CSEF
Centre for Studies in Economics and Finance

WORKING PAPER NO. 794

Human Enhancement Technologies: A Survey Experiment on Private Demand and Governance Preferences.

Giovanni Immordino, Mario Macis, Immacolata Marino, and Fabrizio Panebianco

August 2026

University of Naples Federico II
University of Salerno
Bocconi University, Milan

CSEF - Centre for Studies in Economics and Finance
DEPARTMENT OF ECONOMICS AND STATISTICS – UNIVERSITY OF NAPLES FEDERICO II
80126 NAPLES - ITALY
Tel. and fax +39 081 675372 – e-mail: csef@unina.it
ISSN: 2240-9696

CSEF
Centre for Studies in Economics and Finance

WORKING PAPER NO. 794

Human Enhancement Technologies: A Survey Experiment on Private Demand and Governance Preferences.

Giovanni Immordino[*], Mario Macis[†], Immacolata Marino[‡], and Fabrizio Panebianco[§]

Abstract
When a new technology promises large private benefits but may impose social costs that markets do not price, demand need not reveal how citizens want it governed. We examine this using a nationally representative U.S. survey experiment (N=5,556) on human enhancement technologies (HET). The experiment randomizes benefit domain, mechanism, heritability, purpose, and risk across vignettes; for each respondent's assigned vignette, we elicit stated adoption, preferred regulation, and ethical and societal concerns. Overall, about 53% would adopt. Framing the technology as enhancing rather than restorative lowers adoption by about five percentage points, as much as a severe side-effect profile. About 28% would not adopt at any benefit. This refusal is driven overwhelmingly by the enhancing framing rather than by risk, consistent with a non-compensatory constraint for a substantial subgroup. Most who would adopt still favor strict regulation, and most who would never adopt do not wish to forbid others from doing so. Productivity enhancement generates the most ethical concern of any attribute but attracts the least regulation, and respondents favor subsidizing rather than taxing its adoption, consistent with a concern about access rather than safety. Private demand is therefore an unreliable guide to the governance citizens want, and the divergence we document provides a basis for regulators seeking to align the direction of technical change with societal values and priorities.

Keywords: Human Enhancement Technologies, Survey Experiment, Private–Policy Alignment, Distributional Effects and Inequality, Productivity, Moral Foundations and Ethics.

JEL Classification: I18, O33, D81, D63, J24, L51.

Acknowledgments: This study was reviewed and approved by the "Comitato Etico Per Le Attivit`a Non Biomediche" of the Universit`a degli Studi di Napoli Federico II (PG/2024/0124289) and the Johns Hopkins University Homewood IRB (HIRB00019474). It was pre-registered in the American Economic Association's RCT Registry (AEARCTR-0016972). We would like to thank Alberto Bisin, Ruth Faden, Jeffrey Kahn, Juliana Rebora, and seminar participants at the University of Lugano, the University of North South Wales, the University of Palermo, the Petralia Workshop, the University of Technology Sydney, the Johns Hopkins Berman Institute of Bioethics, the Johns Hopkins Technology and Society Initiative, the Hopkins Business of Health Initiative, and the CEPR Workshop on Frontiers in Measurement and Survey Methods for helpful comments and suggestions. We gratefully acknowledge financial support from the Italian Ministry of University and Research (PRIN PNRR Project No. P2022XJW4T).

[*] University of Naples Federico II and CSEF. E-mail: giovanni.immordino@unina.it

[†] Johns Hopkins University and NBER.

[‡] University of Naples Federico II and CSEF. E-mail: immacolata.marino@unina.it

[§] Università Cattolica del Sacro Cuore di Milano.

1 Introduction

A class of technologies once confined to science fiction is entering clinical and commercial reality. In December 2023, the FDA approved the first therapy based on CRISPR, a technology for targeted DNA editing, for sickle cell disease (U.S. Food and Drug Administration 2023). In January 2024, Neuralink implanted its brain-computer interface (BCI) in a human for the first time, a high-profile entrant in a field where implanted BCIs have been trialed in humans since the 2000s (Guarino 2024; Hochberg et al. 2006). Earlier, in 2018, the controversial CRISPR editing of human embryos had already shown that technologies developed to treat disease can, in principle, be redirected toward inheritable human enhancement, leading to the birth of children carrying genetic edits that did not replicate any known natural variant (Cyranoski and Ledford 2018).[1]

These innovations are leading examples of *human enhancement technologies* (HET), and they share a distinctive economic profile. They promise large private benefits in health, cognition, and well-being, but carry potential costs that private transactions may not fully internalize: irreversibility and tail risks, distributional consequences in education and labor markets, social competition, limits to meaningful consent, and intergenerational spillovers from inheritable interventions.[2] For technologies of this kind, the social problem is not only whether they are safe enough to allow, but which trajectories of technical change should be encouraged, constrained, or prohibited (Acemoglu and Restrepo 2020; Acemoglu and Johnson 2023).

In the standard economic framework, private demand is the natural guide for such choices: willingness to pay reveals how much people value a technology, and a well-designed corrective policy prices any externality and lets the market do the rest. This logic presumes that the considerations driving private adoption and those driving citizens' preferred governance are, if not identical, at least aligned closely enough that demand is a reliable signal of social acceptability. Whether that presumption holds is an empirical question, and one that may be answered differently across human enhancement technologies depending on their specific features. People may bring two quite different kinds of consideration to these technologies, some that enter an ordinary weighing of costs and benefits, and some that operate as constraints standing outside it, and these need not affect private adoption and preferred governance in the same way. The main goal of this paper is to characterize and quantify that potential gap, which we call the *demand–governance wedge*: the divergence, if any, between the attributes that drive private adoption and those that drive citizens' preferred governance.

[1] Several enhancement-adjacent platforms are already used clinically for therapeutic purposes (e.g., non-invasive neurostimulation and implanted stimulators), and gene-editing therapies have entered clinical care; BCIs are also advancing toward clinical use for restoration of communication and motor function.

[2] Public debate over human enhancement as well as opinion polls often emphasizes unequal access, coercive social competition ("arms races"), the impossibility of consent for future generations, and the potential for misuse. See Funk, Kennedy, and Sciupac 2016; Funk and Hefferon 2018; Rainie et al. 2022; Siegrist 2000.

These two kinds of consideration correspond to a key question in the debate over HET: whether support and resistance reflect *performance-based* considerations (safety, reversibility, expected side effects) or instead *principled* constraints not reducible to standard performance-based tradeoffs, even when performance is favorable. The distinction is economically meaningful because it yields different predictions about how adoption and governance respond to changes in technology attributes. If resistance is performance-based, making technologies safer or improving oversight should substantially raise adoption and lower opposition. If resistance is principled, reluctance should persist even under low-risk, high-benefit profiles, for example when a technology is made safer or its benefits increased.[3]

To study these questions, we designed and fielded a large, nationally representative U.S. survey experiment (N = 5,556), quota-matched to the 2024 American Community Survey on gender, age, race and ethnicity, education, and geography. Each respondent evaluates one randomized vignette describing a concrete enhancement technology that orthogonally varies four core attributes: (i) the *benefit domain* (cosmetic, health, or productivity), three broad classes of enhancement that are actively being developed and debated; (ii) the *mechanism and heritability* (non-genetic implant; genetic non-inheritable; genetic inheritable); (iii) the *purpose* (restorative vs. enhancing), which moves the technology across the treatment–enhancement boundary while holding the underlying capability change fixed; and (iv) the *side-effect severity* (mild and temporary vs. severe and long-lasting). The orthogonal assignment of the four core attributes allows us to ask whether the same attribute moves adoption, governance, and ethical and societal concerns in the same or a different way. We elicit private demand (adoption and, conditional on adoption, willingness to pay), governance preferences (preferred legal status and regulatory stringency, including outright prohibition, and whether adoption of the technology should be taxed or subsidized), ethical and societal concerns, and open-ended rationales; in health and productivity scenarios we also elicit reservation-benefit thresholds among initial non-adopters and identify "never adopters" who refuse regardless of benefit. Within a subset of productivity scenarios, we additionally randomize the distributional incidence of the gains (whether they accrue primarily to high-skilled workers, to low-skilled workers, or broadly across skill levels) to test whether adoption, governance, and ethical responses depend on who benefits, in addition to the general effect of productivity enhancement.

Our reliance on a survey experiment reflects two features of these technologies: they are at an early stage, and they are medical, costly, and often irreversible. Because many HET are not yet diffused, evidence on attitudes necessarily relies on hypothetical scenarios rather than realized choices; and because real-stakes experiments are difficult or impossible even where the technologies exist, a survey experiment is the natural tool for measuring attitudes

[3]Related frameworks emphasize protected values, taboo trade-offs, and moral constraints in contested domains (e.g., Friis and Riley 2025; Elías, Lacetera, and Macis 2019; Baron and Spranca 1997; Tetlock et al. 2000; A. E. Roth 2007; Cushman, Young, and Hauser 2006; Fiske and Tetlock 1997).

while diffusion is still limited.[4] Studying attitudes at this upstream stage also provides *ex ante* feedback on which trajectories are seen as legitimate and which trigger distributive or moral resistance, giving policymakers and innovators a chance to shape standards and acceptable use cases before path dependence sets in (Acemoglu and Restrepo 2020; Acemoglu and Johnson 2023).

We organize the results around three main findings. *First*, about 53% of respondents would adopt; demand is highest for health, lowest for cosmetic, and intermediate for productivity. Severe side effects and an enhancing framing each reduce adoption by roughly five percentage points (about 10% of the mean), and the two effects are essentially equal in magnitude. Purpose, a reframing of an otherwise identical capability change, depresses demand as much as a salient prospect of severe harm. Notably, genetic mechanisms do not suppress private adoption; if anything, they raise it modestly, suggesting that for one's own use, genetic mechanism is not a deterrent. On the private adoption margin, domain, risk, and purpose each have larger effects than genetic mechanism or heritability.

Second, about 28% of respondents in the health and productivity scenarios are categorical never-adopters who decline at any benefit. This refusal is driven overwhelmingly by the enhancing framing rather than by risk (a raw gap of roughly 9 percentage points for purpose against 3.5 for risk) and is essentially unmoved by genetic mechanism or heritability. This pattern is consistent with a non-compensatory protected value (Baron and Spranca 1997; Tetlock et al. 2000): for a substantial subgroup, categorical refusal persists regardless of how favorable the performance profile is. Text analysis of open-ended rationales confirms this interpretation: never-adopters invoke moral boundaries, human nature, and dependence far more than adopters, who reason in terms of benefit and usefulness. Unlike on the adoption margin, where purpose and risk have nearly equal effects, categorical refusal is driven overwhelmingly by purpose, consistent with a non-compensatory constraint that no increase in benefit can offset.

Third, the attributes that drive governance preferences are not always the same as those that drive adoption: we document and characterize a demand–governance wedge. The modal governance preference, in every condition, is strict-but-legal regulation (legal but tightly controlled, e.g., clinical trials and ethical oversight; about 44% in the pooled sample), followed by moderate regulation (legal with standards for safety, efficacy, and informed consent; about 32%); both laissez-faire and outright prohibition are rare, and prohibition (around 6%) is barely responsive to any attribute. By the test of A. E. Roth 2007, human enhancement does *not* present as a repugnant transaction: even respondents who would never adopt do not wish to forbid others from doing so, expressing concern through support for oversight within legal access rather than bans. Risk, the enhancing framing, and inheritable interventions all shift

[4] This contrasts with other emerging technologies such as generative AI, which is widely accessible and low-cost to try; this allowed Bursztyn et al. (2025) to embed actual adoption decisions within an experimental survey.

preferences toward stricter oversight, with non-inheritable interventions leaving regulation unmoved while inheritable ones raise it.[5] Ethical and societal concerns help interpret the governance results, particularly for productivity. Productivity generates the largest increase in ethical concern of any attribute, raises beliefs that the technology harms society, lowers perceived social mobility, and is judged unfair when access depends on the ability to pay, yet it reduces desired regulation by 1.052 points relative to the cosmetic baseline (17.6% of the mean), the most permissive stance of any domain. However, respondents oppose taxes on users and favor subsidies for adoption. The distributional-incidence treatment corroborates this interpretation: responses differ by who benefits from the gains, consistent with respondents reading productivity enhancement as a problem of access and fairness rather than safety.

Taken together, these results show that private demand is an unreliable guide to citizens' governance preferences. Adoption responds to several attributes, with domain and purpose each having effects at least as large as risk severity. Purpose, however, plays an outsized role in categorical refusal relative to risk, consistent with a principled rather than purely performance-based constraint for a substantial subgroup. And domain, particularly productivity, raises ethical and societal concerns — including concerns about fairness and access rather than safety — and shifts fiscal preferences far more than it shifts desired regulation. None of these patterns translates into demand for prohibition. As a consequence, policy cannot infer governance preferences from adoption alone: the attributes that drive uptake are not the same as those that drive desired oversight. The results suggest that governance aligned with respondents' preferences would differentiate across applications, applying stricter oversight to higher-risk, enhancing, and inheritable interventions while treating restorative and lower-risk ones more permissively.[6]

Our study contributes to a growing literature on public attitudes toward biomedical and other emerging technologies that bridges economics, ethics, and policy. In economics, work on repugnance and moral constraints shows that deontological reservations can bind even when net benefits are positive (A. E. Roth 2007), and, in related biomedical settings, survey experiments reveal polarization between "never" types and consequentialists who respond to cost–benefit considerations (Elías, Lacetera, and Macis 2019; Hood et al. 2025). Psychology and moral philosophy provide complementary foundations: work on taboo trade-offs and moral foundations explains why some respondents reject "unnatural" enhancements or com-

[5] Inheritable interventions edit germline cells such as sperm, eggs, or embryos, altering DNA in cells whose genetic variants can be passed on to future generations. Non-inheritable interventions, known as somatic editing, affect only the treated individual.

[6] Policy can use various levers to steer the *direction* of innovation. On taxation and incentives, see evidence that the U.S. tax code subsidizes automation and that adjusting tax policy would shift innovative effort toward more labor-augmenting paths (Acemoglu, Manera, and Restrepo 2020). For AI specifically, market forces may foster the "wrong kind of AI," over-automating rather than augmenting workers, motivating standards, public RD, and fiscal instruments (Acemoglu and Restrepo 2020). A broader historical and policy argument for guiding technology toward shared prosperity is developed in *Power and Progress* (Acemoglu and Johnson 2023).

modification on grounds of dignity, sanctity, or fairness (Tetlock et al. 2000; Graham, Haidt, and Nosek 2009; Sandel 2009), and the risk-perception literature predicts outsized sensitivity to severity relative to probability, central to our 1% mild-versus-severe manipulation (Slovic 1987; Sunstein 2002). Distributional and labor-market concerns about who benefits from new technologies connect our productivity scenarios to theories of skill- and task-biased change and inequality (Goldin and Katz 2009; Acemoglu and Autor 2011; Acemoglu and Restrepo 2020; Acemoglu 2025). Complementary to our empirical analysis, Immordino and Panebianco 2026 develop a theoretical model of parental demand for productivity-enhancing HET with heterogeneous moral costs and unequal access. They show that private adoption incentives need not align with preferences over legalization when enhancement generates labor-market externalities, and study the resulting implications for inequality.

Opinion polls document that public support for gene editing is higher for therapeutic than for enhancement uses, varies with religiosity and knowledge, and is shaped by trust in scientists and regulators (Funk, Kennedy, and Sciupac 2016; Funk and Hefferon 2018; Rainie et al. 2022; Siegrist 2000), while governance frameworks emphasize differentiated rules for somatic versus germline interventions (National Academies of Sciences, Engineering, and Medicine 2017; World Health Organization 2021). Our contribution is to use randomized variation in risk, purpose, domain, and heritability (attributes that span performance-based and principled considerations) to identify when technology attributes shift private demand and preferred governance and when resistance persists in ways consistent with non-compensatory moral constraints, and to quantify the demand–governance wedge.

Methodologically, we follow the survey-experimental tradition that uses randomized information and multi-attribute vignettes to identify how beliefs and values map to private adoption and preferred regulation (e.g., Hainmueller, Hopkins, and Yamamoto 2014; Kuziemko et al. 2015; Alesina, Miano, and Stantcheva 2023), extending it to a multi-dimensional human enhancement space that separately identifies the drivers of private demand and of desired regulation. We build on experiments that shift expectations and communication to study downstream attitudes (Alesina, Stantcheva, and Teso 2018; Coibion, Gorodnichenko, and Weber 2022; Haaland, C. Roth, and Wohlfart 2023) and on factorial and conjoint identification strategies (Hainmueller, Hopkins, and Yamamoto 2014; Auspurg and Hinz 2015). Closest in spirit for ethically contested biomedicine, Elías, Lacetera, and Macis 2019 use randomized information with stated choices to distinguish deontological from consequentialist preferences; we adapt this approach to human enhancement technologies.

Finally, our focus on the divergence between private demand and preferred governance complements recent work on AI adoption and regulation (Bursztyn et al. 2025), and we also complement structured outcomes with text-as-data analysis to interpret mechanisms (Gentzkow, Kelly, and Taddy 2019; Ferrario and Stantcheva 2022; Stantcheva 2023).

The remainder of the paper proceeds as follows. Section 2 describes the design, sampling, measurement, and estimation strategy. Section 3 presents the main results on private demand, regulation preferences, and ethical concerns. Section 4 reports auxiliary analyses and robustness checks. Section 5 concludes.

2 Research design, data, and measurements

2.1 Survey structure

The study was implemented as an online survey experiment with embedded randomization modules. The survey consisted of four parts.

Part 1: Socio-demographics and technology orientation. We collected socio-demographic information and also elicited a three-item scale capturing respondents' general attitudes toward new technologies — specifically their tendency toward early adoption, willingness to experiment with unfamiliar technologies, and propensity to seek out information about them.

Part 2: Familiarity, standardized overview, and (randomized) expectations. Respondents reported prior familiarity with human enhancement technologies. All then read a brief standardized description of human enhancement technologies covering potential applications (genetic, neural/biotech), restorative vs. enhancing uses, and the possibility of risks and unintended consequences. A randomized expectations block was shown to 50% of respondents before the scenario module. It elicited prior beliefs on (i) overall well-being impacts, (ii) access and distributional effects, and (iii) ethical implications of human enhancement technologies. This block allows us to test whether prompting ethical and societal considerations affects subsequent responses (Elías, Lacetera, and Macis 2019).

Part 3: Scenario evaluation, attitudes, open text. Each respondent was randomly assigned to one scenario (details below) and then answered modules on adoption and willingness-to-pay (details below), preferred regulatory stringency, and ethical concerns (with randomized placement). Next, we included an open-ended question asking respondents to identify the main considerations driving their previously expressed judgments.

Part 4: Trust, moral foundations, and economic preferences. We concluded with a brief assessment of trust in medical, scientific and political institutions, along with measures of domain-specific risk tolerance and a set of general morality questions capturing the relevance of considerations such as fairness, purity, care for the vulnerable, respect for tradition, disgust/unnaturalness, divine approval, and compassion in respondents' moral judgments.

2.2 Experimental design and sample

2.2.1 Sampling and Data Collection

We fielded our survey to 5,556 adult U.S. respondents recruited by the professional provider Bilendi/Respondi. The questionnaire was programmed and administered in Qualtrics. Data collection occurred between October 9 and October 30, 2025. Respondents provided informed consent on the landing page. The study protocol received ethical approval from the Ethics Board of the University of Naples Federico II and the Johns Hopkins Homewood IRB.

Table 1. Comparison of HET survey sample and ACS 2024 demographics

Group	Category	HET survey (%)	ACS 2024 (%)
Gender	Male	45.3%	49%
	Female	53.8%	51%
	Other	0.87%	
Age	18–24	11.9%	12%
	25–44	30.5%	34%
	45–64	33.5%	31%
	65+	24.1%	23%
Race	White	60.3%	57%
	Black / African American	12.7%	12%
	Hispanic or Latino	15.1%	19%
	Asian	6.2%	6%
	Native American / Alaska Native	1.1%	1%
	Multiracial / Other	4.4%	5%
	Prefer not to answer	0.3%	
Region	Northeast	18.6%	17%
	Midwest	20.9%	20.5%
	South	42.3%	39.0%
	West	18.1%	23.5%
	Other	0.04%	
Education	High school or less	34.9%	37%
	Some college	28.1%	25%
	College degree or higher	36.9%	38%

The sample was quota-matched to the 2024 American Community Survey (ACS) and aligns closely with national benchmarks across gender, age, race/ethnicity, and region (Table 1). Deviations are modest: women are slightly overrepresented, the West somewhat underrepresented, and Hispanic respondents somewhat underrepresented relative to the ACS. Educational attainment is also reasonably similar overall to the ACS distribution, with only modest differences across categories. Overall, the quota design and ACS-based weighting yield a sample that closely approximates the sociodemographic structure of the U.S. adult population.[7]

[7] Appendix Table B1 provides a codebook-style descriptive overview of the variables used throughout the analysis. The table reports summary statistics for the randomized scenario attributes, treatment indicators,

2.2.2 Scenario construction

Each respondent evaluated one vignette drawn at random from a factorial menu varying four core attributes plus a distributional-incidence variation within productivity scenarios.

1. **Benefit domain:** Cosmetic, Health, or Productivity. The vignettes operationalized the *domain* dimension using concrete applications. In the *cosmetic* domain, the scenario concerned an intervention targeting *skin aging* (wrinkles, sun damage, uneven tone). In the *health* domain, the scenario concerned an intervention targeting the *immune system* and resistance to common respiratory infections (e.g., bronchitis, COPD, related illnesses). In the *productivity* domain, the scenario concerned an intervention targeting *memory, concentration, and other cognitive functions*, explicitly linked to study performance, job productivity, and career opportunities. These three domains were selected because they reflect some of the most common and relevant applications of human enhancement technologies discussed in public, policy, and academic debates (National Academies of Sciences, Engineering, and Medicine 2017; Funk, Kennedy, and Sciupac 2016; Funk and Hefferon 2018; Rainie et al. 2022), and because they differ in how need-based versus positional the gain is commonly perceived to be in those debates, allowing us to test whether domain shapes private demand and governance preferences and, if so, through what channels.

2. **Mechanism and heritability:** Genetic (CRISPR-based gene editing), either *non-inheritable* (changes affect the treated individual only) or *inheritable* (changes are inherited by future generations); or Non-genetic (biotechnology implants). In genetic conditions, the scenario described CRISPR as a gene-editing tool that makes precise changes in DNA and applied it to domain-specific targets (genes linked to skin aging, immune resistance to respiratory infections, or cognitive functions such as memory and concentration). The inheritability dimension was conveyed explicitly via standardized wording: *Passes to children? No, changes affect the treated individual only"* (non-inheritable) versus *Passes to children? Yes, inherited by future generations"* (inheritable). In non-genetic conditions, the scenario described biotechnology implants as small devices placed on or within the body that sense internal signals and release medication or stimulation when needed; these scenarios explicitly stated that the technology does not alter genetics and therefore does not pass to children. By varying mechanism and heritability independently of stated risk, the design allows us to identify whether opposition is rooted in the nature of the intervention, its intergenerational reach, or its safety profile.

socio-demographic characteristics, baseline technology attitudes and expectations, outcome variables, and additional preference measures, using survey-question labels whenever available. It therefore complements Table 1, which focuses on representativeness relative to ACS benchmarks, by documenting the full set of experimental and respondent-level variables entering the empirical analysis.

3. **Purpose:** Restorative (return to typical levels) or Enhancing (move beyond typical levels). The *purpose* dimension was implemented through framing rather than by altering the underlying technology. In each domain, respondents were presented with the same technological intervention, but its intended use was described either as restoring a capacity to population-typical functioning (*restorative*) or as increasing that capacity beyond typical levels (*enhancing*). This design choice reflects the reality that the boundary between therapy and enhancement is often conceptually and practically blurred: the same technology may plausibly be used either to remediate deficits or to improve performance above baseline.[8] This contrast is central to the paper's conceptual distinction between performance-based evaluation and principled limits.

4. **Risk:** A 1 in 100 (1%) risk of either *mild, temporary* side effects or *severe, long-lasting* side effects. Risk was communicated as a fixed-probability adverse-event profile. In all scenarios, respondents were told there was a *1 in 100 (1%)* chance of side effects, and the randomized variation concerned the severity and persistence of harm: either *mild, temporary* side effects (e.g., rashes or nausea) or *severe, long-lasting* side effects (e.g., neurological damage or serious immune reactions). Holding the likelihood fixed is especially useful given that small probabilities are often difficult for respondents to interpret consistently, whereas differences in severity are more salient in real-world regulatory debates.

5. **Distributional incidence in productivity scenarios:** Within productivity vignettes, and only for non-genetic HET to limit the number of treatment cells, we also randomize whether cognitive gains are *Education-neutral*, *Complementary to college-level work*, or *Complementary to high-school–level work*. This allows us to test whether private demand and governance preferences respond to the anticipated distribution of gains across skill groups. The education-neutral version captures average perceived gains, while the complements-college and complements-high-school versions make the incidence of benefits explicit, letting us distinguish the intrinsic appeal of productivity enhancement from the perceived distribution of its benefits.

All four core attributes (domain, mechanism/heritability, purpose, and risk) were orthogonally randomized across the full factorial. Because the distributional-incidence sub-randomization within productivity scenarios exists only for non-genetic, enhancing vignettes, the pooled sample contains modest departures from strict orthogonality: the share of non-genetic and enhancing scenarios is slightly higher within the productivity domain than within

[8] In cosmetic scenarios, the restorative framing emphasized correcting accelerated skin aging to achieve a normal appearance, while the enhancing framing emphasized achieving younger-than-typical appearance. In health scenarios, the restorative framing described restoring immune resistance to normal levels, whereas the enhancing framing described increasing resistance well above normal. In productivity scenarios, the restorative framing focused on remediating cognitive impairments due to injury or disability, while the enhancing framing emphasized boosting cognitive performance beyond population norms.

cosmetic or health. Because the regression specifications include domain, mechanism, purpose, and risk together as covariates, the estimated effect of each attribute holds the others fixed, so these estimates are not distorted by the compositional imbalance.

Each vignette includes a concise technology primer (e.g., CRISPR or biotechnology implants), stated function and benefit domain, information on whether changes are transmitted to future generations (for genetic variants), purpose (restoration vs. surpassing typical levels), and risk. Table 2 presents illustrative survey scenarios, and the complete survey instrument is reported in the Online Appendix.

Table 2. Illustrative scenarios used in the survey

Domain	**Mechanism, Purpose, and Risk**	**Scenario description (abridged)**
Cosmetic	Genetic (CRISPR), non-inheritable; restorative; risk of mild, temporary side effects	*CRISPR* repairs genes linked to skin aging that make a person's skin age faster and look worse than normal. Aim: restore appearance (fewer wrinkles, sun damage, uneven tone) to a normal, healthy look. Changes affect only the treated individual (not inherited). There is a 1% chance of mild, temporary side effects (e.g., rashes or nausea).
Health	Genetic (CRISPR), inheritable; enhancing; risk of severe, long-lasting side effects	*CRISPR* enhances genes linked to the immune system. Aim: make a person naturally resistant to common respiratory infections (e.g., bronchitis, COPD) well above normal levels. Change is inherited by future generations. There is a 1% chance of severe, long-lasting side effects (e.g., neurological damage or serious immune reactions).
Productivity (general)	Non-genetic (biotechnology implant); enhancing; risk of severe, long-lasting side effects	Biotechnology implants (sensing and delivering medication/stimulation) enhance memory, concentration, and other cognitive functions. Aim: boost memory, focus, and learning speed above typical population levels, improving study, job productivity, and career opportunities. No genetic alteration; not inherited. There is a 1% chance of severe, long-lasting side effects (e.g., neurological damage or serious immune reactions).
Productivity (education-specific)	Non-genetic (biotechnology implant); enhancing; complementary to college education; risk of mild, temporary side effects	Biotechnology implants that enhance cognitive function (memory, concentration, learning speed) above typical levels. Those whose work or studies rely on skills usually acquired through a college education are expected to gain the greatest advantage. There is a 1% chance of mild, temporary side effects (e.g., rashes or nausea).

