

Institute for Economic Studies, Keio University

Keio-IES Discussion Paper Series

**The Early Spatial Diffusion of Generative AI in Japan:
Retrospective Adoption Cohorts, Composition versus Place,
and the Role of Aging and Labor Shortages**

Hiroyuki Yamada, Atsushi Nakagomi, Takahiro Tabuchi

15 September, 2026

DP2026-020

<https://ies.keio.ac.jp/en/publications/28032/>

Keio University



Institute for Economic Studies, Keio University
2-15-45 Mita, Minato-ku, Tokyo 108-8345, Japan
ies-office-group@keio.jp
15 September, 2026

The Early Spatial Diffusion of Generative AI in Japan: Retrospective Adoption Cohorts, Composition versus Place, and the Role of Aging and Labor Shortages

Hiroyuki Yamada, Atsushi Nakagomi, Takahiro Tabuchi

Keio-IES DP2026-020

15 September, 2026

JEL Classification: O33, R12, J11, J23, L86

Keywords: Generative AI, technology diffusion, regional inequality, population aging, labor shortage, Japan

Abstract

Japan has the lowest rate of generative AI use among major economies, yet little is known about how the technology has spread within the country. This study uses the sixth wave of the Japan COVID-19 and Society Internet Survey (JACSIS, N = 27,630), which asks respondents when they first started using generative AI, to reconstruct cumulative adoption curves for Japanese municipalities from November 2022 to January 2026. Because every respondent became able to adopt at the same moment, the release of ChatGPT, the retrospective cohorts support a discrete-time hazard analysis without left truncation. Linking postal codes to municipal statistics, we document four findings. First, adoption follows the urban hierarchy: by January 2026, 50% of residents of the quintile of municipalities with the highest population density had adopted, compared with 35% in the quintile with the lowest population density, and the ratio between the two rose from 0.55 to 0.71 over the four periods, indicating relative convergence alongside a widening absolute gap. Second, most of the geographic gradient reflects who lives where rather than where they live: individual characteristics explain 64% of the population-density coefficient, and the coefficients on population density and on distance to Tokyo become statistically indistinguishable from zero once individual controls are included. Third, municipal population aging is negatively associated with adoption beyond its compositional effect, but only in the first two years of diffusion. Fourth, the job openings-to-applicants ratio in the respondent's workplace area is positively associated with work-related adoption and unrelated to purely private adoption, a pattern consistent with labor scarcity pulling generative AI into workplaces. The association is statistically significant but economically small: a one-point increase in the ratio is associated with a work-adoption probability less than one percentage point higher. Regional differences are concentrated in whether people adopt, not in how intensively adopters use the technology or whether they stop using it: one in six ever-users has lapsed, and lapsing follows age, education, and job content rather than place. The results are correlational, but they suggest that Japan's low national adoption rate is a consequence of its demographic and occupational composition rather than of a distinct rural disadvantage.

Hiroyuki Yamada

Keio University

hyamada@keio.jp

Atsushi Nakagomi

Chiba University

anakagomi0211@chiba-u.jp

Takahiro Tabuchi

Tohoku University

takahiro.tabuchi.a5@tohoku.ac.jp

The Early Spatial Diffusion of Generative AI in Japan: Retrospective Adoption Cohorts, Composition versus Place, and the Role of Aging and Labor Shortages

Hiroyuki Yamada^a

Atsushi Nakagomi^b

Takahiro Tabuchi^c

* ^a (Corresponding author) Faculty of Economics/Graduate School of Economics

Keio University

2–15–45 Mita, Minato-ku, Tokyo 108–8345, Japan

E-mail: hyamada@keio.jp

<https://orcid.org/0000-0003-3709-1914>

^bDepartment of Social Preventive Medical Sciences, Center for Preventive Medical Sciences

Chiba University

anakagomi0211@chiba-u.jp

^c Division of Epidemiology, School of Public Health, Graduate School of Medicine

