

Generative AI Adoption in India

Kiran Garimella*

April 2026

Abstract

This paper reports the first survey-based estimates of generative AI use among working-age adults in India. We fielded the U.S. Real-Time Population Survey (Bick et al., 2024) instrument across two samples: an online panel ($N=1,065$, January 2026) and a face-to-face survey in rural Uttar Pradesh ($N=730$, February 2026). Pooled and reweighted to national age, gender, and urbanicity targets ($N=1,795$), 63% of Indian working-age adults report using generative AI outside work, 35% use it at work, and it saves them 2.8% of total work hours—well above the U.S. rates of 38%, 32%, and 1.4% from August 2024. Use rises sharply with education (77% among graduates vs. 32% below) and urban residence (75% vs. 54% rural), and falls with age above 35. Women in the rural face-to-face sample report substantially higher use than men, while in the online panel women and men are roughly at parity; the female premium in the pooled data therefore reflects selection into the rural sample rather than a true population-level reversal of the Indian mobile-internet gender gap, and we flag it as a target for future probability-sample waves.

Keywords: Generative AI, India, technology adoption, digital divide, Global South, survey

*School of Communication and Information, Rutgers University. Email: kg766@rutgers.edu. The author thanks the survey firm that administered the rural Uttar Pradesh wave, and Clickworker for the online panel.

1 Introduction

How widely is generative AI actually used, and by whom? Most of the evidence so far comes from the United States and a handful of other high-income countries (Bick et al., 2024; Humlum and Vestergaard, 2024; Fletcher and Nielsen, 2024). India—the world’s most populous country and a major producer and consumer of digital services—has no comparable individual-level estimate.

Measuring adoption in India is harder than in the United States, because no single survey mode is both affordable and representative. Online panels reach large samples cheaply but oversample urban, educated, English-comfortable, and younger respondents (Internet and Mobile Association of India, 2024; GSMA, 2024). Face-to-face surveys are representative but slow and expensive. Phone surveys sit in between and skew toward those who pick up unknown numbers.

We combine two samples. The first is an online Clickworker panel ($N=1,065$, all-India, English, self-complete, January 2026), using the generative-AI module from the U.S. Real-Time Population Survey (Bick and Blandin, 2023; Bick et al., 2024). The second is a face-to-face survey of rural Uttar Pradesh ($N=730$, Hindi, tablet-administered, February 2026) commissioned specifically to correct the online panel’s urban and educational skew. The two samples use the same instrument. We pool them and reweight to national age, gender, and urbanicity targets.

An interactive companion dashboard with all headline numbers and demographic breakdowns is available at <https://genai-india.sci.rutgers.edu/>.

Four findings stand out. First, Indian adoption is high: 63% of working-age adults report using generative AI outside of work and 35% use it at work, well above the 38% and 32% reported for the U.S. in August 2024. Time savings are also higher—2.8% of total work hours versus 1.4% in the U.S. Second, the within-India gradients are steep. Graduates use AI at more than twice the rate of non-graduates (77% vs. 32%); urban adoption exceeds rural by 21 percentage points; the age gradient is sharp above 35. Third, the most common reason

respondents give for not using AI is informational—“not sure how to use it effectively”—rather than structural constraints like access or cost. Fourth, women in our pooled sample report higher use than men, but the pattern is driven almost entirely by the rural face-to-face wave: in the online panel, men and women are roughly at parity; in the rural wave, women adopt at substantially higher rates than men. We read this as positive selection of women into the rural interviewer-administered survey—women who were at home, available, and willing to be interviewed during the day are a more digitally engaged slice of rural women—rather than a true population-level reversal of India’s mobile-internet gender gap.

2 Related Work

2.1 Measuring Generative AI Adoption

The closest precedent for this paper is [Bick et al. \(2024\)](#), who add a generative-AI module to the U.S. Real-Time Population Survey ([Bick and Blandin, 2023](#)). They report 45.5% adoption overall, 32.1% at work, and 1.4% of work hours saved as of August 2024. We use their instrument directly. [Humlum and Vestergaard \(2024\)](#) track ChatGPT adoption in Denmark using web traffic and survey data and find rapid but uneven uptake across occupations. [Fletcher and Nielsen \(2024\)](#) cover generative-AI use in news consumption across six countries. Industry surveys ([Pew Research Center, 2024](#); [NASSCOM, 2024](#)) provide context but not individual-level estimates. Representative individual-level surveys of generative-AI use in the Global South remain scarce; this paper fills that gap for India. In related work, [Choudhury and Garimella \(2026\)](#) analyze actual ChatGPT conversation exports from about 550 Indian users to describe what Indians use the technology for.

2.2 Generative AI and Productivity

Experimental evidence on AI productivity effects has grown quickly. [Brynjolfsson et al. \(2025\)](#) find a 14% gain in customer-service work, concentrated among less-skilled workers.