2.2.3 Outcome measures

Following the vignette, we elicit four classes of outcomes (with module order randomized where noted).

1. **Private demand: Adoption and willingness to pay (WTP).** We first ask: "If you could obtain this technology today, would you adopt it?" (Yes/No).

   - *If Yes:* respondents are asked to state a maximum, one-time WTP (open numeric entry in US dollars).
   - *If No:* we elicit a reservation-benefit threshold using domain-specific sliders, with explicit examples and a non-compensatory opt-out:
     - *Health scenarios:* "Is there an increase in healthy life expectancy that would make you switch to adopting the technology?" Respondents move a slider to the *minimum* number of extra healthy-life years (0–10) that would make them consider adoption (e.g., "if you would adopt with at least 1 extra year, slide to 1; if you'd need at least 5 years, slide to 5"). They must choose *one* option: either use the slider or select a checkbox stating "I would not adopt the technology regardless of the benefit."
     - *Productivity scenarios:* "Is there a salary increase that would make you switch to adopting the technology?" Respondents move a slider to the *minimum* required annual salary increase (0–100%) that would make them consider adoption, with examples spanning small and large gains (e.g., 5%, 50%, or 100% to indicate a doubling of salary). As in the health module, respondents must choose *one* option: either use the slider or select the checkbox "I would not adopt the technology regardless of the benefit."[9]

   The regardless of the benefit option identifies respondents with principled opposition, for whom no level of benefit induces adoption. The domain-specific reservation thresholds measure the minimum gain needed to induce adoption among those who do not initially adopt.

2. **Regulatory preferences.** A single 0–10 slider captures respondents' preferred joint stance on legal status and regulatory stringency for the technology in their assigned scenario. The scale has labeled anchors: 0 (fully legal, no regulation), 1–3 (legal with minimal regulation, with basic safety standards and minimal oversight), 4–6 (legal with moderate regulation, with standards for safety, efficacy, and informed consent rules, but not overly restrictive), 7–9 (legal but tightly controlled, e.g., clinical trials,

[9] This reservation-benefit follow-up was asked only in the health and productivity scenarios and not in the cosmetic domain, where it is less straightforward to express the relevant gains in a common, policy-relevant quantitative metric analogous to healthy-life years or earnings.

ethical oversight, long-term monitoring), and 10 (total prohibition, a complete ban on development or use).

3. **Ethical concerns (randomized placement).** We collect a 0–10 ethical-concern slider and a multi-select checklist of salient concerns.[10] To test priming, we aligned this module with the earlier expectations randomization: the same randomly selected half of respondents who, in Part 2, reported prior expectations about the *societal* consequences of HET saw the ethics module *before* adoption/regulation; the remaining half saw it *after*. This alignment lets us causally estimate whether prompting forward-looking, society-level reasoning shifts private adoption/WTP and desired regulation.

4. **Societal impact, fairness, and fiscal preferences.** Independently of the above, respondents are randomly assigned to one of two parallel five-item modules, containing either positively framed or negatively framed statements about (i) the technology's overall effects on individuals and society, (ii) its implications for social mobility, (iii) the fairness of access based on ability to pay, (iv) the appropriateness of taxing its use, and (v) the desirability of subsidizing it with public funds, measured on 0–10 agreement scales ranging from strong disagreement to strong agreement. Randomizing the valence of the statements helps reduce the risk that responses are mechanically driven by acquiescence or framing artifacts rather than underlying attitudes (Dunsch et al. 2018).

5. **Open-ended rationale.** Additionally, a single open-response item asks respondents to summarize the main factors that influenced their answers. In Section IV, we complement the quantitative analysis with text-as-data methods to classify themes and sentiment in these narratives.

2.2.4 Attention and comprehension checks

To ensure reliability while preserving the internal validity of the randomized treatments, immediately after the scenario we included four brief factual comprehension items. Each item targeted a key feature of the scenario: (i) the benefit domain, (ii) the mechanism and heritability, including whether changes pass to future generations, (iii) the risk level and severity, and (iv) whether the intervention was restorative or enhancing. All four checks were simple true/false statements that required only recalling information explicitly stated in the scenario, with no trick wording or need for additional inference. The correct answer was always "True," because each statement was constructed to match the attributes shown in the respondent's assigned vignette. Importantly, each item referenced only the attributes

[10]The checklist covered concerns related to inequality, moral boundaries, human nature and identity, future generations, eugenic drift, unintended consequences, misuse, privacy, and dependence on technology, and also included an open-ended other option and a no ethical concerns option.

actually shown to the respondent, such as the presented risk type and severity, the specific mechanism and heritability status, and the stated purpose of the intervention. This avoided cross-attribute cues that could contaminate treatment effects while providing a minimal, scenario-specific screen for attentiveness.

Respondents who missed at most one item were classified as passing the comprehension screen; those missing two or more were deemed inattentive and excluded from the analytical sample. Across all respondents before applying exclusion criteria, attention-check performance shows the following distribution of mistakes: 62.83% made zero errors, 24.44% made one error, and the remaining 12.73% made two or more errors, with 1.17% failing all four checks. Consistent with our preregistered quality-control criteria, the final analysis sample includes only respondents who passed the attention checks, defined as making at most one mistake. This retained group represents 87.27% of the original sample. Within the retained analysis sample, 72% made zero mistakes and 28% made exactly one.

Appendix Table B1 reports the corresponding retained-sample comprehension indicator, labelled "Failed at least one scenario comprehension check." This variable equals one for respondents in the analysis sample who missed exactly one of the four scenario-comprehension checks, and zero for respondents who answered all four checks correctly. Thus, while respondents who missed two or more checks are excluded from the main analysis, Table B1 documents the remaining variation in minor comprehension mistakes among retained respondents.

3 Results

3.1 Baseline attitudes

Before presenting the causal estimates for our main outcomes, it is useful to document respondents' baseline orientation toward technology and their prior expectations regarding the future impact of human enhancement technologies.

3.1.1 Attitudes toward technology in general

Responses to the three 0–10 items capturing general attitudes toward technology indicate that, for all three measures, the sample is on average pro-technology but far from uniformly enthusiastic. Each distribution spans the full scale, with nontrivial mass at both extremes. For all three statements—being among the first to try new technologies, liking to experiment with unfamiliar technologies, and actively seeking information about new technological products—the means fall in the mid-to-upper part of the scale, and in each case about 45–50% of respondents choose values of 7 or higher (clear agreement), about 20–30% choose values of 3 or lower (clear disagreement), and the remainder cluster around the middle (4–6). Thus, across all three items, we observe substantial heterogeneity, with a sizable minority of early-adopter–type respondents and a nontrivial group that is relatively cautious or uninterested in new technologies.

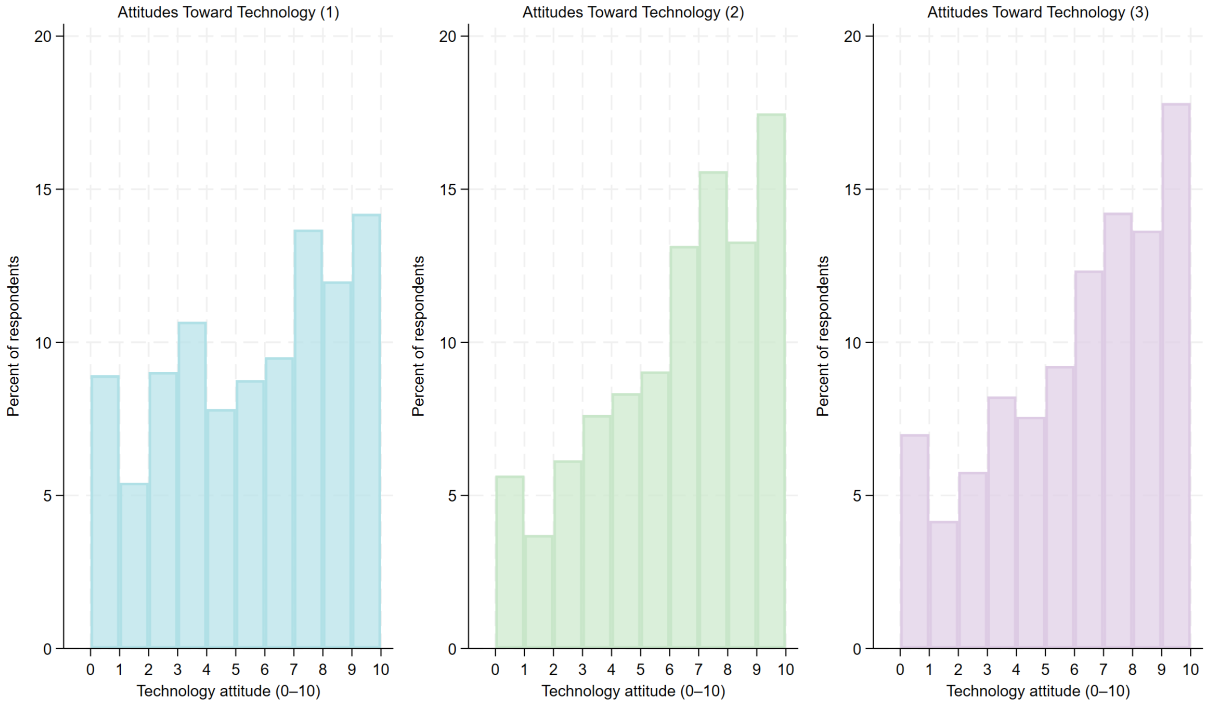

Figure 1. Distribution of responses to the statements: (1) "I am usually among the first of my peers to try new technologies (apps, devices, online services)." (2) "I like to experiment with new technologies even if they are unfamiliar or complex." (3) "I actively seek information about new technological products or services (e.g., from specialized websites or knowledgeable friends)." Responses are measured on an 11-point scale from 0 (strongly disagree) to 10 (strongly agree).

3.1.2 Baseline familiarity with HET and expectations

Responses to the familiarity question show (Figure 2, graph 1) that most people have limited prior exposure to human enhancement technologies. About one third of respondents (34.7%) report that they have never heard of such technologies, and another 36.4% say they have heard of them but do not know much about them. Only 18.1% describe themselves as somewhat familiar, knowing "a few things," while a much smaller group — about 10.8% in total — describe themselves as quite familiar (6.7%) or very familiar (4.1%) with the topic. Overall, fewer than one in five respondents claim more than a basic level of familiarity.

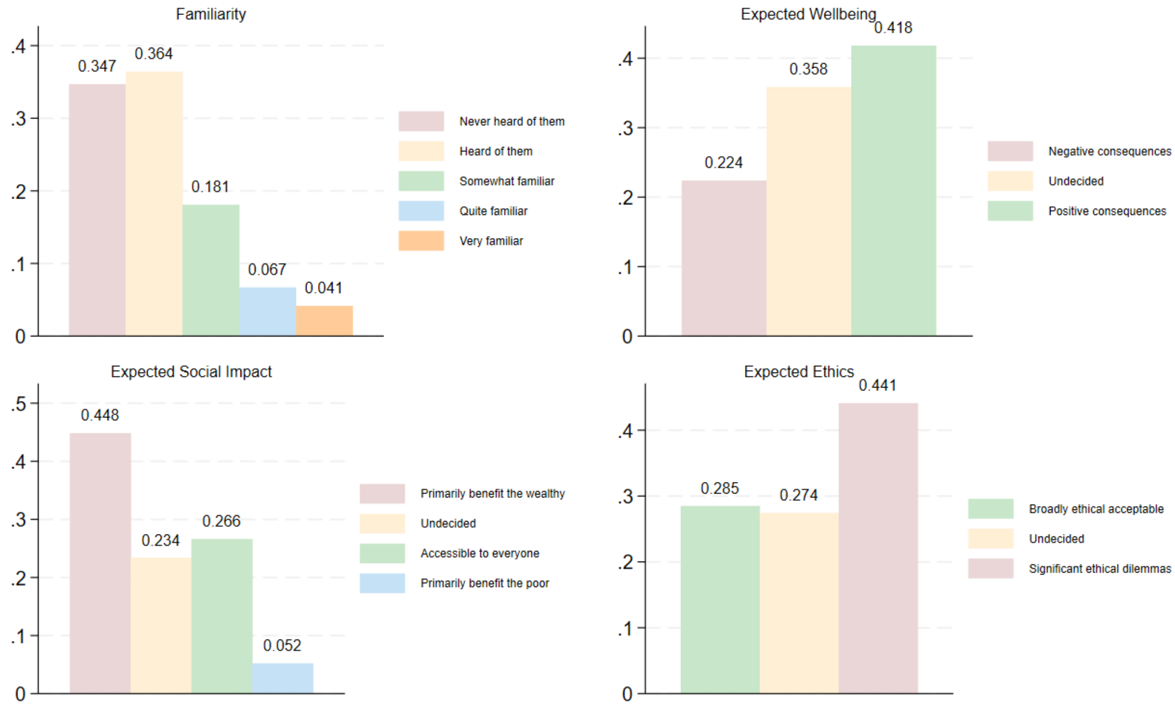

Figure 2. Baseline familiarity and expectations about human enhancement technologies. Graph (1) reports respondents' self-reported familiarity with human enhancement technologies (from "Never heard of them" to "Very familiar"). Graph (2) reports expected consequences of enhancement technologies for personal well-being (negative, undecided, positive). Graph (3) reports perceived social impact and access (primarily benefit the wealthy; undecided; accessible to everyone; primarily benefit the poor). Graph (4) reports perceived ethical acceptability (significant ethical dilemmas; undecided; broadly ethically acceptable).

Figure 2 shows that expectations about the impact of enhancement technologies on people's wellbeing are cautiously optimistic: 41.8% of respondents expect mainly positive consequences, 22.4% anticipate negative consequences, and 35.8% are undecided. Views become markedly more skeptical when we ask about access and distribution: 44.8% believe that enhancement technologies will primarily benefit the wealthy and widen socio-economic gaps, whereas only 26.6% expect them to be accessible and beneficial to everyone, a mere 5.2% think they will mainly benefit the poor, and 23.4% remain uncertain. A similar pattern emerges for ethical implications: 44.1% foresee significant ethical dilemmas, while 28.5% regard enhancement technologies as broadly acceptable and 27.4% are undecided. Overall, these pre-scenario indicators show that expectations about wellbeing benefits are cautiously optimistic, while concerns about distributional fairness and ethical implications are substantial.

3.2 Econometric approach

Because vignette attributes are randomly assigned, differences in outcomes across vignettes identify the causal effect of each attribute on private demand, preferred governance, and ethical concern. The distributional-incidence arm, nested within one domain × mechanism × purpose cell (productivity, non-genetic, enhancing), is not orthogonal to domain, mechanism, or purpose in the pooled sample; we therefore estimate its effects in a separate specification restricted to the relevant subsample (Table 6).

Let Y_i denote an outcome for respondent i. We consider three classes of outcomes. The first captures private demand: willingness to adopt the assigned technology, willingness to pay conditional on adoption, and, among initial non-adopters in the health and productivity domains, whether the respondent would never adopt the technology regardless of the benefit offered. The second captures governance preferences: desired regulatory stringency, support for outright prohibition, and attitudes toward taxing or subsidizing adoption of the technology. The third captures social and moral evaluation: ethical concern, perceived harm, mobility concerns, and fairness of access.

We estimate specifications of the form

$$Y_i = \alpha + D_i'\delta + A_i'\beta + (D_i \times A_i)'\lambda + X_i'\gamma + \varepsilon_i, \tag{1}$$

where D_i is a vector of benefit-domain indicators and A_i is a vector of the other randomized vignette attributes. The domain vector D_i includes indicators for health and productivity applications, with cosmetic applications omitted. The vector A_i includes indicators for genetic inheritable interventions, genetic non-inheritable interventions, enhancing purpose, and severe side effects. The omitted reference category is a cosmetic-domain technology, delivered through a non-genetic mechanism, framed as restorative, and associated with mild temporary side effects.[11]

The interaction term $D_i \times A_i$ allows the effect of risk, enhancement framing, and genetic mechanisms to differ across domains. This is important because the same technical feature may have different economic and ethical meanings depending on the application. For example, severe side effects may be evaluated differently in a health context than in a cosmetic context, and enhancement framing may be especially salient in productivity scenarios, where the technology can affect competition in education and labor markets.

The vector X_i contains pre-treatment respondent characteristics, including socio-demographics, health status, technology orientation, risk preferences, institutional trust, religiosity, and moral-foundation measures. Since vignette attributes are randomly assigned, these controls are not required for identification. We include X_i in robustness specifications to improve

[11] The never-adoption outcome is elicited only in health and productivity scenarios; its regressions therefore exclude cosmetic applications and omit productivity as the reference domain, so the Health coefficient is a health-versus-productivity contrast rather than relative to cosmetic.

precision and to verify that the experimental effects are not driven by chance imbalance in observable characteristics such as income, health status, or religiosity that may independently influence attitudes toward enhancement technologies.

All outcomes are estimated using OLS. For binary outcomes (adoption, never-adoption, and support for prohibition), coefficients represent percentage-point changes relative to the omitted vignette. For continuous outcomes (desired regulation, ethical concern, and related scales), coefficients represent changes in points on the corresponding response scale. We report heteroskedasticity-robust standard errors throughout.

The coefficient vector δ captures how the benefit domain shifts outcomes relative to cosmetic applications. The coefficient vector β captures the effect of technological and moral attributes holding the domain fixed: whether the intervention is genetic, whether it is inheritable, whether it is enhancing rather than restorative, and whether it involves severe rather than mild side effects. The interaction coefficients λ capture whether these attribute effects vary across domains.

3.3 Private demand: adoption and willingness to pay

We begin with private demand: whether respondents say they would adopt the technology described to them, latent demand among those who could be induced to adopt at higher benefits, and categorical refusal.

3.3.1 Descriptive adoption patterns

Figure 3 shows that willingness to adopt is close to an even split overall, but varies meaningfully across domains and attributes. Adoption is lowest in cosmetic scenarios (about 48%), slightly higher in productivity scenarios (roughly 51%), and highest in health scenarios (around 58%). Adoption rates thus differ across domains. Adoption is also higher for genetic technologies than for non-genetic ones (about 54% versus 49%), and within the genetic category, adoption is similar for inheritable and non-inheritable variants (about 55% and 53%). The comparisons also show lower adoption under the severe side-effect profile (49.6% vs. 54.9% under the mild profile) and under enhancing framing (49.6% vs. 55.5% under restorative framing).

Because the vignette attributes are randomized, these cross-group differences are informative about the causal effects of the corresponding attributes. At the same time, the figure reports raw comparisons that average over the full distribution of the other randomized dimensions. We therefore turn next to regression estimates, which provide a more systematic way to recover the marginal effects of each attribute while holding the others fixed and to assess their statistical precision.

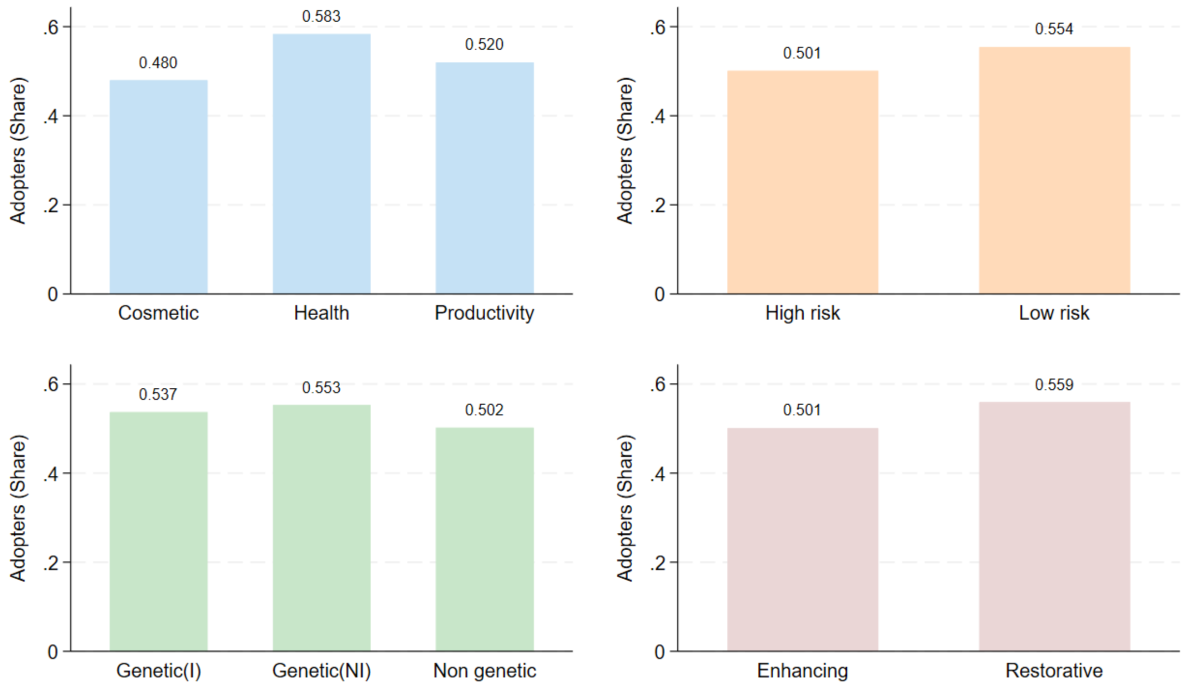

Figure 3. Willingness to adopt the assigned technology. Bars report the share of respondents stating that they would adopt the technology in each treatment condition.

3.3.2 Adoption: regression estimates

Relative to the omitted category (a cosmetic, non-genetic, restorative technology with mild temporary side effects), a health application increases adoption by 10.4 percentage points (about 20% of the 52.8% mean), raising predicted adoption from 50.8% to 61.2%. Productivity applications increase adoption by 5.2 percentage points, about 10% of the mean. Health thus generates the largest adoption premium, productivity a smaller one, and cosmetic use remains the least adopted domain.

Turning to mechanism and heritability: relative to non-genetic implants, inheritable genetic interventions increase adoption by 3.2 percentage points and non-inheritable genetic interventions by 4.5 percentage points, smaller than the health premium but comparable to the productivity premium. Genetic mechanism and heritability therefore do not depress private adoption. Whether the same holds for preferred governance is examined below.

Next, we turn to the comparison between risk and purpose. Increasing side-effect severity from mild and temporary to severe and long-lasting, holding the probability fixed at one percent, reduces adoption by 5.2 percentage points, a 9.8% decline relative to the mean. Framing the same technology as enhancing rather than restorative reduces adoption by 5.5 percentage points, a 10.4% decline. The two effects are nearly identical in magnitude, and we cannot reject their equality. Risk severity is a standard input into a cost–benefit adoption decision, and respondents respond to it as expected. Purpose operates differently:

Table 3. Determinants of HET Adoption

	(1) Adoption	(2) Adoption	(3) Adoption	(4) Adoption
Health	0.104***	0.110***	0.128***	0.085***
	(0.016)	(0.023)	(0.023)	(0.028)
Productivity	0.052***	0.047**	0.083***	0.036
	(0.017)	(0.023)	(0.025)	(0.026)
Genetic Inheritable	0.032**	0.032**	0.030*	0.014
	(0.016)	(0.016)	(0.016)	(0.028)
Genetic Non-Inheritable	0.045***	0.045***	0.043***	0.026
	(0.016)	(0.016)	(0.016)	(0.029)
High Risk	-0.052***	-0.051**	-0.053***	-0.052***
	(0.013)	(0.023)	(0.013)	(0.013)
Enhancing	-0.055***	-0.055***	-0.021	-0.055***
	(0.013)	(0.013)	(0.023)	(0.014)
Health*High Risk		-0.012		
		(0.032)		
Productivity*High Risk		0.009		
		(0.033)		
Health*Enhancing			-0.046	
			(0.032)	
Productivity*Enhancing			-0.058*	
			(0.033)	
Health*Inheritable				0.023
				(0.039)
Productivity*Inheritable				0.035
				(0.040)
Health*Non-Inheritable				0.037
				(0.040)
Productivity*Non-Inheritable				0.017
				(0.040)
Constant	0.508***	0.508***	0.493***	0.520***
	(0.018)	(0.020)	(0.020)	(0.022)
Observations	5,556	5,556	5,556	5,556
Mean	0.528	0.528	0.528	0.528
St. Dev.	0.499	0.499	0.499	0.499

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

the enhancement treatment changes only whether the intervention is described as restoring typical functioning or moving beyond it, while mechanism, domain, and risk probability are held fixed. That a pure reframing reduces adoption by as much as a substantially worse safety profile indicates that the restoration–enhancement boundary carries weight in private demand comparable to that of a salient safety cost.

Column (2) interacts risk with domain. The interactions with health and productivity are small and imprecisely estimated, so the severe-side-effect penalty is similar across cosmetic, health, and productivity applications: risk appears to operate as a general performance penalty.

Column (3) interacts purpose with domain. The main enhancement effect in the cosmetic baseline is negative but imprecise; the additional penalty is -0.046 (s.e. $= 0.032$) in health and -0.058 (s.e. $= 0.033$) in productivity, implying total enhancement penalties of about -6.7 percentage points in health and -7.9 in productivity. The enhancement penalty is thus concentrated in the health and productivity domains rather than in cosmetic use.

Column (4) interacts genetic mechanism with domain. Once genetic mechanism is interacted with domain, the main genetic coefficients are the (cosmetic-domain) genetic effects; the positive health and productivity interactions imply larger genetic effects in those domains. On the personal-use margin, respondents are more responsive to domain, risk, and purpose than to genetic mechanism or heritability.[12]

3.3.3 Willingness to pay

The adoption results identify which attributes affect whether respondents state that they are willing to adopt. We turn, more tentatively, to the intensive margin: among those willing to adopt, how much do they say they would pay? Two cautions apply. Stated willingness to pay for a hypothetical technology is noisy, with an implausibly long upper tail in the raw responses that we trim and cap, so we treat it as exploratory and use distributions rather than means. And because adoption itself responds to the randomized attributes, the set of adopters differs across arms, so any WTP comparison confounds valuation with selection into adoption; we read the patterns as descriptive.

Figure B1 plots inverse demand curves from respondents' stated one-time WTP, conditional on adoption and a positive amount, sorting WTP in descending order against the cumulative number of adopters willing to pay at least each amount. Conditional valuations are highly skewed: most adopters report modest amounts and a small minority report very high ones. Across all four attributes (domain, purpose, risk, and mechanism), two-sample

[12] Appendix Table B4 reports the same specifications from Table 3 with progressively added controls — socio-demographics, health status, religiosity, technology orientation, risk preferences, institutional trust, moral foundations, and the randomized expectations prompt. The main coefficients are stable in sign, magnitude, and significance throughout. Because the attributes are randomly assigned, these controls are not needed for identification; their stability confirms that the pattern is not driven by observable respondent composition.

Kolmogorov–Smirnov tests cannot reject equality of the WTP distributions (all $p \geq 0.6$). The genetic curve lies somewhat above the non-genetic one over part of the range, but this gap is not statistically distinguishable. Subject to the selection caveat, this suggests that domain, purpose, and risk operate mainly on whether respondents adopt rather than on how much adopters would pay.

