



PHOENIX FOR ADVANCED
LEGAL & ECONOMIC
C E N T E R PUBLIC POLICY STUDIES
www.phoenix-center.org

PHOENIX CENTER POLICY BULLETIN NO. 83

George S. Ford, PhD

Riley A. Ford

May 2026

WHO USES ARTIFICIAL INTELLIGENCE AND WHO RESISTS? EVIDENCE FROM THE HOUSEHOLD PULSE SURVEY

Abstract: This paper uses new data from the U.S. Census Bureau’s Household Pulse Survey (Wave 2506) to examine demographic determinants of artificial intelligence (AI) adoption and uses, fear of AI’s career consequences, and active resistance to AI use. Regression analysis on a sample of 7,209 adults finds that AI use rises strongly with education and income and declines with age. In contrast, among those who have adopted AI, the breadth of use is driven primarily by education and age, with income playing little role beyond the adoption decision. Among users, concern about AI’s impact on careers is concentrated among more educated and teleworking individuals, with income providing little explanatory power for career concern. Among non-users, sex and race are the dominant predictors of deliberate abstention. These findings have direct implications for AI literacy programs, workforce policy, and the design of consumer AI regulations.

I. Introduction

Generative artificial intelligence (“AI”) tools, AI-powered search, writing assistants, and AI embedded in workplace software have moved from research concept to daily consumer experience in the span of a few years. The question of who gains access to these tools and on what terms has become a vital concern for policymakers, legislators, and agencies focused on economic equity and workforce readiness. While broadband adoption has long been characterized in terms of a digital divide, AI adoption opens the door for a new sort of digital divide.¹ UNESCO’s Global AI Ethics and Governance Observatory has identified AI literacy as

¹ M. Trucano, *AI and the Next Digital Divide in Education*, BROOKINGS INSTITUTE (July 10, 2023) (available at: <https://www.brookings.edu/articles/ai-and-the-next-digital-divide-in-education>).

PHOENIX CENTER FOR ADVANCED LEGAL & ECONOMIC PUBLIC POLICY STUDIES

5335 Wisconsin Avenue, NW, Suite 440

Washington, D.C. 20015

Tel: (+1) (202) 274-0235 • Fax: (+1) (202) 318-4909

www.phoenix-center.org

an emerging form of digital inequality and has called on governments to intervene before the benefits of AI concentrate further among the already advantaged.² These concerns are well-founded but somewhat underspecified. Whether the household AI adoption gap resembles an access and affordability problem amenable to subsidy and infrastructure investment, or is instead rooted in education and human capital deficit, has direct implications for policy design. The present paper provides some empirical evidence on that question using a nationally representative sample of U.S. adults.

The U.S. Census Bureau's Household Pulse Survey (HPS), Wave 2506, adds a new AI module that enables some descriptive analyses of AI adoption, use, worries about career impacts, and actively opting out of AI use.³ Regression analysis is used to provide an assessment of demographic factors influencing these outcomes. The analysis reveals that AI use is strongly increasing in education and income and strongly decreasing in age. Usage intensity is driven primarily by education and age, with scant evidence income plays a role. Among users, fear of AI's career consequences is driven by age and education but not income, with Blacks and Hispanics less worried about job displacement. Among non-users, active resistance is largely independent of age and income with women more likely to actively opt out. Minorities are less likely to be active resisters.

The paper proceeds as follows. Section II provides a literature review. Section III describes the data and sample construction. Section III presents the analytical framework and specification. Section IV reports results across all three models. Section V discusses draws policy implications, and Section VII concludes.

II. Literature Review

The empirical study of AI adoption is young but growing rapidly. Several recent studies have begun to characterize the demographic profile of AI users. Wang, Zhou, and Liu (2026), using the Pew Research Center's American Trends Panel ($n = 10,087$, 2022 data), examine socioeconomic disparities in AI awareness (defined as the ability to correctly identify AI-enabled technologies in everyday applications) and find that higher education and income are both positively associated, and age negatively associated, with greater AI awareness, AI usage, and perceived familiarity with AI.⁴ The use variable is a categorical variable ranging from "less often" to "almost constantly," and has no non-use response. Also, Wang, *et al.* (2026) use a linear regression model for categorical/ordered outcomes and treat some categorical/ordered variables as continuous

² S. Gonzales, *AI Literacy and the New Digital Divide – A Global Call for Action*, UNESCO (August 6, 2024) (available at: <https://www.unesco.org/ethics-ai/en/articles/ai-literacy-and-new-digital-divide-global-call-action>).

³ Data available at: <https://www.census.gov/data/tables/2025/demo/hhp/2506.html>.

⁴ S. Wang, T. Zhou and X. Liu, *Socioeconomic Disparities in AI Awareness: Examining the Mediating Roles of AI Usage and Familiarity*, INFORMATION, COMMUNICATION & SOCIETY (2026).

PHOENIX CENTER FOR ADVANCED LEGAL & ECONOMIC PUBLIC POLICY STUDIES

5335 Wisconsin Avenue, NW, Suite 440

Washington, D.C. 20015

Tel: (+1) (202) 274-0235 • Fax: (+1) (202) 318-4909

www.phoenix-center.org

(*e.g.*, income), so their findings are questionable, though perhaps the coefficients are directionally valid. Liu and Wang (2025), using global website traffic data from Semrush for the top 40 generative AI tools, find that generative AI users skew young, highly educated, and male across more than 200 economies as of early 2024.⁵ These patterns, however, rest on non-representative web traffic data rather than probability samples, and their generalizability to the U.S. household population is unknown. Humlum and Vestergaard (2024), in a large-scale survey of 100,000 Danish workers in AI-exposed occupations, find that younger and male workers lead ChatGPT adoption, and that employer restrictions and required training are primary barriers among non-users.⁶ Crucially, they find that informing workers about AI's capabilities shifts beliefs but has limited effects on actual adoption, suggesting that some non-adoption reflects deliberate choice rather than ignorance.