Noy and Zhang (2023) find a 40% reduction in writing-task time. Dell’Acqua et al. (2023) report a 12.2% quality gain and 25.1% speed-up for consultants using GPT-4; Peng et al. (2023) report a 55.8% speed-up for developers with GitHub Copilot. Whether these effects scale to aggregate productivity depends on how widely the technology is adopted. Acemoglu (2024) argues that even optimistic scenarios imply a 10-year U.S. GDP effect of 0.5–1.5%; Eloundou et al. (2024) find that 80% of the U.S. workforce has at least 10% of tasks exposed to LLMs. Our survey provides the adoption and time-use inputs needed to repeat these calculations for India.

2.3 Digital Technology Adoption in India and the Global South

India has seen rapid government-scaffolded digital-infrastructure diffusion (UPI, Aadhaar, JAM; Telecom Regulatory Authority of India 2024; Ministry of Electronics and Information Technology 2024), but internet access and digital skills remain unevenly distributed across urban–rural, income, gender, and education lines (Internet and Mobile Association of India, 2024). The GSMA (2024) reports that Indian women are 30% less likely than men to own a smartphone and 40% less likely to use mobile internet. In other Global South settings, technology diffusion is shaped by fixed costs and network externalities (Aker and Mbiti, 2010; Björkegren, 2018), and improved internet access has measurable labor-market effects (Hjort and Poulsen, 2019). Classical frameworks (Rogers, 2003; Venkatesh et al., 2003) were developed in high-income contexts, so their applicability to generative AI in India is an empirical question.

2.4 Survey Modes in LMICs

No single survey mode is both cheap and representative in low- and middle-income countries, and mixed-mode designs have become standard practice (De Leeuw, 2005; Dillman et al., 2014). The World Bank’s Living Standards Measurement Study (World Bank, 2020) combines face-to-face enumeration with phone follow-ups; Pew mixes face-to-face with phone or

web by country ([Pew Research Center, 2018](#)). The design here follows the same logic: pair a cheap mode with a targeted face-to-face survey in the areas the cheap mode under-covers.

3 Data and Methodology

3.1 Survey Instrument

The instrument is adapted from the Real-Time Population Survey module of [Bick et al. \(2024\)](#), with minor localization (Indian rupee income bands; 6-digit PIN code in place of ZIP code). It covers demographics (gender, age, education, household income, PIN code); employment (status, employer type, hours, days); generative-AI use at work (yes/no, products, frequency, time saved, counterfactual hours); generative-AI use outside work (yes/no, products, frequency); and reasons for non-use. Two attention checks ask respondents to select specific answers; respondents failing either check are dropped.

3.2 Wave 1: Online Panel

Wave 1 was fielded on Clickworker in January 2026, targeting Indian residents aged 18–64, in English, self-complete. After dropping respondents who did not finish, failed an attention check, or were missing age or gender, and after geocoding PIN codes to the GHS-SMOD settlement grid ([European Commission Joint Research Centre, 2023](#)), the analytic sample is $N_1 = 1,065$. Online panels in India oversample urban, educated, and English-comfortable respondents; Section 4 quantifies this.

3.3 Wave 2: Face-to-Face Survey in Rural Uttar Pradesh

Wave 2 was commissioned from a commercial survey firm and fielded in February 2026. The mode was interviewer-administered face-to-face on tablets, in Hindi, with paid respondents and convenience sampling. Uttar Pradesh is India’s most populous state (roughly 240 million

residents), Hindi-speaking, and mostly rural; Wave 2 was designed to offset the urban, educated, English-speaking skew of Wave 1, not to represent rural India as a whole. With the same quality filters as Wave 1, $N_2 = 730$. The instrument is identical to Wave 1’s. Translation was managed by the survey firm and is not independently verified.

3.4 Pooling and Weighting

The pooled sample is $N = 1,795$. We construct weights by iterative proportional fitting (raking) to population targets. Age targets (18–24: 18.1%; 25–34: 24.1%; 35–44: 20.7%; 45–54: 15.8%; 55–64: 21.3%) and gender targets (male 51.6%; female 48.4%) come from Indian Census projections.

We report two weighting schemes throughout. The first rakes on age and gender only, matching the specification in [Bick et al. \(2024\)](#) and enabling direct comparison with the U.S. The second adds an urbanicity dimension with targets {urban 36.4%, semi-urban 10.0%, rural 53.6%}; the urban target matches the 2024 World Urbanization Prospects estimate for India ([World Bank, 2024](#)). Urbanicity is assigned by geocoding PIN codes to GHS-SMOD; respondents whose PIN could not be classified are grouped with rural. The resulting weights have mean 0.99 and standard deviation 1.51, with values trimmed at the 1st and 99th percentiles. The urbanicity-adjusted weights are our preferred specification; headline numbers are robust to the exact urban target ([Appendix B](#)).