3.3.4 Never-adopters

Initial non-adoption need not be categorical. A respondent who declines may still be on an ordinary demand curve: the benefit is too small, the risk too high, or the evidence not yet convincing. To distinguish this from non-compensatory refusal, respondents who initially declined in the health and productivity domains were asked whether a larger domain-specific benefit would make them reconsider: additional healthy-life years in health scenarios, an annual salary increase in productivity scenarios. They could either state the minimum benefit required to consider adoption or select that they would not adopt regardless of the benefit. We define the latter as *never-adopters*: respondents who decline to participate even when private benefits are allowed to increase substantially.

Figure 4 shows that this refusal is not marginal. About one quarter of respondents in the health and productivity domains are never-adopters. Categorical refusal is lower in health (about 22%) than in productivity (about 33%), only modestly higher under severe than mild risk (about 29% versus 26%), and sharply higher under an enhancing than a restorative framing (about 32% versus 22%). Genetic mechanism and heritability generate little visible difference. The descriptive pattern thus points to purpose and domain, rather than risk or mechanism, as the main correlates of categorical refusal.

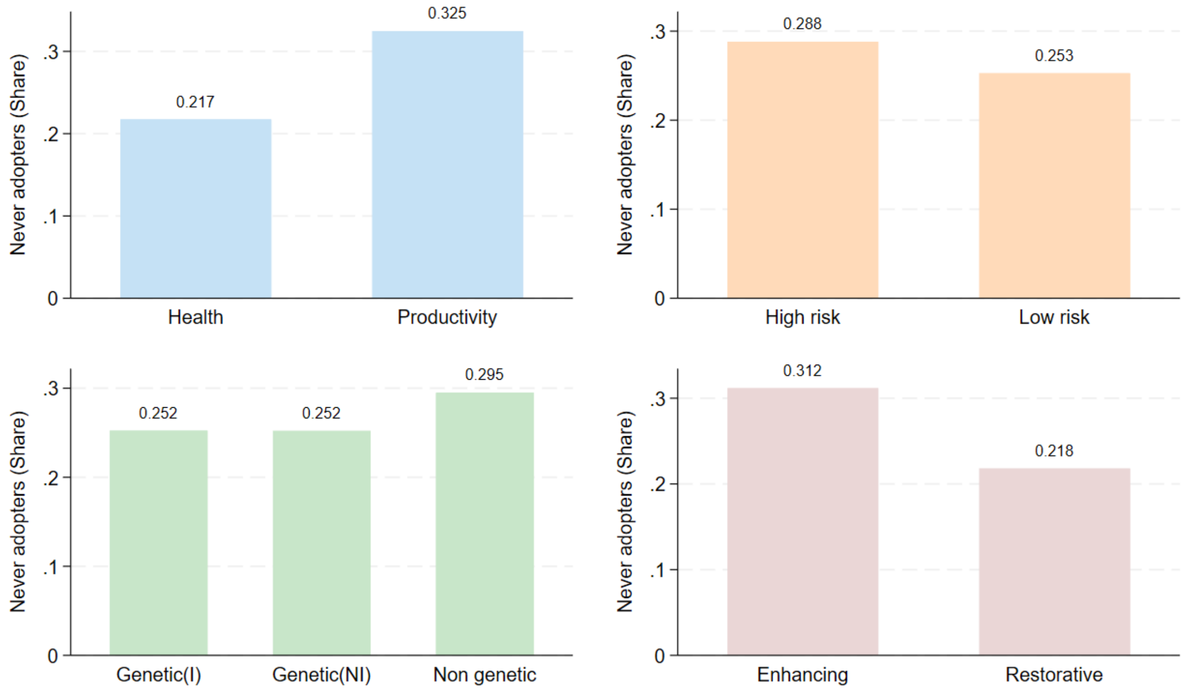

Figure 4. Share of respondents who report that they would never adopt the assigned technology, regardless of the benefits offered.

Table B2 confirms this in the estimates. The mean never-adoption rate is about 28%. Relative to productivity, a health scenario reduces never-adoption by 9.6 percentage points, about 34% of the mean. In addition to raising adoption at the margin, the health domain also moves respondents away from categorical refusal.

Enhancement framing is the strongest positive determinant: describing the technology as enhancing rather than restorative raises never-adoption by 8.2 percentage points, about 30% of the mean, more than twice the effect of severe side effects. Moving from mild to severe side effects raises never-adoption by 3.5 percentage points, about 13% of the mean. Risk matters on this margin, but much less than purpose. The mechanism coefficients are small and insignificant: inheritable and non-inheritable genetic interventions change never-adoption by 1.7 and 1.8 percentage points (both negative), neither distinguishable from zero. Genetic modification does not generate categorical refusal.

The ranking is consistent with a non-compensatory constraint. Severe side effects, the attribute most naturally tied to cost–benefit reasoning, do raise refusal, but modestly. Enhancement framing, which changes the purpose of the same technology rather than the probability of harm, has a much larger effect. Thus, for a substantial minority, categorical refusal persists even when benefits are allowed to increase substantially. The interaction specifications leave this intact. Allowing risk, enhancement, or mechanism effects to vary by domain does not overturn the pattern; the additional high-risk effect in productivity is imprecisely estimated. Categorical refusal is shaped primarily by health versus productivity and by restorative versus enhancing purpose, not by a general aversion to genetic mechanisms or by risk severity.

Finally, we consider the respondents who decline initially but remain persuadable. Figures B2 and B3 plot reservation-benefit distributions for initial non-adopters who did not choose the never-adopt option. The benefit required to reconsider is higher for genetic technologies, especially inheritable ones, and higher under severe risk, while restorative versus enhancing purpose creates little separation. Purpose appears to be central to the participation decision and to categorical refusal; mechanism and heritability matter more for the size of the benefit that would persuade those who remain open.

Whether the attributes that drive private demand also drive what people want for others, through regulation, prohibition, or fiscal treatment, is the question we take up next.

3.4 Desired governance

We now turn from private use to desired governance. We measure desired governance along two margins: regulatory stringency on a legal-to-prohibition scale, where we also examine support for outright prohibition separately as the strongest form of rejection; and fiscal stance, whether respondents would tax or subsidize adoption.

3.4.1 Descriptive regulation patterns

Figure 5 shows the distribution of preferred regulatory regime. Respondents do not concentrate at either extreme. The modal response is legalization with strict regulation, followed by legalization with moderate regulation: in the pooled sample, about 43.5% choose strict regulation (legal but tightly controlled, e.g., clinical trials and ethical oversight) and 32.4% choose moderate regulation (legal with standards for safety, efficacy, and informed consent), while only about 3% favor legalization with no regulation and about 6% favor total prohibition. The dominant response is thus regulated access rather than laissez-faire or prohibition: when respondents express concern, they do so by demanding oversight within legal access rather than by imposing bans.

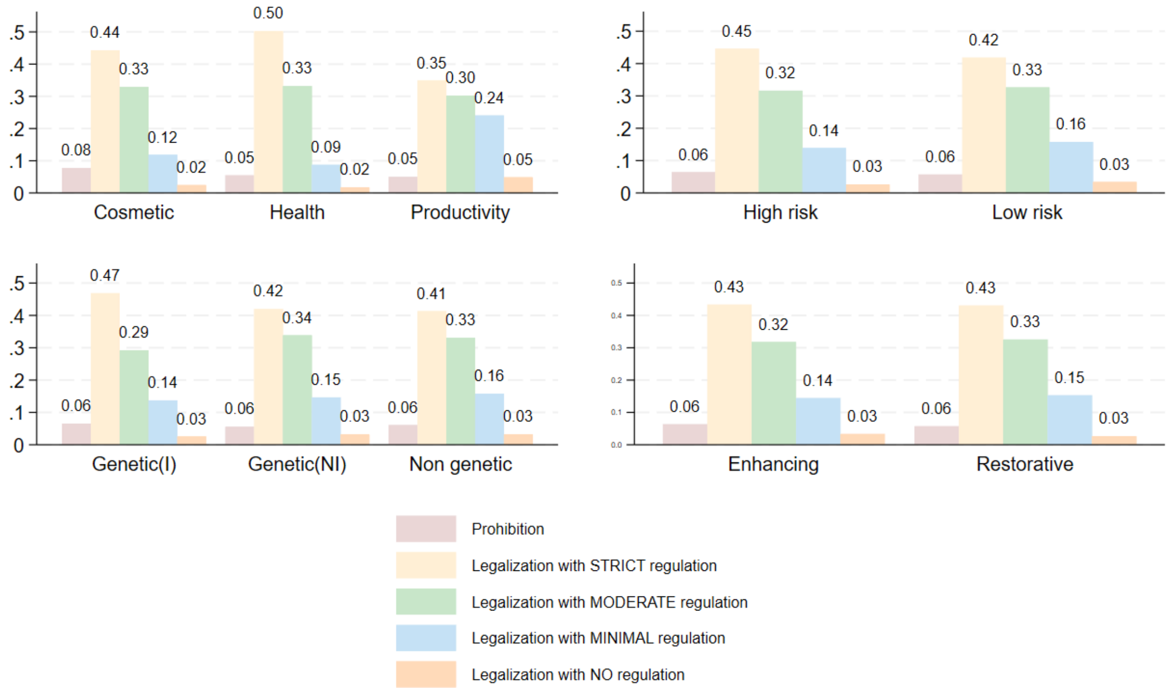

Figure 5. Preferred regulatory regime for the assigned technology. Bars show the distribution of responses across five policy options: prohibition; legalization with strict regulation; legalization with moderate regulation; legalization with minimal regulation; and legalization with no regulation.

3.4.2 Governance: regression estimates

Table 4 reports effects on the 0–10 regulation scale, where higher values indicate stricter preferred governance and 10 corresponds to total prohibition. The outcome mean is 5.98 (s.d. 2.50). Positive coefficients indicate a shift toward tighter oversight, negative coefficients toward more permissive governance.

In Column (1), several attributes raise desired regulation by modest amounts. Health applications increase it by 0.155 points (2.6% of the mean, 6.2% of a standard deviation), and inheritable genetic interventions by 0.151 points; non-inheritable genetic interventions are not distinguishable from non-genetic implants. Risk and purpose move governance in the same direction: severe side effects raise desired regulation by 0.218 points (8.7% of a standard deviation) and an enhancing framing by 0.164 points. These effects are smaller than the same attributes' effects on adoption.

Productivity is the exception, and the largest effect in the table. Productivity applications reduce desired regulation by 1.052 points relative to cosmetic applications, about 17.6% of the mean and several times the magnitude of the risk or enhancement effects, a more permissive stance than even the cosmetic baseline. This permissiveness does not mean respondents are unconcerned about productivity enhancement; we return to its interpretation after presenting the fiscal results in Columns (5)–(6) and the ethics results in the next section.

Columns (2)–(4) let risk, purpose, and mechanism effects vary by domain. In Column (2) the main high-risk coefficient remains positive, and the interactions with health and productivity are imprecisely estimated, so the demand for oversight in response to severe side effects is similar across domains. In Column (3) the enhancement effect varies by domain: the cosmetic-baseline effect is small and insignificant, while the implied effect is about 0.25 points in health and 0.29 in productivity, consistent with the adoption results in placing most weight on the restoration–enhancement boundary in high-stakes domains. In Column (4), the genetic effects also vary by domain. Inheritable interventions show no significant domain contrast; non-inheritable interventions do: the main effect is negative (-0.302), but the health interaction is positive and significant ($+0.422$), implying a net effect of about $+0.120$ for a non-inheritable genetic health technology. Respondents thus demand more oversight for genetic intervention in the health domain even when the change is not inheritable.

Columns (5) and (6) turn to fiscal instruments, which show that governance preferences are not captured by a single ban-to-free-market scale. Column (5) is agreement that it is good to tax users (higher = more support for taxing), and Column (6) is agreement that public funds should subsidize adoption (higher = more support for subsidies).

Health and productivity applications both shift the fiscal stance toward broader access: relative to cosmetic, health makes respondents less willing to tax users (-0.705) and more willing to subsidize (0.841), and productivity does the same (-0.433 and 0.466). Enhancement framing moves the opposite way, toward less willingness to subsidize (-0.306), so restoration attracts an access-expanding fiscal stance while enhancement attracts caution.

These fiscal results help interpret the productivity result. Productivity attracts the least regulation, yet respondents oppose taxing its users and favor subsidizing it. Respondents favor broadening access to productivity enhancement rather than restricting it. Legal permission, oversight, taxation, and access are distinct governance margins, and a respondent can want the same technology legal, lightly regulated, and subsidized at once. Why respondents take this stance toward productivity in particular, and how it reconciles with the concern it raises, we discuss with the ethics and societal-impact results in the next section.

Preferred regulation is therefore not a mechanical extension of private demand. Severe risk, enhancement framing, and inheritable genetics raise desired oversight even where they have limited effects on adoption. Productivity moves the other way: it is treated more permissively, and respondents even favor subsidizing it, despite its raising the most ethical concern of any attribute (Section 3.5).

3.4.3 Prohibition

Prohibition is a different choice from strict regulation: it forbids use rather than setting terms for it. We therefore study it as a separate outcome, since a technology may be viewed

Table 4. Determinants of Desired Regulation of HET

	(1) Regulation	(2) Regulation	(3) Regulation	(4) Regulation	(5) Tax use	(6) Subsidize
Health	0.155**	0.208*	0.012	0.021	-0.705***	0.841***
	(0.076)	(0.109)	(0.107)	(0.130)	(0.095)	(0.100)
Productivity	-1.052***	-0.969***	-1.226***	-1.068***	-0.433***	0.466***
	(0.085)	(0.120)	(0.125)	(0.132)	(0.098)	(0.103)
Genetic Inheritable	0.151*	0.149*	0.163**	0.184	0.050	0.070
	(0.080)	(0.080)	(0.080)	(0.133)	(0.095)	(0.100)
Genetic Non-Inheritable	-0.117	-0.116	-0.105	-0.302**	0.128	0.223**
	(0.079)	(0.079)	(0.080)	(0.141)	(0.095)	(0.100)
High Risk	0.218***	0.309***	0.219***	0.216***	0.016	-0.140*
	(0.065)	(0.112)	(0.065)	(0.065)	(0.078)	(0.082)
Enhancing	0.164**	0.164**	-0.034	0.158**	0.013	-0.306***
	(0.066)	(0.066)	(0.112)	(0.066)	(0.079)	(0.083)
Health*High Risk		-0.109				
		(0.152)				
Productivity*High Risk		-0.166				
		(0.168)				
Health*Enhancing			0.282*			
			(0.152)			
Productivity*Enhancing			0.323*			
			(0.170)			
Health*Inheritable				-0.005		
				(0.184)		
Productivity*Inheritable				-0.101		
				(0.207)		
Health*Non-Inheritable				0.422**		
				(0.187)		
Productivity*Non-Inheritable				0.132		
				(0.204)		
Constant	6.067***	6.022***	6.159***	6.119***	5.199***	4.254***
	(0.087)	(0.097)	(0.096)	(0.107)	(0.104)	(0.111)
Observations	5,556	5,556	5,556	5,556	5,502	5,502
Mean	5.982	5.982	5.982	5.982	4.887	4.543
St. Dev.	2.496	2.496	2.496	2.496	2.903	3.065

Robust standard errors in parentheses
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

as risky or ethically troubling without being viewed as something that should be forbidden. In the language of repugnant transactions, prohibition captures the strongest form of social rejection (A. E. Roth 2007).

We find that support for prohibition is low and weakly responsive to the attributes (Table B3). The mean prohibition rate is only 6.1%: respondents often want meaningful oversight but rarely want to close the market. Human enhancement technologies raise substantial ethical concern (Section 3.5), but do not register as repugnant transactions in the strong sense of eliciting broad support for bans.

Relative to cosmetic applications, health technologies reduce support for prohibition by 2.3 percentage points and productivity technologies by 3.0 percentage points, declines of about 38% and 49% of the mean. When a technology is framed as producing health or productivity benefits, respondents become less willing to ban it, even if they still want oversight. The other attributes behave differently on this margin than on the regulation scale: inheritable and non-inheritable genetic interventions do not raise support for a ban, and severe side effects and enhancement framing each raise prohibition by only 0.7 percentage points, neither distinguishable from zero. The attributes that raise desired regulation, risk, enhancement, and heritability, do not raise demand for prohibition.

The interaction columns in Table B3 reinforce this. In Column (2), high risk raises prohibition in the cosmetic baseline by 2.6 percentage points, but the productivity-high-risk interaction is negative and significant: even with severe side effects, productivity enhancement is no more likely to be banned than its low-risk version, consistent with productivity not being read primarily as a safety problem. Column (3) shows no evidence that enhancement framing generates prohibition in any domain. Column (4) shows that genetic mechanisms do not either; if anything, the health-by-inheritable interaction is negative and marginally significant, implying inheritable genetic interventions in health are less likely to be prohibited than the cosmetic baseline. Respondents may want such technologies tightly regulated, but not banned.

3.5 Ethical concerns

A respondent may be willing to adopt a technology, think it should remain legal, and still regard it as ethically troubling. By identifying where that concern concentrates, we can assess whether it is primarily driven by biomedical risk, intergenerational exposure, distributive considerations, or principled objections to the nature of the intervention itself.

3.5.1 Descriptive ethics patterns

Figure 6 shows that ethical concern is widespread. In every scenario only a small minority (about 4–7%) report no concern, and most report moderate or significant concern. Many respondents who would adopt still report moderate or significant ethical concern. The share reporting significant concern is highest for productivity (41%, versus 35% for cosmetic and 34% for health), somewhat higher under high risk (38% versus 35%) and enhancing framing (38% versus 35%), and somewhat higher for inheritable genetic interventions (40%) than for non-inheritable (35%) or non-genetic (36%) ones.

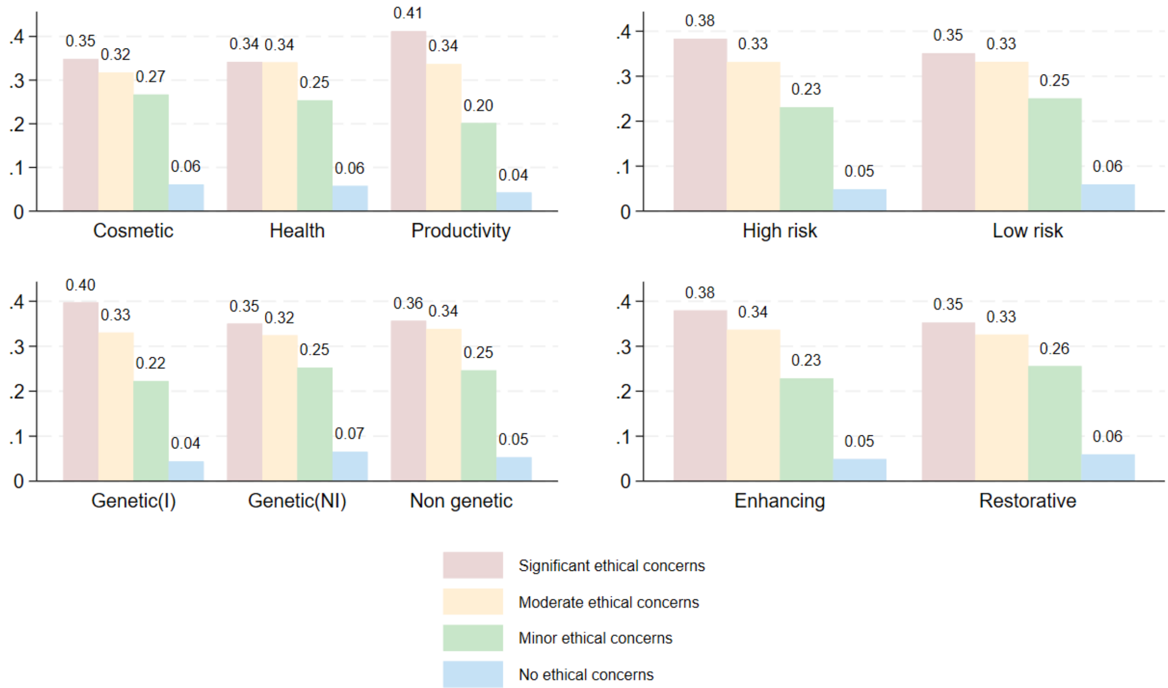

Figure 6. Ethical concerns about the assigned technology. Bars report the fraction of respondents indicating "significant," "moderate," "minor," or "no" ethical concerns. Results are disaggregated by scenario attributes, using the same notation as in Figure 3.

3.5.2 Ethics: regression estimates

Table 5 reports effects on the 0–10 ethical-concern slider (mean 5.35, s.d. 2.79). Higher values indicate greater concern; positive coefficients indicate a shift toward greater concern relative to the omitted scenario (cosmetic, non-genetic, restorative, mild side effects).

Column (1) shows that productivity scenarios generate the largest increase in ethical concern of any attribute. Relative to cosmetic applications, productivity scenarios raise concern by 0.527 points, about 9.9% of the mean and 19% of a standard deviation, larger than the effects of severe side effects, enhancement framing, or inheritable genetic intervention. Health applications do not significantly raise concern relative to cosmetic.

Inheritable genetic interventions raise concern by 0.321 points while non-inheritable ones do not, pointing to an intergenerational channel: respondents are more troubled when changes pass to future generations. Severe side effects raise concern by 0.249 points and enhancement framing by 0.236, both precisely estimated but smaller than the productivity effect.

Table 5. Determinants of HET Ethical concerns

	(1) Ethics	(2) Ethics	(3) Ethics	(4) Ethics	(5) Harm	(6) Mobility	(7) Fairness
Health	0.070	0.202	-0.086	-0.154	-0.342***	0.003	0.751***
	(0.092)	(0.130)	(0.131)	(0.157)	(0.087)	(0.091)	(0.100)
Productivity	0.527***	0.606***	0.233*	0.425***	0.125	0.197**	0.786***
	(0.092)	(0.129)	(0.140)	(0.145)	(0.090)	(0.092)	(0.102)
Genetic Inheritable	0.321***	0.319***	0.346***	0.313**	0.095	-0.016	-0.007
	(0.090)	(0.090)	(0.090)	(0.157)	(0.087)	(0.091)	(0.100)
Genetic Non-Inheritable	-0.049	-0.048	-0.024	-0.383**	-0.145*	-0.154*	-0.223**
	(0.092)	(0.092)	(0.092)	(0.162)	(0.087)	(0.090)	(0.100)
High Risk	0.249***	0.392***	0.250***	0.246***	0.169**	0.182**	0.104
	(0.074)	(0.131)	(0.074)	(0.074)	(0.071)	(0.074)	(0.082)
Enhancing	0.236***	0.235***	-0.039	0.233***	0.361***	0.250***	0.140*
	(0.076)	(0.076)	(0.131)	(0.076)	(0.072)	(0.075)	(0.083)
Health*High Risk		-0.267					
		(0.184)					
Productivity*High Risk		-0.160					
		(0.182)					
Health*Enhancing			0.310*				
			(0.184)				
Productivity*Enhancing			0.537***				
			(0.187)				
Health*Inheritable				0.023			
				(0.221)			
Productivity*Inheritable				-0.008			
				(0.222)			
Health*Non-Inheritable				0.670***			
				(0.227)			
Productivity*Non-Inheritable				0.323			
				(0.224)			
Constant	4.817***	4.748***	4.939***	4.929***	4.444***	4.530***	5.101***
	(0.100)	(0.113)	(0.112)	(0.123)	(0.094)	(0.098)	(0.110)
Observations	5,556	5,556	5,556	5,556	5,502	5,502	5,502
Mean	5.348	5.348	5.348	5.348	4.619	4.784	5.785
St. Dev.	2.788	2.788	2.788	2.788	2.624	2.741	3.069

Robust standard errors in parentheses
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Productivity raises the most concern, yet, as the previous subsection showed, the least desired regulation. Columns (2)–(4) let the risk, enhancement, and genetic effects vary by domain. In Column (2) the high-risk effect is positive in the cosmetic baseline (0.392) and the health and productivity interactions are imprecisely estimated, so severe side effects raise concern without the effect being concentrated in one domain. In Column (3) the enhancement effect varies by domain: small and insignificant in cosmetic (−0.039), positive in health (+0.310), and largest in productivity (+0.537). In Column (4) genetic mechanism also varies by domain: the non-inheritable effect is negative in cosmetic (−0.383) but the health interaction is positive and large (+0.670), a net effect of about +0.29 for a non-inheritable genetic health technology. Genetic intervention becomes ethically salient in a high-stakes biological domain even when the change is not inherited.

Columns (5)–(7) identify the content of the productivity concern. Productivity scenarios raise agreement that the technology will make upward mobility harder (0.197), and that it is unfair if only those who can afford it benefit (0.786). Health scenarios do not raise perceived social harm and in fact lower the concern that the technology will harm people, and they do raise fairness concerns about access (0.751): unequal access is seen as problematic in both domains, but the mobility channel is specific to productivity. Non-inheritable genetic interventions lower perceived harm relative to both non-genetic and genetic inheritable interventions, consistent with a distinct channel: future generations, irreversibility, and consent rather than distributive considerations.

The productivity concern is thus about distribution and access rather than safety, which is why it coexists with low desired regulation and support for subsidies.

3.6 Distributional incidence

Within productivity scenarios, and limited to non-genetic enhancements, we randomized who benefits: the cognitive gain was described as broadly productivity-enhancing (the baseline), as benefiting individuals with college-level skills, or as benefiting individuals with high-school-level skills. This allows us to test whether adoption, governance, and ethical responses depend on who is expected to gain, above and beyond the general effect of productivity enhancement.

Table 6 reports the results. The incidence manipulation has little effect on stated adoption. The education-neutral productivity scenario raises adoption relative to the cosmetic baseline by 5.8 percentage points, although only at the 10 percent level. By contrast, the college- and high-school-complementary productivity variants do not generate statistically significant increases in adoption. The same pattern appears for never-adoption: among respondents in the health and productivity domains, the coefficients for the college- and high-school-complementary variants are small and statistically insignificant. Who benefits from productivity enhancement does not appear to drive the private adoption margin.

Incidence does, however, affect governance and ethical responses. All three productivity variants are associated with substantially lower desired regulation than the cosmetic baseline. The education-neutral, college-complementary, and high-school-complementary versions reduce desired regulation by 1.076, 1.136, and 1.164 points, respectively, on the 0–10 regulation scale. Support for outright prohibition does not increase under any productivity variant; if anything, it falls, especially when the gains are described as benefiting high-school-level skills, where the prohibition coefficient is negative and statistically significant.

The fiscal outcomes show a similar pattern. The education-neutral productivity scenario reduces support for taxing users by 0.754 points and increases support for public subsidies by 0.576 points. The college-complementary version also reduces support for taxing use, although the effect is smaller and only marginally significant, while its effect on subsidies is positive but not statistically significant. The high-school-complementary version generates a strong subsidy response, increasing support for public subsidies by 0.559 points, while its effect on taxing use is negative but not statistically significant.

The final columns show where the ethical concern comes from. Fairness concerns increase across all three productivity conditions: the coefficients are 0.725 for the education-neutral scenario, 0.795 for the college-complementary scenario, and 0.568 for the high-school-complementary scenario, all large and statistically significant. Mobility concerns increase for the education-neutral and college-complementary variants (0.296 and 0.379, respectively) but not when the technology is described as benefiting high-school-level skills. Perceived societal harm does not respond significantly to any productivity variant.

