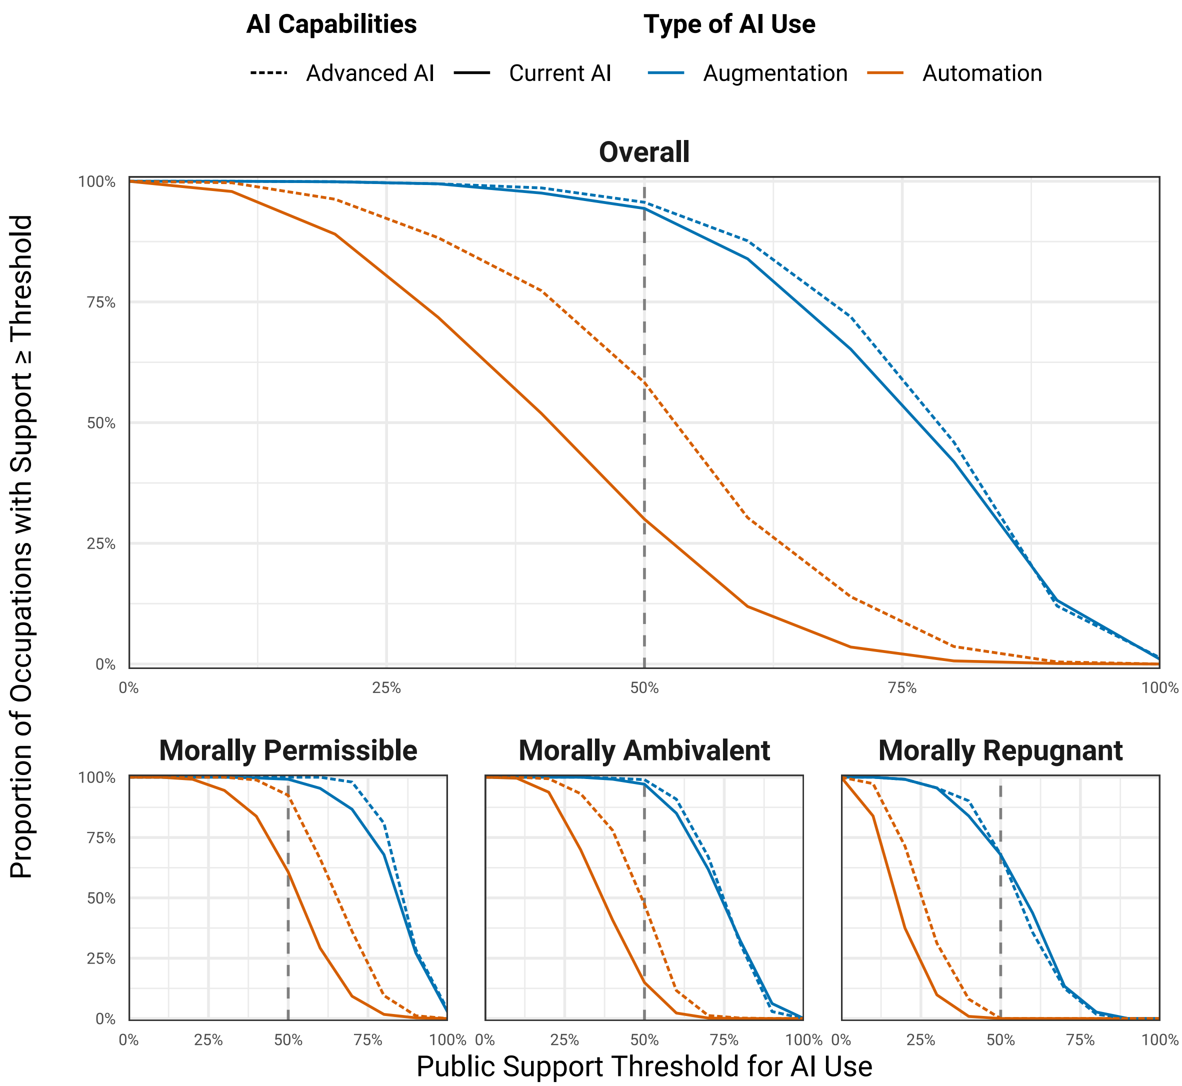

Figure 2. Public Support for AI Augmentation and Automation Under Current vs. Advanced AI. The complementary cumulative distribution functions (CCDFs) illustrate the proportion of occupations where at least a given percentage of the public supports AI augmentation or automation. The top panel shows the distribution for all occupations. The bottom panels show the distribution for occupations in which the use of AI is considered morally permissible, ambivalent, and repugnant.

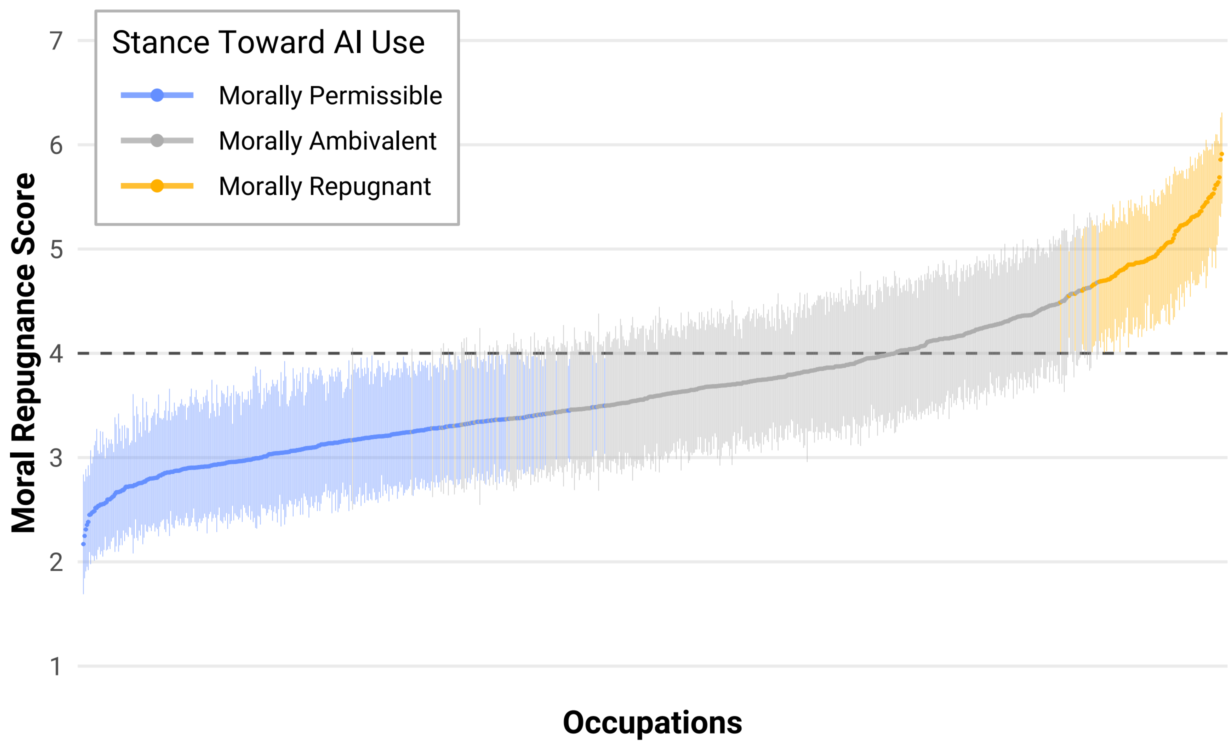

Figure 3. Distribution of Moral Repugnance Scores Across Occupations. Occupations were rated on a scale from 1 (lowest moral repugnance) to 7 (highest). Moral repugnance scores reflect judgments about whether it is inherently wrong to use AI in a given occupation, regardless of performance or cost. 95% confidence intervals were estimated via stratified bootstrapping with 10,000 replicates (see SM3 for details). The use of AI in each occupation is classified as *morally repugnant*, *morally ambivalent*, or *morally permissible* based on whether the confidence interval for its mean score lies entirely above, overlaps with, or lies entirely below the neutral midpoint (4.0), respectively. Overall, the use of AI was classified as morally repugnant in 12% of occupations—indicating categorical rejection even when AI outperforms humans at lower cost—morally permissible in 37%, and morally ambivalent in 51%.

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
Most repugnant		
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]
Least repugnant		
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]

Among morally permissible occupations, AI acceptance is both high at baseline and highly responsive to performance gains (Fig. 2, bottom panel). Virtually all (99.1% [97.5%–99.8%]) of these occupations—such as file clerks, data entry keyers, and cashiers—see support for AI augmentation at current levels, rising to 100.0% [98.9%–100.0%] as AI improves. Automation support is similarly performance-sensitive, increasing from 60.7% [55.3%–65.9%] to 92.5% [89.2%–95.0%] under the advanced AI condition. This pattern suggests that resistance in these occupations is primarily performance-based—as AI becomes more accurate, transparent, and cost-effective, opposition largely dissolves.

In contrast, occupations in which participants are morally ambivalent about the use of AI exhibit moderate acceptance at baseline but only partial responsiveness to AI advancements. These roles—such as nuclear power reactor operators, bakers, and sports medicine physicians—show high support for augmentation (97.1% [95.2%–98.4%]) under current conditions, but automation remains controversial, rising only from 14.9% [11.9%–18.4%] to 47.3% [42.8%–51.9%] under advanced AI conditions. This pattern suggests that while performance improvements can alleviate some concerns, underlying ethical or social hesitations remain. Unlike morally permissible occupations, where AI opposition stems mainly from practical concerns, ambivalent occupations suggest a mix of performance-based and principle-based resistance, meaning that even perfect AI may not fully displace human workers in these roles.

The starkest pattern emerges in occupations where the use of AI is considered morally repugnant. In these occupations, AI is broadly rejected and technological advancements have no impact on acceptance. These occupations—including clergy, child-care workers, marriage and family therapists, and funeral directors—exhibit relatively low and unresponsive augmentation support (67.9% [58.4%–76.4%]) and complete rejection of automation (0.0% [0.0%–3.2%]), both of which remain entirely unchanged under advanced AI. This entrenched resistance suggests absolute moral opposition to AI replacing humans in these roles, aligning with core insights from economic sociology and social psychology (Zelizer 2005; Fiske and Tetlock 1997), in which certain social functions—particularly those rooted in caregiving, emotional labor, or spiritual leadership—are perceived as corrupted by the logic of instrumental rationality (in SM 7 and SM 8 we find that moral repugnance is higher for occupations in such categories). Unlike occupations where AI acceptance is a function of technical progress, these findings suggest fixed moral boundaries that may constrain AI adoption in specific labor sectors, regardless of economic or technological incentives.

When Can-Do and Should-Do Diverge: Misalignment Between Technical Feasibility and Moral Acceptance

Are public attitudes about what AI should do aligned with where AI can perform effectively? To examine this question, we compare our occupation-level moral repugnance scores with external estimates of AI's technical capability (c.f. Shao et al. 2025). Specifically, we draw on Eloundou et al. (2024), who develop a set of occupational exposure metrics to assess the extent to which large language models (LLMs) could automate work tasks. We use one of their measures that estimates the proportion of an occupation's tasks in which an LLM, with or without additional software tools, could reduce task time by at least 50%. Full details on how we harmonize and analyze these independent data sources are provided in SM4, which outlines our procedures for integrating task-based feasibility estimates with survey-based measures of social acceptance.

If repugnance toward AI is primarily driven by concerns about technical feasibility, we would expect a strong positive relationship between the two. Instead, the correlation is weak ($R^2 = 0.134$; see Table S10 for OLS regression results), suggesting that the boundaries of AI's technical aptitude and its social acceptability are often misaligned (Figure 4).

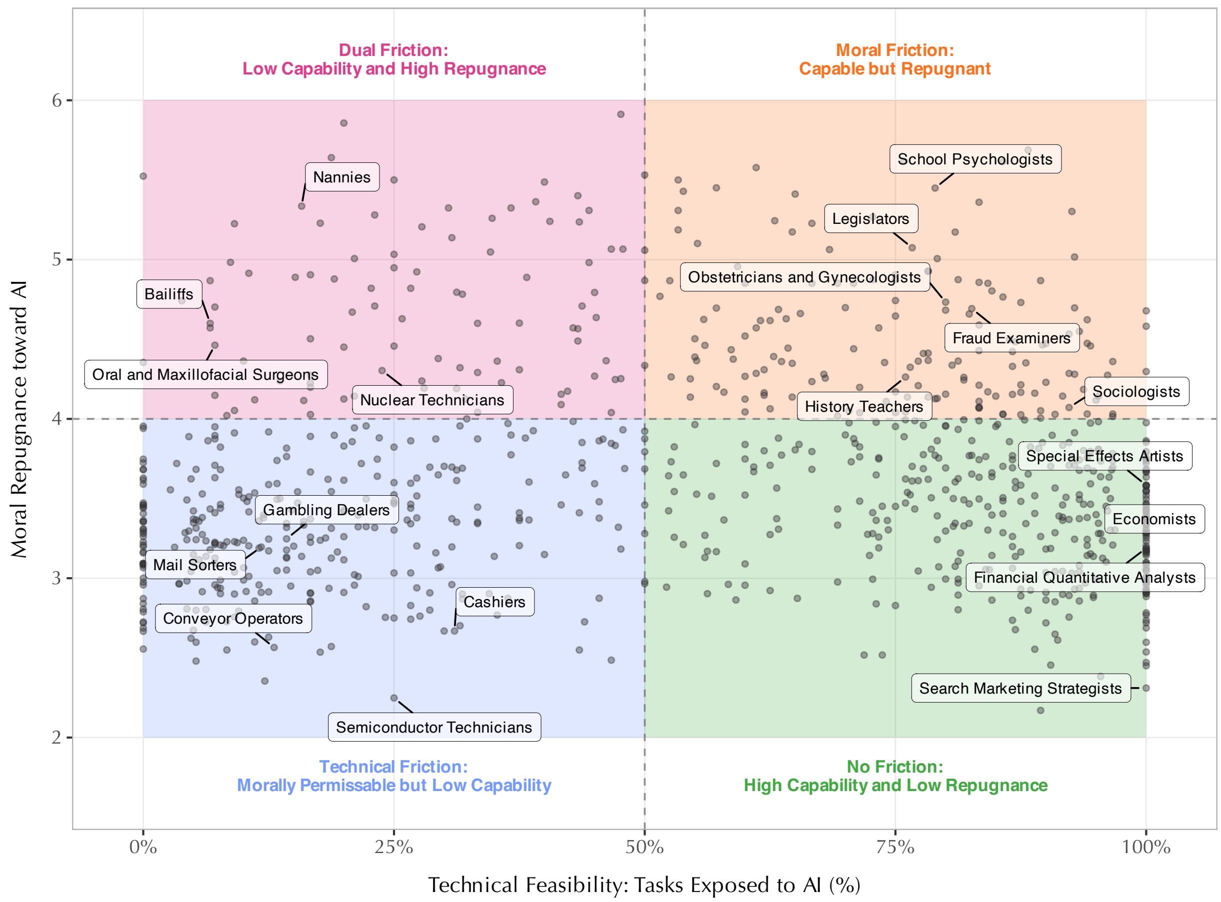

Figure 4. Moral Repugnance vs. Technical Feasibility. Each point represents an occupation. The x-axis measures technical feasibility using Eloundou et al.'s (2024) estimates of the share of tasks where AI could reduce time by at least 50%. The y-axis shows moral repugnance toward AI in that occupation. Occupations in the upper-right quadrant exhibit moral friction: AI is technically capable but socially resisted. The lower-left quadrant reflects a latent zone: AI is not yet feasible but broadly accepted. The top-left and bottom-right quadrants represent cases of social and technical alignment. We note that Shao et al. (2025) also employ a quadrant design to illustrate worker desires vs. technological capabilities, though our axes and conceptual focus differ.