Tohoku University

takahiro.tabuchi.a5@tohoku.ac.jp

Abstract

Japan has the lowest rate of generative AI use among major economies, yet little is known about how the technology has spread within the country. This study uses the sixth wave of the Japan COVID-19 and Society Internet Survey (JACSIS, $N = 27,630$), which asks respondents when they first started using generative AI, to reconstruct cumulative adoption curves for Japanese municipalities from November 2022 to January 2026. Because every respondent became able to adopt at the same moment, the release of ChatGPT, the retrospective cohorts support a discrete-time hazard analysis without left truncation. Linking postal codes to municipal statistics, we document four findings. First, adoption follows the urban hierarchy: by January 2026, 50% of residents of the quintile of municipalities with the highest population density had adopted, compared with 35% in the quintile with the lowest population density, and the ratio between the two rose from 0.55 to 0.71 over the four periods, indicating relative convergence alongside a widening absolute gap. Second, most of the geographic gradient reflects who lives where rather than where they live: individual characteristics explain 64% of the population-density coefficient, and the coefficients on population density and on distance to Tokyo become statistically indistinguishable from zero once individual controls are included. Third, municipal population aging is negatively associated with adoption beyond its compositional effect, but only in the first two years of diffusion. Fourth, the job openings-to-applicants ratio in the respondent's workplace area is positively associated with work-related adoption and unrelated to purely private adoption, a pattern consistent with labor scarcity pulling generative AI into workplaces. The association is statistically significant but economically small: a one-point increase in the ratio is associated with a work-adoption probability less than one percentage point higher. Regional differences are concentrated in whether people adopt, not in how intensively adopters use the technology or whether they stop using it: one in six ever-users has lapsed, and lapsing follows age, education, and job content rather than place. The results are correlational, but they suggest that Japan's low national adoption rate is a consequence of its demographic and occupational composition rather than of a distinct rural disadvantage.

Keywords: Generative AI, technology diffusion, regional inequality, population aging, labor shortage, Japan

JEL Classification Codes: O33, R12, J11, J23, L86

1. Introduction

Three years after the public release of ChatGPT in November 2022, the international dispersion in the use of generative AI is striking. According to the Ministry of Internal Affairs and Communications (MIC, 2026), 52.2% of Japanese adults aged 20 or older had experience with generative AI services in the fiscal 2025 survey, fielded in January and February 2026, compared with 75.6% in the United States, 75.6% in Germany, and 93.6% in China. The Japanese figure a year earlier was 26.7% (MIC, 2025), so use has grown rapidly without closing the gap. Usage data from a major provider point in the same direction: Japan's score on Anthropic's AI Usage Index, the ratio of a country's share of Claude conversations to its share of the world's working-age population, is 1.86, about half of the United States' 3.62 (Anthropic, 2025). Existing explanations of the international gap emphasize national income, skills, and language. However, a country average conceals the geography of adoption within the country, and for Japan, whose regions differ sharply in population density, age structure, and labor market tightness, the internal geography is of direct policy relevance. Regional revitalization programs, most recently the Digital Garden City Nation Initiative (Cabinet Secretariat, 2023), rest on the premise that digital technologies can offset the disadvantages of depopulating and aging regions. Whether generative AI is spreading to such regions, and whether local labor shortages are pulling it into workplaces, are therefore first-order empirical questions.

The economics of technology diffusion offers a clear framework for these questions. Since Griliches (1957), the diffusion of a new technology has been studied as a process in which adoption starts where it is most profitable and spreads with a lag to other locations. Comin and Hobijn (2010) and Comin and Mestieri (2018) distinguish the extensive margin (when a technology arrives) from the intensive margin (how deeply it is used once it has arrived), and show that for recent technologies adoption lags have converged across countries while intensity has diverged. Within countries, Forman, Goldfarb, and Greenstein (2005) ask whether the commercial internet was led by large cities or spread uniformly, and Goolsbee and Klenow (2002) use a single cross-section that records the year of purchase to estimate local spillovers in the diffusion of home computers. A second literature studies the relationship between labor scarcity, demographic change, and technology adoption. Acemoglu (2010) shows formally when scarce labor encourages the adoption of labor-saving technology, Hornbeck and Naidu (2014) and Clemens, Lewis, and Postel (2018) document mechanization after historical labor supply shocks, and Acemoglu and Restrepo (2022) show that the scarcity of middle-aged workers induces the adoption of industrial robots. Whether these forces operate for generative AI, a technology that requires complementary cognitive skills rather than replacing manual labor, is an open question.