Table 6

VARIABLES	(1) Adoption	(2) Never Adopt	(3) Regulation	(4) Prohibition	(5) Tax use	(6) Subsidize	(7) Ethics	(8) Harm	(9) Mobility	(10) Fairness
Health	0.085***	-0.113***	0.017	-0.010	-0.840***	0.893***	-0.154	-0.418***	-0.014	0.732***
	(0.028)	(0.027)	(0.130)	(0.014)	(0.161)	(0.172)	(0.157)	(0.146)	(0.156)	(0.170)
Productivity	0.058*		-1.076***	-0.021	-0.754***	0.576***	0.448***	0.021	0.296*	0.725***
	(0.030)		(0.154)	(0.015)	(0.178)	(0.188)	(0.169)	(0.162)	(0.166)	(0.187)
Productivity (College)	0.014	0.030	-1.136***	-0.027	-0.396*	0.141	0.502**	0.300	0.379*	0.795***
	(0.041)	(0.041)	(0.218)	(0.020)	(0.240)	(0.253)	(0.221)	(0.220)	(0.228)	(0.248)
Productivity (High School)	0.023	-0.003	-1.164***	-0.047***	-0.166	0.559**	0.333	-0.054	-0.089	0.568**
	(0.042)	(0.042)	(0.215)	(0.018)	(0.242)	(0.253)	(0.228)	(0.224)	(0.224)	(0.257)
High Risk	-0.043**	0.056**	0.308***	0.010	0.127	-0.164	0.376***	0.181	0.324***	0.127
	(0.021)	(0.023)	(0.106)	(0.010)	(0.123)	(0.131)	(0.119)	(0.113)	(0.118)	(0.131)
Enhancing	-0.067***	0.063**	0.333***	0.022*	-0.039	-0.047	0.194	0.292**	0.290**	0.177
	(0.024)	(0.026)	(0.115)	(0.012)	(0.138)	(0.147)	(0.134)	(0.126)	(0.132)	(0.146)
Constant	0.522***	0.271***	5.989***	0.059***	5.254***	4.098***	4.883***	4.523***	4.427***	5.106***
	(0.025)	(0.027)	(0.126)	(0.012)	(0.147)	(0.159)	(0.143)	(0.131)	(0.137)	(0.155)
Observations	2,198	1,563	2,198	2,198	2,184	2,184	2,198	2,184	2,184	2,184
Mean	0.502	0.295	5.885	0.0614	4.826	4.446	5.319	4.691	4.850	5.783
St. Dev.	0.500	0.456	2.536	0.240	2.900	3.071	2.801	2.650	2.758	3.064

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

4 Auxiliary analyses and robustness

This section reports three sets of auxiliary results. We first use the adoption and reservation-benefit questions to characterize distinct adoption "types". We then examine heterogeneity across respondent subgroups. We close with robustness checks.

4.1 Characterizing adoption "types"

The adoption and reservation-benefit questions let us classify respondents in the health and productivity domains into three types. *Adopters* would adopt at the benefits described. *Conditional adopters* initially decline but would adopt if the benefit were large enough (additional healthy-life years in health, an annual salary increase in productivity). *Never-adopters* decline and would not adopt at any benefit. The three correspond to distinct margins of demand: current demand, latent demand, and categorical refusal.

After the adoption and reservation-benefit questions, respondents answered a closed-ended question asking which specific ethical concerns, if any, applied to the technology they had evaluated. The list included concerns about playing God or crossing religious boundaries, human nature and identity, dependence on technology, unintended consequences, misuse or abuse, privacy or surveillance, inequality, and an option indicating no ethical concern. Figure 7 reports, for each adoption type, the share of respondents selecting each concern. Because the closed-ended ethical-concern question allowed respondents to select multiple concerns, the entries in Figure 7 are not mutually exclusive and the columns may not add up to 100 percent. Each cell is the percentage of respondents in the corresponding adoption type who selected that concern.

Figure 8 complements the closed-ended checklist with respondents' open-ended explanations of the main factors that influenced their answers. We coded these open-ended responses using a large language model (Claude, Anthropic, model `claude-sonnet-4-5`), which classified each response into a fixed set of non-mutually-exclusive themes following a detailed prompt we designed: benefit or usefulness; safety and side effects; need for more evidence; inequality and access; moral boundaries, human nature, or religion; and dependence, unintended consequences, misuse, or privacy. Appendix C reports the full codebook and coding procedure. For each adoption type, the figure shows the share of respondents whose open-ended rationale was assigned to each theme. Accordingly, the percentages in a column need not sum to 100 percent: some responses do not mention any of the six displayed themes, and, where a response mentions more than one theme, it can be assigned to more than one category.

Figure 7 shows that the closed-ended concern checklist differs systematically across adoption types. Playing God or crossing religious boundaries is cited by 28% of adopters, rising to 40% among conditional adopters and 48% among never-adopters. Human nature and identity follows a similar gradient (30%, 42%, 46%). Unintended consequences is elevated across all three types but higher among conditional adopters and never-adopters (46% each) than among adopters (35%). By contrast, inequality is nearly identical across types (22%, 23%, 22%), indicating that distributional concerns do not distinguish adoption types. Only 12% of adopters report no ethical concerns, falling to 7% among conditional adopters and never-adopters.

Figure 8 shows the open-ended rationales. Adopters overwhelmingly emphasize benefit and usefulness (44%), with concern-related themes all in the single digits. Conditional adopters are more evenly spread, with notably higher rates of safety and side effects (16%), need for more evidence (16%), and dependence, misuse, and privacy (17%). Never-adopters are most likely to cite moral boundaries (26%), more than five times the rate among adopters, and safety and side effects (18%), twice the rate among adopters. Inequality and access is similarly low across all three types (18%, 17%, 13%). Moral boundary concerns are thus the distinguishing feature of never-adopters' open-ended rationales.

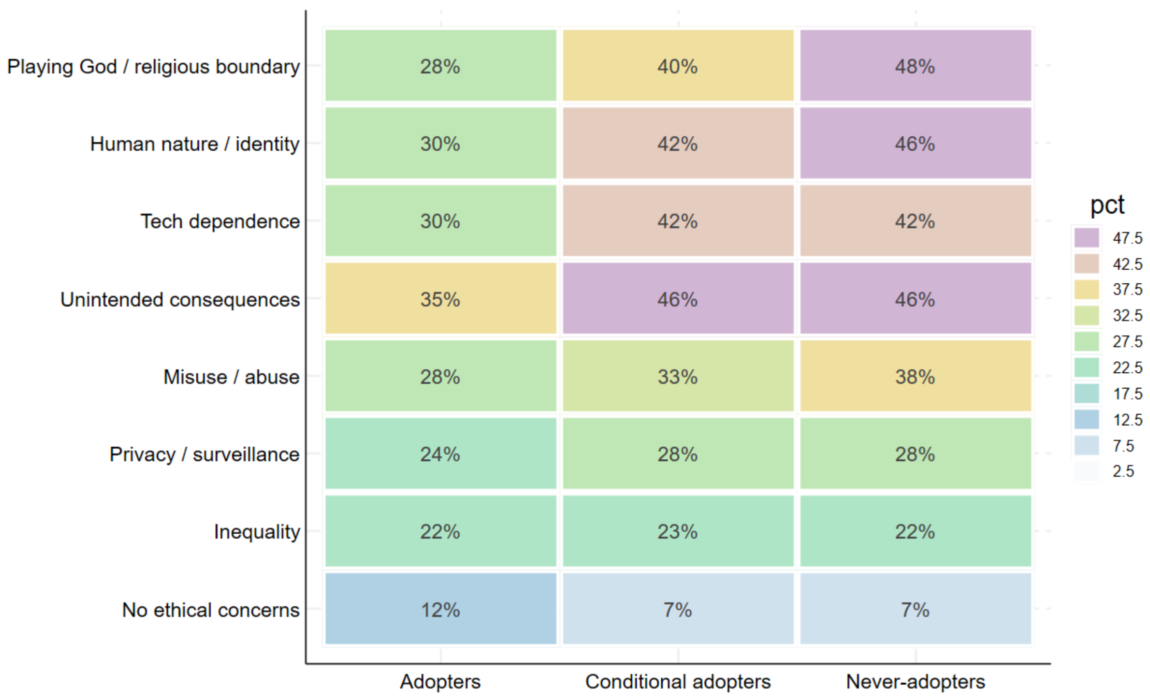

Figure 7. Ethical concerns by adoption type.

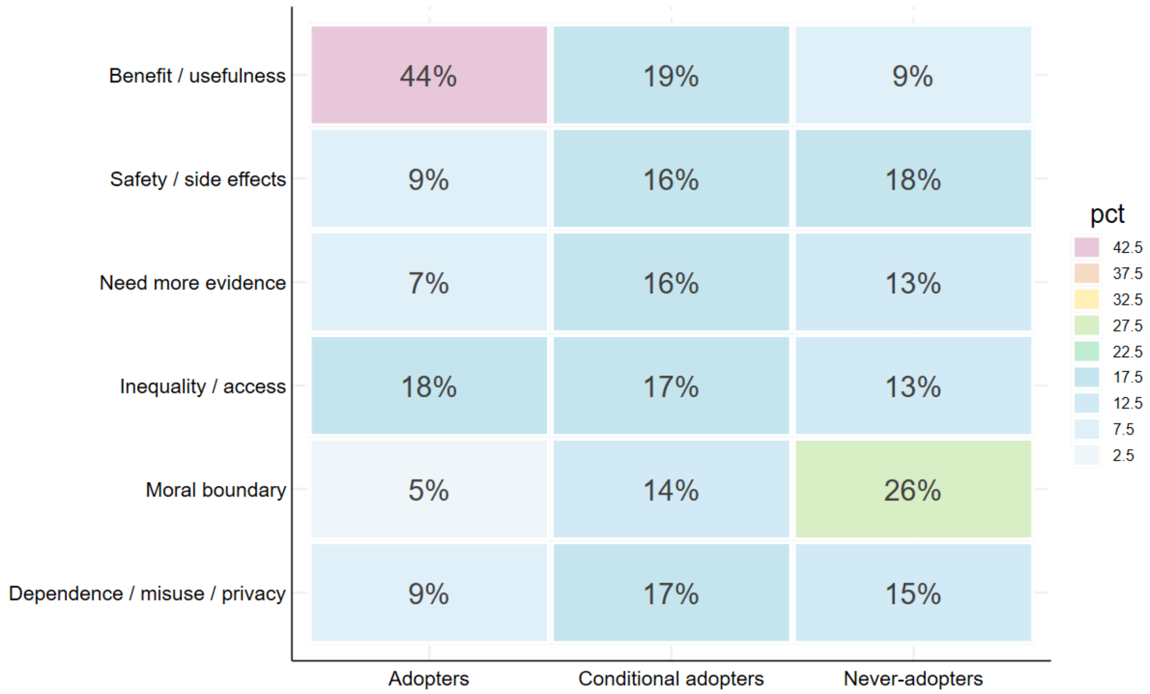

Figure 8. Open-ended rationales by adoption type.

Figures 9 and 10 plot desired regulation and ethical concern by adoption type. Both order the groups the same way: never-adopters report the highest ethical concern and the highest support for strict regulation, including prohibition (about 13%), adopters the lowest, and conditional adopters in between. Nevertheless, the distributions overlap considerably and neither collapses into the extremes: most never-adopters choose interior points on the regulation scale rather than prohibition, and many adopters choose moderate or strict regulation and report meaningful concern.

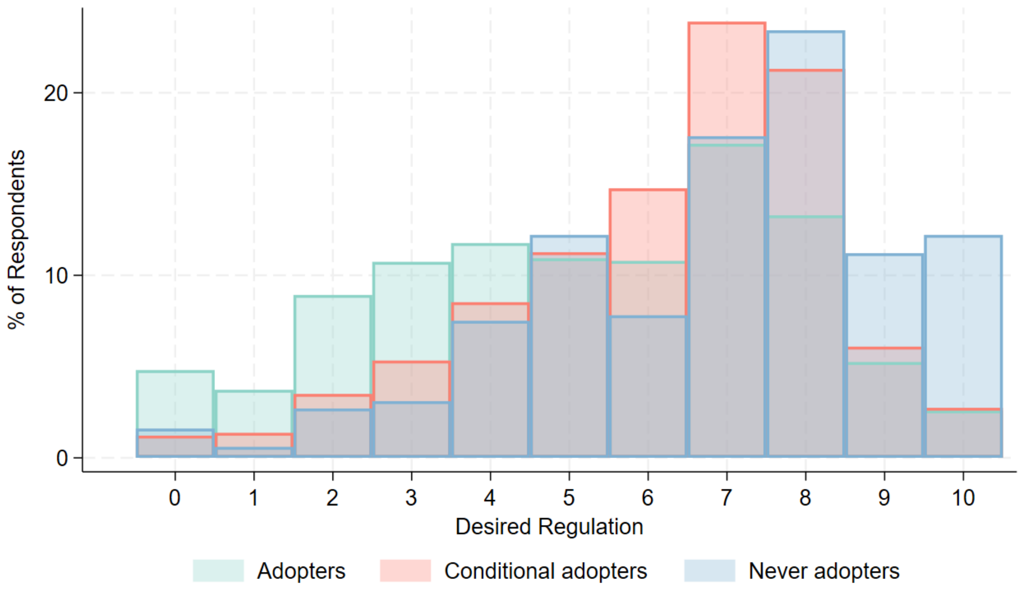

Figure 9. Preferred regulation by adoption type.

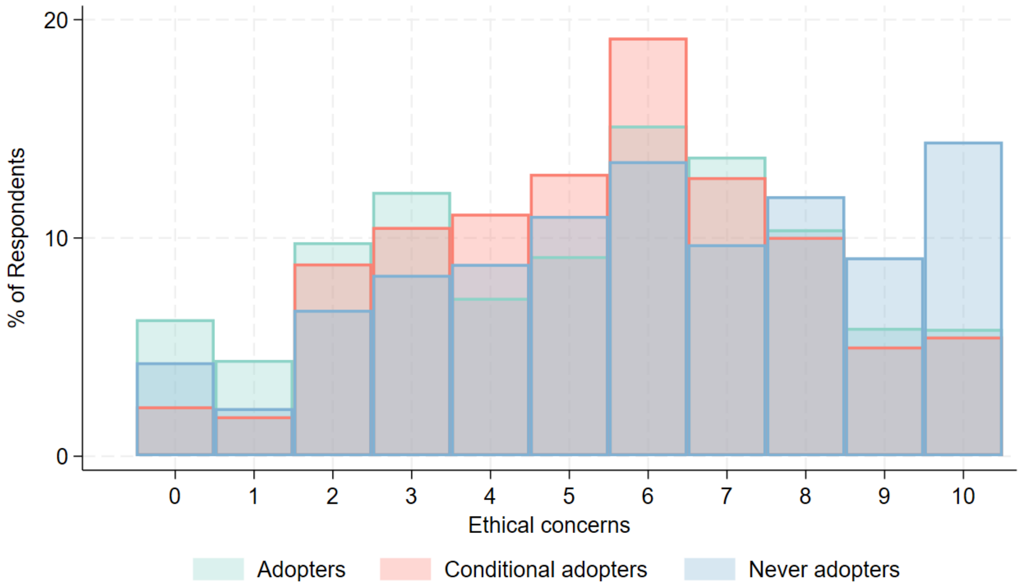

Figure 10. Ethical concern by adoption type.

4.2 Heterogeneity

The main patterns are stable across subgroups. Splitting the sample by prior expectations, gender, education, risk attitudes, and sanctity-related moral views (Figures B4–B8) leaves the ranking largely unchanged: stated willingness to adopt is higher for health applications and lower under severe side effects and enhancement framing, while genetic mechanism and heritability do not reduce stated adoption; inheritable interventions instead raise desired regulation and ethical concern. Full estimates are in the Appendix.

Differences across subgroups are mainly in magnitude, not direction. Respondents with more favorable priors are more supportive (Figure B4); men and women, and more and less educated respondents, differ somewhat in average adoption, regulation, and concern (Figures B5, B6); but in each case respondents still distinguish therapeutic from enhancement uses and react against risk and inheritable interventions. Heterogeneity is more pronounced for risk attitudes and sanctity concerns: more risk-averse respondents adopt less and want tighter regulation, especially under severe side effects or inheritable changes (Figure B7), and respondents with stronger sanctity concerns are less supportive of enhancement, particularly beyond treatment or involving genetic mechanisms (Figure B8). The qualitative ordering of attributes holds throughout.

4.3 Robustness

We assess the robustness of our main findings along three dimensions: (i) inclusion of respondent-level controls, and (ii) inclusion of an indicator for whether respondents were prompted at the start of the survey to state their general expectations about the societal and ethical implications of human enhancement technologies ("expectations prompt"), and (iii) application of a stricter attention/comprehension screen that restricts the sample to respondents with zero errors on the four immediate comprehension items.

Controls. Tables A4-A8 present results from our baseline specifications, adding a rich set of pre-treatment covariates (socio-demographics and baseline measures such as technology orientation, institutional trust, moral foundations, and related background characteristics). Because vignette attributes are randomly assigned, including controls is not necessary for identification, but it can improve precision. The estimated treatment effects on adoption, regulation, and ethical concerns remain similar in sign and magnitude relative to the baseline models, and the ranking of scenario attributes is unchanged.

Expectations prompt. Half of the respondents were randomized to first report their general expectations about the societal and ethical consequences of HET before seeing the specific vignette. We include an indicator for assignment to this expectations-prompt condition in the main specifications (and, where relevant, allow for interactions between the prompt and the randomized vignette attributes). Tables A4-A8, Column 7, show that the prompt reduced the propensity to adopt, but it did not materially alter the estimated causal effects of the vignette attributes. In other words, respondents' reactions to domain, risk, purpose, and mechanism/heritability are not driven by whether general societal/ethical considerations were made salient immediately beforehand.

Stricter attention/comprehension screen. Our main analysis follows the pre-specified rule that respondents pass the attention/comprehension screen if they miss at most one of the four immediate *true/false* comprehension items (and are excluded if they miss two or more). As an additional robustness check, we re-estimate all main regressions restricting the sample to respondents with *zero* errors on all four items (Tables A9-A11). The point estimates and qualitative conclusions are unchanged across adoption, regulation, and ethical-concern outcomes. As expected, some coefficients are estimated less precisely and are therefore less statistically significant in this restricted sample due to the smaller number of observations.

5 Conclusions

Human enhancement technologies are advancing quickly, often first through restorative applications but with clear pathways to supra-typical enhancement. The governance challenge they pose is not only whether they are safe enough, but whether they are seen as socially legitimate. We frame this as a tension between performance-based evaluation, in which attitudes respond to technical attributes such as risk severity and reversibility, and principled constraint, in which opposition reflects considerations not reducible to performance tradeoffs even when performance is favorable. Using a large, nationally representative U.S. survey experiment that randomizes key attributes of hypothetical enhancement technologies, we find that both channels matter and that preferred governance does not follow directly from private demand.

We have three main sets of results. First, about 53% of respondents would adopt the technology described in their vignette, with adoption highest for health applications, lower for productivity, and lowest for cosmetic uses. Two attributes produce the sharpest shifts: severe side effects lower adoption, and framing the same intervention as enhancing rather than restorative lowers it by a comparable amount. Genetic mechanisms do not suppress private adoption; if anything, they modestly raise it.

Second, about 28% of respondents in the health and productivity domains are categorical never-adopters who decline at any benefit. This refusal is driven overwhelmingly by the enhancing framing rather than by risk, consistent with a non-compensatory constraint: for a substantial subgroup, categorical refusal persists regardless of how favorable the performance profile is. Text analysis of open-ended rationales confirms this: never-adopters invoke moral boundaries and human nature far more than adopters, who reason in terms of benefit and usefulness.

Third, we document a demand–governance wedge. Most respondents favor legal access with meaningful oversight; both prohibition and laissez-faire are rare, and support for prohibition is low across all conditions. The same randomized attributes that shape adoption do not map one-for-one into preferred governance. Desired regulation responds to risk severity and heritability, and the response to genetic mechanisms is most pronounced in health:

respondents want more oversight and express greater ethical concern for genetic health interventions even when the change is not inheritable. Productivity shows a different pattern: it raises ethical concern and perceptions of unfairness and reduced mobility, yet attracts more permissive desired regulation than the cosmetic baseline, along with support for subsidizing rather than taxing it. The distributional-incidence treatment corroborates this: responses differ by who benefits from the gains, consistent with respondents reading productivity enhancement as a problem of access and fairness rather than safety. Market uptake is therefore an incomplete signal of social acceptability: governance preferences incorporate intergenerational, distributive, and principled considerations beyond private incentives.

These findings have implications for policy design. First, assessment of HET should extend beyond safety and efficacy. Our results identify three further margins that move public attitudes: intergenerational consequences (heritability), distributional and fairness effects (concern that gains accrue to those who can already afford them, and that mobility may narrow), and the distinction between restoration and enhancement. A governance regime that evaluates these alongside conventional risk assessment will better reflect the range of considerations that shape public attitudes toward these technologies.

Second, the evidence supports oversight rather than bans as the modal form of governance. Support for prohibition is low and weakly responsive to the attributes we randomize. When concern is expressed as demand for oversight rather than prohibition, regulated legality with clear ex ante constraints and credible enforcement is more consistent with observed public preferences than prohibition.

Third, our results are consistent with the cautious approach that existing frameworks already take toward inheritable interventions. Heritability raises ethical concern and desired oversight. That said, our evidence does not establish that the public treats heritability as a categorical, non-tradeable constraint: reservation-benefit responses suggest some respondents would trade it against sufficiently large benefits.

Finally, policy priorities differ across domains. For restorative health applications, governance can prioritize safety, access, and informed consent. For productivity enhancement, the dominant concerns are distributive rather than biomedical: respondents worry about who will be able to afford access and whether the technology will widen existing inequalities. More broadly, where market demand is an incomplete signal of social acceptability, there is a case for steering the direction of technical change toward trajectories that command broad legitimacy, such as restorative, lower-risk, non-inheritable applications, and proceeding more cautiously with those with higher intergenerational consequences or distributional stakes.

These results contribute to an empirically grounded understanding of how the public evaluates contested biomedical innovation. Judgments about one's own use and about permissible use by others appear to be distinct: many who would adopt still want oversight, and many who would not adopt do not want bans. When new technologies generate intergenerational spillovers, distributive consequences, or morally contested uses, the attributes that drive private demand need not be those that drive preferred governance. Future work can examine other institutional settings, track attitudes as these technologies diffuse, and link stated governance preferences to policy designs and political behavior.

References

Acemoglu, Daron (2025). "The simple macroeconomics of AI". In: *Economic Policy* 40.121, pp. 13–58.

Acemoglu, Daron and David Autor (2011). "Skills, Tasks and Technologies: Implications for Employment and Earnings". In: *Handbook of Labor Economics*. Ed. by Orley Ashenfelter and David Card. Vol. 4B. Elsevier, pp. 1043–1171. DOI: 10.1016/S0169-7218(11)02410-5.

Acemoglu, Daron and Simon Johnson (2023). *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*. New York: PublicAffairs. ISBN: 9781541702554.

Acemoglu, Daron, Andrea Manera, and Pascual Restrepo (2020). *Does the U.S. Tax Code Favor Automation?* NBER Working Paper 27052. National Bureau of Economic Research. DOI: 10.3386/w27052. URL: https://www.nber.org/papers/w27052.

Acemoglu, Daron and Pascual Restrepo (2020). "The Wrong Kind of AI? Artificial Intelligence and the Future of Labor Demand". In: *Cambridge Journal of Regions, Economy and Society* 13.1, pp. 25–35. DOI: 10.1093/cjres/rsaa003.

Alesina, Alberto, Alessandro Miano, and Stefanie Stantcheva (2023). "Immigration and Redistribution". In: *Review of Economic Studies* 90.1, pp. 1–39. DOI: 10.1093/restud/rdac058.

Alesina, Alberto, Stefanie Stantcheva, and Edoardo Teso (2018). "Intergenerational Mobility and Preferences for Redistribution". In: *American Economic Review* 108.2, pp. 521–554. DOI: 10.1257/aer.20162015.

Auspurg, Katrin and Thomas Hinz (2015). *Factorial Survey Experiments*. Los Angeles: SAGE Publications. ISBN: 978-1446273102.

Baron, Jonathan and Mark Spranca (1997). "Protected Values". In: *Organizational Behavior and Human Decision Processes* 70.1, pp. 1–16. DOI: 10.1006/obhd.1997.2739.

Bursztyn, Leonardo et al. (2025). *Social Dynamics of AI Adoption*. NBER Working Paper 34488. National Bureau of Economic Research. DOI: 10.3386/w34488. URL: https://www.nber.org/papers/w34488.

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber (2022). "Monetary Policy Communications and Their Effects on Household Inflation Expectations". In: *Journal of Political Economy* 130.6, pp. 1537–1584. DOI: 10.1086/718982.

Cushman, Fiery, Liane Young, and Marc Hauser (2006). "The Role of Conscious Reasoning and Intuition in Moral Judgment: Testing Three Principles of Harm". In: *Psychological Science* 17.12, pp. 1082–1089. DOI: 10.1111/j.1467-9280.2006.01834.x.

Cyranoski, David and Heidi Ledford (2018). "Genome-edited Baby Claim Provokes International Outcry". In: *Nature* 563.7733, pp. 607–608. DOI: 10.1038/d41586-018-07545-0.

Dunsch, Felipe et al. (2018). "Bias in patient satisfaction surveys: a threat to measuring healthcare quality". In: *BMJ Global Health* 3.2, e000694. DOI: 10.1136/bmjgh-2017-000694. URL: https://gh.bmj.com/content/3/2/e000694.

Elías, Julio J., Nicola Lacetera, and Mario Macis (2019). "Paying for Kidneys? A Randomized Survey and Choice Experiment". In: *American Economic Review* 109.8, pp. 2855–2888. DOI: 10.1257/aer.20180282.

Ferrario, Beatrice and Stefanie Stantcheva (2022). "Eliciting people's first-order concerns: Text analysis of open-ended survey questions". In: *AEA Papers and Proceedings*. Vol. 112. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 163–169.

Fiske, Alan P. and Philip E. Tetlock (1997). "Taboo Trade-Offs: Reactions to Transactions That Transgress the Spheres of Justice". In: *Political Psychology* 18.2, pp. 255–297. DOI: 10.1111/0162-895X.00058.

Friis, Simon and James W. Riley (2025). *Performance or Principle: Resistance to Artificial Intelligence in the U.S. Labor Market*. SSRN Working Paper 5560401. Harvard Business School. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5560401.

Funk, Cary and Meg Hefferon (July 2018). *Public Views of Gene Editing for Babies Depend on How It Would Be Used*. Pew Research Center. Report. URL: https://www.pewresearch.org/science/2018/07/26/public-views-of-gene-editing-for-babies-depend-on-how-it-would-be-used/.

Funk, Cary, Brian Kennedy, and Elizabeth Podrebarac Sciupac (July 2016). *U.S. Public Wary of Biomedical Technologies to 'Enhance' Human Abilities*. Pew Research Center. Report. URL: https://www.pewresearch.org/science/2016/07/26/u-s-public-wary-of-biomedical-technologies-to-enhance-human-abilities/.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy (2019). "Text as data". In: *Journal of Economic Literature* 57.3, pp. 535–574.

Goldin, Claudia and Lawrence F. Katz (2009). *The Race between Education and Technology*. Cambridge, MA: Belknap Press.

Graham, Jesse, Jonathan Haidt, and Brian A. Nosek (2009). "Liberals and Conservatives Rely on Different Sets of Moral Foundations". In: *Journal of Personality and Social Psychology* 96.5, pp. 1029–1046. DOI: 10.1037/a0015141.