3.5 Measurement

We follow [Bick et al. \(2024\)](#)’s operational definitions. Work use (QID7) is asked of employed respondents; outside-work use (QID18) is asked of all respondents. “Work use among all adults” treats non-employed respondents as not using AI at work; “work use among employed” restricts to employed respondents. Time savings are computed as the weighted ratio of reported counterfactual hours (QID11; “more than 4” coded as 5) to total work hours last week (QID5), averaged over all employed respondents with non-users contributing zero.

3.6 Limitations

The pooled sample is not a probability sample of India. Wave 1 is a self-selected online panel; Wave 2 is a convenience sample from a single state. The weighting adjusts for age, gender, and urbanicity, but not for education, income, or digital access, so residual selection on these dimensions likely biases the headline adoption numbers upward relative to the true population rate. Wave 2 does not offset gender-based selection into either recruitment channel; we return to this in Section 6. The Wave 2 Hindi translation was managed by the survey firm and not independently verified, so translation drift and interviewer effects are possible. The sample size ($N=1,795$) is much smaller than the U.S. RPS ($N\approx 10,000$), limiting subgroup precision. This is a single cross-section; monthly follow-ups are planned.

4 Selection Diagnostics: What Each Wave Sees

Table 1 compares Wave 1 and Wave 2 on observable demographics, unweighted, to describe the raw selection pattern of each mode.

Table 1: Observable differences between Wave 1 (Clickworker online, all-India, English self-complete) and Wave 2 (UP rural face-to-face, Hindi, interviewer-administered). Percentages are unweighted shares.

Characteristic	Wave 1: Clickworker	Wave 2: UP rural F2F
Female	40.2%	38.5%
Age 18–24	42.7%	28.6%
Age 25–34	37.9%	36.3%
Age 35+	19.3%	35.1%
Urban (GHS-SMOD)	67.3%	46.8%
Rural (GHS-SMOD)	13.3%	44.2%
Graduate or higher	91.4%	61.9%
Income \leq 5 Lakh	51.9%	84.4%
Employed (at work)	57.1%	42.5%
N (unweighted)	1065	730

Three patterns stand out. Wave 1 is younger (42.7% aged 18–24 vs. 28.6% in Wave 2) and more educated (91.4% graduate or higher vs. 61.9%). Wave 1 is also more urban (67.3% vs. 46.8%) and higher-income (only 51.9% below INR 5 lakh vs. 84.4%). On gender, both waves under-represent women (40.2% and 38.5%) relative to the population target of 48.4%—neither mode corrects the other.

The face-to-face wave therefore offsets the online panel’s urban, education, and income skew but leaves the gender skew roughly intact. This is consistent with the well-documented mobile gender gap in India (GSMA, 2024) and with survey-methodology evidence that female respondents are under-represented in both online panels and door-to-door samples, for partly overlapping reasons: device access, availability during interview hours, and household gatekeeping. The substantive numbers in the rest of the paper should be read against this table.

5 Headline Adoption

Figure 1 shows headline adoption rates under both weighting schemes. Under the age \times gender \times urbanicity weighting, 63.3% of Indian working-age adults report using generative AI outside of work, 34.9% report any work-related use, and 67.1% report any use. Under the age \times gender-only weighting that matches [Bick et al. \(2024\)](#), the rates are 66.2%, 38.6%, and 69.4%. Adding urbanicity pulls each number down by 2–4 percentage points, since rural respondents adopt at lower rates.

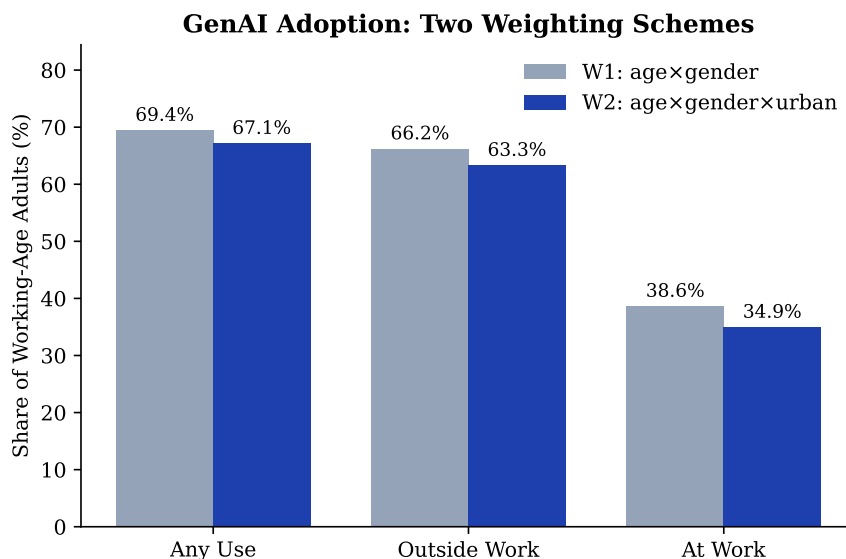


Figure 1: Generative AI adoption rates under two weighting schemes. The first takes on age and gender only (matching [Bick et al. 2024](#)); the second adds urbanicity with urban target 36.4% ([World Bank, 2024](#)). $N = 1,795$.