Alikhani, Harris, and Patnaik (2025), using the AmeriSpeak Omnibus Survey administered by NORC, report that 57% of U.S. adults use generative AI for at least one personal purpose, with higher education levels strongly associated with increased use, where 67% of those with a bachelor's degree or more use AI compared to 46% of high-school graduates.⁷ Their tabulations also document a male professional use advantage and a sharp drop in AI use among those aged 60 and above. At the firm level, McElheran, *et al.* (2024), using the 2018 Annual Business Survey of 850,000 U.S. firms, find that AI adoption is concentrated among larger, younger, and more educated firms, and warn explicitly of a growing "AI divide" if early adoption patterns persist.⁸ Whether the individual-level household patterns mirror the firm-level patterns documented by McElheran, *et al.*, and whether the distinction between deliberate and passive non-adoption matters empirically, are questions the present paper addresses with a larger, more recent probability sample and a regression-based framework that conditions on multiple demographic factors simultaneously.

Understanding who fears AI's career consequences requires first understanding which workers are most exposed to AI's capabilities. Eloundou, *et al.* (2024), applying an LLM-based rubric to O*NET task data, finds that roughly 80% of workers have more than 10% of their job tasks exposed to potential productivity gains from large language models, with exposure peaking

⁵ Y. Liu and H. Wang, *Who on Earth is Using Generative AI?* 199 WORLD DEVELOPMENT 107260 (2026).

⁶ A. Humlum and E. Vestergaard, *The Adoption of ChatGPT*, IZA Institute of Labor Economics, DP No. 16992 (2024) (available at: <https://www.iza.org/publications/dp/16992/the-adoption-of-chatgpt>).

⁷ M. Alikhani, B. Harris, and S. Patnaik, *How are Americans Using AI? Evidence from a Nationwide Survey*, BROOKINGS INSTITUTE (November 25, 2025) (available at: <https://www.brookings.edu/articles/how-are-americans-using-ai-evidence-from-a-nationwide-survey>).

⁸ K. McElheran, J.F. Li, E. Brynjolfsson, Z. Kroff, E. Dinlersoz, L. Foster, and N. Zolas, *AI Adoption in America: Who, What, and Where*, 33 JOURNAL OF ECONOMICS AND MANAGEMENT STRATEGY 375-415 (2024).

around \$90,000 in annual earnings and remaining high into six-figure occupations.⁹ Higher-wage, higher-skill occupations are thus more exposed to AI's transformative potential. Brynjolfsson, Li, and Raymond (2025), studying the staggered rollout of a generative AI assistant to 5,172 customer support agents, find that AI raises productivity by 15% on average, with gains concentrated disproportionately among less experienced and lower-skilled workers.¹⁰ Alikhani, *et al.* (2025) find that workers are broadly skeptical about AI's labor market effects with only 11% anticipating that AI will increase job opportunities in their field over the next five years, and fewer than one in five report that AI has increased their productivity.¹¹ These findings suggest a complex and potentially counter-intuitive relationship between skill, AI exposure, and labor market anxiety.

The digital divide literature, originating in debates over broadband access in the late 1990s and early 2000s, documented that internet access was stratified by income, education, and geography in ways that threatened to compound existing economic inequalities. The conventional policy response—infrastructure subsidies, affordability programs, and buildout mandates—was calibrated to solve an access problem. Whether the emerging AI adoption gap has the same structure, and therefore admits the same policy responses, is not at all obvious. McElheran, *et al.* (2024) raise this concern at the firm level, warning that geographic and organizational clustering of early AI adoption may widen productivity gaps between leading and lagging firms.¹² Liu and Wang (2025) document the same concern internationally, finding that low-income countries account for less than 1% of global ChatGPT traffic while emphasizing the need for targeted investment in digital infrastructure and skills development.¹³ Alikhani, *et al.* (2025) identify education as the dominant demographic factor in U.S. AI adoption, noting that expanded AI take-up is most needed among workers and households with lower levels of education.¹⁴

Bassignana, *et al.* (2025), surveying 1,000 U.S. and U.K. adults and collecting 6,482 real prompts from prior LLM interactions, find that socioeconomic status shapes not only whether

⁹ T. Eloundou, S. Manning, P. Mishkin, and D. Rock, *GPTs are GPTs: An Early Look at the Labor Market Impact Potential of LLMs*, 384 *SCIENCE* 1306-1308 (2024) (available at: <https://www.science.org/doi/10.1126/science.adj0998>).

¹⁰ E. Brynjolfsson, D. Li, and L. Raymond, *Generative AI at Work*, 140 *QUARTERLY JOURNAL OF ECONOMICS* 889-942 (2025).

¹¹ *Supra* n. 7.

¹² *Supra* n. 8.

¹³ *Supra* n. 3.

¹⁴ *Supra* n. 7.

people use AI but how they use it.¹⁵ Higher-SES individuals use AI chatbots more frequently and for more productive purposes (*e.g.*, writing, coding, data analysis, and professional tasks) while lower-SES users engage more with entertainment and general knowledge queries. Higher-SES users also write shorter, more abstract prompts, consistent with greater command of the kind of language that elicits effective AI responses. These differences compound the adoption gap: not only do lower-SES individuals use AI less, but when they do use it they extract less productive value from the interaction. This finding extends Trucano’s (2023) argument that the broadband digital divide is not merely about who has access or who adopts, but about who has the human capital to exploit the technology effectively once inside.¹⁶

As for the digital divide related to AI, a critical question is whether the household AI adoption gap reflects access and awareness barriers amenable to conventional digital divide remedies, or whether it reflects deliberate choice among informed non-users, a distinction with fundamentally different policy implications. Some insights are provided here, but the research in this area will likely expand rapidly in the coming years, becoming increasingly targeted on such concerns.