Guarino, Ben (Jan. 2024). *Elon Musk's Neuralink Has Implanted Its First Chip in a Human Brain. What's Next?* Scientific American. January 30, 2024. URL: https://www.scientificamerican.com/article/elon-musks-neuralink-has-implanted-its-first-chip-in-a-human-brain-whats-next/.

Haaland, Ingar, Christopher Roth, and Johannes Wohlfart (2023). "Designing information provision experiments". In: *Journal of economic literature* 61.1, pp. 3–40.

Hainmueller, Jens, Daniel J. Hopkins, and Teppei Yamamoto (2014). "Causal Inference in Conjoint Analysis: Understanding Multidimensional Choices via Stated Preference Experiments". In: *Political Analysis* 22.1, pp. 1–30. DOI: 10.1093/pan/mpt024.

Hochberg, Leigh R. et al. (2006). "Neuronal Ensemble Control of Prosthetic Devices by a Human with Tetraplegia". In: *Nature* 442.7099, pp. 164–171. DOI: 10.1038/nature04970.

Hood, Amelia J et al. (2025). "Privacy, Policy, and Profits: Survey of Patient Preferences for Research on De-Identified Biosamples". In: *Journal of Empirical Research on Human Research Ethics*, p. 15562646251346940.

Immordino, Giovanni and Fabrizio Panebianco (2026). "From Chance to Choice: The Demand for Human Enhancement Treatments". In: *SSRN Electronic Journal*. SSRN Working Paper. URL: https://ssrn.com/.

Kuziemko, Ilyana et al. (2015). "How Elastic Are Preferences for Redistribution? Evidence from Randomized Survey Experiments". In: *American Economic Review* 105.4, pp. 1478–1508. DOI: 10.1257/aer.20130360.

National Academies of Sciences, Engineering, and Medicine (2017). *Human Genome Editing: Science, Ethics, and Governance*. Washington, DC: National Academies Press. DOI: 10.17226/24623. URL: https://doi.org/10.17226/24623.

Rainie, Lee et al. (Mar. 2022). *AI and Human Enhancement: Americans' Openness Is Tempered by a Range of Concerns*. Pew Research Center. Report. URL: https://www.pewresearch.org/internet/2022/03/17/ai-and-human-enhancement-americans-openness-is-tempered-by-a-range-of-concerns/.

Roth, Alvin E. (2007). "Repugnance as a Constraint on Markets". In: *Journal of Economic Perspectives* 21.3, pp. 37–58. DOI: 10.1257/jep.21.3.37.

Sandel, Michael J. (2009). *The Case Against Perfection: Ethics in the Age of Genetic Engineering*. Cambridge, MA: Harvard University Press.

Siegrist, Michael (2000). "The Influence of Trust and Perceptions of Risks and Benefits on the Acceptance of Gene Technology". In: *Risk Analysis* 20.2, pp. 195–204. DOI: 10.1111/0272-4332.202020.

Slovic, Paul (1987). "Perception of Risk". In: *Science* 236.4799, pp. 280–285. DOI: 10.1126/science.3563507. URL: https://www.science.org/doi/10.1126/science.3563507.

Stantcheva, Stefanie (2023). "How to run surveys: A guide to creating your own identifying variation and revealing the invisible". In: *Annual Review of Economics* 15.1, pp. 205–234.

Sunstein, Cass R. (2002). "Probability Neglect: Emotions, Worst Cases, and Law". In: *The Yale Law Journal* 112.1, pp. 61–107. URL: https://yalelawjournal.org/article/probability-neglect-emotions-worst-cases-and-law.

Tetlock, Philip E. et al. (2000). "The Psychology of the Unthinkable: Taboo Trade-Offs, Forbidden Base Rates, and Heretical Counterfactuals". In: *Journal of Personality and Social Psychology* 78.5, pp. 853–870. DOI: 10.1037/0022-3514.78.5.853.

U.S. Food and Drug Administration (Dec. 2023). *FDA Approves First Gene Therapies to Treat Patients with Sickle Cell Disease*. FDA News Release. December 8, 2023. URL: https://www.fda.gov/news-events/press-announcements/fda-approves-first-gene-therapies-treat-patients-sickle-cell-disease.

World Health Organization (2021). *Human Genome Editing: A Framework for Governance*. Geneva: World Health Organization. URL: https://www.who.int/publications/i/item/9789240030381.

A Survey Instrument

Reader note. Section and subsection titles in this appendix, such as "Common pre-scenario material," are editorial headings and were not shown to respondents. Unless explicitly labeled as a construction note, the text reproduces the wording shown to respondents. Construction notes were not shown to respondents and are included only to explain how randomized modules were combined.

A.1 Common pre-scenario material

Consent screen

This study on "Understanding Attitudes about Human Enhancement Interventions" is conducted by university-based researchers. The study was approved by the Ethics Board of the University of Naples Federico II. In the survey, we will ask you to express your opinions regarding human enhancement interventions, under different hypothetical scenarios, as well as other questions regarding your socio-demographic characteristics and preferences. It should take you about 10 minutes to complete the survey diligently. Withdrawal is possible at any time if you so desire (any data collected will be destroyed). Compensation is conditional on diligently completing the entire survey. By completing this survey, you are consenting to be in this research study. Your participation is voluntary and you can stop at any time. Note that all of the answers that you provide will remain anonymous and treated with confidentiality. The researchers do not know your identity, and they will not be able to match your name with the answers that you provide. The Investigators can be contacted for questions: Prof. Giovanni Immordino (University of Napoli): giovanni.immordino@unina.it; Prof. Mario Macis (Johns Hopkins University): mmacis@jhu.edu; Prof. Immacolata Marino (University of Napoli): immacolata.marino@unina.it; Prof. Fabrizio Panebianco (Catholic University, Milan): fabrizio.panebianco@unicatt.it. If you have questions about your rights as a research participant, you may contact the Ethics Board at the University of Napoli Federico II at guido.durso@unina.it.

Do you agree to participate in this study?

- Yes, I agree to participate in this study.
- No, I do not agree to participate in this study.

Overview of the survey

Overview of the Survey. This survey has three parts: Part 1. Basic demographics. Part 2. Your opinions about human enhancement technology. Part 3. Additional socio-demographics and general preferences.

Initial demographics

What is your age in years?

Please select your age from the scroll down list.

Response options: 18, 19, . . . , 100.

What is your gender?

- Female
- Male
- Non-binary / Third gender
- Prefer to self-describe ________________________
- Prefer not to answer

What is the racial/ethnic group you most identify with?

- African American/Black
- Asian
- Hispanic/Latino
- Native American / Alaska Native
- White
- Multiracial/Other
- Prefer not to answer

What is your highest degree of education attained?

- 8th grade or less
- Some high school
- High school degree/GED
- Some college
- 2-year college degree
- 4-year college degree
- Masters degree

- Doctoral degree
- Professional degree (JD, MD, MBA)

What is your state of residence?

Response options: Alabama, . . . , Wyoming.

Technology orientation

Thinking about your attitudes toward technology in general, to what extent do you agree or disagree with the following statements?

Scale: 0 = strongly disagree, 1, 2, 3, 4, 5 = neutral, 6, 7, 8, 9, 10 = strongly agree.

- I am usually among the first of my peers to try new technologies (apps, devices, online services).
- I like to experiment with new technologies even if they are unfamiliar or complex.
- I actively seek information about new technological products or services (e.g., from specialized websites, magazines, or knowledgeable friends).

Familiarity with human enhancement

How familiar are you with "human enhancement" interventions or technologies?

- I have never heard of human enhancement interventions or technologies before.
- I have heard of them, but I don't know much about them.
- I am somewhat familiar; I know a few things about them.
- I am quite familiar; I have a good understanding of many aspects.
- I am very familiar; I consider myself well-informed or an expert.

Common definition shown before random assignment

Please read the following brief overview of human enhancement technologies.

Human enhancement technologies, ranging from advanced genetic engineering to neural implants, aim to augment human abilities beyond typical levels. These technologies have the potential to significantly enhance physical, mental, and sensory capacities, opening new possibilities for improved health, performance, and well-being. In addition to their use for enhancement, such technologies may also be employed for restorative purposes, helping individuals recover lost functions or address impairments. However, these technologies may also carry risks, including potential unexpected side effects and unintended consequences that could impact individuals and society.

Expectations block

Construction note (not shown to respondents). This block was shown before the scenario module to the respondents assigned to the expectations condition.

What are your expectations regarding the impact of human enhancement technologies on people's wellbeing?

- I expect these technologies will lead to positive physical or psychological consequences for individuals and society as a whole.
- I expect these technologies will lead to negative physical or psychological consequences for individuals and society as a whole.
- I am currently undecided about the impact of these technologies on people's wellbeing.

What are your expectations regarding access and societal impact of human enhancement technologies?

- I expect these technologies to be accessible and beneficial to everyone, regardless of their socio-economic status.
- I expect that these technologies will primarily benefit the wealthy, widening the gap between different socio-economic groups.
- I expect that these technologies will primarily benefit the poor, reducing the gap between different socio-economic groups.
- I am currently undecided about the access and societal impact of these technologies.

What are your expectations regarding the ethical implications of human enhancement technologies?

- I expect that these technologies will create significant ethical dilemmas.
- I expect these technologies to be broadly ethically acceptable and unproblematic.
- I am currently undecided about the ethical implications of these technologies.

Transition to the randomized scenario

Next, we will ask you to consider a specific scenario. Please read it carefully. Immediately after, you will see four brief comprehension (attention) questions about the scenario. To ensure data quality, if you answer two or more of these questions incorrectly, we will not be able to include your responses in the study and the survey may end early. Your thoughtful insights are invaluable to our understanding of public perceptions—thank you for taking the time.

A.2 Scenario construction

Construction note (not shown to respondents). Each respondent was assigned to one vignette. The scenario catalogue and field labels below were not shown to respondents; they reconstruct the randomized instrument without repeating every possible vignette in full. Each scenario family was crossed with one of two adverse-event profiles: mild/temporary or severe/long-lasting, both at a fixed probability of 1 in 100 (1%).

Table A1. Scenario catalogue, not shown to respondents

ID	Domain	Purpose	Mechanism / heritability	Distributional incidence
A1	Appearance	Restoration	Genetic, non-inheritable	–
A2	Appearance	Restoration	Genetic, inheritable	–
A3	Appearance	Restoration	Non-genetic implant	–
A4	Appearance	Enhancement	Genetic, non-inheritable	–
A5	Appearance	Enhancement	Genetic, inheritable	–
A6	Appearance	Enhancement	Non-genetic implant	–
H1	Health	Restoration	Genetic, non-inheritable	–
H2	Health	Restoration	Genetic, inheritable	–
H3	Health	Restoration	Non-genetic implant	–
H4	Health	Enhancement	Genetic, non-inheritable	–
H5	Health	Enhancement	Genetic, inheritable	–
H6	Health	Enhancement	Non-genetic implant	–
P1	Productivity	Restoration	Genetic, non-inheritable	–
P2	Productivity	Restoration	Genetic, inheritable	–
P3	Productivity	Restoration	Non-genetic implant	–
P4	Productivity	Enhancement	Genetic, non-inheritable	–
P5	Productivity	Enhancement	Genetic, inheritable	–
P6	Productivity	Enhancement	Non-genetic implant	–
P7	Productivity	Enhancement	Non-genetic implant	College-skill advantage
P8	Productivity	Enhancement	Non-genetic implant	High-school-skill advantage

Master vignette skeleton

Construction note (not shown to respondents). The bracketed field labels were not displayed. Respondents saw the field insertions combined into a single vignette.

> [1. Mechanism primer]
>
> [2. Application and aim]
>
> How it works: [3. How-it-works line, if included]
>
> What it does: [4. Domain-specific effect]
>
> Main benefit: [5. Benefit wording]
>
> Passes to children? [6. Mechanism / heritability wording]
>
> Goal: [7. Purpose wording]
>
> Risk: There is a 1 in 100 (1%) risk of [8. Adverse-event wording]. Please answer the questions below.

Field 1. Mechanism primer

Mechanism	Text inserted in Field 1
Genetic, non-inheritable or inheritable	CRISPR is a gene-editing tool that functions like molecular scissors to make precise changes in DNA.
Non-genetic implant	Biotechnology implants are small devices placed on or within the body that can sense internal signals and release medication or stimulation when needed.

Fields 2–5 and 7. Scenario-family-specific text

Table A3. Scenario-family-specific respondent-facing text

ID	Insertions
A1	**Restoration; genetic, non-inheritable.** Field 2: Consider a specific application of CRISPR that repairs genes linked to skin aging that cause the person's skin to age faster and look worse than normal. The aim is to restore a person's skin—reducing wrinkles, sun damage, and uneven tone—back to a normal, healthy appearance. Field 3: No line is inserted. Field 4: Restores skin damaged by wrinkles and sun exposure to a normal, healthy appearance. Field 5: Appearance. Field 7: Bring skin back to typical appearance.

ID	Insertions
A2	**Restoration; genetic, inheritable.** Same as A1, with Field 6 set to genetic, inheritable.
A3	**Restoration; non-genetic implant.** Field 2: Consider biotechnology implants designed to restore a person's appearance, specifically to reduce premature or excessive signs of aging and improve overall skin appearance. The aim is to bring skin with more wrinkles, sun damage, and uneven tone than is typical for someone's age back to a normal, healthy look. Field 3: No line is inserted. Field 4: Smooths wrinkles and sun damage, restoring skin to a normal, healthy appearance. Field 5: Appearance. Field 7: Bring skin back to typical appearance.
A4	**Enhancement; genetic, non-inheritable.** Field 2: Consider a specific application of CRISPR that enhances genes linked to skin aging. The aim is to give a person visibly younger-looking skin, with fewer wrinkles and improved tone beyond what is typical for their age. Field 3: No line is inserted. Field 4: Smooths wrinkles and sun damage, giving skin a visibly younger look beyond what is typical for one's age. Field 5: Appearance. Field 7: Look younger than typical aging.
A5	**Enhancement; genetic, inheritable.** Same as A4, with Field 6 set to genetic, inheritable.
A6	**Enhancement; non-genetic implant.** Field 2: Consider biotechnology implants designed to enhance a person's appearance, specifically to reduce visible signs of aging and improve overall skin appearance. The aim is to give a person visibly younger-looking skin, with fewer wrinkles and improved tone beyond what is typical for their age. Field 3: No line is inserted. Field 4: Smooths wrinkles and sun damage, giving skin a visibly younger look beyond what is typical for one's age. Field 5: Appearance. Field 7: Look younger than typical aging.
H1	**Restoration; genetic, non-inheritable.** Field 2: Consider a specific application of CRISPR that repairs genes linked to the immune system that currently leave the person unusually vulnerable to infection. The aim is to bring a person's natural resistance to common respiratory infections—such as bronchitis, COPD, and related illnesses—back up to normal levels. Field 3: One-time CRISPR treatment on blood cells. Field 4: Cuts respiratory infections down to average levels. Field 5: Better health. Field 7: Restore normal immunity.
H2	**Restoration; genetic, inheritable.** Same as H1, except Field 3: No line is inserted; Field 6 set to genetic, inheritable.

ID	Insertions
H3	**Restoration; non-genetic implant.** Field 2: Consider biotechnology implants designed to monitor and adjust body functions in individuals who are unusually vulnerable to infection. The aim is to bring a person's resistance to common respiratory infections—such as bronchitis, COPD, and related illnesses—back up to normal levels. Field 3: No line is inserted. Field 4: Cuts respiratory infections down to average levels. Field 5: Better health. Field 7: Restore normal immunity.
H4	**Enhancement; genetic, non-inheritable.** Field 2: Consider a specific application of CRISPR that enhances genes linked to the immune system. The aim is to make a person naturally resistant to common respiratory infections—such as bronchitis, COPD, and related illnesses—well above normal levels. Field 3: One-time CRISPR treatment on blood cells. Field 4: Gives immunity to common respiratory diseases. Field 5: Better health. Field 7: Immunity above normal.
H5	**Enhancement; genetic, inheritable.** Same as H4, except Field 3: No line is inserted; Field 4: Gives immunity to common respiratory diseases; Field 6 set to genetic, inheritable.
H6	**Enhancement; non-genetic implant.** Field 2: Consider biotechnology implants designed to monitor and adjust body functions; their aim is to make a person resistant to common respiratory infections—such as bronchitis, COPD, and related illnesses—well above normal levels. Field 3: No line is inserted. Field 4: Gives immunity to common respiratory diseases. Field 5: Better health. Field 7: Immunity above normal.
P1	**Restoration; genetic, non-inheritable.** Field 2: Consider a specific application of CRISPR that repairs genes linked to memory, concentration, and other cognitive functions in individuals with cognitive difficulties due to factors such as neurological conditions or past injuries. The aim is to restore cognitive abilities—memory, focus, and learning speed—back to the levels seen in the general population, supporting study, job performance, and career opportunities. Field 3: No line is inserted. Field 4: Restores memory and concentration to normal levels. Field 5: Better cognitive function for study, work, and everyday tasks. Field 7: Bring thinking ability back to average.
P2	**Restoration; genetic, inheritable.** Same as P1, with Field 6 set to genetic, inheritable.

ID	Insertions
P3	**Restoration; non-genetic implant.** Field 2: Consider biotechnology implants that restore normal memory, concentration, and other cognitive functions in individuals with cognitive difficulties due to factors such as neurological conditions or past injuries. The aim is to restore cognitive abilities—memory, focus, and learning speed—back to the levels seen in the general population, supporting study, job performance, and career opportunities. Field 3: No line is inserted. Field 4: Restores memory and concentration to normal levels. Field 5: Better cognitive function for study, work, and everyday tasks. Field 7: Bring thinking ability back to average.
P4	**Enhancement; genetic, non-inheritable.** Field 2: Consider a specific application of CRISPR that enhances genes linked to memory, concentration, and other cognitive functions. The aim is to boost cognitive capacity—memory, focus, and learning speed—well above the levels seen in the general population, potentially improving study performance, job productivity, and career opportunities. Field 3: No line is inserted. Field 4: Raises memory and concentration above normal levels. Field 5: Stronger cognitive function for advanced study, higher workplace performance, and everyday tasks. Field 7: Achieve thinking ability above average.
P5	**Enhancement; genetic, inheritable.** Same as P4, with Field 6 set to genetic, inheritable.
P6	**Enhancement; non-genetic implant.** Field 2: Consider biotechnology implants that enhance memory, concentration, and other cognitive functions. The aim is to boost cognitive capacity—memory, focus, and learning speed—well above the levels seen in the general population, potentially improving study performance, job productivity, and career opportunities. Field 3: No line is inserted. Field 4: Raises memory and concentration above normal levels. Field 5: Stronger cognitive function for advanced study, higher workplace performance, and everyday tasks. Field 7: Achieve thinking ability above average.
P7	**Enhancement; non-genetic implant; college-skill advantage.** Same as P6, with this added to Field 2: People whose work or studies rely on skills usually acquired through a college education are expected to gain the greatest advantage from this technology. Field 5: Stronger cognitive function for advanced study, higher workplace performance, and everyday tasks—especially for individuals with college-level education or higher.

ID	Insertions
P8	**Enhancement; non-genetic implant; high-school-skill advantage.** Same as P6, with this added to Field 2: People whose jobs rely mainly on skills learned through a high-school education or less are expected to gain the greatest advantage from this technology. Field 5: Stronger cognitive function for advanced study, higher workplace performance—especially for individuals whose work relies on high-school-level skills.

Field 6. Mechanism / heritability wording

Mechanism / heritability	Text inserted after "Passes to children?"
Genetic, non-inheritable	No — changes affect the treated individual only.
Genetic, inheritable	Yes — inherited by future generations. In the cognitive/productivity domain, the wording is: "Yes — the change is inherited by future generations."
Non-genetic implant	No — the technology does not alter genetics.

Field 8. Adverse-event wording

Risk profile	Text completing the risk sentence
Mild and temporary	mild, temporary side effects such as rashes or nausea
Severe and long-lasting	severe, long-lasting side effects such as neurological damage or serious immune reactions

A.3 Scenario comprehension checks

Construction note (not shown to respondents). The bracketed text below indicates the wording that changed according to the assigned scenario. Respondents saw the corresponding completed True/False statements, not the brackets.

1. True or False? The main benefit of the technology described in the scenario concerns [appearance / health / productivity at work or in study].
2. True or False? The scenario described [a genetic change affecting only the treated person / an inherited genetic change / an implanted device that does not change genetics].
3. True or False? The goal described was to [restore or enhance] [appearance, health, or cognitive function] [to typical levels or beyond typical levels].

4. True or False? The scenario stated there is a 1% risk of [mild, temporary / severe, long-lasting] side effects.

Each item had the following options:

- True
- False

A.4 Post-scenario outcome modules

Construction note (not shown to respondents). The following modules were asked after the assigned scenario. The order of the ethics, adoption, and regulation modules varied by randomized condition.

Ethical concerns

How do you view the ethical implications of the technology described in the scenario above? Please indicate your stance using the slider below, where: 0 = No Ethical Concerns: No ethical issues identified with the technology 1–2–3 = Minor Ethical Concerns: Some concerns exist but are relatively minor and manageable 4–6 = Moderate Ethical Concerns: Notable ethical issues that require careful consideration 7–10 = Significant Ethical Concerns: Major ethical challenges that may outweigh the benefits

Scale: 0 = No Ethical Concerns, 1, 2–3 = Minor Ethical Concerns, 4–6 = Moderate Ethical Concerns, 7–10 = Significant Ethical Concerns.

In general, which ethical implications of the technology described in the scenario above concern you the most? Please select all that apply.

- Exacerbating social and economic inequality.
- "Playing God" and overstepping moral or religious boundaries.
- Fundamentally altering human nature and identity.
- Limiting the freedom and autonomy of future generations (offspring).
- Risk of eugenic drift leading to controlled or selective breeding practices.
- Risk of unintended consequences affecting health or environment.
- Potential for misuse or abuse in military or commercial applications.
- Privacy concerns related to data collection and surveillance.

- Creating dependencies on technology for basic human functions.
- Other ethical concerns (please specify) ________________________
- I have no ethical concerns.

Regulation, adoption, and willingness to pay

What is your opinion on the regulation and legal status of the technology described in the scenario above? Please indicate your stance using the slider below, where: 0 = No Regulation: Fully legal with no government regulation or oversight 1–3 = Minimal Regulation: Legal with basic safety standards and limited oversight 4–6 = Moderate Regulation: Legal with standard safety, efficacy, and informed consent rules (but not overly restrictive) 7–9 = Strict Regulation: Legal but tightly controlled (e.g., mandatory clinical trials, ethical oversight, long-term monitoring) 10 = Prohibition: Total ban on the technology (no development or use allowed)

Scale: 0 = No Regulation, 1–3 = Minimal Regulation, 4–6 = Moderate Regulation, 7–9 = Strict Regulation, 10 = Prohibition.

If you could obtain this technology today, would you adopt it?

- Yes, I would adopt it.
- No, I would not adopt it.

What is the maximum one-time amount (in US dollars) you would be willing to pay for this technology? Please enter a whole-dollar amount (e.g., 1500). If the amount is zero, write 0. Please do not enter spaces or any other characters, only numbers.

Maximum one-time payment (US $) you would pay ________________________

Reservation-benefit follow-up for initial non-adopters

Construction note (not shown to respondents). The health follow-up was shown to initial non-adopters in health scenarios. The productivity follow-up was shown to initial non-adopters in productivity scenarios. This follow-up was not shown in appearance scenarios.
Is there an increase in healthy life expectancy that would make you switch to adopting the technology? Move the slider to the minimum number of extra healthy-life years that would make you consider adopting the technology. For example: If you would adopt the technology if it gave you at least 1 extra year, slide to 1. If you'd need at least 5 years, slide to 5. Etc. Please choose only one option: Use the slider if any amount of life extension would make the technology

worth considering. Select the checkbox if you would not adopt the technology regardless of the life extension benefit.

Options: I would not adopt the technology regardless of the benefit; or scale: 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10 life-extension years required.

Is there a salary increase that would make you switch to adopting the technology? Move the slider to the minimum annual salary increase that would make you consider adopting the technology. For example: If you would adopt the technology if it increased your salary by at least 5%, slide to 5. If you'd need a 50% salary increase, slide to 50. If you'd need your salary to double, slide to 100. Please choose only one option: Use the slider if any salary increase would make the technology worth considering. Select the checkbox if you would not adopt the technology regardless of the salary increase benefit.

Options: I would not adopt the technology regardless of the benefit; or scale: 0, 5, 10, . . . , 100% salary increase required.

Additional social-evaluation modules

Construction note (not shown to respondents). Respondents were randomly assigned to one of the two versions below.

Scale for each item: 0 = strongly disagree, 1, 2, 3, 4, 5 = neutral, 6, 7, 8, 9, 10 = strongly agree.

Negatively framed version.

- This technology is likely to harm individuals and society overall.
- This technology will make it harder for most people to move up in society.
- It is unfair if only those who can afford this technology are able to benefit from it.
- It is wrong to impose taxes on people who choose to use this technology.
- Public funds should not be used to subsidize this technology.

Positively framed version.

- This technology will contribute positively to individuals and society overall.
- This technology will help most people improve their social and economic status.
- It is fair if access to this technology depends on the ability to pay for it.
- Taxing the use of this technology is a good way to support those who may be negatively affected by it.

- Subsidizing this technology with public funds would help promote its fair and widespread use.

Open-ended rationale

What were the main factors that influenced your answers regarding the technology considered above? Please provide a brief summary in the space below.

A.5 Final background and preference measures

Additional socio-demographics

Now, we would like you to please answer a few additional socio-demographic questions.

What is your current relationship status? (select the option that best applies to you)

- Single
- Unmarried but in a relationship
- Married/Domestic partnership
- Separated/divorced
- Widow(er)
- Other (specify) ______________________

What is your parental status? (select all that apply)

- I have children under age 18
- I have children older than 18
- I do not have children

Which of the following best describes your current labor market status?

- Employed full time
- Employed part time
- Self-employed/Entrepreneur
- Unemployed
- Homemaker

- Student
- Retired
- Other (specify) ________________________

Approximately, what was your total household income, before taxes, in 2024?

- $0–$9,999
- $10,000–$19,999
- $20,000–$29,999
- $30,000–$39,999
- $40,000–$49,999
- $50,000–$74,999
- $75,000–$99,999
- $100,000–$149,999
- $150,000–$199,999
- $200,000+
- Prefer not to answer

What is your current religion, if any?

- Atheist/Agnostic
- Christian
- Jewish
- Muslim
- Other (please specify below) ________________________
- Prefer not to answer

How often do you typically attend religious or worship services (not including weddings and funerals) when able (e.g. when in good health)?

- More than once a week
- Once a week

- Once or twice a month
- A few times a year
- Seldom
- Never

Which of the following best describes the area where you currently live?

- City with a population over 1 million
- City with a population of 500,000 to 1 million
- City with a population of 100,000 to 499,999
- City or town with a population of 50,000 to 99,999
- Town with a population of 10,000 to 49,999
- Small town or village with a population of less than 10,000
- Rural area (not in a town, village, or city)

Health and medication background

How would you rate your current health status? Would you say that it is

- Excellent
- Very good
- Good
- Fair
- Poor

Have you ever been diagnosed with any of the following? Please select all that apply.