Evidence on who adopts generative AI is accumulating quickly. Bick, Blandin, and Deming (2025) document steep gradients by education, income, and occupation in the United States. Humlum and Vestergaard (2024) show unequal adoption within exposed occupations in Denmark, and Chatterji et al. (2025) show that use rises with national income. For Japan, Nakagomi, Abe, and Tabuchi (2025) find that generative AI use is higher among younger, better-educated respondents living in urban and less deprived areas. These studies are cross-sectional descriptions of who uses the technology at one point in time. None uses information on when individuals adopted, none decomposes the urban–rural gap into the part attributable to the

composition of residents and the part attributable to place, and none relates adoption to local aging or labor market tightness. This study fills these gaps.

We use the sixth wave of the Japan COVID-19 and Society Internet Survey (JACSIS), a nationwide internet survey administered between December 2025 and February 2026 to 27,630 respondents. The wave contains a question on when the respondent first started using generative AI, with four retrospective cohorts (November 2022 to December 2023, 2024, the first half of 2025, and July 2025 onward), together with home and workplace postal codes. Two features make these data unusually well suited to a diffusion study. First, because generative AI became available to everyone at the same moment, every respondent entered the risk set on the same date and there is no left truncation, a problem that has complicated diffusion studies of most other technologies. Second, the cohort boundaries coincide with well-known product events (the release of GPT-4, the renaming of Bing Chat and Bard, and the 2025 wave of new models), which limits recall error. We convert postal codes into municipality codes and link them to the 2020 Population Census, the annual Basic Resident Register (2022 to 2026), the 2024 Economic Census, urban employment area definitions, and office-level job openings-to-applicants ratios that we collected from the publications of 45 prefectural labor bureaus.

Our empirical analysis proceeds in four steps. We first reconstruct cumulative adoption curves by municipality type and test whether the gap between high- and low-population density areas has narrowed or widened. We then estimate discrete-time hazard models of adoption with period-specific coefficients on population density and distance to Tokyo, and decompose the gradient into composition and place using Oaxaca–Blinder and Gelbach (2016) decompositions. Third, we add municipal aging and labor market tightness, the latter measured at the workplace rather than the residence. Finally, we examine spatial autocorrelation with global and local Moran statistics and spatial lag and error models, and we compare regional gradients on the extensive margin (adoption) with those on the intensive margin (frequency of use, work use, and paid subscriptions).

This study makes three contributions. First, it provides the first reconstruction of the within-country spatial diffusion of generative AI from individual retrospective adoption dates, and shows that the diffusion of generative AI in Japan has been relatively convergent across the urban hierarchy while the absolute gap has widened. Second, it shows that the bulk of the urban–rural gap in adoption is compositional: after controlling for age, education, employment, remote work, and job content, the coefficients on population density and distance are close to zero, and the Gelbach decomposition attributes 64% of the raw population-density gradient to observable individual characteristics. Third, it brings worker-level evidence to bear on two hypotheses drawn from the literature on labor scarcity, demographic change, and technology adoption (Acemoglu, 2010; Hornbeck and Naidu, 2014; Clemens, Lewis, and Postel, 2018; Acemoglu and Restrepo, 2022): that population aging slows the arrival of a new technology, and that scarce labor induces employers to adopt technologies that economize on it. We find that local aging slows adoption beyond its compositional effect only in the earliest phase of diffusion, and that local labor shortages are associated with work-related but not with private adoption, although the association is small in magnitude. Beyond these three contributions, we extend the analysis from the extensive margin to the retention margin of diffusion: one in six respondents who ever used generative AI has stopped using it, and we show that this lapsing, like the intensity of use, carries

no regional disadvantage and is determined by age, education, and the tasks of the job. All of our estimates are correlational, and we are careful throughout to describe them as associations.