III. Data

The data are drawn from Wave 2506 of the U.S. Census Bureau’s Household Pulse Survey (“HPS”). The HPS is a nationally representative online survey of U.S. adults designed to measure household economic and social conditions. Respondents are sampled from the Census Bureau’s Master Address File; data are weighted to align with American Community Survey (“ACS”) population estimates. Wave 2506 introduced a new topical module on artificial intelligence—the first such module in the HPS series. The module covers AI adoption and usage, attitudes toward AI among users, and whether respondents have actively opted out. All four elements are used in this paper.

The raw PUF contains 7,485 respondent records. Missing values are addressed by pairwise deletion ($n = 276$). A binary indicator equal to 1 if the household used AI for any purpose in the past two months. Respondents who selected any of seven AI use purposes (see Table 1) are coded as users ($USE = 1$). Respondents who selected the explicit non-user flag (“No one in my household has used an AI tool”) are coded $USE = 0$. The 154 observations where neither use nor non-use was indicated are dropped from the sample. As a measure of scope of use, the responses to the individual use questions are summed ($UNEN$). The analytical sample contains 7,209 observations. Users ($USE = 1$) number 4,564 (63.3%); confirmed non-users ($USE = 0$) number

¹⁵ E. Bassingnana, A. Curry, and D. Hovy, *The AI Gap: How Socioeconomic Status Affects Language Technology Interactions*, Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 18647–18664 (July 27 - August 1, 2025, Association for Computational Linguistics).

¹⁶ *Supra* n. 1.

2,645 (36.7%), with weighted mean shares of 57.1% users and 42.9% non-users. On average, a user responds affirmatively to 2.4 of the uses.

A second binary indicator (*WORRY*) equals 1 if the respondent selected “I am worried about AI’s impact on my career.” *WORRY* is coded as missing for non-users (only users responded to this question) and the *WORRY* model is estimated on the user subsample ($n = 4,564$). In that subsample, 24.4% express career worry from AI. Finally, AI non-use is not a unitary phenomenon. Passive non-use may reflect lack of occasion, limited awareness, or absence of a relevant use case. Active resistance reflects an affirmative decision against AI, presumably driven by privacy concerns, distrust of AI outputs, ideological objection, or deliberate preference assessment. These mechanisms differ in their policy implications and in their expected demographic correlates. A binary indicator (*OPTOUT*) equals 1 if the respondent reports having *actively* opted out of AI use, and 0 if the respondent has not. Respondents answering “Not sure” are excluded from the *OPTOUT* model to maintain a clean binary outcome. The *OPTOUT* model is estimated on the non-user subsample with binary opt-out responses ($n = 2,062$). Among these non-users, 32.4% have actively opted out.

Table 1. Summary of AI Use Types, Worry, and Opting Out

Item	Unweighted Count	Unweighted Mean	Weighted Mean
Finding factual information	2,946	0.409	0.360
Assisting with schoolwork	641	0.089	0.119
Assisting with work projects	1,869	0.259	0.214
Performing a task ... otherwise hired someone to do	370	0.051	0.046
Assisting with creative tasks	1,538	0.213	0.201
Brainstorming or idea generation	1,901	0.264	0.234
Integrated into another product you use	1,702	0.236	0.195
No one in my household has used an AI tool	2,645	0.367	0.429
Mean Adoption (<i>USE</i>)	7,209	0.633	0.571
Mean Answer Count (<i>USEN</i>)	4,564	2.40	2.40
Mean Worried about AI’s impact on my career (<i>WORRY</i>)	4,564	0.262	0.244
Mean Actively opt out of Using AI (<i>OPTOUT</i>)	2,062	0.317	0.324

Table 1 reports use prevalence by purpose. Finding factual information is the most common use (36%), followed by brainstorming (23.4%), assisting with work projects (21.4%), creative tasks (20.1%), and integrated AI products (19.5%). Schoolwork assistance (11.9%) and replacing a hired task (4.6%) are less common. The mean adoption rate is 57.7% and the mean use count (*USEN*) is 2.4 uses. Of users, 24.4% express concerns about job security, and for non-users 32.4% actively opt out of using AI.

PHOENIX CENTER FOR ADVANCED LEGAL & ECONOMIC PUBLIC POLICY STUDIES

5335 Wisconsin Avenue, NW, Suite 440

Washington, D.C. 20015

Tel: (+1) (202) 274-0235 • Fax: (+1) (202) 318-4909

www.phoenix-center.org

Several covariates are used in the analysis; descriptive statistics are summarized in Table 2 along with the observed *USE* rate for each demographic. Age is grouped into six categories: 18–29 (reference), 30–39, 40–49, 50–59, 60–69, 70+. Income is a seven-category household income variable (reference: Category 1, less than \$25,000), ranging through Category 7 (\$150,000 or more). Education is coded as three binary indicators relative to a less-than-some-college reference: some college or an associate’s degree; bachelor’s degree; and graduate or professional degree. Sex is a binary indicator (*FEMALE* = 1). Race/Ethnicity indicators are created for Black non-Hispanic, Hispanic, Asian non-Hispanic, and Other races non-Hispanic, with White non-Hispanic as the reference group. Also, there is a binary indicator for telework (any number of days). Finally, a binary indicator is created for metropolitan statistical area residents.