- Asthma, chronic bronchitis, or COPD
- Cancer
- Diabetes or pre-diabetes
- Heart attack, heart disease, or other heart condition
- High blood pressure or cholesterol

- Kidney disease
- Stroke
- Muscular dystrophy
- Skin conditions
- Inherited genetic condition (one that can be passed on to your children)
- Anxiety and/or depression
- Covid-19
- Other ______________________
- None

Do you regularly take any medications for your health conditions? If yes, please select all that apply

- Prescription medications for a chronic condition
- Over-the-counter medications for symptom relief
- Vitamins or dietary supplements
- Herbal supplements or natural remedies
- Other ______________________
- No, I do not take any medications regularly

Trust and risk preferences

Finally, just a few more questions about your general perceptions and preferences.

How much do you TRUST each of the following:

Scale: very low trust, low trust, moderate trust, high trust, very high trust.

- The President of the USA
- Medical doctors
- Scientists
- Members of Congress
- Other people in general

- Federal regulatory agencies
- The media

People can behave differently in different situations. On a scale where the value 0 means: "risk averse" and the value 10 means: "fully prepared to take risks", how would you rate your willingness to take risks in the following areas?

Scale: 0 = Totally risk averse, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10 = Fully prepared to take risks.

- while driving
- in financial matters
- during leisure and sport
- in your occupation
- with your health

Moral-foundation items

When you decide if something is right or wrong, how relevant is each of the following? Please rate each item on the scale shown.

Scale: Not at all relevant to my view of right/wrong, not very relevant, slightly relevant, somewhat relevant, very relevant, Extremely relevant—key to my view of right/wrong.

- Some people were treated differently from others.
- Standards of purity and decency were violated.
- Someone cared for someone weak or vulnerable.
- Traditions of society were followed.
- Something disgusting was done, even if no one was harmed.
- Someone acted in a way God would approve.
- Compassion was shown for the suffering.
- Unnatural acts were performed.

B Additional Figures and Tables

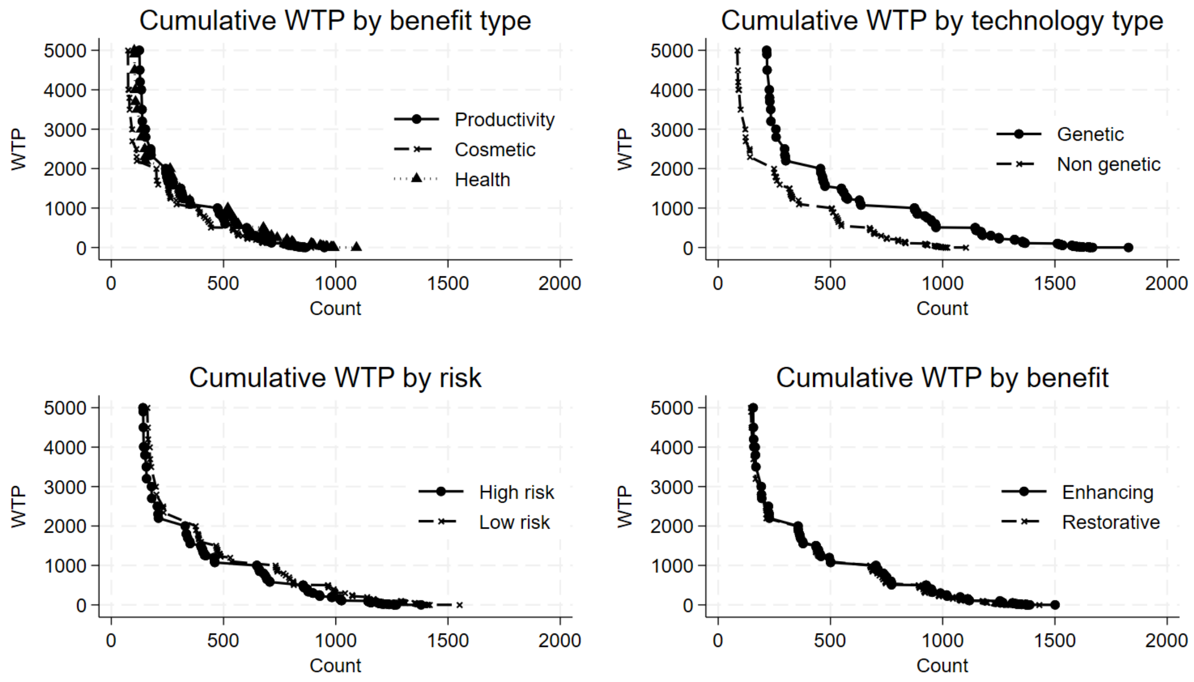

Figure B1. Demand Curves

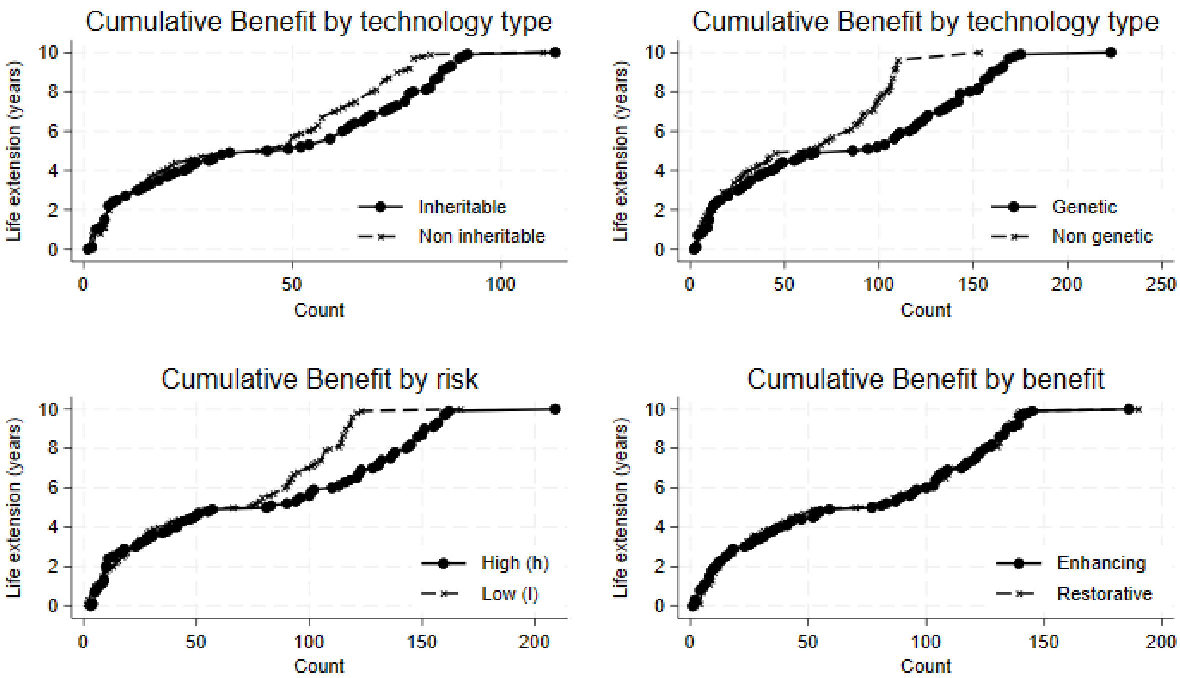

Figure B2. Reservation Benefits among Initial Non-Adopters - Health

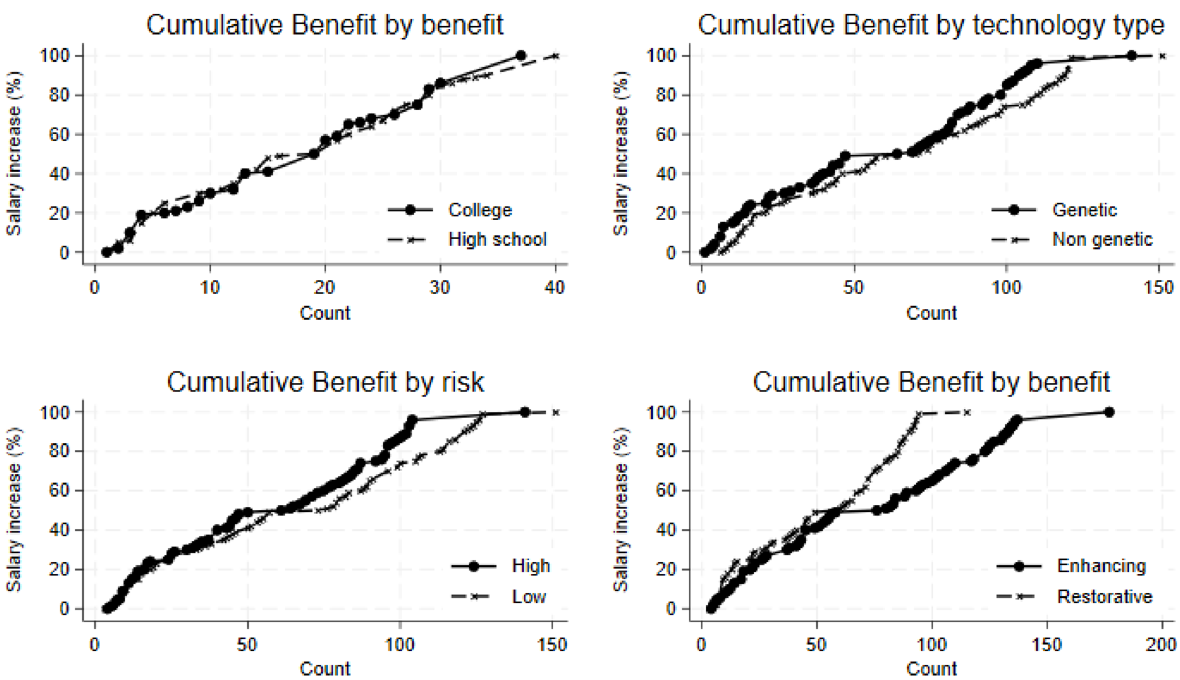

Figure B3. Reservation Benefits among Initial Non-Adopters - Productivity

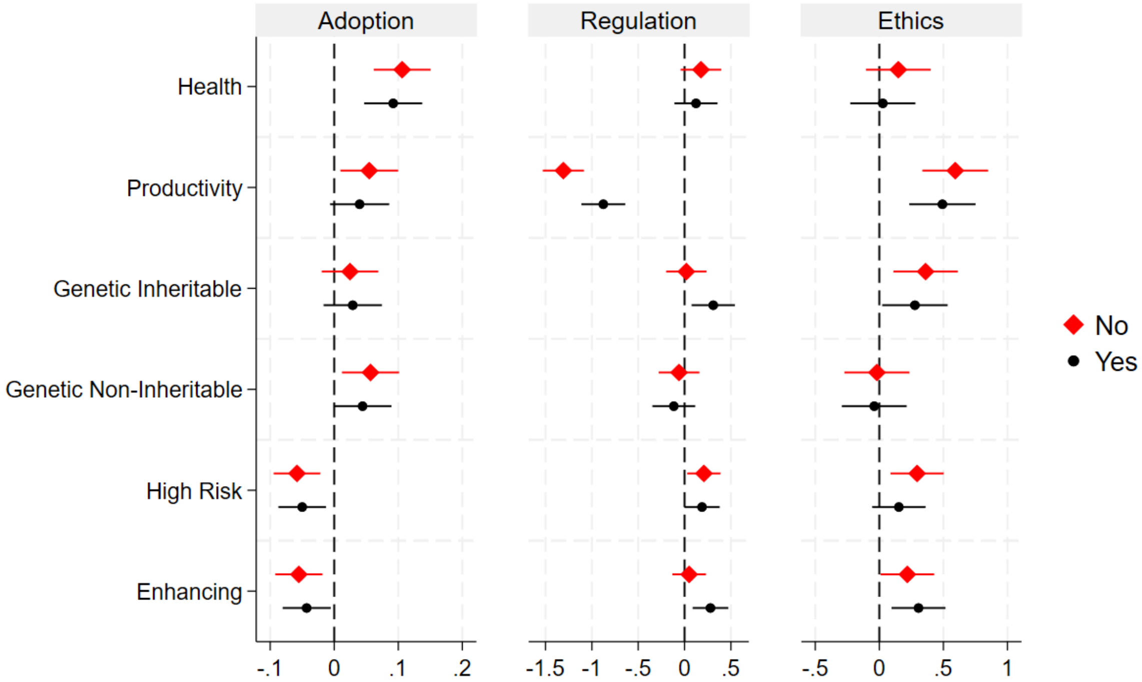

Figure B4

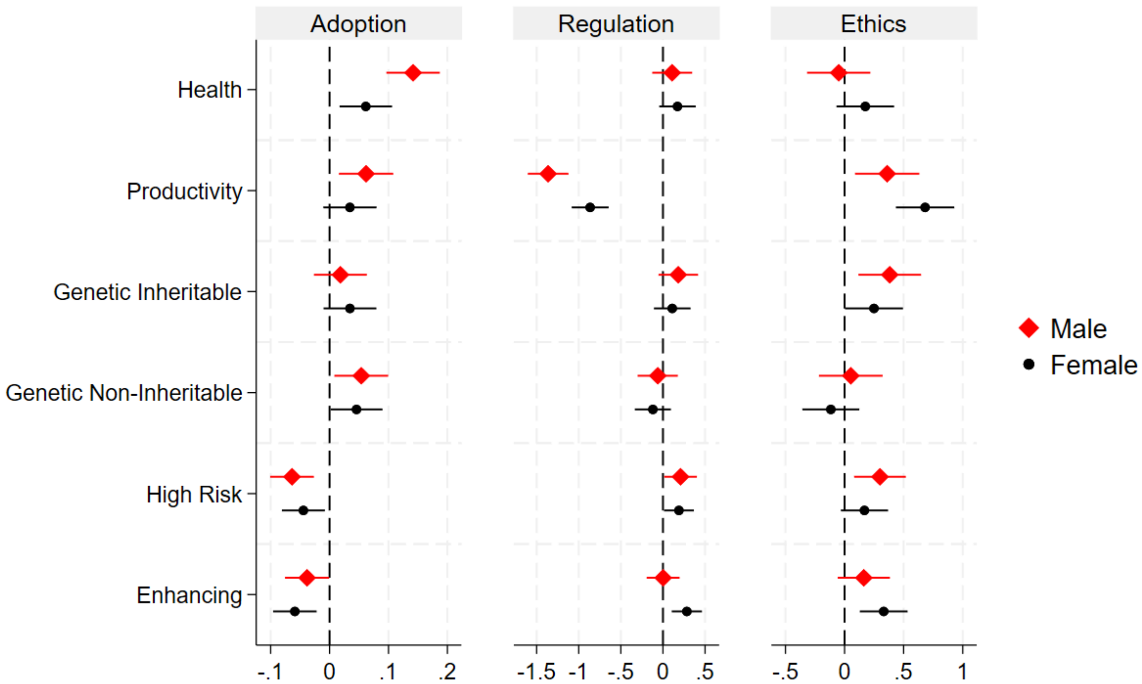

Figure B5

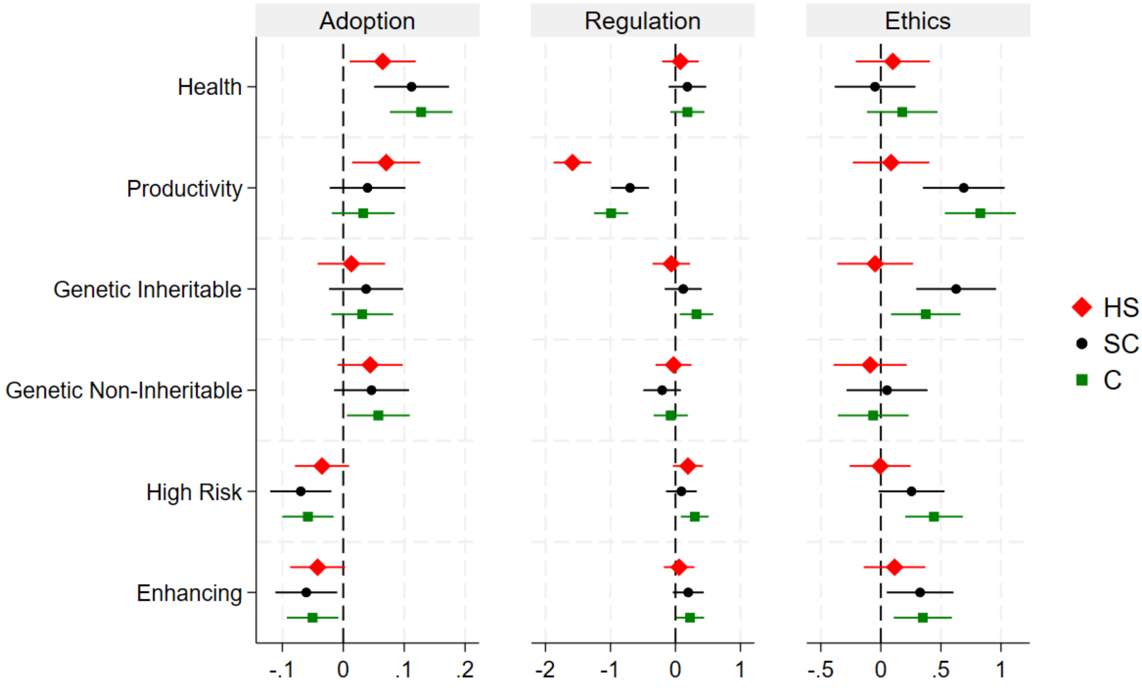

Figure B6

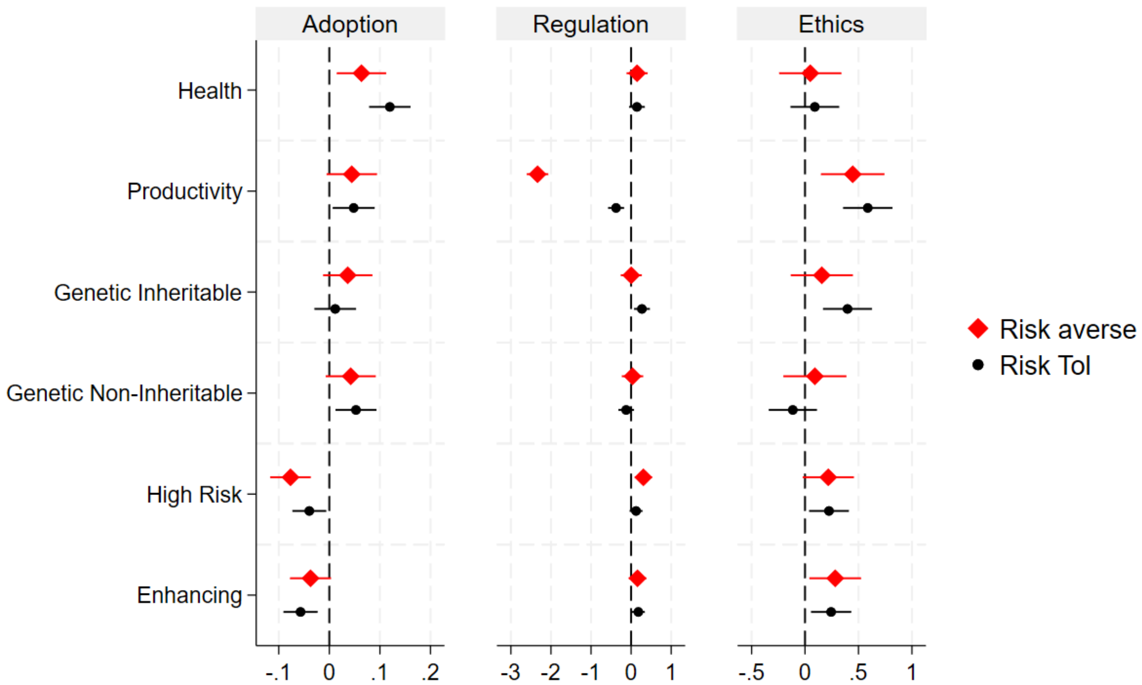

Figure B7

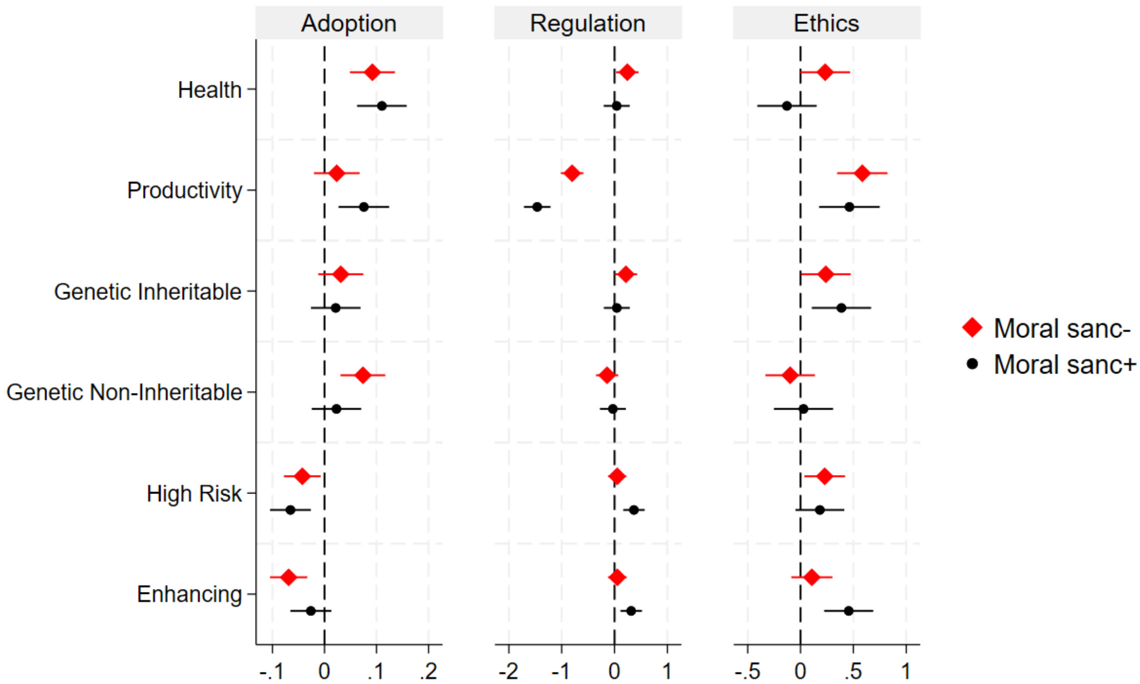

Adoption
Regulation
Ethics
Health
Productivity
Genetic Inheritable
Genetic Non-Inheritable
High Risk
Enhancing
-.1
0
.1
.2
-2
-1
0
1
-.5
0
.5
1
Moral sanc-
Moral sanc+

Figure B8

Table B1. Descriptive statistics with survey labels

Variable	**Mean**	**SD**	**Min**	**Max**	**N**
Scenario domain: appearance/cosmetic	0.335	0.472	0.000	1.000	5556
Scenario domain: health	0.337	0.473	0.000	1.000	5556
Scenario domain: productivity/cognition	0.328	0.470	0.000	1.000	5556
Technology type: CRISPR gene editing	0.604	0.489	0.000	1.000	5556
Technology type: biotechnology implant/non-genetic	0.396	0.489	0.000	1.000	5556
Risk: severe, long-lasting side effects	0.496	0.500	0.000	1.000	5556
Risk: mild, temporary side effects	0.504	0.500	0.000	1.000	5556
Genetic change inherited by future generations	0.503	0.500	0.000	1.000	3358
Not inherited; affects only the treated individual	0.497	0.500	0.000	1.000	3358
Goal: enhancement beyond typical levels	0.540	0.498	0.000	1.000	5556
Goal: restoration to typical levels	0.460	0.498	0.000	1.000	5556
Scenario detail: appearance/cosmetic	0.335	0.472	0.000	1.000	5556
Scenario detail: health	0.337	0.473	0.000	1.000	5556
Scenario detail: productivity/cognition	0.250	0.433	0.000	1.000	5556
Scenario detail: productivity, college-level skills benefit most	0.040	0.196	0.000	1.000	5556
Scenario detail: productivity, high-school-level skills benefit most	0.038	0.191	0.000	1.000	5556
Technology/inheritance: CRISPR, inherited	0.304	0.460	0.000	1.000	5556
Technology/inheritance: CRISPR, not inherited	0.300	0.458	0.000	1.000	5556
Technology/inheritance: biotechnology implant/non-genetic	0.396	0.489	0.000	1.000	5556
Completion time in minutes	13.110	60.142	3.367	3705.417	5556
What is your age in years?	48.845	17.994	18.000	96.000	5556
Gender: female	0.536	0.499	0.000	1.000	5556
Gender: male	0.455	0.498	0.000	1.000	5556
Gender: non-binary/third gender, self-described, or prefer not to answer	0.009	0.094	0.000	1.000	5556
Race/ethnicity: African American/Black	0.126	0.332	0.000	1.000	5556
Race/ethnicity: Asian	0.063	0.242	0.000	1.000	5556
Race/ethnicity: Hispanic/Latino	0.151	0.358	0.000	1.000	5556
Race/ethnicity: Multiracial/Other	0.044	0.205	0.000	1.000	5556
Race/ethnicity: Native American/Alaska Native	0.010	0.101	0.000	1.000	5556
Race/ethnicity: prefer not to answer	0.003	0.055	0.000	1.000	5556
Race/ethnicity: White	0.603	0.489	0.000	1.000	5556
Education category (ordered numeric)	2.024	0.847	1.000	3.000	5556
Education: high school or less	0.347	0.476	0.000	1.000	5556
Education: some college	0.282	0.450	0.000	1.000	5556
Education: college degree or higher	0.371	0.483	0.000	1.000	5556
Relationship status: married/domestic partnership or relationship	0.507	0.500	0.000	1.000	5556
Relationship status: separated/divorced, widow(er), or other	0.193	0.395	0.000	1.000	5556
Relationship status: single	0.300	0.458	0.000	1.000	5556

Continued on next page

Table B1. Descriptive statistics with survey labels (continued)

Variable	Mean	SD	Min	Max	N
Parental status: no children	0.383	0.486	0.000	1.000	5406
Parental status: children older than 18	0.337	0.473	0.000	1.000	5406
Parental status: children under age 18	0.237	0.425	0.000	1.000	5406
Parental status: children under and older than 18	0.040	0.196	0.000	1.000	5406
Parental status: other	0.003	0.054	0.000	1.000	5406
Employment status: employed full time or part time	0.491	0.500	0.000	1.000	5406
Employment status: homemaker, student, or other	0.110	0.313	0.000	1.000	5406
Employment status: retired	0.247	0.431	0.000	1.000	5406
Employment status: self-employed/entrepreneur	0.058	0.233	0.000	1.000	5406
Household income: category midpoint/average	73836.240	58734.812	5000.000	250000.000	5276
Household income: lower bound of category	60837.756	48929.870	0.000	200000.000	5276
Household income: upper bound of category	87807.992	69771.223	9999.000	300000.000	5276
Religion: atheist/agnostic	0.123	0.328	0.000	1.000	5406
Religion: Christian	0.648	0.478	0.000	1.000	5406
Religion: Jewish	0.030	0.171	0.000	1.000	5406
Religion: Muslim	0.022	0.146	0.000	1.000	5406
Religion: other	0.090	0.286	0.000	1.000	5406
Religion: prefer not to answer	0.088	0.283	0.000	1.000	5406
Religious/worship service attendance: never	0.290	0.454	0.000	1.000	5406
Religious/worship service attendance: seldom	0.330	0.470	0.000	1.000	5406
Religious/worship service attendance: once or twice a month	0.097	0.296	0.000	1.000	5406
Religious/worship service attendance: once a week	0.215	0.411	0.000	1.000	5406
Religious/worship service attendance: more than once a week	0.068	0.251	0.000	1.000	5406
Urbanization level (1 = rural, 7 = largest city)	4.268	1.919	1.000	7.000	5406
Area: rural area	0.109	0.312	0.000	1.000	5406
Area: small town/village, under 10,000	0.093	0.290	0.000	1.000	5406
Area: town, 10,000-49,999	0.162	0.368	0.000	1.000	5406
Area: city/town, 50,000-99,999	0.165	0.371	0.000	1.000	5406
Area: city, 100,000-499,999	0.174	0.380	0.000	1.000	5406
Area: city, 500,000-1,000,000	0.122	0.327	0.000	1.000	5406
Area: city, over 1,000,000	0.175	0.380	0.000	1.000	5406
Self-rated health status (1 = poor, 5 = excellent)	3.392	0.984	1.000	5.000	5406
Current health status: poor	0.027	0.161	0.000	1.000	5406
Current health status: fair	0.150	0.357	0.000	1.000	5406
Current health status: good	0.362	0.481	0.000	1.000	5406
Current health status: very good	0.326	0.469	0.000	1.000	5406
Current health status: excellent	0.135	0.342	0.000	1.000	5406
Number of self-reported diagnosed health conditions	1.119	0.843	0.000	2.000	5406
Diagnosed health conditions: none	0.303	0.460	0.000	1.000	5406