The rest of this paper is organized as follows. Section 2 describes the data, the postal-code linkage, and the construction of adoption cohorts. Section 3 presents the empirical framework. Section 4 reports the results, and Section 5 reports the checks on recall, mobility, lapsed users, and representativeness that we regard as mandatory for retrospective adoption data. Section 6 discusses the findings and Section 7 concludes.

2. Data

2.1. The JACSIS survey and the analysis sample

We utilize data from the sixth wave of the Japan COVID-19 and Society Internet Survey (JACSIS). The JACSIS is a population-based online questionnaire survey that recruits participants from a panel maintained by Rakuten Insight, Inc., which comprises approximately 2.2 million individuals aged 15 or above. The sixth wave was administered between December 2025 and February 2026. A follow-up survey targeted 24,626 individuals who had participated in prior waves and consented to further contact, of whom 16,518 provided valid responses (a response rate of 67.1%), and a supplementary survey using the same questionnaire targeted new panelists aged 16–79 and yielded 13,828 respondents. Following the exclusion rules applied to the same data set by Yamada et al. (2026), we drop respondents who completed the questionnaire in less than ten minutes, who failed the attention check, who reported an implausible pattern of chronic diseases, or who reported a household of more than fifteen members. The analysis sample consists of 27,630 respondents. Inverse probability weights that align the sample with the 2022 Comprehensive Survey of Living Conditions on sex, age, region, education, and marital status are provided with the data. We report weighted estimates as our main results and unweighted estimates alongside them.

The sixth wave is the first JACSIS wave to include a module on generative AI. Unlike most of the module, which was administered only to current users, the question on when the respondent started using generative AI was asked of every respondent, and it is the basis of this study.

2.2. Adoption cohorts and the person-period file

The adoption question reads, in translation, "When did you start using generative AI (for example, ChatGPT, Gemini, Claude, or Microsoft Copilot)? Trial use counts." The six answer categories are "never used", "used in the past but not currently", and four categories for respondents who still use generative AI and started in one of four periods: November 2022 to December 2023 (described in the question as the release of ChatGPT, GPT-4, Bard, Claude, and Bing Chat), 2024 (the renaming to Copilot and Gemini and the spread of general use), January to June 2025, and July 2025 or later. In the unweighted sample, 51.1% had never used generative AI, 7.7% had lapsed, and 7.2%, 10.5%, 12.5%, and 11.0% had started in the four successive periods. With the survey weights, 36.2% were current users and 43.9% had ever used the technology.

We treat the four cohorts as periods of a discrete-time hazard model. Period 1 covers November 2022 to December 2023 (14 months), period 2 covers 2024 (12 months), period 3

covers January to June 2025 (6 months), and period 4 covers July 2025 to the survey date (about 7 months). Each current user contributes one observation per period up to and including the period of adoption, and each never-user contributes four right-censored observations. Lapsed users, whose adoption date is not recorded, are excluded from the main hazard sample and are treated in Section 5 with bounds. The person-period file contains 86,629 observations and 11,385 adoption events. Because the periods differ in length, all hazard models include period dummies and an offset for the logarithm of the number of months in the period.

2.3. Geographic linkage

Every respondent reported a seven-digit home postal code, and every employed respondent reported a workplace postal code. We convert postal codes into five-digit municipality codes using the Japan Post address file, resolving large-establishment postal codes with the separate establishment file. Home postal codes map directly to a municipality for 98.1% of respondents. For the remaining 1.9%, whose postal codes are obsolete or invalid, we assign the modal municipality of the first five or three digits and flag the assignment. Workplace postal codes map to a municipality for 91% of employed respondents. Because Hamamatsu City reorganized its wards in 2024, we treat Hamamatsu as a single unit in all sources. Respondents are located in 1,478 of Japan's 1,910 municipalities (77%), which contain about 96% of the population, and 862 municipalities have at least five respondents.