The divergence between social and technical feasibility has two implications. Methodologically, it validates that our measure of moral repugnance captures a distinct form of resistance, not merely disguised skepticism about AI's capabilities. If participants were expressing performance concerns in moral terms, we would observe much stronger alignment between technical feasibility and repugnance. That we do not underscores that moral boundaries operate independently from performance-based judgments.

Substantively, the misalignment suggests distinct patterns of resistance and acceleration across the labor market, where existing exposure models may either overestimate or underestimate the likelihood of AI adoption. Occupations in the no-friction quadrant (bottom-right)—where technical feasibility and moral acceptability are aligned—represent the cases where we should expect the most rapid adoption (e.g., data entry keyers, file clerks). At the other extreme, occupations in the dual-friction quadrant (upper-left)—where both technical infeasibility and moral opposition align—represent the cases where we should expect the slowest or effectively no adoption (e.g., actors, clergy).

Of particular interest are the off-diagonal quadrants, where social and technical feasibility as misaligned. These quadrants indicate where moral attitudes may act as a catalyst or constraint relative to technical feasibility alone. Occupations in the moral friction quadrant (upper-right)—such as caregiving, counseling, and spiritual roles—are technically suitable for automation, yet face principled resistance. Here, moral repugnance may inhibit adoption even when AI performance is strong. By contrast, the technical friction quadrant (bottom-left) includes jobs like biomedical engineers, where AI is not yet technically capable but faces little social opposition.[1]

[1] While our primary measure focuses on moral resistance, we conducted exploratory analyses (not reported here) using an independent "moral embrace" scale that assessed whether people actively want AI used in a given occupation. These early tests suggest that, in practice, the absence of moral repugnance often coincides with moral openness or even support. This implies that occupations in the moral acceleration quadrant, where AI is not yet capable but faces little normative resistance, may represent sites of anticipatory adoption, where social acceptance could outpace technical feasibility.

In these cases, public openness may accelerate adoption ahead of technical maturity. Moral considerations may thus serve both an enabling and constraining role, shaping where and how quickly AI actually diffuses, above and beyond what exposure scores alone would predict. See Table S12 for representative examples from each quadrant, and Table S11 for the overall distribution of occupations across these four categories.

These "off-diagonal" occupations—where technical capacity and moral acceptability diverge—highlight the importance of going beyond exposure models rooted in technical feasibility to understand AI's real-world impact. Moral boundaries may shield technically automatable roles from replacement, while enabling adoption in sectors that are not yet technically tractable. Anticipating where AI will transform work therefore requires attention not only to what machines can do, but also to what society is willing to let them do.

DISCUSSION

Our findings highlight a core tension in how society views AI in the workforce. On one hand, once AI is assumed to surpass human capabilities, support for replacing human labor expands dramatically, rising from 30% to 58% of occupations. This indicates that many concerns undergirding algorithmic aversion, such as accuracy, reliability, and bias, can be engineered away. If AI convincingly proves itself superior, many people appear ready to endorse automation in roles as diverse as data entry, transportation planning, and even aspects of healthcare diagnostics.

On the other hand, moral repugnance creates a moral floor for a smaller but meaningful subset of occupations, approximately 12% of those examined. These roles—frequently associated with caregiving, spiritual well-being, or deeply interpersonal connections—elicit stable opposition that does not diminish when AI is positioned as flawless. In these cases, participants' moral repugnance scores suggest a sense of sacredness or irreducible humanity. For these occupations, cost and performance considerations fail to lower this moral threshold.

AI and the Future of Work

These insights challenge technologically and economically deterministic narratives that forecast a linear march of AI adoption as the technology improves. A more accurate model is that public acceptance involves both a performance threshold—where better, cheaper AI can convert skeptics—and a moral threshold—where no amount of improvement resolves objections. In practical terms, this model implies that AI's labor-market impact will be uneven: many jobs may see rapid automation as capabilities advance, whereas others could remain effectively untouched by AI for reasons of principle or identity.

For policymakers, recognizing moral constraints is essential. If certain occupations are regarded as intrinsically human, attempts to automate them may face backlash, reputational damage, or outright prohibition. Policy debates over AI licensing, certification, and liability should account for this moral dimension—particularly in high-stakes fields such as eldercare, therapy, or spiritual counsel. Legislation that blindly incentivizes AI adoption without attending to these deep-seated beliefs may trigger public controversies similar to those surrounding genetically modified crops or embryonic stem cell research (Frewer et al. 2013; Bauer 1995; *Ethical Issues in Human Stem Cell Research Volume I: Report and Recommendations of the National Bioethics Advisory Commission* 1999).

For businesses seeking to implement AI, these findings caution that superior performance alone will not always guarantee public acceptance. Companies may instead need hybrid solutions that preserve some human touch in morally sensitive tasks or carefully reframe the role of AI as an augmenting tool rather than a replacement. Even in professions generally open to automation, organizations must communicate transparency and address ethical concerns, as consumer or employee pushback can stall promising innovations.

Limitations

While our survey is large and quota-matched to U.S. census demographics, it remains a non-probability internet panel. Biases could arise, for example, if AI skeptics are over- or underrepresented in our sample. We focused on the U.S. labor market, but cultural norms surrounding technology and labor differ across societies. Future cross-national studies could reveal whether morally protected roles vary internationally or if similar moral boundaries arise elsewhere. Moreover, norms can shift over time. Practices once considered taboo—such as in vitro fertilization—often gain acceptance through reframing, generational change, and institutional endorsement. The moral repugnance we observe may likewise erode over time if cultural perceptions of AI evolve, or it may harden if controversies reinforce the sense that certain roles should remain exclusively human. Moreover, the time horizon for AI's improvement is unspecified; acceptance of AI "surpassing humans" presumably depends on demonstrable successes that may or may not materialize soon.

Self-reported attitudes do not always translate into real-world behavior. People might overstate their moral objections when there is no personal cost in a survey context. Conversely, real-life concerns about AI safety or bias could intensify resistance beyond what we see under hypothetical scenarios. We also measured moral repugnance at the occupation level. This approach helps capture the holistic identity of an occupation but may obscure important nuances of specific tasks within that role. Future work might examine which particular tasks evoke the strongest moral boundaries—caregiving might be “off-limits,” for instance, while scheduling or record-keeping within the same job is deemed acceptable for AI.

CONCLUSION

Far from a simple tale of cost and feasibility, our data reveal that even when AI is apt for a task, people may vigorously resist its use. For most occupations, better performance suffices to convert skeptics, suggesting a broad openness to automation as AI improves. But for a notable subset, defined by their moral or symbolic significance, performance is irrelevant: the use of AI is categorically rejected, even under ideal conditions.

This dynamic implies a future in which AI advances rapidly through permissive sectors of the economy, even as a limited but resilient frontier of human-exclusive work persists. Understanding the contours of that frontier—who and what it protects—is essential for anticipating the social consequences of automation. As AI capabilities evolve, the moral boundaries around work may shift, but they are unlikely to vanish entirely. Whether they endure or erode will depend not only on what AI can do, but on what society chooses to let it do.

REFERENCES

Acemoglu, Daron, and David Autor. 2011. *Skills, Tasks and Technologies: Implications for Employment and Earnings*. Edited by David Card and Orley Ashenfelter. Vol. 4b. Handbook of Labor Economics. https://doi.org/10.1016/s0169-7218(11)02410-5.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *The Quarterly Journal of Economics* 118 (4): 1279–333. https://doi.org/10.1162/003355303322552801.

Bansak, Kirk, and Elisabeth Paulson. 2024. "Public Attitudes on Performance for Algorithmic and Human Decision-Makers." *PNAS Nexus* 3 (12): pgae520. https://doi.org/10.1093/pnasnexus/pgae520.

Bauer, Martin, ed. 1995. *Resistance to New Technology*. Cambridge University Press. https://doi.org/10.1017/cbo9780511563706.020.

Becker, Howard Saul. 1982. *Art Worlds*. University of California Press.

Bigman, Yochanan E., and Kurt Gray. 2018. "People Are Averse to Machines Making Moral Decisions." *Cognition* 181: 21–34. https://doi.org/10.1016/j.cognition.2018.08.003.

Boateng, Godfred O., Torsten B. Neilands, Edward A. Frongillo, Hugo R. Melgar-Quiñonez, and Sera L. Young. 2018. "Best Practices for Developing and Validating Scales for Health, Social, and Behavioral Research: A Primer." *Frontiers in Public Health* 6: 149. https://doi.org/10.3389/fpubh.2018.00149.

Boussioux, Léonard, Jacqueline N Lane, Miaomiao Zhang, Vladimir Jacimovic, and Karim R Lakhani. 2024. "The Crowdless Future? Generative AI and Creative Problem-Solving." *Organization Science* 35 (5): 1589–607. https://doi.org/10.1287/orsc.2023.18430.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. 2018. "What Can Machines Learn and What Does It Mean for Occupations and the Economy?" *AEA Papers and Proceedings* 108: 43–47. https://doi.org/10.1257/pandp.20181019.

Castelo, Noah, Maarten W. Bos, and Donald R. Lehmann. 2019. "Task-Dependent Algorithm Aversion." *Journal of Marketing Research* 56 (5): 809–25. https://doi.org/10.1177/0022243719851788.

Dell'Acqua, Fabrizio, Edward McFowland, Ethan R. Mollick, et al. 2023. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality." *SSRN Electronic Journal*, ahead of print. https://doi.org/10.2139/ssrn.4573321.

National Center for O*NET Development. n.d. "O*NET OnLine." Accessed January 29, 2025. www.onetonline.org/.

Dietvorst, Berkeley J., Joseph P. Simmons, and Cade Massey. 2015. "Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err." *Journal of Experimental Psychology: General* 144 (1): 114–26. https://doi.org/10.1037/xge0000033.

Dong, Mengchen, Jane Rebecca Conway, Jean-François Bonnefon, Azim Shariff, and Iyad Rahwan. 2024. "Fears About Artificial Intelligence Across 20 Countries and Six Domains of Application." *American Psychologist*, ahead of print. https://doi.org/10.1037/amp0001454.

Douglas, Mary. 1966. *Purity and Danger: An Analysis of Concepts of Pollution and Taboo*. Routledge & Kegan Paul.

Elías, Julio J, Nicola Lacetera, and Mario Macis. 2019. "Paying for Kidneys? A Randomized Survey and Choice Experiment." *American Economic Review* 109 (8): 2855–88. https://doi.org/10.1257/aer.20180568.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2024. "GPTs Are GPTs: Labor Market Impact Potential of LLMs." *Science* 384 (6702): 1306–8. https://doi.org/10.1126/science.adj0998.

Ethical Issues in Human Stem Cell Research Volume I: Report and Recommendations of the National Bioethics Advisory Commission. 1999. The National Bioethics Advisory Commission.

Felten, Edward, Manav Raj, and Robert Seamans. 2021. "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses." *Strategic Management Journal* 42 (12): 2195–217. https://doi.org/10.1002/smj.3286.

Fiske, Alan Page, and Philip E. Tetlock. 1997. "Taboo Trade-offs: Reactions to Transactions That Transgress the Spheres of Justice." *Political Psychology* 18 (2): 255–97. https://doi.org/10.1111/0162-895x.00058.

Freitas, Julian De, Stuti Agarwal, Bernd Schmitt, and Nick Haslam. 2023. "Psychological Factors Underlying Attitudes toward AI Tools." *Nature Human Behaviour* 7 (11): 1845–54. https://doi.org/10.1038/s41562-023-01734-2.

Frewer, Lynn J., Ivo A. van der Lans, Arnout R.H. Fischer, et al. 2013. "Public Perceptions of Agri-Food Applications of Genetic Modification – A Systematic Review and Meta-Analysis." *Trends in Food Science & Technology* 30 (2): 142–52. https://doi.org/10.1016/j.tifs.2013.01.003.

Frey, Carl Benedikt, and Michael A. Osborne. 2017. "The Future of Employment: How Susceptible Are Jobs to Computerisation?" *Technological Forecasting and Social Change* 114: 254–80. https://doi.org/10.1016/j.techfore.2016.08.019.

Glikson, Ella, and Anita Williams Woolley. 2020. "Human Trust in Artificial Intelligence: Review of Empirical Research." *Academy of Management Annals* 14 (2): 627–60. https://doi.org/10.5465/annals.2018.0057.

Globig, Laura K, Rachel Xu, Steve Rathje, and Jay Joseph Van Bavel. 2024. *Perceived (Mis)Alignment in Generative Artificial Intelligence Varies Across Cultures*. https://doi.org/10.31234/osf.io/suqa2.

Hochschild, Arlie Russell. 1983. *The Managed Heart: Commercialization of Human Feeling*. University of California Press.

Jackson, Summer R. 2023. "(Not) Paying for Diversity: Repugnant Market Concerns Associated with Transactional Approaches to Diversity Recruitment." *Administrative Science Quarterly* 68 (3): 824–66. https://doi.org/10.1177/00018392231183649.

Lamont, Michèle. 1992. *Money, Morals, and Manners: The Culture of the French and the American Upper-Middle Class*. University of Chicago Press.