Table 2 places these numbers alongside the U.S. estimates of [Bick et al. \(2024\)](#). On an apples-to-apples basis (the age \times gender-only column), India’s outside-work rate is 28 percentage points above the U.S.; adding urbanicity narrows the gap only slightly. Our sample is not a probability sample, and the gap should not be read as a direct population-level claim. A more defensible reading is that, among the Indians actually reachable via an online panel or a rural-UP survey firm—each with its own self-selection—generative AI use is strikingly high.

Table 2: Generative AI adoption: India (this paper) vs. United States (Bick et al., 2024). The first India column is weighted on age and gender only; the second adds urbanicity.

Metric	India (age×gender)	India (+urban)	US, Aug 2024
Any generative AI use (%)	69.4	67.1	45.5
Outside-work use (%)	66.2	63.3	37.7
Work use, all adults (%)	38.6	34.9	32.1
Time savings (% of all work hours)	3.1	2.8	1.4

Figure 2 shows the frequency of use. Among all working-age adults, 12.7% use generative AI every day outside work, 29.9% on more than one day last week, and 20.7% on one day. At work the corresponding rates are 6.5%, 13.4%, and 9.6%, with a further 3.8% reporting work use but not in the reference week; a small residual gap between these shares and the 34.9% work-use headline reflects item non-response on the follow-up frequency question. Use is concentrated among adopters—most users use it frequently rather than occasionally, mirroring the U.S. pattern reported by Bick et al. (2024).




Figure 2: Frequency of generative-AI use last week, as a share of all working-age adults.

6 Demographic Heterogeneity

6.1 Education and Urbanicity

Education and urbanicity are the strongest predictors of adoption in the data. Graduates use generative AI outside work at 77.0%, compared with 32.1% among non-graduates—a 45-percentage-point gap (Figure 3). Urban respondents adopt at 75.2%, semi-urban at 61.0%, and rural at 54.0% (Figure 4), a 21-point urban–rural spread. Both effects remain large and statistically significant when included in the regression together (Table 3).




Figure 3: Outside-work generative-AI adoption by education (W2 weighting).




Figure 4: Outside-work generative-AI adoption by urban/rural classification (GHS-SMOD; W2 weighting).

6.2 Age

Figure 5 shows outside-work adoption is high and essentially flat for respondents aged 18–24 (80.8%) and 25–34 (82.5%), then drops to 49.5% among those 35 and older. Work-related adoption peaks in the 25–34 group (40.0%), consistent with higher employment rates in that bracket.




Figure 5: Generative-AI adoption by age group.

6.3 Income

The income gap is modest: 66.0% among respondents above INR 5 lakh vs. 61.8% below (Figure 6). Once education and urbanicity are controlled for, income is not statistically distinguishable from zero (Table 3). Since most generative-AI tools have free tiers, adoption turns more on device access, digital literacy, and task fit—dimensions captured by education and urbanicity—than on willingness to pay.




Figure 6: Outside-work generative-AI adoption by household income.

6.4 Gender

The pooled gender pattern is striking. Women in our sample report substantially higher outside-work adoption (78.3%) than men (49.5%), a 29-percentage-point gap (left panel of Figure 7). Taken at face value, this would reverse the mobile-internet gender gap documented for India (GSMA, 2024). We do not think that reading is correct.

The right panel of Figure 7 shows outside-work adoption by gender within each wave. In the online Clickworker panel, men and women are essentially at parity (85.4% male vs. 84.8% female); the large female premium only appears in the rural Uttar Pradesh face-to-face wave (52.8% male vs. 77.9% female, a 25-point gap). The pooled female premium is

therefore almost entirely a Wave-2 phenomenon. The most plausible explanation is positive selection into the rural interviewer-administered survey. Daytime face-to-face interviewer visits systematically reach women who are at home, available, and willing to be interviewed; in rural Uttar Pradesh, that subset is likely more digitally engaged than rural women as a whole, in a way that “men reachable by the same interviewer” are not more digitally engaged than rural men overall.