We attach the following municipal statistics. From the 2020 Population Census we take population, area, population density, and the share of the population living in densely inhabited districts (DIDs). From the Basic Resident Register we take the share of residents aged 65 or older on January 1, 2022, at the start of the diffusion window, and the log change in the population aged 15–64 between January 2022 and January 2026. From the 2024 Economic Census we take the number of private establishments and employees. We classify municipalities by the 2020 urban employment area (UEA) definitions of Kanemoto and Tokuoka (2002), distinguishing center cities, suburbs, and municipalities outside any UEA, and we compute the great-circle distance from each municipality's main polygon centroid to Tokyo Station using the administrative boundary data of the Ministry of Land, Infrastructure, Transport and Tourism. For maps and for the measure of local exposure in Section 3.6 we use secondary medical areas as the aggregation unit. A secondary medical area is a group of municipalities that a prefecture defines under the Medical Care Act as the unit for planning general hospital care, so that it approximates the area within which residents travel for everyday services. The 2024 list has 335 such areas, with a median of five municipalities each, and 333 of them contain respondents to the survey.

Labor market tightness is measured by the job openings-to-applicants ratio at the level of the public employment security office (Hello Work) that has jurisdiction over the respondent's workplace municipality. Office-level ratios are not published centrally. We collected them from the monthly and annual reports of 45 prefectural labor bureaus and averaged the fiscal-year values for 2023 to 2025 across roughly 500 offices. Two prefectures, Gunma and Yamanashi, do not publish office-level ratios. For these we use the prefectural ratio and flag the substitution. Municipalities are assigned to offices using the jurisdiction lists of the Ministry of Health, Labour and Welfare.

2.4. Control variables

Individual controls are age and its square, sex, education (junior high or high school, vocational or two-year college, university or above, and other), employment status (regular employee or executive, self-employed or freelance, non-regular employee, and not employed), firm size, household income (six categories including "prefer not to answer" and "do not know"), household size, frequency of working from home, and whether the job is mainly desk work. Firm size, working from home, and desk work are asked only of the employed, and we code the non-employed as an explicit category so that they remain in the sample. Prefecture fixed effects are included in all individual-level models.

2.5. Descriptive statistics

Table 1 reports weighted means by UEA type. Current use is 37.2% in center cities, 35.8% in suburbs, and 25.4% outside UEAs. The share who started in the first period is 6.5%, 5.2%, and 3.1%, respectively. The share of lapsed users is nearly identical across the three groups (7.6%, 8.0%, and 6.0%), so the lower use outside UEAs reflects fewer trials rather than more abandonment. Work-related use is 23.3%, 22.1%, and 16.7%.

3. Empirical framework

3.1. Cumulative adoption curves and convergence

Let $S_r(t)$ denote the share of residents of area r who had adopted by the end of period t . We compute $S_r(t)$ for quintiles of municipal population density and for UEA types, and we summarize convergence with the ratio $S_1(t)/S_5(t)$ between the quintiles of municipalities with the lowest and the highest population density. A rising ratio indicates relative convergence, whereas a rising difference $S_5(t) - S_1(t)$ indicates absolute divergence. Both can occur at once during the early phase of a logistic diffusion process, and we report both.

3.2. Discrete-time hazard of adoption

Our main specification is a complementary log-log model of the hazard of adoption in period t , the discrete-time counterpart of a proportional hazards model (Jenkins, 1995):

$$\log(-\log(1 - h_{it})) = \alpha_t + \sum_{s=1}^4 \beta_s Z_{r(i)} D_{st} + X_i \gamma + \delta_{p(i)} + \ln m_t, \quad (1)$$

where h_{it} is the probability that respondent i , who has not adopted before period t , adopts in period t . The vector $Z_{r(i)}$ contains municipal characteristics of the residence $r(i)$ (log population density, log distance to Tokyo, the DID share, the share aged 65 or older, and the log change in the working-age population), and D_{st} is an indicator for period s , so that the coefficient β_s on the municipal characteristics is specific to each period. The vector X_i contains the individual controls, $\delta_{p(i)}$ are prefecture fixed effects, and the log of the period length in months, $\ln m_t$, enters as an offset. Standard errors are clustered by municipality. The period-specific coefficients β_s on population density and distance are the objects of interest: a decline in their magnitude across periods indicates that later adopters are drawn increasingly from less densely populated and more remote places. We estimate the model without individual controls, with individual controls, with municipal characteristics, with survey weights, and with UEA type replacing population density.