Leuker, Christina, Lasare Samartzidis, and Ralph Hertwig. 2021. "What Makes a Market Transaction Morally Repugnant?" *Cognition* 212: 104644. https://doi.org/10.1016/j.cognition.2021.104644.

Loaiza, Isabella, and Roberto Rigobon. 2024. *The EPOCH of AI: Human-Machine Complementarities at Work*. https://doi.org/10.2139/ssrn.5028371.

Longoni, Chiara, Andrea Bonezzi, and Carey K Morewedge. 2019. "Resistance to Medical Artificial Intelligence." *Journal of Consumer Research* 46 (4): 629–50. https://doi.org/10.1093/jcr/ucz013.

Manning, Benjamin S, Kehang Zhu, and John J Horton. 2024. "Automated Social Science: Language Models as Scientist and Subjects." *arXiv*, ahead of print. https://doi.org/10.48550/arxiv.2404.11794.
March, James G., and Johan P. Olsen. 1996. "Institutional Perspectives on Political Institutions." *Governance* 9 (3): 247–64. https://doi.org/10.1111/j.1468-0491.1996.tb00242.x.
Nussbaum, Martha C. 2010. *From Disgust to Humanity: Sexual Orientation and Constitutional Law*. Oxford University Press.
Phillips, Damon J, Catherine J Turco, and Ezra W Zuckerman. 2013. "Betrayal as Market Barrier: Identity-Based Limits to Diversification among High-Status Corporate Law Firms." *American Journal of Sociology* 118 (4): 1023–54. https://doi.org/10.1086/668412.
Qin, Xin, Xiang Zhou, Chen Chen, et al. 2025. "AI Aversion or Appreciation? A Capability–Personalization Framework and a Meta-Analytic Review." *Psychological Bulletin* 151 (5): 580–99. https://doi.org/10.1037/bul0000477.
Roth, Alvin E. 2007. "Repugnance as a Constraint on Markets." *Journal of Economic Perspectives* 21 (3): 37–58. https://doi.org/10.1257/jep.21.3.37.
Satz, Debra. 2010. *Why Some Things Should Not Be for Sale*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195311594.001.0001.
Shao, Y., Zope, H., Jiang, Y., Pei, J., Nguyen, D., Brynjolfsson, E. and Yang, D., 2025. Future of Work with AI Agents: Auditing Automation and Augmentation Potential across the US Workforce. *arXiv preprint arXiv:2506.06576*.
Tu, Tao, Anil Palepu, Mike Schaekermann, et al. 2024. "Towards Conversational Diagnostic AI." *arXiv*, ahead of print. https://doi.org/10.48550/arxiv.2401.05654.
Velthuis, Olav. 2005. *Talking Prices: Symbolic Meanings of Prices on the Market for Contemporary Art*. Princeton University Press.
Wallach, Wendell, and Colin Allen. 2008. *Moral Machines: Teaching Robots Right from Wrong*. Oxford University Press.
Webb, Michael. 2019. "The Impact of Artificial Intelligence on the Labor Market." *SSRN Electronic Journal*, ahead of print. https://doi.org/10.2139/ssrn.3482150.
Zelizer, Viviana A. 2005. *The Purchase of Intimacy*. Princeton University Press.
Zelizer, Viviana A. Rotman. 1979. *Morals and Markets: The Development of Life Insurance in the United States*. Columbia University Press.

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

Table of contents

S1 Materials and Methods

S1.1 Research Objectives

This study aimed to systematically map resistance to artificial intelligence (AI) across the U.S. labor market by evaluating a comprehensive set of occupations. Our primary objective was to characterize the overall distribution and prevalence of resistance to AI adoption, rather than to precisely estimate attitudes toward individual occupations or detect specific pairwise differences. This emphasis on breadth over precision informed key methodological decisions, from sample size determination to scale construction.

S1.2 Ethics

This research was approved by the Harvard University-Area Committee on the Use of Human Subjects under protocol numbers IRB24-0098 (main study) and IRB24-0756 (scale development and validation study). All participants provided informed consent before beginning the study and received monetary compensation for their involvement. No deception was used.

S1.3 Sample

S1.3.1 Target Sample and Size Determination

Our target sample was 23,500 occupation ratings (25 ratings $\times$ 940 occupations) from 2,350 participants, with each participant evaluating 10 randomly assigned occupations. Sample size calculations were informed by our scale validation study (Supplementary Materials S2) of nine occupations spanning low, medium, and high moral repugnance to AI. This study yielded moral repugnance scores ranging from 3.09 to 5.12 (M = 3.88, SD = 0.78) with standard deviations ranging from 1.31 to 1.72 (M = 1.58, SD = 0.15). Using this mean standard deviation, 25 responses per occupation provided a margin of error of 0.65 at a 95% confidence interval level. This provides sufficient precision for characterizing distributional patterns while enabling broad occupational coverage.

Participants were recruited through Prolific, a non-probability (opt-in) Internet panel. To enhance the generalizability of our distributional findings to the U.S. population and enable exploration of demographic heterogeneity in resistance to AI, we used Prolific's representative sampling service[1] to implement census-matched demographic quotas (age, sex, ethnicity, and political affiliation). While the limited ratings per occupation constrain occupation-specific inferences, the combination of demographic representativeness and randomized occupation assignment should help attenuate sampling biases when examining distributional patterns.

[1] https://researcher-help.prolific.com/en/article/e6555f

S1.3.2 Final Sample and Exclusion Criteria

Of 2,436 participants who completed the survey, we excluded 79 based on two exclusion criteria designed to filter out potential bots and inattentive respondents. First, we excluded participants who completed the survey in less than one-third of the median completion time (15.1 minutes, corresponding to a threshold of 302 seconds; n = 22). Second, we excluded participants with reCaptcha scores below 0.5 (n = 57). These scores, provided by Qualtrics[2] based on Google's bot detection algorithm, range from 0 to 1, with scores ≥ 0.5 indicating likely human respondents and scores < 0.5 indicating likely automated behavior. No participants met both exclusion criteria.

The final sample comprised 2,357 participants who provided 23,570 ratings across 940 occupations. Participants were 51.2% women with a mean age of 45.2 years (SD = 16.0, range = 18-85). Tables Table S1, Table S2 show simple and detailed demographic characteristics of our final sample. Tables Table S3, Table S4 compare these demographic characteristics to the U.S. population. Compared to the U.S. population, our sample is slightly younger and less White, with more respondents identifying as Mixed ethnicity (Table S3).

S1.4 Procedure

The complete survey instrument is available in the Open Science Framework project repository[3].

To systematically map attitudes toward AI across the U.S. labor market, we drew occupations from the O*NET database (version 29.0), a comprehensive and standardized catalog of occupational information sponsored by the U.S. Department of Labor. From the complete set of 1,016 O*NET occupations, we included all except 76 residual "All Other" categories (e.g., "Personal Service Managers, All Other" and "Education Administrators, All Other"). These residual categories comprise occupations "not listed separately" in O*NET, with descriptions taking the form of "All education administrators not listed separately." These categories were excluded because they are challenging for participants to meaningfully evaluate since they were not presented with the specific occupations being referenced.

Each participant evaluated 10 randomly assigned occupations. Each evaluation began with the O*NET job title and description, followed by questions about AI use under two scenarios: current AI capabilities and advanced AI capabilities. These scenarios were designed to distinguish between all-cause resistance to AI under current circumstances and enduring forms of social resistance that are not alleviated through technical advancement. The advanced AI scenario helps isolate enduring forms of social resistance by minimizing concerns about AI's technical capabilities and economic viability. By "turning off" resistance mechanisms rooted in quality concerns, we can better identify fundamental social limits to the acceptance of new technologies.

For current capabilities, participants evaluated their support for AI augmentation and automation of core tasks after reading:

[2] https://www.qualtrics.com/support/survey-platform/survey-module/survey-checker/fraud-detection/

> "The following questions ask about AI's role in performing the core tasks of a specific job, given the current level of AI technology. Remember: Core tasks are the essential responsibilities that define a job, even if they do not take up the most time in a typical workday."

Under the current AI scenario, participants first evaluated whether AI *can* augment or automate core tasks before assessing whether it *should* be used in these ways. This design choice deliberately contrasted technical feasibility with social acceptability, implicitly prompting participants to distinguish between practical limitations and normative concerns about AI use.

Participants were then instructed to consider a hypothetical scenario involving advanced AI that outperforms humans and does so at lower cost:

> "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. Given this scenario, rate your level of agreement with the following statements:"

Participants then rated their support for AI augmentation and automation under the advanced AI scenario, followed by their moral repugnance toward AI use in the occupation.

S1.5 Measures

Our analysis focused on two primary outcomes: support for AI and moral repugnance toward AI. Support for AI was measured through four variants (augmentation and automation under both current and advanced scenarios), while moral repugnance captured enduring opposition to AI use. All items used 7-point Likert scales (1 = Strongly Disagree, 7 = Strongly Agree) unless noted otherwise. Analyses use unweighted measures except where specified.

Support for AI measures were based on single items:

1. **Current AI - augmentation support:** "Given current capabilities, AI should be used to assist workers with the core tasks of this job."
2. **Current AI - automation support**: "Given current capabilities, AI should be used to fully automate the core tasks done by workers in this job."
3. **Advanced AI - augmentation support**: "If AI outperforms humans, we should use it to significantly assist workers with the core tasks of this job."
4. **Advanced AI - automation support**: "If AI outperforms humans, we should use it to fully automate the core tasks of this job."

Responses were dichotomized, with scores ≥ 4 (neutral or positive agreement) coded as 1, reflecting our operational definition of support as the absence of resistance.

Moral repugnance toward AI was measured using a validated 7-item scale (for scale construction and validation details, see Supplementary Materials S2) administered during the advanced capabilities scenario. The scale assessed enduring moral opposition to AI use, independent of practical benefits:

1. **Human dignity**: "AI should never be used in core tasks of this job because it violates human dignity."
2. **Off-limits**: "No matter how advanced AI becomes, this job should remain off-limits to machines."
3. **Moral responsibility**: "This job is too important to be left to AI, as it involves deep moral responsibilities."
4. **Imperfect human preference**: "I believe it is better for a human to do this job imperfectly than for AI to do it perfectly."
5. **Human-only decisions**: "This job involves decisions that should never be handed over to AI."
6. **Avoidance**: "Even if the quality is high, I would avoid using services or products in this field if I knew AI was involved in their creation."
7. **Betrayal**: "I would feel betrayed if I found out AI was being used in this job."

Individual moral repugnance scores were calculated by averaging responses across the seven items. Occupation-level scores represent the mean of individual scores.

S2 Supplementary Text: Scale Development and Validation

In this section, we introduce and validate a new, parsimonious scale of moral repugnance toward AI in occupational contexts—that is, we measure the extent to which individuals view the use of AI in certain jobs as inherently wrong or off-limits, irrespective of any demonstrated benefits in efficiency or performance. Existing discussions of moral repugnance often center on market transactions (e.g., organ sales, commercial surrogacy) and frequently blend consequentialist objections (e.g., potential harm, reduced quality) with deeper principle-based stances. By contrast, we focus on the deontological core of moral repugnance: the idea that some uses of AI violate fundamental moral boundaries and therefore remain objectionable even if AI is more capable or cost-effective than humans.

Our goal in developing this scale stems from both theoretical and practical considerations. Theoretically, prior accounts of moral repugnance often conflate outcome-driven concerns (e.g., risk of harm, safety issues) with absolute moral prohibitions (e.g., "certain tasks must not be done by machines, as a matter of principle"). We seek to disentangle the two by explicitly framing our items to capture enduring moral objections—those that would remain even if AI were flawless and cost-effective. Practically, we need a measure concise enough for use across multiple occupations, including those involving physical labor and those requiring high-level moral or emotional judgment. Existing instruments—developed primarily for taboo trades or commodification—are not readily adaptable to AI adoption contexts or fail to isolate strict deontological objections from broader concerns about technology's outcomes. The scale we introduce here aims to fill this methodological gap and support the rigorous study of moral limits to AI in diverse professional domains.

S2.1 Theoretical Foundations and Conceptual Framework

S2.1.1 Moral Repugnance in Prior Literature

Moral repugnance has long been studied in the context of controversial market transactions and taboo trades (Roth 2007; V. A. Zelizer 2005; Fiske and Tetlock 1997; Jackson 2023; Satz 2010; Leuker, Samartzidis, and Hertwig 2021; Elías, Lacetera, and Macis 2019). In these settings, *repugnance* typically refers to the sense that buying or selling certain goods or services is beyond the pale, even if the exchange could be mutually beneficial. Many studies highlight concerns about commodification, human dignity, or fairness and justice—yet often blend these principled stances with consequentialist fears about potential societal harm or exploitation.

A key takeaway from this literature is that moral repugnance can stem from more than just practical or utilitarian objections. For instance, Elías, Lacetera, and Macis (2019) identify appeals to human dignity as a driving factor in repugnance to certain transactions, distinct from perceived economic inefficiencies. Phillips, Turco, and Zuckerman (2013) similarly show how a sense of betrayal can trigger resistance to market activities that violate identity-based norms. However, these and other existing studies do not consistently differentiate whether the repugnance is contingent on negative outcomes (e.g., harm to third parties) or stands as a blanket prohibition on moral or cultural grounds.

S2.1.2 Applying Moral Repugnance to Technology Use

Despite its longstanding presence in discussions of taboo trades, moral repugnance has received comparatively little attention as a driver of technological resistance. Indeed, most empirical work on public attitudes to AI centers on performance-related factors—such as potential job loss, privacy breaches, or biased outcomes—implicitly assuming that if AI becomes sufficiently accurate, trustworthy, and cost-effective, these sources of opposition will recede. However, if some occupations are considered intrinsically off-limits to machines, then no level of technical proficiency or safety will fully mitigate resistance.

Yet the possibility that AI can do a job "better" than humans does not necessarily alleviate all concerns. Attitudes toward AI in such roles may rest not on whether the AI is capable, but whether these tasks ought to be performed by machines at all. This point becomes clearer when we consider certain roles perceived to involve unique moral or emotional dimensions (e.g., clergy, end-of-life care, mental health counseling). If people believe these tasks should never be entrusted to machines—regardless of how "good" the machines become—then a strictly performance- or risk-based analysis overlooks a more fundamental type of objection. While improved engineering may address reliability or safety concerns, it cannot eliminate principled reservations that equate automation with a violation of human dignity, sanctity, or moral responsibility.