GenAI Adoption by Gender: Pooled vs. Wave-split

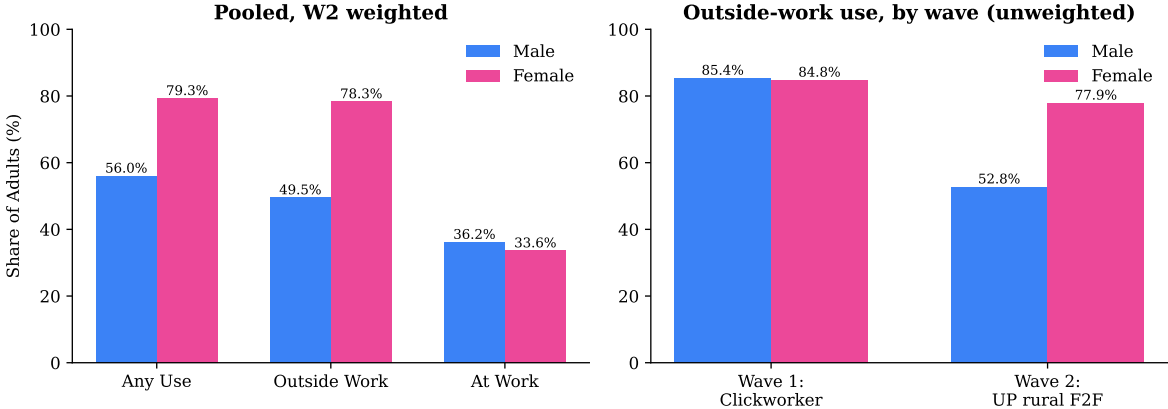


Figure 7: Generative-AI adoption by gender. Left: pooled weighted rates. Right: unweighted outside-work adoption by gender, split by wave.

The regression (Table 3) is consistent with this reading. The main effect of female is large and positive (+0.233, $p < 0.01$), and the female×urban interaction is negative (−0.158, $p = 0.011$). Conditional on other demographics, the female premium is concentrated in rural respondents—exactly where Wave 2 was fielded and where selection into a face-to-face survey is tightest. We therefore report the female premium as a feature of our sample, driven by who is reachable for a daytime rural interviewer visit, and not as evidence that Indian women in the population have overtaken Indian men in generative-AI use. Testing the population-level claim requires a probability-sample wave in which women are enumerated from a sampling frame rather than from willingness to participate.

Table 3: Linear probability model of outside-work generative-AI use on the pooled sample, weighted to national age, gender, and urbanicity targets. Heteroskedasticity-robust standard errors (HC1). Reference categories: male, 18–24, below graduate, income below 5 lakh, rural.

Variable	Coefficient	Std. Error
Constant	0.347***	(0.041)
Female	0.233***	(0.053)
Age 25–34	0.010	(0.029)
Age 35+	-0.177***	(0.037)
Graduate/Postgraduate	0.343***	(0.042)
Income > 5 Lakh	0.018	(0.037)
Urban	0.176***	(0.047)
Semi-urban	-0.018	(0.052)
Female × Urban	-0.158**	(0.062)
N	1795	
R^2	0.283	

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

7 Products, Time Savings, Non-adoption, and Work Context

7.1 Products

Figure 8 shows which products respondents use. ChatGPT dominates both contexts, consistent with its global brand recognition. Gemini is a clear second, reflecting Android’s near-universal market share in India. DeepSeek has a non-trivial presence, particularly at work. GitHub Copilot and Midjourney appear as specialist tools.

GenAI Products Used (W2 weighted)

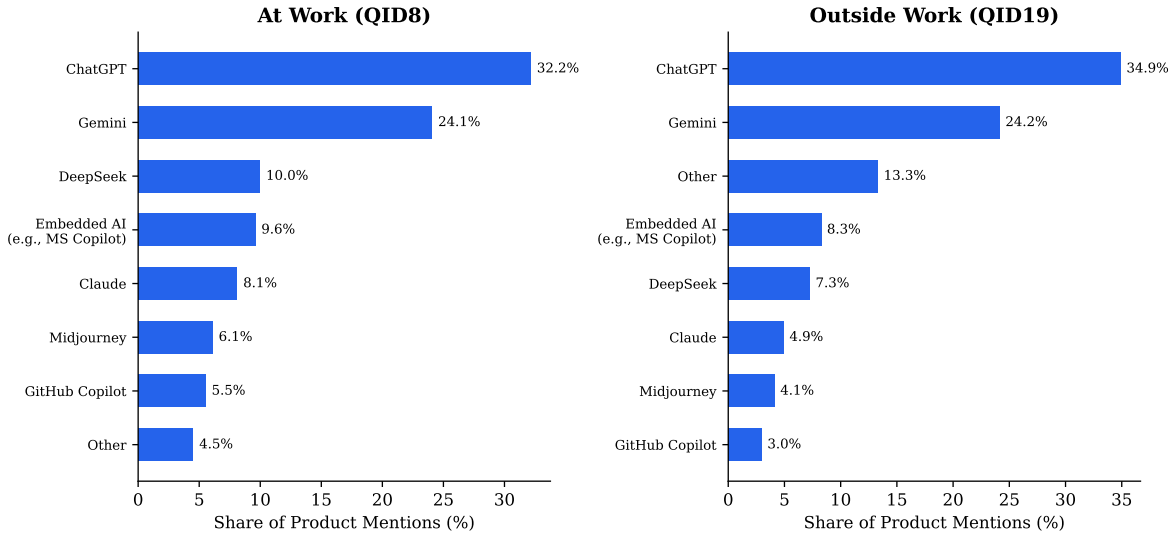


Figure 8: Generative-AI products used at work (left) and outside work (right). Respondents could select multiple products.