Table 2. Covariate Descriptive Statistics
(*n* = 7,209)

Covariate	Mean	<i>USE</i>	Covariate	Mean	<i>USE</i>
Age 18-29	0.186	0.686	No College	0.367	0.419
Age 30-39	0.180	0.602	Some College	0.291	0.554
Age 40-49	0.167	0.651	Bach. Degree	0.185	0.706
Age 50-59	0.146	0.621	Graduate Degree	0.157	0.797
Age 60-69	0.176	0.464	White	0.592	0.571
Age 70+	0.146	0.374	Black	0.142	0.522
Income < 24,000	0.186	0.392	Hispanic	0.117	0.452
Income 25,000 to 34,999	0.069	0.429	Asian	0.065	0.782
Income 35,000 to 49,999	0.092	0.449	Other Races	0.084	0.655
Income 50,000 to 74,999	0.171	0.567	Telework	0.198	0.797
Income 75,000 to 99,999	0.129	0.592	Not Telework	0.802	0.515
Income 100,000 to 149,999	0.160	0.664	Metro	0.522	0.595
Income ≥ 150,000	0.194	0.763	Not Metro	0.478	0.545
Female	0.506	0.531			
Male	0.494	0.611			

Several patterns in AI adoption (*USE*) are apparent. Adoption declines monotonically with age, from 68.6% among the youngest cohort (18–29) to 37.4% among those aged 70 and above. The income gradient is strong and approximately monotone: adoption rises from 39.2% among households earning less than \$24,000 to 76.3% among those earning \$150,000 or more, with a notable step-up beginning around the \$50,000 threshold. Education shows the steepest gradient of any covariate: adoption is 41.9% among those without a college degree, 55.4% among those with some college, 70.6% among college graduates, and 79.7% among those with a graduate or professional degree. Men adopt at higher rates than women (61.1% vs. 53.1%), a gap that the regression analysis below attributes primarily to differences among less-educated respondents. Among racial and ethnic groups, Asian respondents show the highest adoption rate (78.2%),

PHOENIX CENTER FOR ADVANCED LEGAL & ECONOMIC PUBLIC POLICY STUDIES

5335 Wisconsin Avenue, NW, Suite 440

Washington, D.C. 20015

Tel: (+1) (202) 274-0235 • Fax: (+1) (202) 318-4909

www.phoenix-center.org

followed by White non-Hispanic (57.1%), Other races (65.5%), Black non-Hispanic (52.2%), and Hispanic (45.2%). Teleworkers adopt at substantially higher rates than non-teleworkers (79.7% vs. 51.5%), and metropolitan residents adopt at modestly higher rates than non-metropolitan residents (59.5% vs. 54.5%).

IV. Empirical Analysis

Binary outcomes are estimated using a Logit Model with survey weights and heteroskedasticity-robust standard errors. The number of AI uses—a count variable—is estimated by Poisson regression with robust standard errors to address the potential for over-dispersion in the distribution. The models are estimated independently. As the Logit coefficients are not directly interpretable, the average marginal effects (“AMEs”) are reported. For the Logit models, the AMEs may be interpreted as the percentage point difference in the probability of the outcome associated with each covariate (all are binary). With Poisson, the coefficients are semi-elasticities and are interpreted as the percent change in the count of AI uses associated with each covariate.

A. AI Adoption and Uses

Table 3 reports average marginal effects from the AI adoption (*USE*) and AI uses models with 7,209 observations (the full sample) for *USE* and 4,564 for *USEN* (only users). The mean of the *USE* outcome is 0.57 and the *USEN* outcome is 2.4. For the *USE* model, the reported estimates are AMEs (percentage point change) and for the *USEN* model they are semi-elasticities (percentage change).

Table 3. Summary Results of AI Use Model (*USE* and *USEN*)

Covariate	<i>USE</i> (Δ pp)	<i>USEN</i> (Δ %)	Covariate	<i>USE</i> (Δ pp)	<i>USEN</i> (Δ %)
Age 30-39	-0.152*** (-3.93)	-0.151*** (-2.80)	Some College	0.081*** (2.83)	0.161** (2.54)
Age 40-49	-0.089** (-2.15)	-0.238*** (-4.17)	Bach. Degree	0.147*** (5.07)	0.190*** (3.26)
Age 50-59	-0.133*** (-3.41)	-0.343*** (-5.41)	Graduate Degree	0.258*** (8.96)	0.228*** (4.00)
Age 60-69	-0.235*** (-6.21)	-0.540*** (-8.88)	Female	-0.050** (-2.35)	-0.075** (-2.54)
Age 70+	-0.326*** (-8.49)	-0.549*** (-8.05)	Black	-0.021 (-0.33)	-0.103* (-1.67)
Income 25,000 to 34,999	0.039 (0.71)	0.151 (1.49)	Hispanic	-0.042 (-1.19)	-0.112 (-1.31)
Income 35,000 to 49,999	0.050 (1.10)	0.015 (0.16)	Asian	0.101** (3.07)	0.059 (1.11)
Income 50,000 to 74,999	0.149*** (3.79)	-0.002 (-0.03)	Other Races	0.043 (1.01)	0.037 (0.63)
Income 75,000 to 99,999	0.128*** (2.83)	0.051 (0.56)	Telework	0.127*** (5.71)	0.044 (1.44)
Income 100,000 to 149,999	0.160*** (4.09)	0.026** (0.38)	Metro	0.015 (0.73)	0.096*** (3.15)
Income \geq 150,000	0.209*** (5.35)	0.148** (2.13)			
Observations	7,209	4,564			
Wald χ^2	564.2***	397.3***			
Pseudo R ²	0.135	0.043			
Stars *** 1% ** 5% * 10%					

For the *USE* model (adoption), most of the coefficients are well estimated and exhibit sensible patterns. AI use declines nearly monotonically with age; all coefficients are statistically different from zero at the 5% level or better. Relative to the 18–29 reference group, respondents aged 30–39 are 15.2 percentage points (pp) less likely to use AI; those aged 60–69 are 23.5pp less likely; and those aged 70+ are 32.6pp less likely.