To clarify this distinction, it is helpful to refer to two major frameworks in moral philosophy: consequentialism and deontology. Under consequentialism, the moral acceptability of an action depends primarily on its outcomes—if AI provides net benefits with minimal harm, then adopting it is justified. By contrast, deontology holds that some actions are inherently right or wrong, independent of their consequences (March and Olsen 1996; Wallach and Allen 2008).

In technology contexts, a consequentialist objection to AI can often be "engineered away" through better performance or safeguards, whereas a deontological objection persists even when those practical concerns are resolved. Addressing deontological opposition requires cultural, social, or normative change—processes that tend to unfold more slowly and unpredictably than purely technical improvements. Understanding these distinct foundations of moral repugnance is therefore crucial for analyzing why certain occupations are viewed as simply incompatible with AI, no matter its capabilities.

In light of this theoretical framing, our central aim is to isolate these purely principle-based concerns—what we term the "deontological core" of moral repugnance—by building a scale that explicitly references the idea that "certain tasks should remain off-limits to AI, no matter the efficiency or cost savings."

S2.2 Scale Development Process

S2.2.1 Item Generation

To develop a concise yet conceptually robust set of items, we synthesized insights from (1) moral philosophy on deontological prohibitions, (2) empirical studies of moral repugnance in market contexts, and (3) emerging research on AI ethics and resistance. We started with a working definition of moral repugnance toward AI as principled, categorical opposition to AI adoption in particular occupations, even when AI performance clearly exceeds human abilities on consequential metrics.

Guided by earlier literature on moral repugnance (Leuker, Samartzidis, and Hertwig 2021), we drafted items designed to capture this deontological stance. Specifically, we avoided references to harmful or negative outcomes (e.g., "AI would pose a risk to the user") and instead emphasized fundamental moral or existential concerns (e.g., "It violates human dignity," "This job should remain off-limits to machines," "I would feel betrayed if AI were doing this job"). We also included items tapping into emotional and behavioral dimensions, such as disgust or refusal to use AI-made products or services, to capture repugnance's visceral and motivational qualities (Phillips, Turco, and Zuckerman 2013; Nussbaum 2010).

In total, this process yielded a pool of 13 candidate items, each framed such that the participant would imagine a future scenario where AI outperforms humans on the core tasks of a given job. We deliberately emphasized this advanced capability to ensure that any stated opposition could be attributed to moral repugnance rather than concerns about AI's current technical limitations. Because we aimed to develop a final scale with approximately six items, we followed best practices in scale development by initially generating a pool of roughly twice that number. This approach allowed us to systematically prune items based on empirical performance while ensuring the final scale retained conceptual breadth (Boateng et al. 2018).

The full set of 13 candidate items is presented below. The bold labels, which were not shown to participants, serve as shorthand for reference throughout the analysis:

1. **Disgust:** I would feel a sense of disgust at the idea of using AI to perform core tasks in this job, even if it outperforms humans.
2. **Principled objection:** AI should never be used for the core tasks of this job as a matter of principle.
3. **Human dignity violation:** AI should never be used in core tasks of this job because it violates humanity dignity.
4. **Inherent wrongness:** I find the use of AI in this job inherently wrong, regardless of its performance.
5. **Job off-limits:** No matter how advanced AI becomes, this job should remain off-limits to machines.
6. **Deep moral responsibility:** This job is too important to be left to AI, as it involves deep moral responsibilities.
7. **Efficiency irrelevant:** Even if AI can perform this job more efficiently, it is still morally wrong to use it.

8. **Human imperfection preferred:** I believe it is better for a human to do this job imperfectly than for AI to do it perfectly.
9. **Human-only decisions:** This job involves decisions that should never be handed over to AI.
10. **Avoidance:** Even if the quality is high, I would avoid using services or products in this field if I knew AI was involved in their creation.
11. **Alternative seeking:** I would actively seek out alternatives that don't use AI for this job, even if they were less convenient or more expensive.
12. **Reduced interest:** The involvement of AI in creating a product or service would make me less interested in buying it, regardless of the price.
13. **Betrayal:** I would feel betrayed if I found out AI was being used in this job.

S2.2.2 Pre-Test and Cognitive Testing

Before deploying our scale in a larger validation study, we conducted a small pre-test with 30 U.S.-based adults recruited from Prolific. Participants answered the 13 candidate items and provided open-ended feedback on clarity, wording, and relevance. The design of the pre-test survey was identical to the subsequent scale validation study except participants were asked to provide open-ended feedback after each set of questions.

Feedback indicated the scenario was generally well-understood, and none of the items elicited persistent confusion. Consequently, we proceeded with all 13 items unchanged into the main validation study.

S2.2.3 Survey Design

The full survey instruments for the pre-test and final validation studies are available in the project repository on OSF[4].

Participants were randomly assigned to rate one of nine occupations. These occupations were selected based on two dimensions: (1) levels of moral repugnance toward AI (based on a previous pilot survey), and (2) the technical suitability of AI for these professions, as derived from Brynjolfsson, Mitchell, and Rock (2018). We selected occupations representing low, medium, and high levels on both dimensions to ensure that the scale could be applied to a broad range of occupations. The selected occupations are shown in Table S6.

Participants were shown the title of the occupation as well as the O*NET description of that occupation. Participants were then asked a series of questions about their familiarity with the occupation and their support for AI automation augmentation of the core tasks of the occupation under current AI capabilities. Participants were told that "Core tasks are the essential responsibilities that define a job, even if they do not take up the most time in a typical workday."

Finally, participants were asked to consider advanced AI capabilities, with the instructions prompting them as follows (emphasis in original): "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." Participants were then shown the 13 moral repugnance items, which were randomly ordered. All response options were on a 7-point Likert scale (1 = strongly disagree, 7 = strongly agree). Two attention check questions were included throughout the survey to ensure data quality.

S2.2.4 Sample Recruitment

We recruited 300 participants from Prolific, using quotas for sex, ethnicity, age, and political affiliation to ensure the sample was broadly representative of the US population and similar to the representative sample we recruited for the main study.

S2.2.5 Exclusion criteria

We excluded participants who failed one or both attention check questions, resulting in a final sample of 283 participants.

S2.3 Results

S2.3.1 Descriptive statistics

Descriptive statistics are provided in Table S7. Fig. S1 shows the distribution of moral repugnance item scores by occupation.

S2.3.2 Initial item validation

Our analysis provides strong evidence that the 13 items measuring moral repugnance toward AI represent a single, cohesive construct.

Factor Analysis. Parallel analysis suggested retaining a single factor, supporting a unidimensional structure. Exploratory factor analyses (EFA) using Principal Axis Factoring with Promax rotation supported this unidimensional structure. These methods were chosen because they allow for factor intercorrelations, which we expected given the nature of our items. All items loaded strongly onto the single factor (loadings: 0.79-0.91). The single factor explained 74.3% of the total variance, indicating that the items collectively capture a dominant underlying construct of moral repugnance toward AI use.

Reliability Analysis. The 13-item scale measuring moral repugnance toward AI use demonstrated exceptionally high internal consistency (Cronbach's α = 0.97). Item deletion analysis revealed minimal changes in alpha when individual items were removed (α range: 0.971-0.973), confirming that all items contribute similarly to the scale. Item-total correlations (range: 0.78-0.90) indicate that each item is strongly related to the overall construct.

Item Redundancy and Scale Refinement. Several high item-total correlations (>0.85) and inter-item correlations (range: 0.59-0.84, with some approaching 0.85) suggest redundancy among the items. This indicates that certain items may be overlapping in what they measure. Given the evidence of item redundancy, we implemented a systematic pruning process to enhance the scale's parsimony while maintaining its reliability and validity. This process aimed to reduce redundancy while retaining the scale's conceptual breadth.

S2.3.3 Item pool reduction

We employed the ICLUST clustering algorithm to identify groups of conceptually similar and statistically redundant items. Subsequently, we pruned the item pool by eliminating items with low discrimination between different occupations, as indicated by their eta-squared values. This approach allowed us to remove overlapping items without significantly compromising the construct's breadth. The item "Less interested in AI-made products" was the first to be eliminated due to its low eta-squared value.

Through iterative pruning of redundant items, we arrived at a final 7-item scale. These items were selected based on their capacity to capture key dimensions of moral repugnance and effectively discriminate between different occupational contexts. Detailed documentation of the pruning process is available in the replication notebooks, accessible via the authors' GitHub repository.

S2.3.4 Confirmatory Factor Analysis

We conducted a Confirmatory Factor Analysis (CFA) on the 7-item scale. The model demonstrated excellent fit, with Comparative Fit Index (CFI) = 0.97, Tucker-Lewis Index (TLI) = 0.96, and Standardized Root Mean Square Residual (SRMR) = 0.03. The Root Mean Square Error of Approximation (RMSEA) value of 0.11, while slightly higher than the conventional threshold of 0.08, is acceptable given the unidimensional nature of the construct and the strong fit indices overall. The final 7-item scale exhibited excellent reliability, with Cronbach's alpha = 0.95, indicating high internal consistency.

S2.4 Final Scale and Discussion

Table S8 The resulting 7-item scale provides a parsimonious, reliable, and valid measure of moral repugnance toward AI in professional contexts. Each retained item reflects key ethical concerns about AI's role in occupations where human dignity, moral judgment, and ethical responsibility are central. The scale successfully captures the broad spectrum of public moral opposition to AI use, ranging from emotional responses such as disgust and betrayal to principled objections based on human dignity.

S3 Supplementary Text: Categorizing Moral Repugnance Toward AI Use in Occupations

We classify each occupation into one of three categories—morally repugnant, morally ambivalent, or morally permissible—based on its occupation-level moral repugnance score. This score is calculated as the mean of individual moral repugnance ratings of that occupation.

Ratings of moral repugnance items were given on a 7-point Likert scale (1 = strongly disagree, 7 = strongly agree), with 4 as the neutral midpoint. The scale items assess negative moral evaluations of AI use, creating a bipolar measurement where high scores indicate strong moral repugnance toward AI in an occupation, while low scores reflect disagreement with this negative moral assessment. Importantly, low scores do not imply active endorsement of AI integration; rather, they are best interpreted as reflecting moral permissibiliy—a lack of perceived moral constraint. Midpoint scores indicate moral ambivalence or evaluative uncertainty, where judgments are mixed or indeterminate.

We then classified the 940 distinct occupations into one of three moral categories based on whether the 95% confidence interval for the occupation-level repugnance score included the neutral midpoint:

- **Morally Repugnant:** Occupations for which the entire confidence interval lay above 4, indicating strong moral repugnance toward the use of AI.
- **Morally Ambivalent:** Occupations for which the confidence interval included 4, indicating mixed or indeterminate moral judgments toward the use of AI.
- **Morally Permissible:** Occupations for which the confidence interval lay entirely below 4, suggesting an absence of perceived moral constraint toward the use of AI.

S3.1 Bootstrap Estimation of Occupation-Level Confidence Intervals

To estimate confidence intervals for occupation-level moral repugnance scores while accounting for the repeated-measures structure of the data, we implemented a cluster bootstrap at the participant level. This approach preserves the partially crossed design in which each participant rated multiple occupations, and each occupation was rated by multiple participants.

Resampling within occupations would incorrectly treat occupation-level scores as independent, ignoring the fact that each participant contributes multiple, correlated ratings. By resampling at the participant level, we preserve the dependence among responses and avoid underestimating uncertainty due to shared variance across ratings from the same individual.

We generated 10,000 bootstrap resamples by sampling participants with replacement and, for each resample, computed occupation-level mean moral repugnance scores. We then calculated bias-corrected and accelerated (BCa) bootstrap confidence intervals, which adjust for both bias and skewness in the resampled distribution.

The classification results are robust to the choice of bootstrap confidence interval method; percentile, studentized, and BCa approaches produced substantively consistent outcomes (Table S9).

S4 Supplementary Text: Analysis of Moral Repugnance and Technical Suitability for AI Adoption

S4.1 Technical Suitability Measure

We obtained our primary measure of technical suitability for AI from Eloundou et al. (2024), who assess the extent to which occupational tasks can be accelerated using large language models (LLMs). The authors define occupation-level exposure scores based on the proportion of tasks where AI can reduce task completion time by at least 50%. Their approach classifies tasks into three tiers of exposure to AI and aggregates these scores to the occupation level. They define three measures (p. 1306):

- *E1*: "The share of an occupation's tasks where access to an LLM alone or with a simple interface would lead to 50% time savings."
- *E2*: "The share of an occupation's tasks where additional software is needed on top of an LLM to realize 50% time savings."
- *E1 + E2*: The share of an occupation's tasks in which an LLM (with or without additional software) can reduce task time by 50%.

We use *E1 + E2* as our primary measure because it represents the broadest conception of LLM-based AI capability, thereby aligning most closely with our paper's focus on a future where AI outperforms humans.

We retrieved these data from the authors' public GitHub repository (@openai/GPTs-are-GPTs) and ran their published script to generate the occupation-level exposure measures.[5] Eloundou et al. (2024) construct exposure measures using two approaches: one based on human annotations and another based on LLM-generated labels. We use the LLM-coded scores, which the authors find are closely aligned with human annotation scores.

We merged these scores with our dataset of 940 occupations for which we collected moral repugnance scores, using O*NET-SOC occupation codes. Out of 940 occupations, 17 occupations (1.8%)—all military roles—had no corresponding score in Eloundou et al. (2024); we excluded these from our analyses of technical suitability vs. moral repugnance, leaving a final sample of 923 occupations for this comparison. Across these 923 occupations, the technical suitability scores ranged from 0 to 1 (mean = 0.55, SD = 0.34).

While this dataset represents the most recent, comprehensive, and granular estimate of occupational exposure to AI, one limitation is that the exposure scores are specific to LLM-based AI (specifically, ChatGPT). As a result, occupations that involve manual labor or physical dexterity receive low exposure scores, even though future AI developments—such as humanoid robotics—may be capable of automating such tasks. Our study takes a broader perspective on resistance to AI, considering not only

LLMs but also advances in robotics and embodied AI, which may shift the exposure landscape beyond what is captured in this dataset.

S4.2 Regression Analyses

To formally test the relationship between moral repugnance and technical suitability, we fit the following regression using OLS:

$$repugnance\_score_i = \beta_0 + \beta_1\, technical\_suitability_i + \beta_2\, technical\_suitability_i^2 + \epsilon_i$$

We included a second-degree term after visual inspection of a LOESS regression (Fig. 4, main text) suggested a curvilinear (inverted-U) relationship. Results are presented in Table S10.