7.2 Time Savings

Following [Bick et al. \(2024\)](#), we compute time savings as the weighted sum of reported counterfactual hours saved (QID11) divided by the weighted sum of total work hours last week (QID5), across all employed respondents, with non-users contributing zero to the numerator but their hours to the denominator. Under this definition, generative AI saves 2.8% of total work hours. The modal reported saving among work users is 1–2 hours per week, with a non-trivial tail above 4 hours (Figure 9).




Figure 9: Distribution of reported time savings among generative-AI work users.

Our 2.8% aggregate is roughly double the 1.4% U.S. estimate in [Bick et al. \(2024\)](#), reflecting both higher adoption and higher per-user savings in our sample. Self-reported counterfactual-hours questions are noisy, so the number should be read as indicative. Even a large discount would leave India’s figure above the U.S. baseline.

7.3 Reasons for Non-Adoption

Figure 10 shows the reasons non-users give for not using generative AI. At work, the leading reasons are privacy or security concerns (32%), the view that AI cannot help with the respondent’s job (26%), not knowing how to use AI effectively (13%), and employer restrictions (13%). Outside work, the most common reason is the view that AI cannot help with the respondent’s tasks (28%), followed by privacy concerns (25%) and uncertainty about effective use (23%). Two findings stand out. First, privacy and security are a first-order concern in both contexts—more so than in U.S. surveys—and likely reflect both data-protection worries and a still-uneven regulatory environment. Second, uncertainty about how to use AI effectively is a substantial barrier in both contexts, suggesting that training and awareness programs could meaningfully expand adoption among non-users who do not face structural constraints.

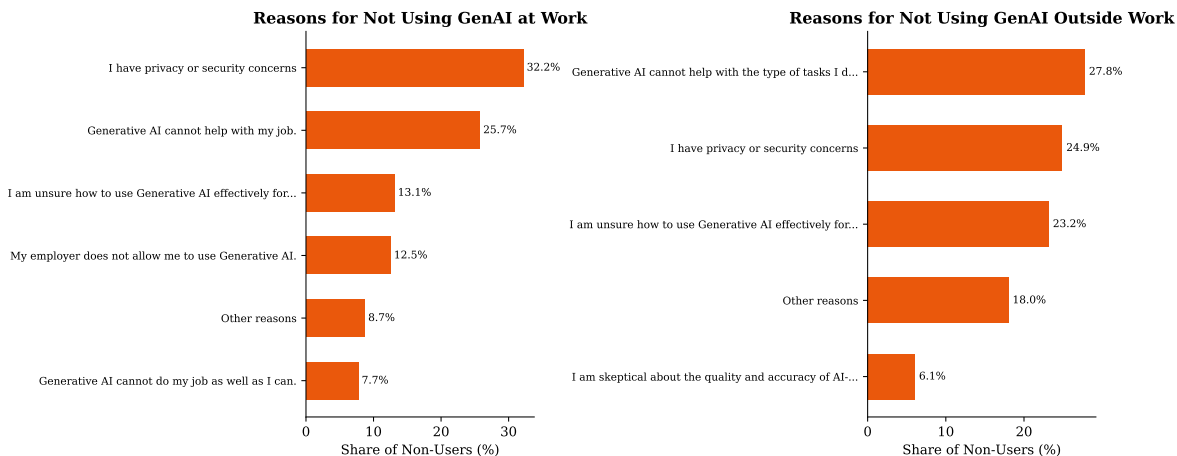


Figure 10: Reasons for not using generative AI at work (left) and outside work (right), among non-users.

Appendix C breaks these reasons down by education and urbanicity. The informational barrier is more common among non-graduates and rural respondents than among graduates and urban respondents, suggesting that part of the education and urban–rural adoption gap reflects awareness rather than underlying need.

7.4 Employer Type

Table 4 reports work-AI adoption among employed respondents by employer type. Private-sector, government, and self-employed respondents adopt at 61–64%. Non-profit employees adopt at 34%, but the cell is small ($N = 32$) and this should be read as suggestive.

Table 4: Work generative-AI adoption by employer type. W2-weighted among employed respondents.

Employer type	N	Work adoption (%)
Work in a business owned by someone else in this household	74	63.7
Government (including state or local governments, a public school, university or hospital)	83	63.2
Private-sector, for-profit company	598	61.6
Self-employed	233	61.0
Non-profit organization (including charitable organizations)	32	34.2

8 Discussion

8.1 India vs. the United States

Indian adoption rates exceed the U.S. rates of [Bick et al. \(2024\)](#) on every comparable measure. On the age-and-gender-only comparison, the outside-work gap is 28 percentage points; with urbanicity added, it is 26 points. Two factors plausibly contribute. Our sample is not a probability sample, and residual selection on education, digital literacy, and willingness to participate likely pushes our estimates up. At the same time, structural features of India—a large English-literate and IT-trained population, widespread cheap mobile data, a young median age, and a population already comfortable with text-based digital services (UPI, WhatsApp, government portals)—plausibly lower the cost of adopting a text-based AI chatbot. Both channels are likely at play; separating them is a task for probability-sample follow-ups.