The income-adoption gradient is monotonic and non-linear with larger effects concentrated in the upper half of the distribution. The \$25,000–\$49,999 group are statistically indistinguishable from the lowest income group. Beginning at the \$50,000–\$74,999 group, adoption begins to rise nearly monotonically: +14.9pp, +12.8pp, +16.0pp, and +20.9pp, all significant at the 1% level or better. Like age and income, education is a strong and monotone predictor of adoption. Relative to the less-than-some-college reference, respondents with some college education are 8.1pp more likely to use AI, those with a college degree are 14.7pp more likely, and those with graduate degrees are 25.8pp more likely, other things constant.

PHOENIX CENTER FOR ADVANCED LEGAL & ECONOMIC PUBLIC POLICY STUDIES

5335 Wisconsin Avenue, NW, Suite 440

Washington, D.C. 20015

Tel: (+1) (202) 274-0235 • Fax: (+1) (202) 318-4909

www.phoenix-center.org

Females have a slight (conditional) adoption deficit (-5pp , $p = 0.019$); a statistically significant effect. The raw difference in AI adoption rates between men and women is 8pp , where men adopt at a rate of 61.1% versus 53.1% for women. Thus, the observed raw adoption gap is explained only in part by the covariates. As for race, Asian respondents are 10.1pp ($p = 0.002$) more likely to use AI than White non-Hispanic respondents; Black, Hispanic and Other race respondents are show no significant difference from White non-Hispanic respondents in overall adoption. Teleworkers are 12.7pp ($p < 0.001$) more likely to use AI, which is unsurprising. Metro residents do not appear more likely to adopt AI than more rural respondents.

In large part, age, income, and education are the dominant factors affecting AI use, with females using AI slightly less than males. Race has mixed effects and does not appear to be a substantial factor. Though AI is somewhat in its infancy, even if adoption is high, these results contribute to the policy discussion for a new type of digital divide where AI use is concentrated in higher incomes, higher education, and greater youth.

The Poisson model of the number of AI uses (*USEN*) is estimated on the 4,564 user subsample. The dependent variable ranges from 1 to 7, with a mean of 2.4 for users, representing the count of distinct AI purposes selected by the respondent. The model is estimated with robust standard errors, so the standard errors are valid regardless of overdispersion in the count distribution. The reported results are the coefficients (semi-elasticities) and measure the percent change in the number of uses when the binary indicator is 1 and not 0. The sizes of the coefficient are not directly comparable to the *USE* model.

The results reveal an important asymmetry in adoption. Age continues to predict increasing AI use strongly and monotonically among users. Relative to users aged 18–29, those aged 30–39 use AI for 15% fewer purposes, those aged 50–59 for 34% fewer purposes, and those aged 60+ for approximately 54% fewer purposes. Older users not only adopt AI at lower rates, but they also use it less broadly when they do adopt. The effects of income are non-monotonic and statistically significant effects are limited to the highest income groups, with a large 14.8% increase in use intensity for the highest-income group. Education remains a positive and significant predictor of use intensity across all degree levels. Users with some college education use AI for 16% more purposes than the less-than-some-college reference group; those with a bachelor's degree for 19% more; and graduate degree holders for 23% more.

The female coefficient is negative and statistically significant at the 5% level (-7.5%). Women who adopt AI use it for fewer purposes than comparable men, though the effect is relatively small. Among racial and ethnic groups, Black users use AI for 10.3% fewer purposes than White non-Hispanic users, a result that is marginally significant. The effects for Hispanics (-11.2%), Asians ($+5.9\%$), and Other races (3.7%) are less precisely estimated and not statistically different from zero. Metro residents use AI for nearly 10% more purposes than non-metro residents—a statistically significant result, suggesting that urban location predicts intensity of use and the

PHOENIX CENTER FOR ADVANCED LEGAL & ECONOMIC PUBLIC POLICY STUDIES

5335 Wisconsin Avenue, NW, Suite 440

Washington, D.C. 20015

Tel: (+1) (202) 274-0235 • Fax: (+1) (202) 318-4909

www.phoenix-center.org

adoption decision itself (though the adoption effect is small). Teleworkers' intensity (+4.4%) has a small effect statistically no different from zero.

Taken together, the results characterize AI use as a two-stage process: the decision to adopt AI and the breadth of AI use conditional on adoption are governed by different factors. Age and education predict both margins consistently, but income—a strong predictor of adoption—largely disappears as a predictor of use intensity except at the highest incomes. Crossing the adoption threshold is primarily a function of age, income, and education. Once that threshold is crossed, income recedes as a determinant and the breadth of AI use is shaped primarily by age and education. Thus, while income gets households in the AI door, it does not determine how broadly they use AI once inside. This asymmetry indicates the adoption decision and the intensity decision are distinct; treating them as a single process would obscure a meaningful difference in how income operates across the two margins. The analysis also raises an important social question regarding the relative importance of adoption versus the intensity of use, since the latter is conditional on the former.