S4.3 Quadrant Classification

We classified occupations as High vs. Low in moral repugnance using a threshold of 4.0 on the 7-point scale and High vs. Low technical suitability using a threshold of 0.5 on technical suitability. Table S11 reports the number of occupations in each category. Table S12 provides examples of occupations in each category.

S5 Supplementary Text: Estimating Employment Bounds for Morally Repugnant AI Occupations

We estimate bounds on total employment in the 12% of occupations where AI use is classified as morally repugnant. Employment data are drawn from the U.S. Bureau of Labor Statistics Occupational Employment and Wage Statistics (OEWS) database. OEWS data are mapped to O*NET occupations through standard crosswalks, which create a many-to-one mapping of O*NET occupations to OEWS occupations. To avoid double-counting within OEWS occupations, we calculate lower and upper bounds on employment, as well as a weighted estimate. The detailed mapping procedure and calculation of bounds is available in the replication materials on the OSF project repository.

The lower bound estimate assumes that all employees in mixed repugnance classifications are in non-repugnant roles, while the upper bound assumes all are in repugnant roles. We also calculate a weighted employment measure, allocating employment proportionally based on the percentage of SOC codes classified as repugnant for each OEWS title.

The estimated employment in occupations where AI use is considered morally repugnant ranges from 17.3 million to 21.9 million workers, representing 11.5% to 14.6% of the total workforce (149.8 million). The weighted estimate places employment in these occupations at 19.8 million, or 13.2% of the workforce. This range closely aligns with the 12% of occupations classified as morally repugnant for AI use, suggesting that these occupations are proportionally distributed across the workforce. This proportionality indicates that both small and large occupations are represented, without either extreme dominating the repugnant classification.

S6 Supplementary Text: Relationship Between Moral Repugnance and Occupational Demographics

How does moral resistance to AI shape labor market inequality along demographic lines? To examine the relationship between occupation-level moral repugnance scores and key occupational demographics, we used three main data sources: the 2022 Occupational Employment and Wage Statistics (OEWS) and the 2023 Employment Projections (EP) datasets from the U.S. Bureau of Labor Statistics (BLS), along with demographic data from the 2023 Current Population Survey (CPS).

Figs Fig. S2, Fig. S3, Fig. S4, Fig. S5 display the relationships between moral repugnance scores and occupational characteristics. For robustness, we include alternative measures from multiple data sources where available, finding similar patterns across datasets. These visualizations suggest consistent associations across alternative specifications of employment and wage variables.

Table S13 presents the results of ordinary least squares (OLS) regressions that analyze the influence of occupational characteristics on moral repugnance scores. Occupations with higher wages, fewer employees, higher proportions of women, and higher proportions of White workers are associated with significantly higher moral repugnance scores. These patterns suggest that public perceptions of AI's moral suitability vary systematically along demographic and economic lines, hinting at broader implications for AI's role in shaping labor market inequalities. Insofar as moral repugnance toward AI insulates occupations from AI, this protection is not distributed randomly but instead map onto existing social and economic divisions, with implications for how AI could reinforce or disrupt existing inequalities in the labor market.

S7 Supplementary Text: LLM-Based Occupation Classification

What kinds of work more likely to trigger moral repugnance when replaced or augmented by AI? In this section, we outline a theory-driven approach to identifying occupational categories where societal resistance to AI may be particularly pronounced. This framework draws on sociological insights into how labor acquires moral and cultural worth (Hochschild, n.d.; V. A. R. Zelizer, n.d.; Lamont, n.d.). It complements our empirically driven approach (see Supplementary Materials Section S8), in which we use LASSO regressions on occupational attributes to identify bottom-up predictors of moral repugnance.

To analyze patterns of societal resistance to AI across different types of work, we categorized occupations into six thematic domains based on the nature of their tasks and societal roles. These domains reflect conceptual categories of work that, according to sociological research, are likely to elicit moral, cultural, or ethical resistance to automation, whether because they involve emotional labor, sacred duties, or claims to authenticity (Velthuis, n.d.; Douglas, n.d.; Becker, n.d.). Specifically, we identified:

1. Care Work and Emotional Labor
2. Religious and Sacred Professions
3. Cultural and Artistic Professions
4. Legal and Ethical Decision-Making
5. Education and Moral Development
6. Healthcare (Especially Life-and-Death Decisions)

Although these six categories capture key domains in which AI may encounter significant cultural or ethical resistance, they are neither exhaustive nor perfectly discrete. Certain occupations may straddle multiple categories, and novel domains of moral repugnance could emerge as technology continues to evolve. Nevertheless, this theory-driven typology provides a clear framework for investigating how AI clashes with culturally charged forms of work, and it generates testable expectations about where resistance should be strongest. By confirming that occupations in these categories indeed exhibit higher scores on our moral repugnance measure, we also gain confidence in the validity of our methodological approach. In the next section, we complement this top-down strategy with a bottom-up analysis that uses a comprehensive set of occupational characteristics (from O*NET) to identify unexpected or hidden drivers of resistance to AI.

S7.1 Classification Approach

To classify occupations, we used ChatGPT (gpt-4o-2024-08-06) with a structured prompt to assess each occupation's description and assign membership scores across the six predefined categories. Each occupation received a continuous score ranging from 0 to 1 for each category, indicating the extent to which it fit the classification. A score of 0 meant the occupation did not belong to the category at all, while a score of 1 indicated full membership. Intermediate values (e.g., 0.5) represented partial relevance.

Code to reproduce the classification is provided in the OSF repository `scripts/chatgpt_occupation_classification.ipynb`. The full dataset of classified occupations is also available in the OSF repository (`data/processed/gpt_occupation_classification.csv`).

The classification was conducted using the following prompt:

Listing 1: Prompt Instructions for Classifying Occupations

```
# **Instructions for Classifying Occupations and Assigning Scores**

You have been provided an occupation that includes the title and
description. Your task is to classify the occupation into one or
more of the categories defined below, based on the nature of the
work as described. The occupation may belong to multiple categories.

## **Step 1: Classification into Categories**

Please assess whether the occupation fits into any of the 6
categories listed below. Each category has a short description to
help guide your classification.

**For each category**, assign a score between **0 and 1** that
reflects the 'occupations degree of membership in that category. A
score of **0** means the occupation does not belong in the category
at all, and a score of **1** means it fully belongs in the
category. You may assign any value in between (e.g., 0.5 if it
likely fits the category).

# **Categories for Classification**

## **Category 1: Care Work and Emotional Labor**
- **Definition**: This category includes roles that involve
caregiving, emotional support, and nurturing. These occupations
typically focus on providing care to individuals who are vulnerable
or in need of support. The roles often involve emotional labor and
relational work, requiring a deep connection between the caregiver
and the recipient.

## **Category 2: Religious and Sacred Professions**
- **Definition**: This category includes occupations that are
deeply tied to spiritual or religious authority and rituals. These
roles are often considered sacred or culturally significant, and
they play an important part in guiding people through spiritual or
moral aspects of life.

## **Category 3: Cultural and Artistic Professions**
```

```
- **Definition**: This category includes occupations centered
around the creation and expression of culturally significant and
symbolic works. These roles emphasize originality, emotional
resonance, and the personal involvement of the creator, where
authenticity and the human touch are considered vital to the value
of the work. The products or performances from these professions
are seen as essential expressions of identity, tradition, and human
experience.

## **Category 4: Legal and Ethical Decision-Making**
- **Definition**: This category includes occupations that involve
making complex legal or ethical decisions. These roles are focused
on upholding justice, ensuring fairness, and applying moral and
ethical reasoning in matters of legal and institutional importance.

## **Category 5: Education and Moral Development**
- **Definition**: This category includes occupations that are
involved in teaching and the development of moral, ethical, and
civic values. These roles focus on shaping future generations and
instilling knowledge and responsibility through education and
mentoring.

## **Category 6: Healthcare (Especially Life-and-Death Decisions)**
- **Definition**: This category includes occupations in healthcare,
particularly those involving life-and-death decision-making and the
direct care of patients. These roles require critical moral
judgment, compassion, and an understanding of complex medical or
ethical dilemmas.

# **Example of How to Classify and Rate**

1. **Classify the occupation**: Determine how well the occupation
fits into any of the 12 categories above. For each relevant
category, assign a score from **0 to 1** reflecting the strength of
its membership in that category.

2. **Output your response as valid JSON object**: Your output
should be in valid JSON format. Include the verbatim title of the
occupation. Here is an example:

<jsonl>
{
  "occupation_title": "Chief Executives",
  "care_work_and_emotional_labor": 0.0,
  "religious_and_sacred_professions": 0.0,
  "creative_and_artistic_professions": 0.0,
  "legal_and_ethical_decision_making": 0.0,
  "education_and_moral_development": 0.0,
  "healthcare_life_and_death_decisions": 0.0
}
```

```
</jsonl>

# **Final Notes**

Please review the occupation title and description carefully before
assigning scores. There are no right or wrong answers, but your
assessment should be based on a thorough reading of the
occupation's description and your understanding of the
responsibilities involved in each role. Output just the JSON and
nothing else.
```

S7.1.1 Binary Classification (Threshold at 0.5)

For analytical purposes, we dichotomized each category by setting a threshold of 0.5. Occupations scoring ≥ 0.5 were classified as belonging to the category. Occupations scoring < 0.5 were classified as not belonging.

S7.2 Descriptive Statistics

Table S14 presents summary statistics of raw scores for each category across the 940 O*NET occupations. This table includes the number and percentage of occupations meeting the binary classification threshold for membership. Fig. S6 shows the distribution of raw scores.

S7.3 Relationship Between Classification Scores and Moral Repugnance

To assess the relationship between moral repugnance and occupational classification, we computed the mean repugnance score for occupations classified within each category.

Table S15 presents descriptive statistics for repugnance scores for occupations classified as belonging to each category (membership score ≥ 0.5). Fig. S7 plots the mean repugnance scores by occupational category.

To visualize the relationship between category scores and moral repugnance, Fig. S8 presents scatterplots of moral repugnance vs. classification scores across the six categories.

S8 Supplementary Text: LASSO Selection of Occupational Characteristics Associated with Moral Repugnance Toward AI

To identify occupational characteristics predictive of AI resistance, we employed LASSO (Least Absolute Shrinkage and Selection Operator) regression with stability selection. We used stability selection with a cutoff of 0.75 (requiring features to appear in at least 75% of resampling iterations) and Per Family Error Rate (PFER) of 2 (limiting expected false positives to 2). All variables were standardized prior to analysis (mean = 0, SD = 1).

From 112 O*NET occupational characteristics, we identified 10 occupational characteristics predictive of moral repugnance toward AI (Table S16). We then fitted an OLS regression using these 10 selected predictors (Table S17). The model explained 49.2% of the variance in moral repugnance scores (adjusted $R^2 = 0.492$, $F(10,868) = 86.07$, $p < .001$). Eight predictors showed significant associations ($p < .01$): assisting and caring for others ($\beta = 0.116$), public speaking ($\beta = 0.107$), exposed to disease or infections ($\beta = 0.109$), deal with physically aggressive people ($\beta = 0.090$), stress tolerance ($\beta = 0.078$), and negative associations with degree of automation ($\beta = -0.104$), repairing electronic equipment ($\beta = -0.099$), and pace determined by equipment ($\beta = -0.070$).

List of Figures

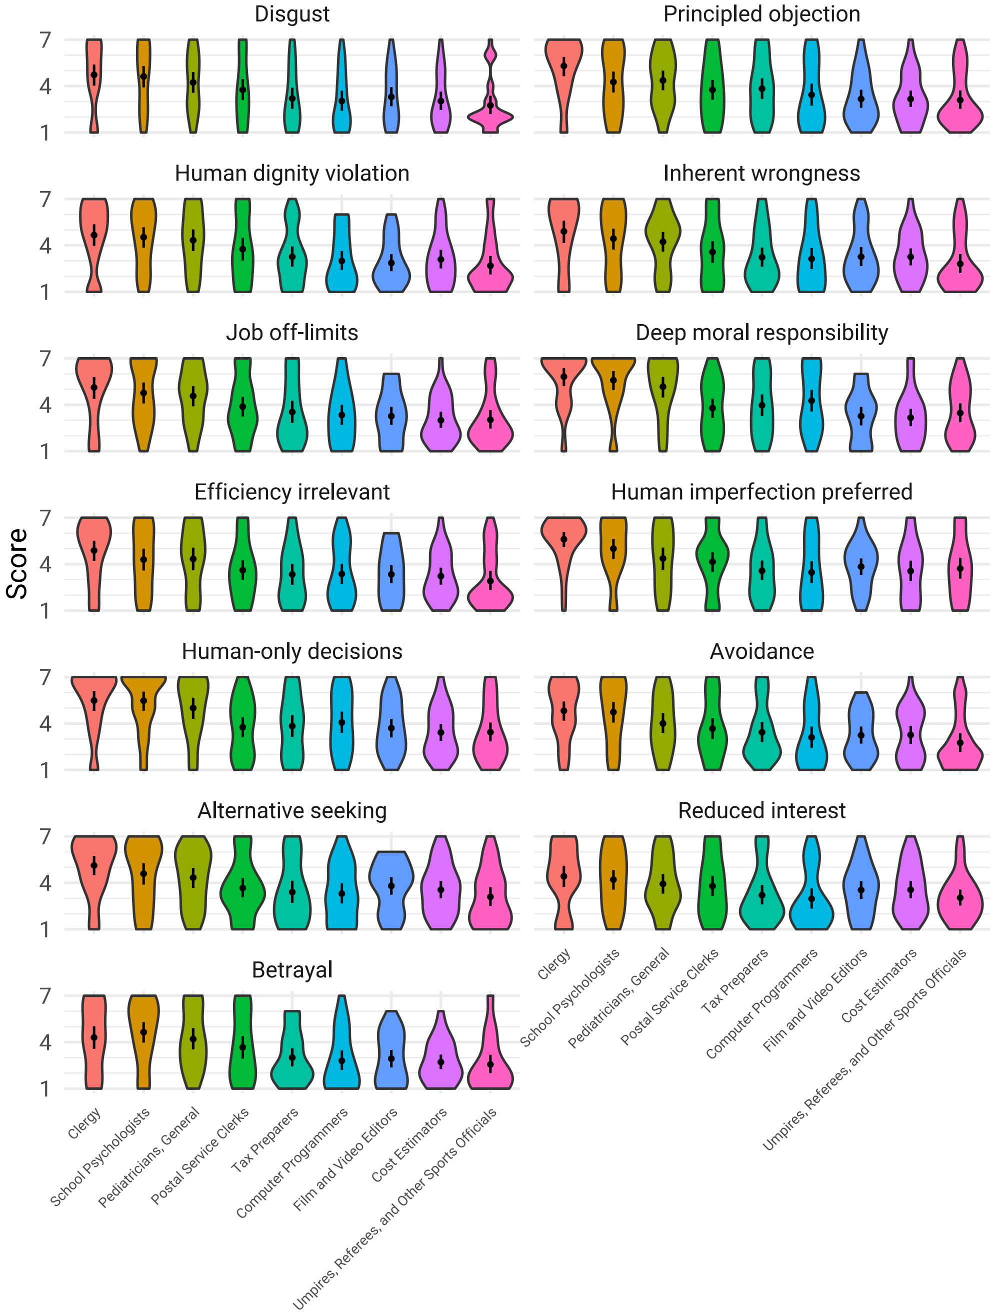

Fig. S1: Moral repugnance toward AI scale item score distributions, by occupation. Dots and vertical lines indicate means with bootstrapped 95% confidence intervals (10,000 iterations).