8.2 The Digital Divide in AI Use

The data show four gradients in generative-AI adoption. Education is the largest: graduates are 45 percentage points more likely to use AI than non-graduates, likely reflecting both digital literacy and task fit. Urbanicity is next: urban respondents are 21 points ahead of rural respondents, consistent with broader patterns in digital infrastructure. Age is sharp above 35, with a 30+ point gap between the under-35 and 35+ groups. Income, by contrast, is weak once education and urbanicity are controlled for, consistent with the low marginal monetary cost of free-tier AI tools. Gender does not appear on this list as a substantive gradient: the female premium in our sample looks like selection into both recruitment modes rather than a population pattern.

8.3 Policy Implications

Three implications follow for Indian policy. First, privacy and security concerns and uncertainty about effective use together account for roughly half of the reasons non-users give for not using generative AI. Both are plausibly addressable—the first through clearer data-protection regulation and plainer-language disclosures by AI providers, the second through training and awareness programs as extensions of existing Digital India initiatives. Informational barriers are consistently more prominent among non-graduates and rural respondents (Appendix C). Second, the 21-point urban–rural gap combined with India’s still-majority-rural population means that most Indians likely have not yet benefited from generative AI; vernacular-language and low-bandwidth tools are the natural bridge. Third, 61% of employed respondents report work-related use, and the 2.8% time-savings estimate, even heavily discounted for self-report noise, points to a productivity contribution that firms and policymakers should plan around.

8.4 What Indians Use ChatGPT For

This paper focuses on *whether and how often* Indians use generative AI. The complementary question of *what they use it for* is taken up by [Choudhury and Garimella \(2026\)](#), who analyze ChatGPT conversation exports from about 550 Indian users—drawn from the same Clickworker panel as Wave 1 of this survey—using the classification approach of [Chatterji et al. \(2025\)](#). They find that practical guidance and information-seeking dominate Indian ChatGPT use, with less writing assistance than in global benchmarks. This is broadly consistent with our finding that outside-work use substantially exceeds work use.

8.5 Future Work

The natural next step is a probability-sample benchmark wave. Combining the current online-plus-face-to-face infrastructure with a periodic full-probability-sample calibration wave—

along the lines of the World Bank LSMS (World Bank, 2020)—would let us estimate and correct the residual mode-specific biases documented here, particularly the gender selection. We also plan monthly follow-ups to track adoption dynamics, and parallel fieldings in other Global South countries (Pakistan, Nigeria, Brazil) to enable cross-country comparison.

9 Conclusion

This paper provides the first survey-based estimates of generative-AI adoption in India. Combining an online panel with a face-to-face survey in rural Uttar Pradesh, and reweighting to national age, gender, and urbanicity targets, 63% of Indian working-age adults use generative AI outside of work, 35% use it at work, and it saves them 2.8% of total work hours. These rates exceed the U.S. benchmarks of Bick et al. (2024) on every dimension. Adoption is strongly gradient in education, urbanicity, and age; income matters little once the other three are controlled for. Women in our sample report higher use than men, but we interpret this as selection into both recruitment modes rather than a population-level reversal of the Indian mobile-internet gender gap.

Our estimates are not a substitute for a probability-sample estimate, and the paper is explicit about what the current design does and does not correct. The natural next step is a probability-sample calibration wave to pin down the residual biases, particularly in gender, and monthly follow-ups to track adoption as it evolves.

References

- Acemoglu, D. (2024). The simple macroeconomics of AI. *NBER Working Paper*, (32487).
- Aker, J. C. and Mbiti, I. M. (2010). Mobile phones and economic development in africa. *Journal of Economic Perspectives*, 24(3):207–232.