B. *Worry About Job Displacement*

Table 5 reports average marginal effects from the logistic model of career fear, estimated on the 4,564 user observations (non-users omitted). Fear of AI's career impact declines sharply with age among users. Relative to users aged 18–29, users aged 50–59 are 14.2pp less likely to express career fear ($p < 0.001$), those aged 60–69 are 22.9pp less likely ($p < 0.001$), those aged 70 or more are 24.8pp less likely ($p < 0.001$). Older users who have adopted AI are substantially less concerned about its career consequences, perhaps because retirement or late-career status reduces the relevance of labor market displacement.

Income is irrelevant to career fear among users. None of the six income dummy coefficients achieves statistical significance at conventional levels, and the point estimates show no monotone pattern. A respondent's income level, conditional on AI use, provides no information about their likelihood of fearing AI's labor market impact.

Table 5. Summary Results of Worry About Job Displacement (WORRY)

Covariate	WORRY (Δ pp)	Covariate	WORRY (Δ pp)
Age 30-39	-0.027 (-0.59)	Some College	0.017 (0.41)
Age 40-49	-0.073 (-1.60)	Bach. Degree	0.092** (2.25)
Age 50-59	-0.142*** (-3.23)	Graduate Degree	0.090** (2.25)
Age 60-69	-0.229*** (-5.16)	Female	-0.019 (-0.83)
Age 70+	-0.248*** (-4.96)	Black	-0.130*** (-2.93)
Income 25,000 to 34,999	0.103 (1.20)	Hispanic	-0.112** (-2.29)
Income 35,000 to 49,999	0.021 (0.32)	Asian	0.019 (0.57)
Income 50,000 to 74,999	0.038 (0.76)	Other Races	0.013 (0.28)
Income 75,000 to 99,999	-0.024 (-0.41)	Telework	0.099*** (5.00)
Income 100,000 to 149,999-	-0.031 (-0.70)	Metro	0.045** (2.23)
Income \geq 150,000	-0.011 (-0.23)		
Observations	4,564		
Wald χ^2	192.4***		
Pseudo R ²	0.104		
Stars *** 1% ** 5% * 10%			

Education is a positive and significant predictor of career fear. College graduates are 9.2pp ($p = 0.025$) more likely than the less-than-some-college reference group to fear AI's career impact, and graduate degree holders are 9.0pp more likely ($p = 0.024$). Having some college education or an associate's degree shows no significant difference from the reference (+1.7pp, $p = 0.683$). Highly educated users worry more about AI displacing them, a result that is consistent with the observation that higher-skilled occupations are increasingly exposed to generative AI tools. Teleworkers express deep concern with 9.9pp ($p < 0.001$) with job displacement.

Women are 1.9pp ($p = 0.41$) less likely than men to fear AI's career impact among users, but the difference is small and not statistically significant. Black users are 13.0pp ($p = 0.003$) less likely than White non-Hispanic users to fear AI's career impact. Hispanic users are 11.2pp less likely ($p = 0.022$). These are large and significant effects. Asian users show no significant difference from White non-Hispanic users (+1.9pp, $p = 0.570$). Other races are no different than the base group. Respondents living in metropolitan areas are 4.5pp ($p = 0.026$) more likely to fear AI's career impact, consistent with greater exposure to AI-adopting employers in urban labor markets.

C. *Active Opt-Out Among Non-Users*

Table 6 reports the AMEs from the Logit model of active opt-out, estimated on 2,062 non-users with binary opt-out responses. The mean of the dependent variable is 32.4%. Among non-users, age groups 30–39, 40–49, 50–59, and 60–69 are all statistically indistinguishable from the 18–29 reference group. Only the 70+ group show significant differences: respondents aged 70 or more are 24pp less likely to opt out than the youngest non-users ($p = 0.004$). This result may be an artifact of knowledge. Across the bulk of the working-age non-user population, age does not predict whether abstention is deliberate or passive. This finding contrasts sharply with Model 1, where every age group above 30 shows significantly lower adoption. Non-use rises monotonically with age; active resistance does not. The middle-aged non-user is no more a deliberate resister than the young non-user.

Table 6. Summary Results of Active Opt-Out of AI Use (OPTOUT)

Covariate	OPTOUT (Δ pp)	Covariate	OPTOUT (Δ pp)
Age 30-39	-0.003 (-0.03)	Some College	0.011 (0.29)
Age 40-49	-0.011 (0.12)	Bach. Degree	0.042 (0.96)
Age 50-59	-0.075 (-0.88)	Graduate Degree	0.043 (0.89)
Age 60-69	-0.037 (-0.45)	Female	0.075** (2.20)
Age 70+	-0.223*** (-2.89)	Black	-0.106* (-1.74)
Income 25,000 to 34,999	0.144* (1.89)	Hispanic	-0.107** (-2.17)
Income 35,000 to 49,999	0.032 (0.49)	Asian	-0.350*** (-4.62)
Income 50,000 to 74,999	0.094* (1.70)	Other Races	-0.057 (-0.75)
Income 75,000 to 99,999	0.043 (0.72)	Telework	0.102*** (2.54)
Income 100,000 to 149,999-	0.043 (0.77)	Metro	-0.020 (-0.59)
Income \geq 150,000	0.002 (0.04)		
Observations	2,062		
Wald χ^2	89.5***		
Pseudo R ²	0.066		
Stars *** 1% ** 5% * 10%			

Income shows no systematic relationship with active opt-out with few statistically significant coefficients and then only at the 10% level. Two income dummies achieve marginal

PHOENIX CENTER FOR ADVANCED LEGAL & ECONOMIC PUBLIC POLICY STUDIES

5335 Wisconsin Avenue, NW, Suite 440

Washington, D.C. 20015

Tel: (+1) (202) 274-0235 • Fax: (+1) (202) 318-4909

www.phoenix-center.org

significance—\$25,000-\$34,000 (+14.4pp, $p = 0.058$) and \$50,000-\$74,999 (+9.4pp, $p = 0.090$)—but the pattern has discernible gradient. Income predicts adoption strongly; it does not predict resistance among those who have not adopted. The effect on the \$50,000-\$74,000 group may reflect Eloundou, *et al.*'s (2024) finding that LLM exposure peaks around \$90,000 in annual earnings, a deliberate opt out related to employment concerns. Education shows no statistically significant relationship with active opt-out among non-users. Actively opting out of AI use is nearly unrelated to age, income, and education.