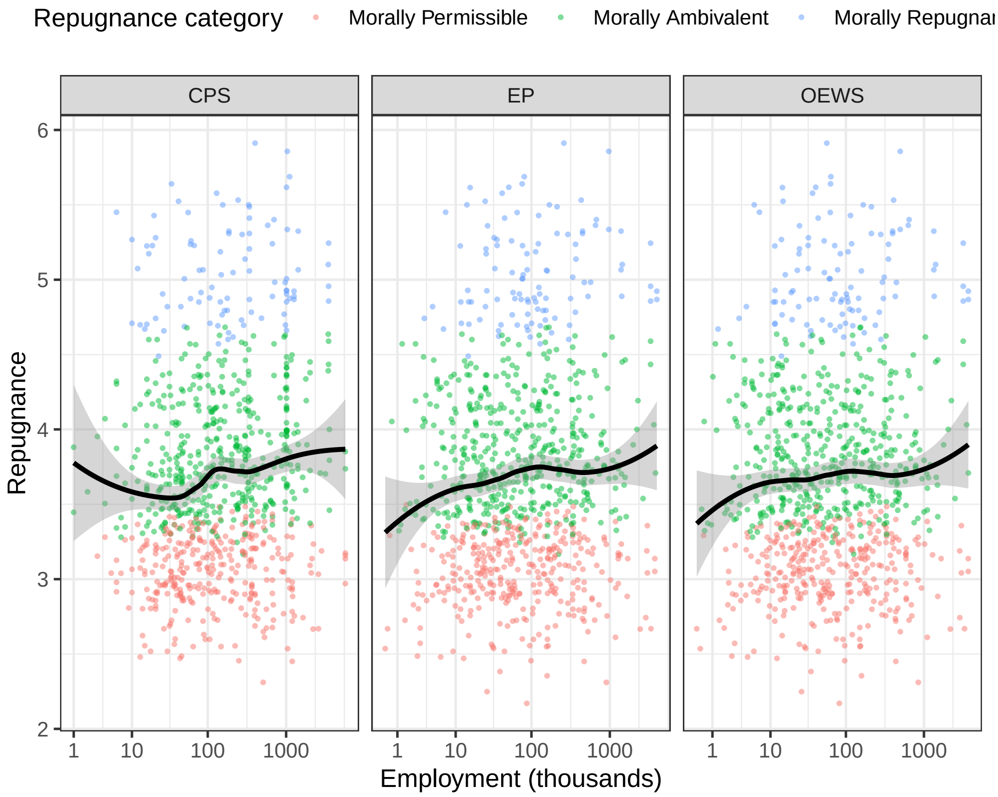

Fig. S2: Moral repugnance scores vs. employment. Moral repugnance scores are calculated at the occupational level by taking the mean individual moral repugnance score.

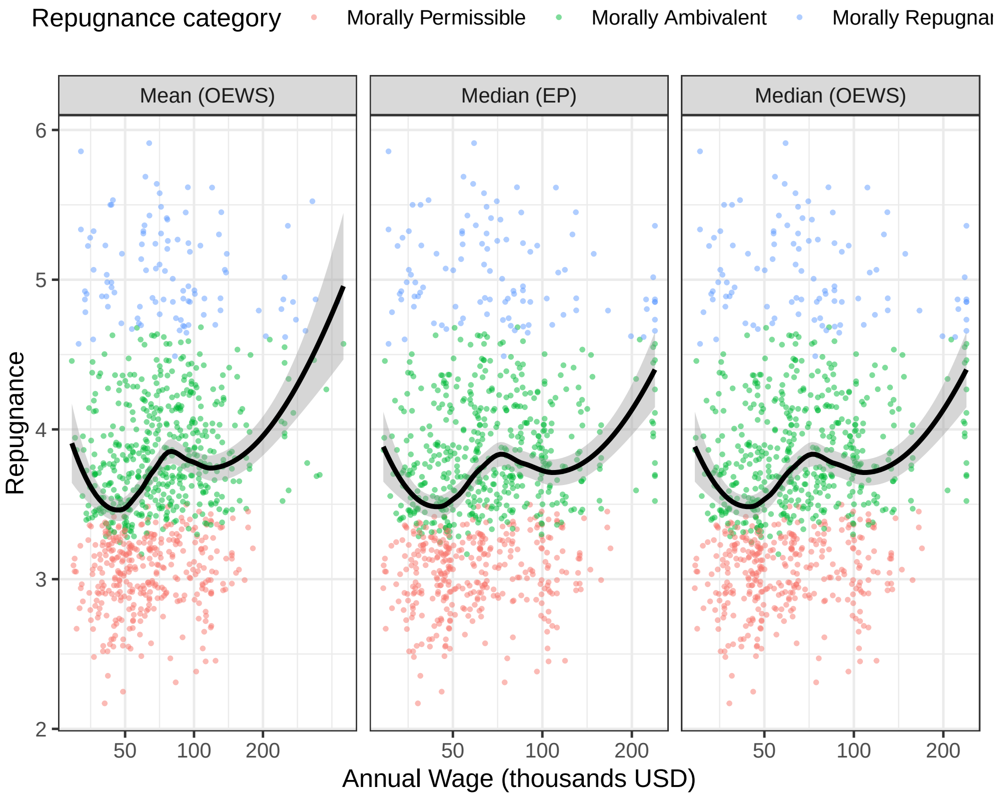

Fig. S3: Moral repugnance scores vs. annual wages. Moral repugnance scores are calculated at the occupational level by taking the mean individual moral repugnance score.

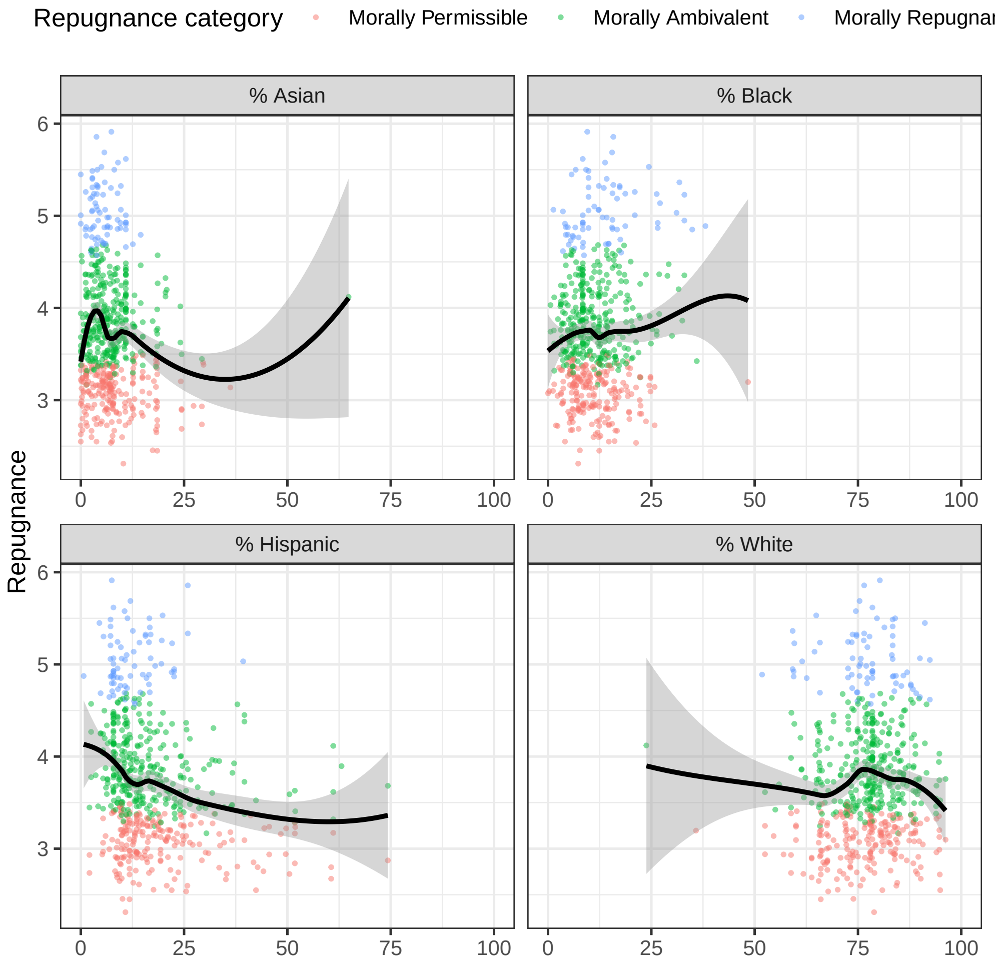

Fig. S4: Moral repugnance scores vs. ethnicity. Moral repugnance scores are calculated at the occupational level by taking the mean individual moral repugnance score.

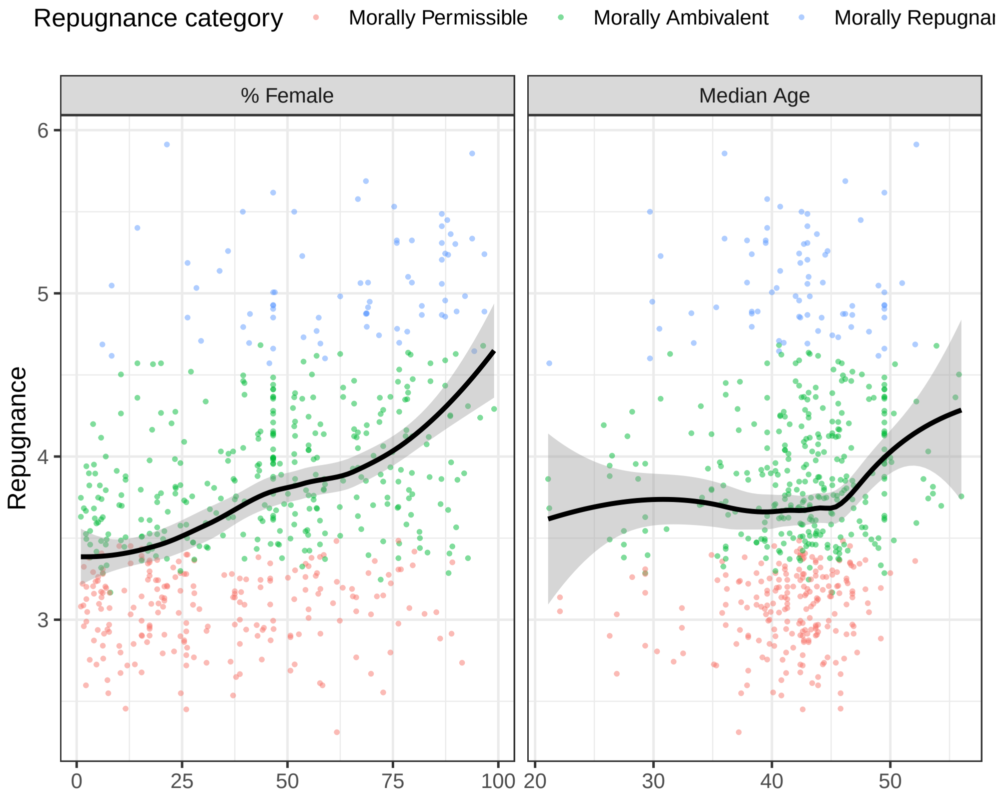

Fig. S5: Moral repugnance scores vs. sex and age. Moral repugnance scores are calculated at the occupational level by taking the mean individual moral repugnance score.