Continued on next page

Table B1. Descriptive statistics with survey labels (continued)

Variable	**Mean**	**SD**	**Min**	**Max**	**N**
Diagnosed health conditions: one condition	0.276	0.447	0.000	1.000	5406
Diagnosed health conditions: two or more conditions	0.422	0.494	0.000	1.000	5406
Medication/supplement intensity (0-2)	1.032	0.790	0.000	2.000	5406
Medication/supplement use: none regularly	0.297	0.457	0.000	1.000	5406
Medication/supplement use: one type	0.375	0.484	0.000	1.000	5406
Medication/supplement use: two or more types	0.329	0.470	0.000	1.000	5406
Employment status: unemployed	0.094	0.292	0.000	1.000	5406
Region: Northeast	0.187	0.390	0.000	1.000	5556
Region: Midwest	0.208	0.406	0.000	1.000	5556
Region: South	0.422	0.494	0.000	1.000	5556
Region: West	0.182	0.386	0.000	1.000	5556
Region: other/non-state territory	0.000	0.019	0.000	1.000	5556
Technology attitude: first among peers to try new technologies (1)	5.174	3.029	0.000	10.000	5556
Technology attitude: likes experimenting with unfamiliar or complex technologies (2)	5.800	2.815	0.000	10.000	5556
Technology attitude: actively seeks information about new technologies (3)	5.715	2.922	0.000	10.000	5556
Familiarity with human enhancement interventions or technologies	2.095	1.078	1.000	5.000	5556
Familiarity: never heard of human enhancement technologies before	0.346	0.476	0.000	1.000	5556
Familiarity: heard of them, but don't know much	0.364	0.481	0.000	1.000	5556
Familiarity: somewhat familiar	0.181	0.385	0.000	1.000	5556
Familiarity: quite familiar	0.067	0.251	0.000	1.000	5556
Familiarity: very familiar/well-informed	0.042	0.200	0.000	1.000	5556
Expected wellbeing impact of human enhancement technologies	2.197	0.778	1.000	3.000	2772
Expected wellbeing impact: negative consequences	0.223	0.417	0.000	1.000	2772
Expected wellbeing impact: undecided	0.357	0.479	0.000	1.000	2772
Expected wellbeing impact: positive consequences	0.420	0.494	0.000	1.000	2772
Expected access and societal impact of human enhancement technologies	1.925	0.958	1.000	4.000	2772
Expected societal impact: primarily benefit the wealthy	0.447	0.497	0.000	1.000	2772
Expected societal impact: undecided	0.233	0.423	0.000	1.000	2772
Expected societal impact: accessible and beneficial to everyone	0.268	0.443	0.000	1.000	2772
Expected societal impact: primarily benefit the poor	0.052	0.223	0.000	1.000	2772
Expected ethical implications of human enhancement technologies	2.167	0.830	1.000	3.000	2772
Expected ethics: broadly acceptable and unproblematic	0.275	0.447	0.000	1.000	2772
Expected ethics: undecided	0.283	0.450	0.000	1.000	2772
Expected ethics: significant ethical dilemmas	0.442	0.497	0.000	1.000	2772
Preferred legal status: no regulation	0.030	0.171	0.000	1.000	5556
Preferred legal status: legalization with minimal regulation	0.149	0.357	0.000	1.000	5556
Preferred legal status: legalization with moderate regulation	0.324	0.468	0.000	1.000	5556
Preferred legal status: legalization with strict regulation	0.435	0.496	0.000	1.000	5556
Preferred legal status: prohibition	0.061	0.239	0.000	1.000	5556

Continued on next page

Table B1. Descriptive statistics with survey labels (continued)

Variable	**Mean**	**SD**	**Min**	**Max**	**N**
Maximum one-time willingness to pay for the technology (US dollars)	725901.437	26186976.178	0.000	1.000e+09	2930
Maximum willingness to pay as a percentage of household income	595.077	22851.965	0.000	1142857.125	2820
Maximum one-time willingness to pay, top 10% trimmed (US dollars)	1235.540	1494.942	0.000	5000.000	2930
Attitude: technology likely to harm individuals and society overall	6.378	2.310	0.000	10.000	5502
Attitude: technology will make upward mobility harder	6.252	2.464	0.000	10.000	5502
Attitude: unfair if only those who can afford the technology benefit from it	7.372	2.055	0.000	10.000	5502
Attitude: wrong to impose taxes on people who use the technology	6.625	2.408	0.000	10.000	5502
Attitude: public funds should not subsidize the technology	7.115	2.265	0.000	10.000	5502
Attitude: technology will contribute positively to individuals and society overall	7.421	1.304	0.000	10.000	2824
Attitude: technology will help most people improve social/economic status	7.451	1.332	0.000	10.000	2824
Attitude: fair if access depends on ability to pay	7.774	1.396	0.000	10.000	2824
Attitude: taxing use can support those negatively affected by the technology	7.524	1.361	0.000	10.000	2824
Attitude: public subsidies would promote fair and widespread use	7.638	1.381	0.000	10.000	2824
Household income category (ordered numeric)	5.532	2.492	1.000	10.000	5276
Household income: $0-$9,999	0.075	0.264	0.000	1.000	5276
Household income: $10,000-$19,999	0.075	0.264	0.000	1.000	5276
Household income: $20,000-$29,999	0.103	0.304	0.000	1.000	5276
Household income: $30,000-$39,999	0.094	0.292	0.000	1.000	5276
Household income: $40,000-$49,999	0.075	0.264	0.000	1.000	5276
Household income: $50,000-$74,999	0.189	0.392	0.000	1.000	5276
Household income: $75,000-$99,999	0.147	0.355	0.000	1.000	5276
Household income: $100,000-$149,999	0.134	0.341	0.000	1.000	5276
Household income: $150,000-$199,999	0.060	0.238	0.000	1.000	5276
Household income: $200,000+	0.046	0.209	0.000	1.000	5276
Trust in the President of the USA	2.337	1.395	1.000	5.000	5384
Trust in medical doctors	3.534	0.970	1.000	5.000	5384
Trust in scientists	3.432	1.021	1.000	5.000	5384
Trust in Members of Congress	2.104	0.997	1.000	5.000	5384
Trust in other people in general	2.759	0.848	1.000	5.000	5384
Trust in federal regulatory agencies	2.530	0.993	1.000	5.000	5384
Trust in the media	2.247	1.041	1.000	5.000	5384
Trust in science/medical experts index	3.483	0.886	1.000	5.000	5384
Trust in political/regulatory/media institutions index	2.396	0.724	1.000	5.000	5384
Risk preference: willingness to take risks while driving	3.979	3.223	0.000	10.000	5360
Risk preference: willingness to take risks in financial matters	4.542	2.770	0.000	10.000	5360
Risk preference: willingness to take risks during leisure and sport	4.979	2.751	0.000	10.000	5360
Risk preference: willingness to take risks in your occupation	4.600	2.949	0.000	10.000	5360
Risk preference: willingness to take risks with your health	3.975	3.007	0.000	10.000	5360
Risk preference index	4.415	2.425	0.000	10.000	5360

Continued on next page

Table B1. Descriptive statistics with survey labels (continued)

Variable	**Mean**	**SD**	**Min**	**Max**	**N**
Moral relevance: some people were treated differently from others	4.248	1.303	1.000	6.000	5326
Moral relevance: standards of purity and decency were violated	4.089	1.383	1.000	6.000	5326
Moral relevance: someone cared for someone weak or vulnerable	4.368	1.311	1.000	6.000	5326
Moral relevance: traditions of society were followed	3.721	1.354	1.000	6.000	5326
Moral relevance: something disgusting was done, even if no one was harmed	4.129	1.456	1.000	6.000	5326
Moral relevance: someone acted in a way God would approve	3.876	1.651	1.000	6.000	5326
Moral relevance: compassion was shown for the suffering	4.651	1.267	1.000	6.000	5326
Moral relevance: unnatural acts were performed	3.912	1.555	1.000	6.000	5326
Moral care foundation index	4.422	1.065	1.000	6.000	5326
Moral sanctity foundation index	3.945	1.073	1.000	6.000	5326
Failed at least one scenario comprehension check	0.268	0.443	0.000	1.000	5556
Religious attendance: at least once a month indicator	0.283	0.450	0.000	1.000	5406
Good/very good/excellent self-rated health indicator	0.823	0.382	0.000	1.000	5406
Risk-averse indicator	0.627	0.484	0.000	1.000	5360
High moral care foundation indicator	0.609	0.488	0.000	1.000	5326
High moral sanctity foundation indicator	0.447	0.497	0.000	1.000	5326
Health scenario: would adopt for some healthy-life-expectancy benefit	0.478	0.500	0.000	1.000	780
Health scenario: would not adopt regardless of life-extension benefit	0.522	0.500	0.000	1.000	780
Productivity scenario: would adopt for some salary-increase benefit	0.324	0.468	0.000	1.000	876
Productivity scenario: would not adopt regardless of salary increase	0.676	0.468	0.000	1.000	876
Would adopt for some additional benefit	0.397	0.489	0.000	1.000	1656
Would not adopt regardless of additional benefit	0.603	0.489	0.000	1.000	1656
Randomized to expectations questions before scenario	0.499	0.500	0.000	1.000	5556
Post-scenario ethical concerns category	2.278	0.654	1.000	3.000	5556
Change in ethical concerns from pre-scenario expectations to post-scenario rating	0.071	1.016	-2.000	2.000	2772

Notes:

Table B2. Determinants of Never HET Adoption

	(1) Never Adopt	(2) Never Adopt	(3) Never Adopt	(4) Never Adopt
Health	-0.096***	-0.081***	-0.091***	-0.116***
	(0.015)	(0.020)	(0.021)	(0.023)
Genetic Inheritable	-0.017	-0.017	-0.016	0.010
	(0.018)	(0.018)	(0.018)	(0.023)
Genetic Non-Inheritable	-0.018	-0.018	-0.017	-0.010
	(0.018)	(0.018)	(0.018)	(0.023)
High Risk	0.035**	0.019	0.035**	0.035**
	(0.014)	(0.019)	(0.014)	(0.014)
Enhancing	0.082***	0.081***	0.087***	0.080***
	(0.015)	(0.015)	(0.022)	(0.015)
Productivity*High Risk		0.032		
		(0.029)		
Health*Enhancing			-0.010	
			(0.029)	
Productivity*Inheritable				-0.056
				(0.035)
Productivity*Non-Inheritable				-0.012
				(0.035)
Constant	0.266***	0.258***	0.263***	0.276***
	(0.018)	(0.019)	(0.021)	(0.020)
Observations	3,696	3,696	3,696	3,696
Mean	0.270	0.270	0.270	0.270
St. Dev.	0.444	0.444	0.444	0.444

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table B3. Determinants of the Demand for Prohibition

	(1) Prohibition	(2) Prohibition	(3) Prohibition	(4) Prohibition
Health	-0.023***	-0.011	-0.025**	-0.010
	(0.008)	(0.011)	(0.012)	(0.014)
Productivity	-0.030***	-0.014	-0.040***	-0.025*
	(0.008)	(0.011)	(0.011)	(0.013)
Genetic Inheritable	0.001	0.001	0.002	0.014
	(0.008)	(0.008)	(0.008)	(0.015)
Genetic Non-Inheritable	-0.007	-0.007	-0.006	-0.003
	(0.008)	(0.008)	(0.008)	(0.015)
High Risk	0.007	0.026**	0.007	0.007
	(0.006)	(0.012)	(0.006)	(0.006)
Enhancing	0.007	0.007	-0.000	0.007
	(0.006)	(0.006)	(0.012)	(0.007)
Health*High Risk		-0.024		
		(0.016)		
Productivity*High Risk		-0.032**		
		(0.016)		
Health*Enhancing			0.004	
			(0.016)	
Productivity*Enhancing			0.018	
			(0.016)	
Health*Inheritable				-0.033*
				(0.020)
Productivity*Inheritable				-0.003
				(0.020)
Health*Non-Inheritable				-0.006
				(0.020)
Productivity*Non-Inheritable				-0.010
				(0.019)
Constant	0.073***	0.064***	0.076***	0.067***
	(0.009)	(0.010)	(0.010)	(0.011)
Observations	5,556	5,556	5,556	5,556
Mean	0.0608	0.0608	0.0608	0.0608
St. Dev.	0.239	0.239	0.239	0.239

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table B4. Determinants of HET Adoption: Robustness

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Adoption	Adoption	Adoption	Adoption	Adoption	Adoption	Adoption
Health	0.100***	0.096***	0.097***	0.096***	0.096***	0.097***	0.098***
	(0.016)	(0.016)	(0.016)	(0.015)	(0.015)	(0.015)	(0.015)
Productivity	0.047***	0.044***	0.045***	0.038**	0.037**	0.036**	0.037**
	(0.017)	(0.017)	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)
Genetic Inheritable	0.027*	0.027*	0.028*	0.025	0.022	0.024	0.023
	(0.016)	(0.016)	(0.016)	(0.015)	(0.015)	(0.015)	(0.015)
Genetic Non-Inheritable	0.050***	0.051***	0.050***	0.048***	0.048***	0.048***	0.048***
	(0.016)	(0.016)	(0.016)	(0.016)	(0.015)	(0.015)	(0.015)
High Risk	-0.054***	-0.053***	-0.052***	-0.043***	-0.042***	-0.041***	-0.040***
	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)
Enhancing	-0.049***	-0.049***	-0.050***	-0.051***	-0.053***	-0.054***	-0.054***
	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)
Age	-0.004***	-0.004***	-0.004***	-0.001**	-0.001	-0.001	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Male	0.092***	0.099***	0.100***	0.054***	0.044***	0.042***	0.043***
	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)
African American/Black	0.022	0.019	0.011	0.001	0.005	0.004	0.002
	(0.023)	(0.022)	(0.022)	(0.022)	(0.022)	(0.022)	(0.021)
White	-0.049***	-0.054***	-0.059***	-0.030*	-0.037**	-0.036**	-0.038**
	(0.017)	(0.017)	(0.017)	(0.016)	(0.016)	(0.016)	(0.016)
Education	-0.001	-0.006	-0.006	-0.008	-0.008	-0.007	-0.008
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
Employed	0.068***	0.062***	0.057***	0.033*	0.030	0.032	0.033*
	(0.021)	(0.021)	(0.021)	(0.020)	(0.020)	(0.020)	(0.020)
Other	-0.027	-0.033	-0.035	-0.017	-0.017	-0.011	-0.009
	(0.028)	(0.028)	(0.028)	(0.027)	(0.027)	(0.027)	(0.027)
Retired	0.030	0.016	0.013	0.036	0.025	0.025	0.027

	(0.026)	(0.026)	(0.026)	(0.025)	(0.025)	(0.025)	(0.025)
Income	0.000	-0.000	-0.000	-0.000	-0.000*	-0.000*	-0.000*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Urbanization level (1 = rural, 7 = largest city)	0.017***	0.016***	0.016***	0.012***	0.010***	0.010***	0.010***
	(0.004)	(0.004)	(0.004)	(0.003)	(0.003)	(0.003)	(0.003)
Self-rated health status (1 = poor, 5 = excellent)		0.044***	0.041***	0.018**	0.005	0.005	0.005
		(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
Number of self-reported health conditions		0.032***	0.032***	0.030***	0.032***	0.034***	0.034***
		(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
Medication intensity (0–2)		0.031***	0.030***	0.023**	0.019**	0.021**	0.020**
		(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Religiosity: More than once a week			0.020	-0.016	-0.036	-0.033	-0.030
			(0.025)	(0.023)	(0.023)	(0.023)	(0.023)
Christian			0.045***	0.020	0.005	0.003	0.002
			(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
Tech Attitudes (1)				0.017***	0.013***	0.012***	0.012***
				(0.004)	(0.004)	(0.004)	(0.004)
Tech Attitudes (2)				0.018***	0.016***	0.018***	0.018***
				(0.004)	(0.004)	(0.004)	(0.004)
Tech Attitudes (3)				0.025***	0.021***	0.021***	0.021***
				(0.004)	(0.004)	(0.004)	(0.004)
Risk Lover					0.018***	0.018***	0.018***
					(0.003)	(0.003)	(0.003)
Trust: Science					0.022***	0.029***	0.029***
					(0.008)	(0.008)	(0.008)
Trust: Politics					0.070***	0.067***	0.067***
					(0.010)	(0.010)	(0.010)
Moral Care						-0.023***	-0.022***
						(0.008)	(0.008)
Moral Sanctity						0.005	0.004
						(0.008)	(0.008)

Randomized expectations							-0.068***
							(0.013)
Constant	0.466***	0.284***	0.273***	0.004	-0.193***	-0.142**	-0.112**
	(0.043)	(0.051)	(0.051)	(0.050)	(0.052)	(0.056)	(0.056)
Observations	5,276	5,276	5,276	5,276	5,234	5,202	5,202
Mean	0.534	0.534	0.534	0.534	0.535	0.535	0.535
St. Dev.	0.499	0.499	0.499	0.499	0.499	0.499	0.499

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table B5. Determinants of HET Never Adoption: Robustness

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Never Adopt	Never Adopt	Never Adopt	Never Adopt	Never Adopt	Never Adopt	Never Adopt
Health	-0.094***	-0.093***	-0.093***	-0.097***	-0.097***	-0.098***	-0.099***
	(0.015)	(0.015)	(0.015)	(0.014)	(0.014)	(0.014)	(0.014)
Genetic Inheritable	-0.005	-0.005	-0.005	-0.000	0.004	0.004	0.005
	(0.018)	(0.018)	(0.018)	(0.017)	(0.017)	(0.017)	(0.017)
Genetic Non-Inheritable	-0.015	-0.014	-0.013	-0.010	-0.004	-0.004	-0.004
	(0.018)	(0.018)	(0.018)	(0.017)	(0.017)	(0.017)	(0.017)
High Risk	0.036**	0.034**	0.034**	0.028**	0.029**	0.028**	0.028**
	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
Enhancing	0.074***	0.073***	0.074***	0.073***	0.075***	0.075***	0.075***
	(0.015)	(0.015)	(0.015)	(0.014)	(0.014)	(0.014)	(0.014)
Age	0.003***	0.004***	0.004***	0.002***	0.002***	0.002***	0.002***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Male	-0.078***	-0.083***	-0.083***	-0.049***	-0.041***	-0.040***	-0.042***
	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)
African American/Black	-0.024	-0.024	-0.020	-0.010	-0.017	-0.016	-0.013
	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)
White	0.029	0.033*	0.036**	0.014	0.016	0.015	0.018
	(0.018)	(0.018)	(0.018)	(0.017)	(0.017)	(0.017)	(0.017)
Education	0.003	0.004	0.003	0.005	0.008	0.009	0.010
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Employed	-0.080***	-0.081***	-0.078***	-0.058***	-0.055**	-0.057**	-0.059***
	(0.023)	(0.023)	(0.023)	(0.023)	(0.022)	(0.022)	(0.022)
Other	0.000	0.006	0.008	0.002	-0.002	-0.004	-0.006
	(0.031)	(0.031)	(0.031)	(0.031)	(0.030)	(0.030)	(0.030)
Retired	-0.042	-0.036	-0.034	-0.050*	-0.041	-0.042	-0.044
	(0.029)	(0.029)	(0.029)	(0.029)	(0.029)	(0.029)	(0.029)
Income	-0.000**	-0.000**	-0.000**	-0.000	-0.000	-0.000	-0.000

	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Urbanization level (1 = rural, 7 = largest city)	-0.010**	-0.010**	-0.010**	-0.008**	-0.005	-0.005	-0.005
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Self-rated health status (1 = poor, 5 = excellent)		-0.010	-0.010	0.004	0.011	0.011	0.010
		(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Number of self-reported health conditions		-0.027***	-0.028***	-0.027***	-0.027***	-0.026***	-0.026***
		(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Medication intensity (0–2)		-0.019*	-0.019*	-0.015	-0.013	-0.013	-0.013
		(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)
Religiosity: More than once a week			0.013	0.039	0.042	0.042	0.041
			(0.029)	(0.028)	(0.028)	(0.028)	(0.028)
Christian			-0.028*	-0.018	-0.009	-0.011	-0.011
			(0.015)	(0.015)	(0.015)	(0.016)	(0.016)
Tech Attitudes (1)				-0.009**	-0.007	-0.007	-0.007
				(0.004)	(0.004)	(0.004)	(0.004)
Tech Attitudes (2)				-0.013***	-0.011**	-0.011**	-0.011**
				(0.005)	(0.005)	(0.005)	(0.005)
Tech Attitudes (3)				-0.016***	-0.013***	-0.014***	-0.014***
				(0.005)	(0.005)	(0.005)	(0.005)
Risk lover					-0.020***	-0.020***	-0.020***
					(0.003)	(0.003)	(0.003)
Trust: Science					-0.037***	-0.038***	-0.039***
					(0.009)	(0.009)	(0.009)
Trust: Politics					-0.025**	-0.024**	-0.025**
					(0.011)	(0.011)	(0.011)
Moral Care						0.004	0.003
						(0.009)	(0.009)
Moral Sanctity						0.004	0.005
						(0.009)	(0.009)
Randomized expectations							0.059***
							(0.014)

Constant	0.316***	0.384***	0.392***	0.568***	0.743***	0.722***	0.696***
	(0.046)	(0.055)	(0.055)	(0.056)	(0.059)	(0.064)	(0.064)
Observations	3,508	3,508	3,508	3,508	3,489	3,474	3,474
Mean	0.267	0.267	0.267	0.267	0.267	0.267	0.267
St. Dev.	0.442	0.442	0.442	0.442	0.443	0.442	0.442

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table B6. Determinants of HET Desired Regulation: Robustness

	(1) Regulation	(2) Regulation	(3) Regulation	(4) Regulation	(5) Regulation	(6) Regulation	(7) Regulation
Health	0.139*	0.136*	0.138*	0.141*	0.149*	0.160**	0.162**
	(0.078)	(0.078)	(0.078)	(0.079)	(0.080)	(0.080)	(0.080)
Productivity	-1.103***	-1.105***	-1.107***	-1.096***	-1.095***	-1.076***	-1.075***
	(0.086)	(0.086)	(0.086)	(0.085)	(0.085)	(0.085)	(0.085)
Genetic Inheritable	0.153*	0.154*	0.150*	0.155*	0.158*	0.146*	0.144*
	(0.082)	(0.082)	(0.082)	(0.081)	(0.082)	(0.081)	(0.081)
Genetic Non-Inheritable	-0.094	-0.094	-0.092	-0.091	-0.090	-0.073	-0.074
	(0.081)	(0.081)	(0.081)	(0.081)	(0.081)	(0.081)	(0.081)
High Risk	0.193***	0.193***	0.191***	0.176***	0.174***	0.169**	0.170**
	(0.067)	(0.067)	(0.067)	(0.066)	(0.067)	(0.066)	(0.066)
Enhancing	0.155**	0.155**	0.158**	0.161**	0.170**	0.186***	0.187***
	(0.067)	(0.067)	(0.067)	(0.067)	(0.067)	(0.067)	(0.067)
Age	0.004	0.003	0.004	-0.000	-0.002	-0.004	-0.004
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Male	-0.284***	-0.272***	-0.272***	-0.201***	-0.187***	-0.138**	-0.135*
	(0.068)	(0.068)	(0.068)	(0.068)	(0.069)	(0.069)	(0.069)
African American/Black	-0.197*	-0.193*	-0.174	-0.157	-0.192	-0.197*	-0.202*
	(0.117)	(0.117)	(0.117)	(0.117)	(0.117)	(0.117)	(0.117)
White	0.192**	0.188**	0.206**	0.159*	0.171*	0.173*	0.171*
	(0.088)	(0.088)	(0.089)	(0.089)	(0.089)	(0.089)	(0.089)
Education	0.162***	0.156***	0.152***	0.155***	0.163***	0.148***	0.147***
	(0.047)	(0.047)	(0.047)	(0.047)	(0.047)	(0.047)	(0.047)
Employed	-0.049	-0.052	-0.039	0.002	0.002	-0.003	0.000
	(0.103)	(0.103)	(0.104)	(0.103)	(0.104)	(0.103)	(0.103)
Other	-0.077	-0.092	-0.086	-0.119	-0.128	-0.168	-0.165
	(0.138)	(0.138)	(0.138)	(0.139)	(0.140)	(0.139)	(0.139)
Retired	-0.017	-0.036	-0.030	-0.063	-0.028	-0.050	-0.045

	(0.128)	(0.128)	(0.128)	(0.129)	(0.130)	(0.129)	(0.129)
Income	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Urbanization level (1 = rural, 7 = largest city)	0.002	0.001	0.000	0.006	0.010	0.012	0.013
	(0.018)	(0.018)	(0.018)	(0.018)	(0.018)	(0.018)	(0.018)
Self-rated health status (1 = poor, 5 = excellent)		0.019	0.020	0.059	0.080**	0.067*	0.067*
		(0.039)	(0.039)	(0.039)	(0.040)	(0.039)	(0.039)
Number of self-reported health conditions		0.008	0.004	0.006	0.009	-0.009	-0.009
		(0.047)	(0.047)	(0.047)	(0.047)	(0.047)	(0.047)
Medication intensity (0–2)		0.092*	0.094*	0.106**	0.114**	0.099*	0.097*
		(0.052)	(0.051)	(0.051)	(0.052)	(0.052)	(0.052)
Religiosity: More than once a week			0.191	0.249*	0.280*	0.214	0.220
			(0.145)	(0.144)	(0.146)	(0.147)	(0.147)
Christian			-0.162**	-0.121*	-0.105	-0.148**	-0.149**
			(0.071)	(0.071)	(0.072)	(0.074)	(0.074)
Tech Attitudes (1)				-0.046**	-0.042**	-0.030	-0.029
				(0.021)	(0.021)	(0.021)	(0.021)
Tech Attitudes (2)				-0.013	-0.008	-0.024	-0.024
				(0.023)	(0.023)	(0.023)	(0.023)
Tech Attitudes (3)				-0.033	-0.023	-0.027	-0.028
				(0.021)	(0.021)	(0.021)	(0.021)
Risk lover					-0.039**	-0.047***	-0.047***
					(0.018)	(0.018)	(0.018)
Trust: Science					-0.083*	-0.154***	-0.154***
					(0.046)	(0.047)	(0.047)
Trust: Politics					-0.106*	-0.097	-0.096
					(0.061)	(0.062)	(0.062)
Moral Care						0.241***	0.243***
						(0.043)	(0.043)
Moral Sanctity						0.076*	0.074*
						(0.043)	(0.043)