Table 6 shows that women are 7.5pp ($p = 0.028$) more likely than men to actively opt out among non-users. Women who do not use AI are disproportionately deliberate in their abstention, while men are passive non-users. Whether that rejection reflects privacy concerns, distrust of AI accuracy, lack of relevant use cases, or other factors is not identified in the HPS data. Asian non-users are 34.6pp ($p < 0.001$) less likely than White non-Hispanic non-users to have actively opted out. This is the largest and most precisely estimated finding in the model. Asian respondents are significantly more likely than average to use AI (+4.2pp in the *USE* model). Among those who do not use AI, they are significantly less likely to have actively opted out. Asian non-users are disproportionately passive, and their non-use reflects circumstance rather than deliberate choice. Black non-Hispanic non-users and Hispanics and are both about 10pp less likely to opt out ($p = 0.081, 0.030$).

Teleworkers are more likely (+10.8pp, $p = 0.011$) to actively opt-out of AI use, and they express greater worry about job displacement. The effect size is relatively large. These results might suggest some type of (perhaps futile) rebellion against a technology threat. Respondents living in metro areas are no more likely to actively opt-out of AI use than their rural counterparts.

V. Policy Implications

The results reported above have several implications for policy discussions about AI adoption, workforce readiness, and the emerging AI divide. The AI divide is a human capital problem, not an access problem. The adoption gradient is dominated by age and education, not geography or income at the lower end of the distribution. Income matters, but only above the \$50,000 threshold with the bottom half of the income distribution shows little differentiation in AI adoption rates. Females adopt AI at a slightly lower rate than males. Metropolitan location, the traditional proxy for digital infrastructure advantage, has a small but statistically insignificant in the adoption model. These patterns place the AI divide closer to Trucano's (2023) second and third divides (skills and human capital) than to the first divide of access and affordability.¹⁷ The conventional digital divide policy toolkit of subsidized devices, broadband buildout, and

¹⁷ *Supra* n. 1.

affordability programs addresses a problem that does not appear to be the binding constraint here.

The intensity margin compounds the adoption gap. The *USEN* model results show that the same groups disadvantaged in adoption—older workers, less-educated respondents, women—also use AI less broadly conditional on adoption. The divide is not just about who enters the AI economy; it is about how deeply they engage once inside. A household that adopts AI for a single purpose presumably extracts far less value than one that integrates it across multiple dimensions of work and daily life. Policies that promote adoption without also addressing engagement depth will close only part of the gap in AI-derived economic benefit.

Fear of AI displacement is concentrated where exposure is highest. Higher-educated and teleworking users are the most likely to worry about AI's career consequences. Income is irrelevant to that worry. This is counterintuitive relative to narratives that frame AI displacement risk as primarily a threat to lower-wage workers. Despite Eloundou, *et al.*'s (2024) finding that LLM exposure peaks around \$90,000 in annual earnings, the present results do not suggest that workers whose incomes are most exposed to generative AI are the ones expressing concern; more educated persons are, however, more worried. Black and Hispanic workers are significantly less worried. Whether that reflects differential occupational exposure, different assessments of AI's capabilities, or other factors is not identified here, but the pattern warrants attention in workforce policy discussions that assume AI anxiety is concentrated among vulnerable low-wage workers.

Active resistance to AI is not the same as passive non-use, and policy should treat them differently. Among non-users, 32.4% have made an affirmative decision to abstain from AI. That decision is largely independent of age, income, and education, which are the factors that drive adoption. It is predicted instead by sex, race and ethnicity, and telework status. Women are disproportionately active resisters; Asian, Hispanic, and Black non-users are disproportionately passive. AI literacy programs aimed at non-users conflate two fundamentally different populations. Passive non-users—those who have not opted out and may simply lack occasion or awareness—are plausible targets for outreach and demonstration programs. Active resisters have made a considered choice and are unlikely to be moved by awareness campaigns. The finding that teleworkers are more likely to opt out despite being more likely to use AI, and more likely to fear displacement, hints at a deliberate resistance among workers who feel AI is being imposed on their working lives, which is a concern more responsive to workplace governance and worker voice policies than to literacy programs.

Finally, the distinction between adoption and intensity—between crossing the threshold into AI use and the breadth of use conditional on adoption—has direct policy implications that the digital divide literature has not fully confronted. The broadband analogy is instructive but cuts both ways. The original digital divide concern was that non-adopters were being left behind economically and socially, and the policy remedy was to close the adoption gap through infrastructure investment, affordability programs, and digital literacy initiatives. But broadband

PHOENIX CENTER FOR ADVANCED LEGAL & ECONOMIC PUBLIC POLICY STUDIES

5335 Wisconsin Avenue, NW, Suite 440

Washington, D.C. 20015

Tel: (+1) (202) 274-0235 • Fax: (+1) (202) 318-4909

www.phoenix-center.org

policy quickly learned that adoption alone was an insufficient target. A household that uses broadband only for entertainment derives far less economic benefit than one that uses it for job search, education, remote work, and commerce. More consequentially, broadband access enabled platforms (e.g., social media) that have been associated with mental health harms, political polarization, and misinformation diffusion. The same infrastructure that enables telemedicine and remote work also enables doom-scrolling and radicalization. Broadband policy never seriously grappled with that duality, treating access as unambiguously good and leaving use-quality questions to other regulatory domains. The result was a policy framework calibrated to the adoption margin while remaining largely silent on the intensity and use-case margins where the most consequential effects, whether positive and negative, were operating.