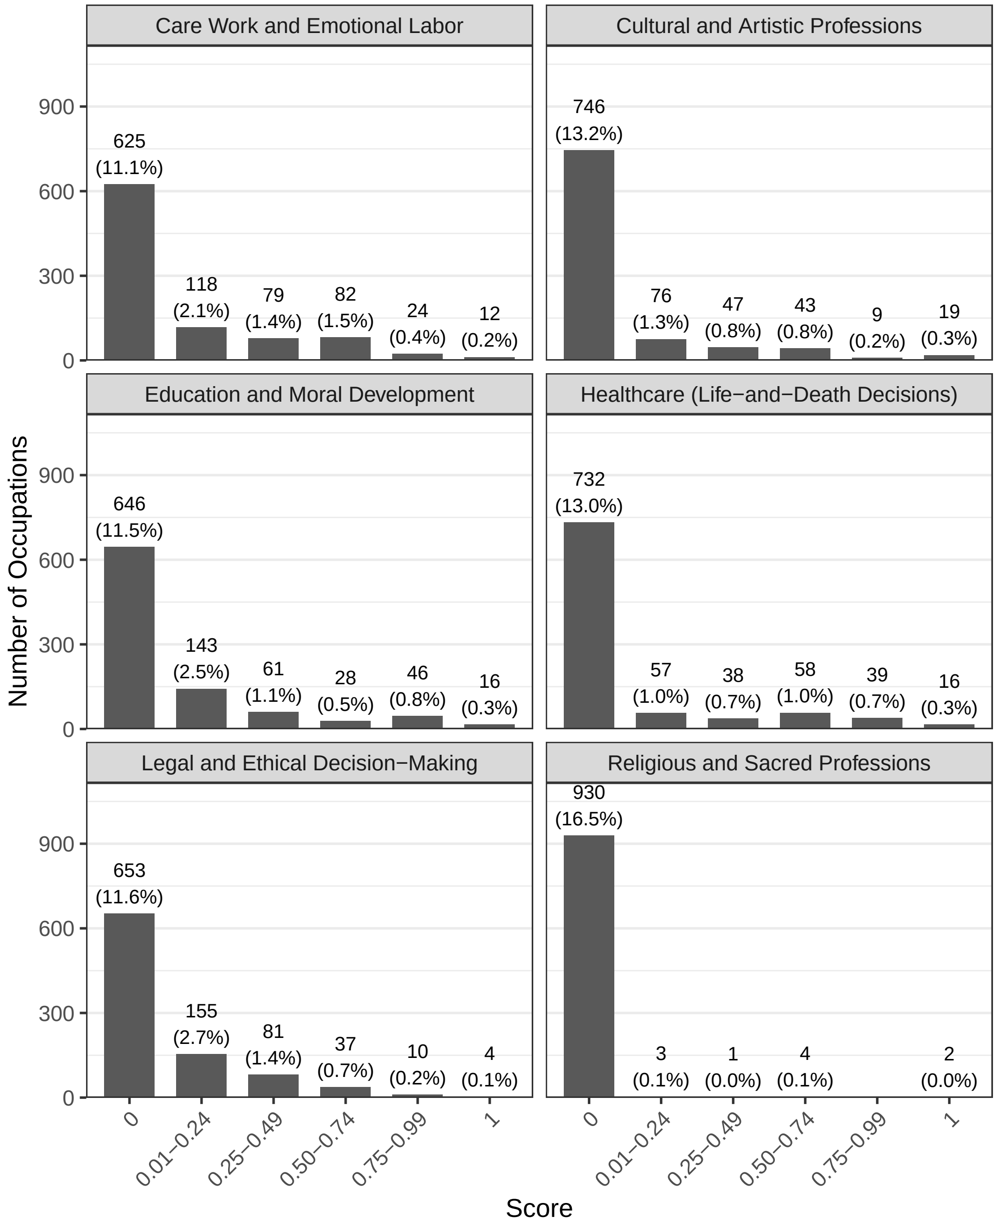

Fig. S6: Distribution of Occupational Membership Scores in Morally and Culturally Significant Work Categories

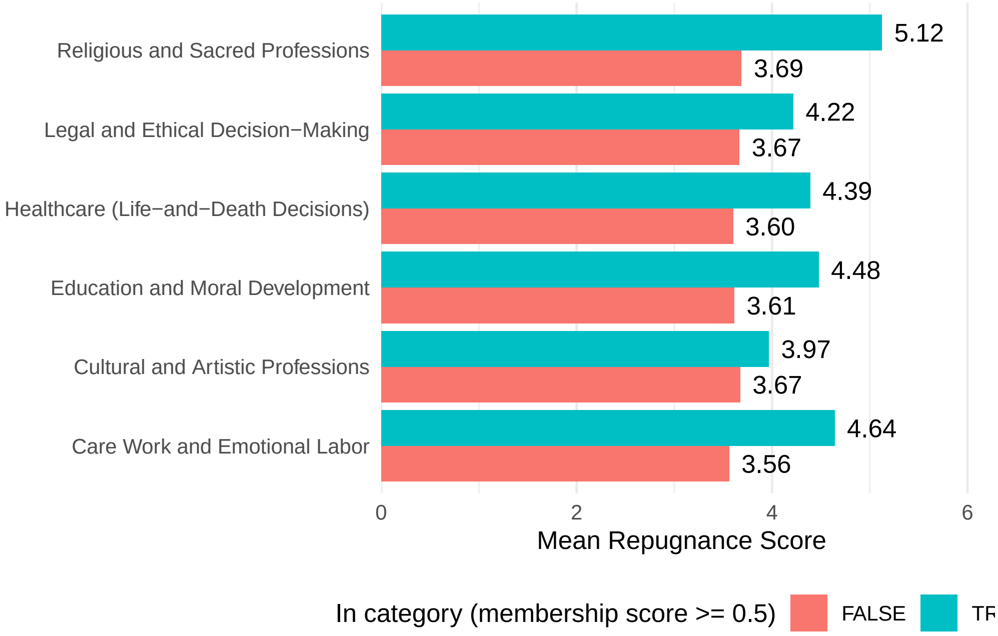

Fig. S7: Mean Moral Repugnance by Occupational Category.

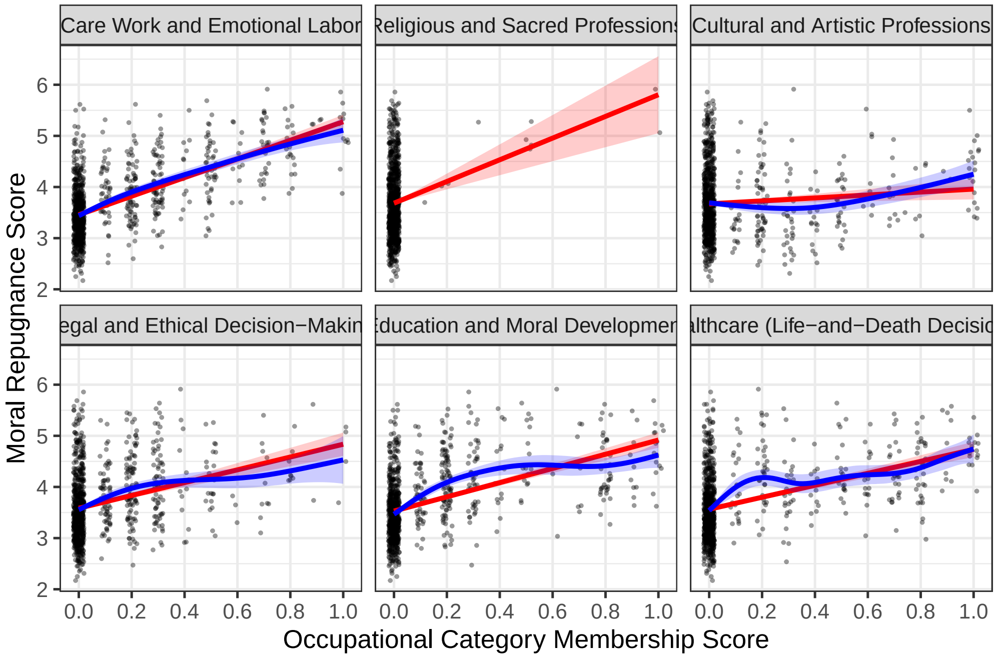

Fig. S8: Relationship Between Occupational Membership and Moral Repugnance. Red lines indicate linear regression fit with 95% confidence interval; blue lines indicate GAM fit with 95% confidence interval. Points are horizontally jittered for clarity.

List of Tables

Table S1: Main Study Participant Demographics.

	N	%
Age		
18-24	304	12.9
25-34	436	18.5
35-44	412	17.5
45-54	370	15.7
55+	835	35.4
Sex		
Male	1151	48.8
Female	1206	51.2
Ethnicity		
White	1482	62.9
Asian	150	6.4
Black	299	12.7
Mixed	261	11.1
Other	165	7.0
Political Affiliation		
Republican	668	28.3
Democrat	674	28.6
Independent	1015	43.1
Education		
Less than high school	20	0.8
High school graduate	304	12.9
Some college	569	24.1
2 year degree	298	12.6
4 year degree	774	32.8
Professional degree	352	14.9
Doctorate	40	1.7
AI Use		
Never	344	14.6
Less than once a month	537	22.8
A few times a month	481	20.4
About once per week	172	7.3
A few times per week	433	18.4
About once per day	157	6.7
Several times per day	233	9.9

Note: Values represent demographics of participants included in the final sample (N = 2,357) after applying exclusion criteria. Age was provided as an integer: mean = 45.2, SD = 16.0, median = 45, range = 18–85.

Table S2: Full distribution of participant demographic characteristics in the main study.

			Republican		Democrat		Independent		
Ethnicity	Age		Male	Female	Male	Female	Male	Female	All
White	18-24	N	15	10	20	23	40	40	148
		%	0.6	0.4	0.8	1.0	1.7	1.7	6.3
	25-34	N	22	24	31	33	62	59	231
		%	0.9	1.0	1.3	1.4	2.6	2.5	9.8
	35-44	N	24	24	32	29	62	58	229
		%	1.0	1.0	1.4	1.2	2.6	2.5	9.7
	45-54	N	33	30	29	28	50	47	217
		%	1.4	1.3	1.2	1.2	2.1	2.0	9.2
	55+	N	110	127	98	112	98	112	657
		%	4.7	5.4	4.2	4.8	4.2	4.8	27.9
Asian	18-24	N	6	2	5	3	8	3	27
		%	0.3	0.1	0.2	0.1	0.3	0.1	1.1
	25-34	N	3	10	2	2	6	9	32
		%	0.1	0.4	0.1	0.1	0.3	0.4	1.4
	35-44	N	3	4	3	3	6	9	28
		%	0.1	0.2	0.1	0.1	0.3	0.4	1.2
	45-54	N	3	12	3	2	7	6	33
		%	0.1	0.5	0.1	0.1	0.3	0.3	1.4
	55+	N	3	1	6	9	3	8	30
		%	0.1	0.0	0.3	0.4	0.1	0.3	1.3
Black	18-24	N	5	6	5	6	11	9	42
		%	0.2	0.3	0.2	0.3	0.5	0.4	1.8
	25-34	N	5	16	9	10	23	15	78
		%	0.2	0.7	0.4	0.4	1.0	0.6	3.3
	35-44	N	6	11	6	6	16	15	60
		%	0.3	0.5	0.3	0.3	0.7	0.6	2.5
	45-54	N	6	12	6	6	9	9	48
		%	0.3	0.5	0.3	0.3	0.4	0.4	2.0
	55+	N	13	9	11	17	5	16	71
		%	0.6	0.4	0.5	0.7	0.2	0.7	3.0
Mixed	18-24	N	18	2	6	6	11	12	55
		%	0.8	0.1	0.3	0.3	0.5	0.5	2.3
	25-34	N	8	6	8	6	15	15	58
		%	0.3	0.3	0.3	0.3	0.6	0.6	2.5
	35-44	N	10	6	6	6	15	12	55
		%	0.4	0.3	0.3	0.3	0.6	0.5	2.3
	45-54	N	6	6	6	6	9	10	43
		%	0.3	0.3	0.3	0.3	0.4	0.4	1.8
	55+	N	3	12	5	11	8	11	50
		%	0.1	0.5	0.2	0.5	0.3	0.5	2.1
Other	18-24	N	5	3	6	3	9	6	32
		%	0.2	0.1	0.3	0.1	0.4	0.3	1.4
	25-34	N	3	3	6	6	12	7	37
		%	0.1	0.1	0.3	0.3	0.5	0.3	1.6
	35-44	N	3	3	8	6	12	8	40
		%	0.1	0.1	0.3	0.3	0.5	0.3	1.7
	45-54	N	3	6	6	3	5	6	29
		%	0.1	0.3	0.3	0.1	0.2	0.3	1.2
	55+	N	2	5	3	6	5	6	27
		%	0.1	0.2	0.1	0.3	0.2	0.3	1.1
	All	N	318	350	326	348	507	508	2357
		%	13.5	14.8	13.8	14.8	21.5	21.6	100.0

Table S3: Comparison of main study participant demographics with U.S. population estimates (marginal distribution).

	Survey %	U.S. %	Diff. (pp)
Age			
18-24	12.9	11.9	1.0
25-34	18.5	17.8	0.7
35-44	17.5	16.6	0.9
45-54	15.7	16.3	-0.6
55+	35.4	37.5	-2.0
Sex			
Male	48.8	49.0	-0.2
Female	51.2	51.0	0.2
Ethnicity			
White	62.9	70.1	-7.2
Asian	6.4	5.9	0.5
Black	12.7	12.2	0.5
Mixed	11.1	5.7	5.4
Other	7.0	6.1	0.9

Table S4: Comparison of main study participant demographics with U.S. population estimates (joint distribution). U.S. population estimates are based on data from the American Community Survey (ACS), 2021.

Age	Ethnicity	Survey N	Survey %	U.S. %	Diff. (pp)
Male					
18-24	White	75	3.18	3.86	-0.68
	Black	21	0.89	0.86	0.03
	Asian	19	0.81	0.33	0.48
	Mixed	35	1.48	0.51	0.97
	Other	20	0.85	0.50	0.35
25-34	White	115	4.88	5.78	-0.90
	Black	37	1.57	1.23	0.34
	Asian	11	0.47	0.59	-0.12
	Mixed	31	1.32	0.65	0.67
	Other	21	0.89	0.73	0.16
35-44	White	118	5.01	5.53	-0.52
	Black	28	1.19	1.02	0.17
	Asian	12	0.51	0.56	-0.05
	Mixed	31	1.32	0.57	0.75
	Other	23	0.98	0.67	0.31
45-54	White	112	4.75	5.69	-0.94
	Black	21	0.89	0.96	-0.07
	Asian	13	0.55	0.48	0.07
	Mixed	21	0.89	0.45	0.44
	Other	14	0.59	0.54	0.05
55+	White	306	12.98	13.69	-0.71
	Black	29	1.23	1.69	-0.46
	Asian	12	0.51	0.79	-0.28
	Mixed	16	0.68	0.63	0.05
	Other	10	0.42	0.69	-0.27
Female					
18-24	White	73	3.10	3.69	-0.59
	Black	21	0.89	0.84	0.05
	Asian	8	0.34	0.33	0.01
	Mixed	20	0.85	0.49	0.36
	Other	12	0.51	0.46	0.05
25-34	White	116	4.92	5.60	-0.68
	Black	41	1.74	1.27	0.47
	Asian	21	0.89	0.62	0.27
	Mixed	27	1.15	0.63	0.52
	Other	16	0.68	0.65	0.03

35-44	White	111	4.71	5.37	-0.66
	Black	32	1.36	1.11	0.25
	Asian	16	0.68	0.63	0.05
	Mixed	24	1.02	0.55	0.47
	Other	17	0.72	0.60	0.12
45-54	White	105	4.45	5.59	-1.14
	Black	27	1.15	1.07	0.08
	Asian	20	0.85	0.55	0.30
	Mixed	22	0.93	0.46	0.47
	Other	15	0.64	0.51	0.13
55+	White	351	14.89	15.33	-0.44
	Black	42	1.78	2.17	-0.39
	Asian	18	0.76	0.99	-0.23
	Mixed	34	1.44	0.73	0.71
	Other	17	0.72	0.76	-0.04

Table S5: Descriptive Statistics at the Occupation Level in the Main Study.