- Bick, A. and Blandin, A. (2023). Real-time labor market estimates during the 2020 coronavirus outbreak. *Journal of Monetary Economics*.
- Bick, A., Blandin, A., and Deming, D. J. (2024). The rapid adoption of generative AI. *Federal Reserve Bank of St. Louis Working Paper*, (2024-027). Revised November 2024.
- Björkegren, D. (2018). The adoption of network goods: Evidence from the spread of mobile phones in Rwanda. *The Review of Economic Studies*, 86(3):1033–1060.
- Brynjolfsson, E., Li, D., and Raymond, L. R. (2025). Generative AI at work. *Quarterly Journal of Economics*. Forthcoming.
- Chatterji, A., Dahal, A., Deshmukh, A., Green, T., Handa, K., Kon, M., Maheshwari, A., Rieser, V., and Singhal, K. (2025). How people use ChatGPT. *OpenAI Research*.
- Choudhury, P. and Garimella, K. (2026). How people use ChatGPT: Evidence from Indian users. Working paper. https://gvrkiran.github.io/content/How_people_use_ChatGPT.pdf.
- De Leeuw, E. D. (2005). To mix or not to mix data collection modes in surveys. *Journal of Official Statistics*, 21(2):233–255.
- Dell’Acqua, F., McFowland III, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraymer, L., Candelon, F., and Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. *Harvard Business School Technology & Operations Mgt. Unit Working Paper*, (24-013).
- Dillman, D. A., Smyth, J. D., and Christian, L. M. (2014). Internet, phone, mail, and mixed-mode surveys: The tailored design method.
- Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). GPTs are GPTs: An early look at the labor market impact potential of large language models. *Science*, 384(6702).

European Commission Joint Research Centre (2023). GHS-SMOD – GHS settlement model grid.

Fletcher, R. and Nielsen, R. K. (2024). Generative AI and news: Reuters Institute Digital News Report. *Reuters Institute for the Study of Journalism*.

GSMA (2024). The mobile gender gap report 2024. Technical report, GSMA.

Hjort, J. and Poulsen, J. (2019). The arrival of fast internet and employment in Africa. *American Economic Review*, 109(3):1032–1079.

Humlum, A. and Vestergaard, E. (2024). The adoption of ChatGPT. *University of Chicago Working Paper*.

Internet and Mobile Association of India (2024). Internet in India report.

Ministry of Electronics and Information Technology (2024). Digital India programme. <https://www.digitalindia.gov.in/>.

NASSCOM (2024). AI adoption index: India.

Noy, S. and Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654):187–192.

Peng, S., Kalliamvakou, E., Cihon, P., and Demirer, M. (2023). The impact of AI on developer productivity: Evidence from GitHub Copilot. *arXiv preprint arXiv:2302.06590*.

Pew Research Center (2018). Methods: Survey research in emerging and developing nations. Technical report, Pew Research Center. <https://www.pewresearch.org/methods/>.

Pew Research Center (2024). AI in daily life. *Pew Research Center Report*.

Rogers, E. M. (2003). Diffusion of innovations.

Telecom Regulatory Authority of India (2024). Indian telecom services performance indicators.

Venkatesh, V., Morris, M. G., Davis, G. B., and Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3):425–478.

World Bank (2020). Living standards measurement study: Guidelines on survey design. Technical report, The World Bank.

World Bank (2024). Urban population (% of total population) — India. Technical report, World Bank Open Data (World Urbanization Prospects). <https://data.worldbank.org/indicator/SP.URB.TOTL.IN.ZS?locations=IN>.

A Survey Instrument

Table 5 lists the key questions from the survey instrument, which is otherwise identical to the generative-AI module of the U.S. Real-Time Population Survey (Bick and Blandin, 2023; Bick et al., 2024).

Table 5: Key Survey Questions

ID	Question Text
Q27	What is your gender?
Q28	What is your age?
Q29	What is your highest level of education?
Q30	What is your estimated total annual household income?
Q34	What is your 6-digit postal code (PIN code)?
QID3	What is your current employment status?
QID4	What type of employer do you work for?
QID5	How many hours did you work last week?
QID7	Do you use Generative AI as part of your job?
QID8	Which Generative AI products do you use for your job?
QID9	Did you use Generative AI for your job last week?
QID11	Without Generative AI, how many additional hours would you have needed?
QID16, QID17	Reasons for not using Generative AI at work.
QID18	Do you use Generative AI outside of your job?
QID19	Which Generative AI products do you use outside of your job?
QID20	Did you use Generative AI outside of your job last week?
QID25, QID26	Reasons for not using Generative AI outside of work.

B Urbanicity-Target Sensitivity

Table 6 shows headline adoption under three urbanicity targets: 33%, 36.4% (baseline, from World Bank 2024 (World Bank, 2024)), and 37%. Across all three scenarios we hold the semi-urban target fixed at 10% and set the rural target to $1 - \text{urban} - 0.10$, so that the three shares sum to 100%. The span of plausible targets moves the headline numbers by at most 0.6 percentage points.

Table 6: Headline adoption under alternative urbanicity targets (all W2 raking, $\text{age} \times \text{gender} \times \text{urbanicity}$).

Urban target	Outside-work (%)	Work, all (%)	Any use (%)
0.330	62.8	34.3	66.8
0.364	63.3	34.9	67.1
0.370	63.4	35.1	67.2

C Non-adoption Reasons by Demographics

The CSV file `tables/nonadoption_by_demo.csv` reports the W2-weighted share of non-users citing each reason, broken out by context (work / outside work), demographic dimension (education: graduate vs. below graduate; area: urban vs. rural), and reason. We summarize the qualitative patterns in Section 7 of the main text; the full table is omitted here for space and available in the replication archive.