AI presents the same duality in sharper form. Intensity of AI use is not unambiguously beneficial, and the relationship between intensity and benefit depends entirely on the use case. For a skilled worker using AI to accelerate research, draft communications, or analyze data, broader use likely represents genuine productivity enhancement. For a student using AI to complete assignments, broader use may represent substitution of AI output for the cognitive effort that produces learning, undermining the human capital accumulation that drives long-term outcomes. The findings in this paper that younger respondents use AI most broadly, and that schoolwork assistance is among the commonly reported use cases, make this concern immediate rather than prospective. A student using AI for seven purposes including homework completion is not obviously better off than one using it for none, and may be worse off in ways that compound over time as the learning deficit accumulates.

These findings suggest that AI adoption policy needs to be use-case specific in a way that broadband policy never was. Broadband access is a relatively neutral infrastructure input across use cases – the pipe does not know whether it is carrying a job application or a social media feed, and policy treated it accordingly. In contrast, *AI is not neutral*: the same tool that raises a skilled worker’s productivity may undermine a student’s learning. Closing the AI adoption and intensity gaps is therefore not an unambiguous policy goal. Programs aimed at expanding AI use should specify which uses they are promoting and for whom. Education-specific AI policy may need to run in the opposite direction from general AI literacy policy by not just expanding use but shaping, structuring, and in some contexts limiting it. The broadband experience suggests that a technology policy framework built entirely around closing an adoption gap, without confronting the use-quality question, will eventually face a reckoning with the harms it enabled by treating access as an end in itself.

VI. Conclusion

This paper uses new data from the U.S. Census Bureau’s Household Pulse Survey, Wave 2506, to examine the demographic determinants of AI adoption, use intensity, career worry, and active resistance among U.S. adults. The analysis produces three principal findings.

PHOENIX CENTER FOR ADVANCED LEGAL & ECONOMIC PUBLIC POLICY STUDIES

5335 Wisconsin Avenue, NW, Suite 440

Washington, D.C. 20015

Tel: (+1) (202) 274-0235 • Fax: (+1) (202) 318-4909

www.phoenix-center.org

First, AI adoption is strongly increasing in education and income and strongly decreasing in age, with a threshold structure in income that concentrates the gradient in the upper half of the distribution. Telework is a strong positive predictor of adoption, reflecting the occupational and workplace channels through which AI enters daily life. The dominant cleavages are age and education—human capital dimensions that operate on generational timescales. The implication is that the conventional policy toolkit for closing digital divides (*e.g.*, infrastructure investment, device subsidies, affordability programs) addresses the wrong binding constraint.

Second, among users, fear of AI's career consequences is driven by education, age, and labor market context, not income. Higher-educated and teleworking users are the most worried about AI displacing them, consistent with evidence that generative AI exposure is concentrated in higher-skill, higher-wage occupations. Black and Hispanic users are significantly less likely to express career worry, a finding that deserves further investigation. The income-neutrality of career fear challenges policy narratives that frame AI displacement risk as primarily a threat to lower-wage workers.

Third, active resistance to AI among non-users is predicted by a different demographic profile than non-use itself. Age, income, and education are largely irrelevant to whether a non-user is a deliberate resister or a passive abstainer. Women are significantly more likely to actively opt out; Asian, Hispanic, and Black non-users are significantly less likely. Teleworkers opt out at higher rates despite being more likely to adopt and more worried about displacement, suggesting that AI resistance in some cases reflects workplace grievance rather than ignorance. These patterns imply that AI literacy programs aimed at non-users may conflate two distinct populations with different policy needs.

Taken together, these results characterize AI non-use as a heterogeneous phenomenon with distinct demographic profiles for distinct mechanisms. The adoption gap, the fear gap, the intensity gap, and the resistance gap are not the same problem in different magnitudes—they are different problems with different demographic signatures and different policy remedies. Progress on AI equity requires first distinguishing which problem is being addressed and for whom. The Household Pulse Survey's new AI module, the first of its kind in a nationally representative government survey, provides the empirical foundation for that distinction and a baseline against which future waves can measure how these patterns evolve as AI continues to diffuse through the economy.

Finally, the distinction between adoption and intensity carries a policy caution that the broadband experience makes concrete. Broadband policy treated access as unambiguously good and left use-quality questions unaddressed; the benefits and harms of social media are the most visible consequence of that omission. AI presents the same duality in sharper form. Intensity of use is not uniformly beneficial: broader AI use by skilled workers likely represents genuine productivity enhancement, while broader use by students (and even adults) completing

PHOENIX CENTER FOR ADVANCED LEGAL & ECONOMIC PUBLIC POLICY STUDIES

5335 Wisconsin Avenue, NW, Suite 440

Washington, D.C. 20015

Tel: (+1) (202) 274-0235 • Fax: (+1) (202) 318-4909

www.phoenix-center.org

assignments may represent substitution for the cognitive effort that produces learning. Closing the AI adoption and intensity gaps is therefore not an unambiguous policy goal.

PHOENIX CENTER FOR ADVANCED LEGAL & ECONOMIC PUBLIC POLICY STUDIES

5335 Wisconsin Avenue, NW, Suite 440

Washington, D.C. 20015

Tel: (+1) (202) 274-0235 • Fax: (+1) (202) 318-4909

www.phoenix-center.org