	Mean	SD	Min	Median	Max
Responses	25.07	1.55	18.00	25.00	28.00
Moral repugnance toward AI	3.69	0.69	2.17	3.58	5.91
Moral Repugnance Scale Items					
3. Human dignity violation	3.47	0.69	2.00	3.36	5.80
5. Job off-limits	3.74	0.74	2.08	3.60	6.23
6. Deep moral responsibility	3.82	0.90	1.96	3.68	6.32
8. Human imperfection preferred	3.97	0.70	2.42	3.89	6.29
9. Human-only decisions	3.99	0.78	2.15	3.88	6.28
10. Avoidance	3.51	0.64	2.00	3.41	5.92
13. Betrayal	3.36	0.66	1.89	3.24	5.65
Support for AI					
Current AI: Augmentation	4.53	0.62	2.50	4.55	6.35
Current AI: Automation	3.21	0.66	1.39	3.20	5.46
Advanced AI: Augmentation	4.69	0.60	2.29	4.74	6.12
Advanced AI: Automation	3.74	0.70	1.82	3.77	5.87
Participant Demographics					
Age	45.21	3.11	33.42	45.21	53.78
Female	0.51	0.10	0.22	0.52	0.80
Bachelor's degree	0.49	0.10	0.22	0.50	0.84
Daily AI user	0.17	0.08	0.00	0.16	0.42
Ethnicity					
White	0.63	0.09	0.38	0.63	0.89
Asian	0.06	0.05	0.00	0.05	0.26
Black	0.13	0.07	0.00	0.12	0.35
Mixed	0.11	0.06	0.00	0.11	0.33
Other	0.07	0.05	0.00	0.07	0.30
Political Affiliation					
Democrat	0.29	0.09	0.00	0.28	0.64
Independent	0.43	0.10	0.17	0.43	0.72
Republican	0.28	0.09	0.00	0.28	0.61

All values represent averages at occupation level. Categorical variables were transformed into binary variables. Bachelor's degree = 1 for participants with a four-year degree, professional degree, or doctorate. Daily AI user = 1 for participants who reported using AI ”about once per day” or ”several times per day.”

Table S6: Occupations Selected for Moral Repugnance Toward AI Scale Development and Validation Study. Occupations selected to represent a range of public moral repugnance and machine learning suitability scores, based on data from Brynjolfsson, Mitchell, and Rock (2018). Occupations are categorized by quintiles of repugnance and suitability for machine learning, illustrating diverse roles across the spectrum of public attitudes and technical feasibility for AI automation.

Repugnance quintile	Suitability for Machine Learning quintile	Occupation
1	1	Cost Estimators
	3	Computer Programmers
	5	Tax Preparers
3	1	Umpires, Referees, and Other Sports Officials
	3	Film and Video Editors
	5	Postal Service Clerks
5	1	Pediatricians, General
	3	School Psychologists
	5	Clergy

Table S7: Descriptive Statistics for Moral Repugnance Toward AI Scale Items, by Occupation.

Item	Mean	SD
Clergy (N=33)		
Disgust	4.73	2.05
Principled objection	5.30	1.83
Human dignity violation	4.67	2.09
Inherent wrongness	4.91	2.14
Job off-limits	5.12	2.07
Deep moral responsibility	5.82	1.72
Efficiency irrelevant	4.88	1.92
Human imperfection preferred	5.61	1.48
Human-only decisions	5.48	1.89
Avoidance	4.82	1.84
Alternative seeking	5.12	1.87
Reduced interest	4.42	2.00
Betrayal	4.30	2.13
Computer Programmers (N=30)		
Disgust	4.62	2.06
Principled objection	4.26	2.06
Human dignity violation	4.53	1.99
Inherent wrongness	4.44	2.09
Job off-limits	4.76	2.03
Deep moral responsibility	5.59	1.84
Efficiency irrelevant	4.29	2.08
Human imperfection preferred	5.00	1.91
Human-only decisions	5.47	1.86
Avoidance	4.74	2.03
Alternative seeking	4.59	2.08
Reduced interest	4.21	1.92
Betrayal	4.65	2.04
Cost Estimators (N=31)		
Disgust	4.23	1.87
Principled objection	4.37	1.79
Human dignity violation	4.33	1.99
Inherent wrongness	4.23	1.81
Job off-limits	4.57	1.83
Deep moral responsibility	5.17	1.90
Efficiency irrelevant	4.33	2.06
Human imperfection preferred	4.37	2.03
Human-only decisions	5.00	1.97
Avoidance	4.00	1.86
Alternative seeking	4.33	1.86

Reduced interest	3.93	1.80
Betrayal	4.20	1.95
Film and Video Editors (N=30)		
Disgust	3.76	2.09
Principled objection	3.76	1.95
Human dignity violation	3.76	2.15
Inherent wrongness	3.58	2.08
Job off-limits	3.88	1.92
Deep moral responsibility	3.79	1.85
Efficiency irrelevant	3.61	1.89
Human imperfection preferred	4.15	1.89
Human-only decisions	3.76	1.87
Avoidance	3.67	1.95
Alternative seeking	3.67	1.85
Reduced interest	3.79	1.98
Betrayal	3.67	2.13
Pediatricians, General (N=30)		
Disgust	3.20	1.90
Principled objection	3.83	1.93
Human dignity violation	3.27	1.82
Inherent wrongness	3.23	1.79
Job off-limits	3.53	2.06
Deep moral responsibility	3.97	1.97
Efficiency irrelevant	3.33	1.83
Human imperfection preferred	3.57	1.79
Human-only decisions	3.83	2.00
Avoidance	3.43	1.83
Alternative seeking	3.40	2.01
Reduced interest	3.20	1.85
Betrayal	3.00	1.66
Postal Service Clerks (N=33)		
Disgust	3.03	1.81
Principled objection	3.43	2.06
Human dignity violation	3.00	1.74
Inherent wrongness	3.13	1.93
Job off-limits	3.33	1.83
Deep moral responsibility	4.27	2.05
Efficiency irrelevant	3.37	1.88
Human imperfection preferred	3.47	2.08
Human-only decisions	4.07	1.96
Avoidance	3.10	1.92
Alternative seeking	3.30	1.86
Reduced interest	2.97	1.85
Betrayal	2.80	1.83

School Psychologists (N=34)		
Disgust	3.30	1.78
Principled objection	3.17	1.64
Human dignity violation	2.87	1.53
Inherent wrongness	3.27	1.72
Job off-limits	3.27	1.66
Deep moral responsibility	3.27	1.68
Efficiency irrelevant	3.33	1.67
Human imperfection preferred	3.83	1.51
Human-only decisions	3.70	1.68
Avoidance	3.23	1.52
Alternative seeking	3.80	1.63
Reduced interest	3.53	1.66
Betrayal	2.93	1.62
Tax Preparers (N=30)		
Disgust	3.03	1.70
Principled objection	3.16	1.59
Human dignity violation	3.10	1.70
Inherent wrongness	3.26	1.57
Job off-limits	3.00	1.51
Deep moral responsibility	3.16	1.66
Efficiency irrelevant	3.23	1.56
Human imperfection preferred	3.55	1.89
Human-only decisions	3.42	1.59
Avoidance	3.26	1.67
Alternative seeking	3.55	1.65
Reduced interest	3.55	1.65
Betrayal	2.71	1.35
Umpires, Referees, ... (N=32)		
Disgust	2.75	1.68
Principled objection	3.09	1.71
Human dignity violation	2.69	1.73
Inherent wrongness	2.81	1.77
Job off-limits	3.03	1.75
Deep moral responsibility	3.47	1.76
Efficiency irrelevant	2.91	1.82
Human imperfection preferred	3.72	1.97
Human-only decisions	3.44	1.85
Avoidance	2.75	1.74
Alternative seeking	3.09	1.77
Reduced interest	3.03	1.49
Betrayal	2.56	1.72

Table S8: Moral Repugnance Toward AI Scale Development and Validation: Final Items and Factor Loadings.

Item	Factor Loading
Human Dignity Violation	0.87
Job Off Limits	0.90
Deep Moral Responsibility	0.84
Human Imperfection Preferred	0.80
Human Only Decisions	0.83
Avoidance	0.89
Betrayal	0.84

Table S9: Distribution of occupations across moral categories (count and percentage) by bootstrap CI method.

Repugnance classification	BCa	Percentile	Student-t
Morally Permissible	346 (36.8%)	361 (38.4%)	328 (34.9%)
Morally Ambivalent	482 (51.3%)	463 (49.3%)	508 (54.0%)
Morally Repugnant	112 (11.9%)	116 (12.3%)	104 (11.1%)

Table S10: OLS Regression of Moral Repugnance on Technical Suitability for AI.

	(1)
Technical suitability	3.254***
	(0.271)
Technical suitability^2	−3.041***
	(0.257)
(Intercept)	3.169***
	(0.054)
Num.Obs.	923
R2	0.136
R2 Adj.	0.134

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table S11: Distribution of Occupations by Technical Suitability and Moral Repugnance

Technical Suitability	Moral Repugnance	N	%
Low	Low	308	32.8%
	High	93	9.9%
High	Low	358	38.1%
	High	164	17.4%
(Missing)	Low	3	0.3%
	High	14	1.5%

Table S12: Examples of Occupations by Technical Suitability and Moral Repugnance Classification

Occupation	Technical Suitability Score	Repugnance Score
Low Technical Suitability, Low Moral Repugnance		
Paper Goods Machine Setters, Operators, and Tenders	0.00	2.56
Highway Maintenance Workers	0.05	2.60
Shoe Machine Operators and Tenders	0.05	2.48
Segmental Pavers	0.08	2.55
Extruding and Forming Machine Setters, Operators, and Tenders, Synthetic...	0.05	2.62
Low Technical Suitability, High Moral Repugnance		
Ambulance Drivers and Attendants, Except Emergency Medical Technicians	0.09	5.23
Embalmers	0.04	4.74
Hairdressers, Hairstylists, and Cosmetologists	0.09	4.98
Home Health Aides	0.07	4.87
Athletes and Sports Competitors	0.00	5.52
High Technical Suitability, Low Moral Repugnance		
Radio Frequency Identification Device Specialists	0.90	2.46
Search Marketing Strategists	1.00	2.31
Geographic Information Systems Technologists and Technicians	1.00	2.45
Transportation Planners	0.95	2.38
Cartographers and Photogrammetrists	1.00	2.47
High Technical Suitability, High Moral Repugnance		
Education and Childcare Administrators, Preschool and Daycare	0.88	5.69
Psychology Teachers, Postsecondary	0.76	5.62
Administrative Law Judges, Adjudicators, and Hearing Officers	0.86	5.62
School Psychologists	0.79	5.45
Psychiatrists	0.83	5.36

Table S13: OLS Regression of Occupational Demographics Predicting Moral Repugnance Toward AI.

	(1)
Log Mean Wage	0.200*
	(0.078)
Log Total Employment	−0.031*
	(0.015)
Log Median Age	0.118
	(0.190)
Pct. Female	0.010***
	(0.001)
Pct. Black	0.005
	(0.004)
Pct. Asian	−0.030***
	(0.004)
Pct. Hispanic	−0.010***
	(0.003)
Intercept	1.328
	(0.958)
Num.Obs.	601
R2	0.257
R2 Adj.	0.248

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Standard errors in parentheses. *p* < 0.05, **p** < 0.01, ***p*** < 0.001.

Table S14: Descriptive Statistics of Occupational Membership in Morally and Culturally Significant Work Categories

	Mean	SD	Min	Median	Max	In category	% in category
Care Work and Emotional Labor	0.13	0.23	0.00	0.00	1.00	118	12.6%
Religious and Sacred Professions	0.01	0.06	0.00	0.00	1.00	6	0.6%
Cultural and Artistic Professions	0.08	0.21	0.00	0.00	1.00	71	7.6%
Legal and Ethical Decision-Making	0.09	0.17	0.00	0.00	1.00	51	5.4%
Education and Moral Development	0.12	0.24	0.00	0.00	1.00	90	9.6%
Healthcare (Life-and-Death Decisions)	0.11	0.25	0.00	0.00	1.00	113	12.0%

Table S15: Descriptive Statistics of Repugnance Scores by Occupational Category

Occupational Category	Mean	SD	Min	Median	Max
In category					
Legal and Ethical Decision-Making	4.22	0.63	3.04	4.17	5.62
Education and Moral Development	4.48	0.63	3.04	4.37	5.91
Care Work and Emotional Labor	4.64	0.61	3.05	4.68	5.91
Healthcare (Life-and-Death Decisions)	4.39	0.59	2.98	4.33	5.64
Cultural and Artistic Professions	3.97	0.62	2.80	3.91	5.52
Religious and Sacred Professions	5.12	0.43	4.74	4.99	5.91
Not in category					
Care Work and Emotional Labor	3.56	0.59	2.17	3.47	5.62
Religious and Sacred Professions	3.69	0.69	2.17	3.58	5.86
Cultural and Artistic Professions	3.67	0.70	2.17	3.54	5.91
Legal and Ethical Decision-Making	3.67	0.69	2.17	3.53	5.91
Education and Moral Development	3.61	0.65	2.17	3.49	5.86
Healthcare (Life-and-Death Decisions)	3.60	0.65	2.17	3.47	5.91

Table S16: O*NET Occupational Characteristics Selected via LASSO Stability Selection. Selection probabilities indicate the proportion of resampling iterations in which each characteristic was selected.

Occupational Characteristic	Selection Probability
Assisting and caring for others	1.00
Deal with physically aggressive people	1.00
Pace determined by speed of equipment	1.00
Exposed to disease or infections	0.98
Public speaking	0.97
Degree of automation	0.94
Repairing and maintaining electronic equipment	0.90
Concern for others	0.86
Self control	0.85
Stress tolerance	0.82

Table S17: OLS Regression of LASSO-Selected O*NET Occupational Characteristics Predicting Moral Repugnance Toward AI. All coefficients are standardized.

	(1)
Assisting and caring for others	0.116***
	(0.031)
Deal with physically aggressive people	0.090***
	(0.021)
Pace determined by speed of equipment	−0.070**
	(0.021)
Exposed to disease or infections	0.109***
	(0.026)
Public speaking	0.107***
	(0.019)
Degree of automation	−0.104***
	(0.018)
Repairing and maintaining electronic equipment	−0.099***
	(0.019)
Concern for others	0.001
	(0.033)
Self control	0.044
	(0.033)
Stress tolerance	0.078**
	(0.027)
(Intercept)	3.670***
	(0.016)
Num.Obs.	879
R2	0.498
R2 Adj.	0.492

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$