Randomized expectations							-0.166**
							(0.066)
Constant	5.917***	5.814***	5.866***	6.264***	6.746***	5.843***	5.917***
	(0.211)	(0.253)	(0.254)	(0.268)	(0.294)	(0.314)	(0.315)
Observations	5,276	5,276	5,276	5,276	5,234	5,202	5,202
Mean	5.983	5.983	5.983	5.983	5.982	5.980	5.980
St. Dev.	2.503	2.503	2.503	2.503	2.502	2.903	3.065

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table B7. Determinants of HET Prohibition: Robustness

	(1) Prohibition	(2) Prohibition	(3) Prohibition	(4) Prohibition	(5) Prohibition	(6) Prohibition	(7) Prohibition
Health	-0.022***	-0.023***	-0.022***	-0.022***	-0.023***	-0.023***	-0.023***
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Productivity	-0.033***	-0.034***	-0.034***	-0.034***	-0.034***	-0.033***	-0.033***
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Genetic Inheritable	0.003	0.004	0.003	0.003	0.003	0.003	0.003
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Genetic Non-Inheritable	-0.005	-0.004	-0.005	-0.004	-0.004	-0.003	-0.003
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
High Risk	0.005	0.005	0.005	0.005	0.005	0.005	0.004
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Enhancing	0.007	0.007	0.007	0.007	0.007	0.008	0.008
	(0.007)	(0.006)	(0.006)	(0.006)	(0.007)	(0.006)	(0.006)
Age	0.000	0.001*	0.001**	0.001**	0.001*	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Male	-0.013**	-0.013*	-0.013*	-0.012*	-0.012*	-0.008	-0.008
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
African American/Black	0.008	0.006	0.006	0.007	0.005	0.005	0.006
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)
White	0.007	0.008	0.009	0.008	0.009	0.010	0.010
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Education	-0.016***	-0.018***	-0.018***	-0.018***	-0.016***	-0.017***	-0.017***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Employed	-0.004	-0.008	-0.008	-0.008	-0.006	-0.008	-0.008
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Other	-0.002	-0.001	-0.001	-0.001	0.003	-0.003	-0.003
	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
Retired	-0.004	-0.007	-0.007	-0.008	-0.002	-0.005	-0.006

	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)
Income	-0.000	-0.000	-0.000	-0.000	-0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Urbanization level (1 = rural, 7 = largest city)	0.004**	0.004**	0.003*	0.003*	0.004**	0.004**	0.004**
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Self-rated health status (1 = poor, 5 = excellent)		0.019***	0.018***	0.018***	0.018***	0.017***	0.017***
		(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Number of self-reported health conditions		-0.000	-0.000	-0.000	0.000	0.000	0.000
		(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Medication intensity (0–2)		0.001	0.001	0.001	0.003	0.001	0.001
		(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Religiosity: More than once a week			0.028*	0.029*	0.025	0.019	0.019
			(0.015)	(0.015)	(0.015)	(0.016)	(0.016)
Christian			-0.008	-0.008	-0.011	-0.017**	-0.017**
			(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
Tech Attitudes (1)				0.001	0.000	0.001	0.001
				(0.002)	(0.002)	(0.002)	(0.002)
Tech Attitudes (2)				-0.001	-0.000	-0.001	-0.001
				(0.002)	(0.002)	(0.002)	(0.002)
Tech Attitudes (3)				-0.001	-0.001	-0.001	-0.001
				(0.002)	(0.002)	(0.002)	(0.002)
Risk lover					0.002	0.001	0.001
					(0.002)	(0.002)	(0.002)
Trust: Science					-0.017***	-0.019***	-0.019***
					(0.005)	(0.005)	(0.005)
Trust: Politics					0.008	0.007	0.007
					(0.007)	(0.007)	(0.007)
Moral Care						0.008*	0.008*
						(0.005)	(0.005)
Moral Sanctity						0.013***	0.013***
						(0.004)	(0.004)

Randomized expectations							0.006
							(0.006)
Constant	0.089***	0.028	0.032	0.035	0.061**	0.001	-0.001
	(0.020)	(0.023)	(0.023)	(0.025)	(0.029)	(0.031)	(0.031)
Observations	5,276	5,276	5,276	5,276	5,234	5,202	5,202
Mean	0.0595	0.0595	0.0595	0.0595	0.0594	0.0588	0.0588
St. Dev.	0.237	0.237	0.237	0.237	0.236	0.235	0.235

Robust standard errors in parentheses

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

Table B8. Determinants of HET Ethics: Robustness

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Ethics	Ethics	Ethics	Ethics	Ethics	Ethics	Ethics
Health	0.071	0.070	0.080	0.073	0.077	0.070	0.075
	(0.093)	(0.092)	(0.092)	(0.092)	(0.091)	(0.091)	(0.091)
Productivity	0.534***	0.525***	0.534***	0.526***	0.536***	0.534***	0.537***
	(0.093)	(0.093)	(0.092)	(0.092)	(0.092)	(0.092)	(0.092)
Genetic Inheritable	0.319***	0.333***	0.328***	0.323***	0.333***	0.334***	0.329***
	(0.091)	(0.090)	(0.090)	(0.090)	(0.090)	(0.090)	(0.090)
Genetic Non-Inheritable	-0.036	-0.023	-0.028	-0.025	-0.021	-0.024	-0.025
	(0.093)	(0.092)	(0.092)	(0.091)	(0.090)	(0.090)	(0.090)
High Risk	0.224***	0.221***	0.222***	0.238***	0.235***	0.233***	0.235***
	(0.075)	(0.075)	(0.075)	(0.074)	(0.074)	(0.074)	(0.074)
Enhancing	0.256***	0.256***	0.251***	0.244***	0.239***	0.250***	0.251***
	(0.076)	(0.076)	(0.076)	(0.076)	(0.075)	(0.075)	(0.075)
Age	-0.023***	-0.021***	-0.021***	-0.017***	-0.015***	-0.017***	-0.017***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Male	-0.085	-0.069	-0.057	-0.122	-0.165**	-0.135*	-0.128*
	(0.075)	(0.075)	(0.075)	(0.075)	(0.075)	(0.076)	(0.076)
African American/Black	0.321**	0.305**	0.270**	0.250**	0.231*	0.228*	0.215*
	(0.127)	(0.126)	(0.127)	(0.125)	(0.124)	(0.125)	(0.125)
White	0.019	0.031	0.023	0.073	0.062	0.083	0.076
	(0.096)	(0.096)	(0.096)	(0.095)	(0.095)	(0.095)	(0.095)
Education	0.058	0.018	0.018	0.014	0.062	0.081	0.077
	(0.052)	(0.052)	(0.052)	(0.052)	(0.052)	(0.052)	(0.052)
Employed	0.640***	0.565***	0.548***	0.497***	0.490***	0.475***	0.482***
	(0.114)	(0.113)	(0.113)	(0.113)	(0.113)	(0.112)	(0.112)
Other	0.174	0.173	0.156	0.204	0.244*	0.227	0.235
	(0.149)	(0.148)	(0.148)	(0.147)	(0.147)	(0.148)	(0.148)
Retired	0.380***	0.310**	0.294**	0.307**	0.389***	0.362**	0.373**

	(0.146)	(0.146)	(0.146)	(0.145)	(0.145)	(0.145)	(0.145)
Income	-0.000*	-0.000***	-0.000***	-0.000***	-0.000***	-0.000***	-0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Urbanization level (1 = rural, 7 = largest city)	0.035*	0.024	0.021	0.015	0.018	0.019	0.021
	(0.021)	(0.021)	(0.021)	(0.021)	(0.021)	(0.021)	(0.021)
Self-rated health status (1 = poor, 5 = excellent)		0.280***	0.253***	0.203***	0.170***	0.154***	0.156***
		(0.043)	(0.043)	(0.044)	(0.044)	(0.044)	(0.044)
Number of self-reported health conditions		-0.067	-0.074	-0.071	-0.042	-0.037	-0.036
		(0.053)	(0.053)	(0.053)	(0.053)	(0.053)	(0.053)
Medication intensity (0–2)		0.152***	0.145**	0.125**	0.141**	0.139**	0.136**
		(0.058)	(0.058)	(0.058)	(0.058)	(0.058)	(0.058)
Religiosity: More than once a week			0.771***	0.709***	0.569***	0.488***	0.501***
			(0.160)	(0.159)	(0.158)	(0.160)	(0.159)
Christian			0.083	0.035	-0.071	-0.173**	-0.174**
			(0.080)	(0.080)	(0.080)	(0.081)	(0.081)
Tech Attitudes (1)				0.099***	0.077***	0.073***	0.075***
				(0.024)	(0.024)	(0.024)	(0.024)
Tech Attitudes (2)				-0.071***	-0.074***	-0.075***	-0.075***
				(0.025)	(0.025)	(0.026)	(0.026)
Tech Attitudes (3)				0.057**	0.051**	0.049**	0.046*
				(0.024)	(0.024)	(0.024)	(0.024)
Risk lover					0.095***	0.089***	0.088***
					(0.019)	(0.019)	(0.019)
Trust: Science					-0.349***	-0.323***	-0.322***
					(0.050)	(0.052)	(0.052)
Trust: Politics					0.318***	0.264***	0.266***
					(0.064)	(0.065)	(0.065)
Moral Care						-0.104**	-0.100**
						(0.049)	(0.049)
Moral Sanctity						0.274***	0.270***
						(0.050)	(0.050)

Randomized expectations							-0.378***
							(0.074)
Constant	5.402***	4.493***	4.500***	4.217***	4.418***	4.040***	4.206***
	(0.233)	(0.276)	(0.275)	(0.287)	(0.311)	(0.331)	(0.332)
Observations	5,276	5,276	5,276	5,276	5,234	5,202	5,202
Mean	5.351	5.351	5.351	5.351	5.355	5.353	5.353
St. Dev.	2.791	2.791	2.791	2.791	2.790	2.792	2.792

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table B9. Determinants of HET Adoption: Stricter attention

	(1) Adoption	(2) Adoption	(3) Adoption	(4) Adoption
Health	0.133***	0.130***	0.143***	0.122***
	(0.018)	(0.026)	(0.026)	(0.032)
Productivity	0.078***	0.065**	0.106***	0.068**
	(0.020)	(0.028)	(0.031)	(0.031)
Genetic Inheritable	0.041**	0.041**	0.039**	0.029
	(0.019)	(0.019)	(0.019)	(0.032)
Genetic Non-Inheritable	0.039**	0.039**	0.037*	0.031
	(0.019)	(0.019)	(0.019)	(0.033)
High Risk	-0.058***	-0.068***	-0.058***	-0.058***
	(0.016)	(0.026)	(0.016)	(0.016)
Enhancing	-0.072***	-0.072***	-0.051*	-0.072***
	(0.016)	(0.016)	(0.026)	(0.016)
Health*High Risk		0.006		
		(0.037)		
Productivity*High Risk		0.026		
		(0.039)		
Health*Enhancing			-0.022	
			(0.037)	
Productivity*Enhancing			-0.048	
			(0.040)	
Health*Inheritable				0.004
				(0.045)
Productivity*Inheritable				0.043
				(0.048)
Health*Non-Inheritable				0.028
				(0.045)
Productivity*Non-Inheritable				-0.012
				(0.048)
Constant	0.481***	0.485***	0.471***	0.487***
	(0.020)	(0.023)	(0.023)	(0.025)
Observations	4,065	4,065	4,065	4,065
Mean	0.507	0.507	0.507	0.507
St. Dev.	0.500	0.500	0.500	0.500

Robust standard errors in parentheses
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table B10. Determinants of HET Regulation: Stricter attention

	(1)	(2)	(3)	(4)	(5)	(6)
	Regulation	Regulation	Regulation	Regulation	Tax use	Subsidize
Health	0.195**	0.324***	0.081	0.045	-0.788***	1.046***
	(0.087)	(0.125)	(0.122)	(0.150)	(0.108)	(0.114)
Productivity	-0.799***	-0.644***	-1.114***	-0.912***	-0.483***	0.564***
	(0.102)	(0.145)	(0.158)	(0.158)	(0.117)	(0.124)
Genetic Inheritable	0.146	0.142	0.170*	0.145	0.059	0.143
	(0.093)	(0.093)	(0.093)	(0.156)	(0.111)	(0.117)
Genetic Non-Inheritable	-0.045	-0.045	-0.022	-0.333**	0.139	0.216*
	(0.092)	(0.092)	(0.093)	(0.162)	(0.112)	(0.118)
High Risk	0.195**	0.380***	0.196**	0.193**	0.024	-0.205**
	(0.076)	(0.129)	(0.076)	(0.076)	(0.091)	(0.096)
Enhancing	0.219***	0.217***	-0.014	0.220***	0.038	-0.342***
	(0.077)	(0.077)	(0.129)	(0.077)	(0.092)	(0.097)
Health*High Risk		-0.261				
		(0.173)				
Productivity*High Risk		-0.315				
		(0.201)				
Health*Enhancing			0.231			
			(0.173)			
Productivity*Enhancing			0.544***			
			(0.207)			
Health*Inheritable				0.004		
				(0.211)		
Productivity*Inheritable				-0.045		
				(0.250)		
Health*Non-Inheritable				0.465**		
				(0.214)		
Productivity*Non-Inheritable				0.418*		
				(0.244)		
Constant	5.983***	5.893***	6.083***	6.075***	5.200***	4.110***
	(0.101)	(0.114)	(0.112)	(0.126)	(0.122)	(0.129)
Observations	4,065	4,065	4,065	4,065	4,027	4,027
Mean	6.064	6.064	6.064	6.064	3	2
St. Dev.	2.463	2.463	2.463	2.463	3	2

Robust standard errors in parentheses
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table B11. Determinants of HET Ethical concerns: Stricter attention

	(1) Ethics	(2) Ethics	(3) Ethics	(4) Ethics	(5) Harm	(6) Mobility	(7) Fairness
Health	0.085	0.309**	-0.069	-0.126	-0.089	-0.136	0.355***
	(0.105)	(0.147)	(0.149)	(0.180)	(0.088)	(0.096)	(0.078)
Productivity	0.659***	0.805***	0.316*	0.588***	0.206**	0.261***	0.429***
	(0.111)	(0.156)	(0.174)	(0.173)	(0.092)	(0.097)	(0.083)
Genetic Inheritable	0.392***	0.387***	0.417***	0.350*	0.171*	0.080	0.109
	(0.105)	(0.105)	(0.106)	(0.181)	(0.088)	(0.095)	(0.079)
Genetic Non-Inheritable	-0.057	-0.057	-0.033	-0.313*	-0.058	-0.073	0.037
	(0.109)	(0.109)	(0.109)	(0.186)	(0.091)	(0.097)	(0.079)
High Risk	0.269***	0.518***	0.270***	0.268***	0.058	0.022	-0.079
	(0.088)	(0.150)	(0.087)	(0.088)	(0.073)	(0.078)	(0.065)
Enhancing	0.227**	0.225**	-0.049	0.220**	0.147*	0.213***	0.113*
	(0.089)	(0.089)	(0.150)	(0.090)	(0.075)	(0.081)	(0.066)
Health*High Risk		-0.453**					
		(0.210)					
Productivity*High Risk		-0.296					
		(0.218)					
Health*Enhancing			0.309				
			(0.210)				
Productivity*Enhancing			0.598***				
			(0.226)				
Health*Inheritable				0.058			
				(0.252)			
Productivity*Inheritable				0.084			
				(0.265)			
Health*Non-Inheritable				0.594**			
				(0.260)			
Productivity*Non-Inheritable				0.121			
				(0.271)			
Constant	4.590***	4.470***	4.711***	4.690***	6.096***	6.000***	7.080***
	(0.116)	(0.129)	(0.130)	(0.143)	(0.096)	(0.104)	(0.088)
Observations	4,065	4,065	4,065	4,065	4,027	4,027	4,027
Mean	5.174	5.174	5.174	5.174	4.634	4.769	5.669
St. Dev.	2.815	2.815	2.815	2.815	2.665	2.754	3.066

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

C Textual analysis: Coding methodology

C.1 Overview and unit of analysis

To analyze the open-ended responses in which respondents described the main factors influencing their answers about adoption, regulation, and ethical concerns regarding the human enhancement technology presented to them, we designed a fixed taxonomy of eight themes (T0–T7), together with explicit decision rules for assigning each theme, a normalization step for short or generic responses, tie-breaker rules for very short responses, a validation check before assigning the residual category, and a set of binding example classifications (the full coding prompt is reproduced in Section C.7). The classification of each response according to this taxonomy was then carried out by a large language model (Claude, Anthropic, model `claude-sonnet-4-5`): the model received the fixed prompt together with each response and returned, for each response, a binary (0/1) indicator for each of the eight themes, with no further input or post-hoc adjustment from us.

The unit of analysis is a single open-ended response. Responses were processed in batches of 40. A response could receive a value of 1 on multiple themes if it raised multiple distinct considerations.

C.2 Coding taxonomy

Table C1 summarizes the eight themes used to code the open-ended responses. T0 captures uninformative or non-substantive responses (e.g., "I don't know," blank responses, or vacuous statements with no identifiable reasoning). T1–T6 capture substantive considerations that map onto the main dimensions of adoption and regulation preferences emphasized elsewhere in the paper: perceived benefits, safety concerns, demand for evidence, distributional/access concerns, moral and religious boundaries, and concerns about dependence, misuse, or privacy. T7 is a strict residual category for informative responses that do not fit any of T1–T6.

Table C1. Coding taxonomy for open-ended responses

Code	Theme	Definition
T0	Uninformative / non-substantive	Blank, "I don't know," "no opinion," or vacuous responses with no identifiable reason (e.g., "personal choice" with no elaboration).
T1	Benefit / usefulness	Practical value, effectiveness, improvements to health, appearance, or productivity, convenience, or quality of life ("if it works, why not").
T2	Safety / side effects	Physical risk, side effects, medical safety, reversibility, or long-term health consequences, including reference to the risk probability presented in the vignette.
T3	Need for more evidence / uncertainty	Desire for more research, testing, long-term data, proof of efficacy, or regulatory approval prior to forming a judgment; general statements of uncertainty ("too new," "not enough studies").
T4	Inequality / access	Cost, affordability, fairness of access, competitive advantage (e.g., in school or work), social mobility, or concerns that the technology would widen gaps between rich and poor.
T5	Moral boundary / human nature / religion	"Playing God," religious or spiritual objections, naturalness, human identity, "going against nature," dignity, or the distinction between treatment and enhancement.
T6	Dependence, unintended consequences, misuse, or privacy	Becoming dependent on the technology, loss of control, unintended or unforeseen consequences, misuse or abuse by individuals or institutions, surveillance, or data privacy.
T7	Other informative	Informative content that does not fit T1–T6. Strict residual category, assigned only after explicitly checking and rejecting T1–T6.

C.3 Decision rules

The coding prompt instructs the model to follow four steps for each response:

1. **Normalize short or generic responses.** Before assigning themes, short or generic phrases are mapped to the corresponding theme based on surface wording alone (e.g., "useful," "helps people" → T1; "risky," "side effects" → T2; "expensive," "unfair access" → T4; "unnatural," "against religion" → T5; "addiction," "privacy" → T6).

2. **Determine T0.** A response is coded T0 if, after normalization, it contains no identifiable reason, concern, or consideration.

3. **Determine T1–T6.** The model checks each substantive theme definition against the response and assigns a value of 1 to every theme for which the response provides a relevant cue. Multiple themes may be assigned.

4. **Assign T7 only as a residual.** T7 is assigned only after the model has explicitly checked and rejected T1–T6 for a response that is nonetheless informative.

C.3.1 Handling short and ambiguous responses

For responses of one to four words, the model applies the following priority order: (1) nonresponse/ignorance cues → T0; (2) vacuous or non-specific phrasing → T0; (3) benefit/usefulness cue → T1; (4) safety/risk cue → T2; (5) evidence/uncertainty cue → T3; (6) inequality/access cue → T4; (7) moral/religious/naturalness cue → T5; (8) dependence/misuse/privacy cue → T6; (9) T7 only if the response is informative but clearly outside T1–T6. When a response could plausibly be coded as either T7 or one of T1–T6, the named theme (T1–T6) is preferred.

C.3.2 Validation check

Before finalizing a T7 assignment, the model is required to re-check the response against each of T1–T6 and to assign the corresponding theme(s) instead of (or in addition to) T7 if any relevant cue is present. This step is intended to keep T7 a genuinely residual category.

C.4 Limitations

As with any automated coding procedure, this approach is subject to several limitations. First, theme assignment retains an element of subjectivity for responses that touch on multiple overlapping considerations (e.g., a response citing both cost and fairness may be coded under T4 alone, even though it could also be read as touching on T1 if framed around access to benefits). Second, very short or terse responses are coded primarily via lexical pattern-matching rather than inferred intent, which may occasionally misclassify idiosyncratic phrasing. Third, the lexical normalization step, while improving consistency, could in principle bias short responses toward the themes whose example phrases most closely match the respondent's wording. Finally, we did not conduct a separate human-coded validation sample; the reported results reflect the model's classifications under the fixed prompt described in Section C.7.

C.5 Prevalence of themes

Table C2 reports the share of open-ended responses (N = 5,556) coded under each theme. Because themes are not mutually exclusive, the column does not sum to 100%.

Table C2. Prevalence of coded themes in open-ended responses

Code	Theme	% of responses
T0	Uninformative / non-substantive	25.0
T1	Benefit / usefulness	27.0
T2	Safety / side effects	12.7
T3	Need for more evidence / uncertainty	9.3
T4	Inequality / access	16.4
T5	Moral boundary / human nature / religion	13.2
T6	Dependence, unintended consequences, misuse, or privacy	10.7
T7	Other informative	7.9
N		5,556

C.6 Coding codebook

Table C3 summarizes, for each output variable, the allowed values and the key decision rule used by the model.

C.7 Full coding prompt

The following prompt (system message) was submitted to the model together with a JSON array of {`id`, `comment`} pairs for each batch of responses.

Role: You are an expert in quantitative coding of open-ended text responses.

Task: For each response to the survey item: "What are the main factors that influenced your answers about adoption, regulation, and ethical concerns regarding this technology?"—create reproducible theme indicators using a fixed taxonomy.

Theme coding (T0–T7)

Hard constraints

Table C3. Coding codebook

Variable	Allowed values	Key decision rule
`responseid`	string	Unique respondent identifier.
`comments`	string	Verbatim open-ended response.
T0	0 / 1	1 if response is uninformative, non-substantive, or vacuous.
T1	0 / 1	1 if response references benefit, usefulness, or quality-of-life improvements.
T2	0 / 1	1 if response references safety, side effects, or health risk.
T3	0 / 1	1 if response references need for more evidence, testing, or uncertainty.
T4	0 / 1	1 if response references cost, access, fairness, or inequality.
T5	0 / 1	1 if response references moral, religious, or naturalness objections.
T6	0 / 1	1 if response references dependence, unintended consequences, misuse, or privacy.
T7	0 / 1	1 if response is informative but does not fit T1–T6 (strict residual).

- Use ONLY these 8 themes (T0–T7). Do NOT create new themes or sub-themes.
- For each response, create T0. . . T7 indicators (0/1).
- Themes are non-mutually-exclusive: a response may receive multiple 1s if it raises multiple distinct considerations.

Theme definitions

T0 – Uninformative / non-substantive / nonsense. Blank or near-blank responses; "I don't know," "no opinion," "N/A"; irrelevant or unintelligible text; vacuous responses with no identifiable reason (e.g., "just my opinion," "personal choice," with no elaboration). A response is informative only if it contains at least one identifiable reason, concern, or consideration.

T1 – Benefit / usefulness. References to the technology's practical value, effectiveness, improvement to health/appearance/productivity, convenience, quality of life, or "if it works, why not." Includes short generic positive-utility terms (e.g., "it helps people," "useful," "good for health").

T2 – Safety / side effects. References to physical risk, side effects, medical safety, reversibility, long-term health consequences, or the risk probability presented in the vignette. Includes short terms like "risky," "dangerous," "side effects."

T3 – Need for more evidence / uncertainty. References to wanting more research, testing, long-term data, proof of efficacy, regulatory approval/trials before forming a judgment, or general statements of uncertainty about consequences ("too new," "not enough studies," "wait and see").

T4 – Inequality / access. References to cost, affordability, who can pay, fairness of access, competition (e.g., in school/work), social mobility, or the technology widening gaps between rich and poor.

T5 – Moral boundary / human nature / religion. References to "playing God," religious or spiritual objections, naturalness/unnaturalness, human identity, what it means to be human, "going against nature," dignity, or enhancement vs. treatment as a moral line.

T6 – Dependence, unintended consequences, misuse, or privacy. References to becoming dependent on technology, loss of control, unintended/unforeseen consequences, misuse or abuse by individuals/institutions, surveillance, or data privacy.

T7 – Other informative. Informative content not fitting T1–T6. Strict residual category—assign only after explicitly checking and rejecting T1–T6.

Theme assignment rules (very important)

- A response may have multiple themes (multiple indicators can be 1).
- T7 is a strict residual category. Use it only AFTER explicitly checking and rejecting T1–T6.
- Do NOT assign T7 to short/generic phrases that plausibly fit T1–T6.
- When in doubt between T7 and a named theme (T1–T6), prefer the named theme.

Required pre-step: normalize short/generic responses

Before assigning themes, map short or generic wording to concepts using surface wording only (no deep inference about intent):

- "benefits," "useful," "helps people," "improves life" → T1
- "risky," "side effects," "dangerous," "health risks" → T2
- "needs more research," "too new," "not tested enough" → T3
- "expensive," "only rich people," "unfair access," "inequality" → T4
- "unnatural," "playing god," "against religion," "not human anymore" → T5
- "addiction," "lose control," "abuse," "privacy," "surveillance," "unintended effects" → T6

For very short responses (1–4 words), prioritize direct lexical mapping to T1–T6 before using T7. If no identifiable concept remains after normalization, code T0.

Short comment tie-breakers (1–4 words; required)

1. Ignorance/nonresponse ("I don't know," "none," "N/A") → T0
2. Vacuous/non-specific ("personal opinion," "just because") → T0
3. Benefit/usefulness cue ("helps," "good," "useful") → T1
4. Safety/risk cue ("risky," "side effects," "dangerous") → T2
5. Evidence/uncertainty cue ("too new," "untested") → T3
6. Inequality/access cue ("expensive," "unfair," "only for rich") → T4
7. Moral/religious/naturalness cue ("unnatural," "against god/nature") → T5
8. Dependence/misuse/privacy cue ("addiction," "abuse," "privacy") → T6
9. T7 only if informative but clearly outside T1–T6

Validation check before finalizing T7 (required)

If T7 = 1, re-check the response for any cue of: usefulness/benefit (T1); safety/side effects (T2); evidence/testing/uncertainty (T3); cost/access/inequality (T4); naturalness/morality/religion (T5); dependence/misuse/privacy (T6). If any cue is present, assign the corresponding theme(s) and remove/reduce T7.

Mandatory coding examples (binding precedents)

- "It could really help people with disabilities" → T1 = 1, not T7
- "1% risk of serious side effects is too high" → T2 = 1, not T7
- "We don't know the long-term effects yet" → T3 = 1, not T7
- "Only rich people will be able to afford this" → T4 = 1, not T7
- "This goes against human nature" → T5 = 1, not T7
- "People could become dependent on it" → T6 = 1, not T7
- "I don't know" → T0 = 1
- "Just my opinion" → T0 = 1

Output (one row per response): `responseid`, `comments`, T0…T7.