3.3. Composition versus place

To separate the part of the urban–rural gap that reflects the characteristics of residents from the part that reflects the place itself, we use two decompositions. The Oaxaca–Blinder decomposition of the difference in current use between residents of center cities and residents of municipalities outside UEAs attributes the gap to differences in observed characteristics and to differences in the returns to those characteristics. The Gelbach (2016) decomposition attributes the change in the coefficient on log population density between a specification with no controls and the full specification to each group of controls, in a way that does not depend on the order in which controls are added. We form five groups of controls: age and sex, education, employment status and firm size, income and household size, and working from home and desk work. The last group deserves comment: working from home and desk work may be mechanisms through which place affects adoption rather than pre-determined characteristics, and we report them separately for this reason.

3.4. Aging and labor shortages

We enter the municipal share aged 65 or older with and without the individual age controls. The comparison separates the compositional effect of aging (older people adopt less) from a place effect (people of a given age adopt less in older municipalities). We also interact aging with period dummies. For labor shortages the analysis is guided by a hypothesis, although the design cannot test it formally. We have no exogenous source of variation in labor market tightness, so the most the data can show is whether the pattern of associations is consistent with the hypothesis. The literature on induced innovation holds that scarce labor raises the return to technologies that economize on it. Acemoglu (2010) formalizes the conditions under which labor scarcity encourages the adoption of labor-saving technology, Hornbeck and Naidu (2014) and Clemens, Lewis, and Postel (2018) document mechanization after historical labor supply shocks, and Acemoglu and Restrepo (2022) show that the scarcity of middle-aged workers induces the adoption of industrial robots. Generative AI economizes on labor in tasks such as drafting, summarizing, coding, and answering routine enquiries, which are among the tasks that establishments in a tight labor market find hardest to staff. The hypothesis is therefore that employers facing a tight local labor market introduce the technology into work and encourage their employees to use it. Two implications follow and determine the design. The relevant labor market is that of the workplace rather than of the residence, so we measure tightness by the job openings-to-applicants ratio of the public employment security office with jurisdiction over the workplace municipality. And because the mechanism operates through the employer, it should affect adoption for work and not adoption for private purposes, which provides a consistency check rather than a test. An association of the same size with private adoption would speak against the hypothesis. We estimate probit models of work-related adoption and of purely private adoption on the workplace ratio, with the individual controls and prefecture fixed effects, and we report average marginal effects, that is, the change in the probability of the outcome per unit change in a regressor averaged over the sample, so that the magnitudes can be read directly. The second outcome serves as a placebo in this sense. Neither outcome identifies a causal effect of labor scarcity, and Section 4.5 discusses the remaining threats, in particular the joint determination of tightness and adoption by unobserved establishment characteristics.

3.5. Spatial statistics

The regressions above treat municipalities as independent observations. Adoption may, however, be spatially clustered, either because neighboring municipalities share unobserved characteristics or because adoption in one place raises adoption in its neighbors. We examine this at the municipal level with two standard tools of spatial statistics, which we describe briefly for readers outside the field.

The first is Moran's I (Moran, 1950), the spatial analogue of an autocorrelation coefficient. For a variable y observed in n municipalities, let z be the $n \times 1$ vector of deviations $z_i = y_i - \bar{y}$, and let W be the $n \times n$ matrix whose element w_{ij} equals one if municipalities i and j lie within 50 km of each other and zero otherwise. Then

$$I = \frac{n}{S_0} \cdot \frac{z'Wz}{z'z}, \quad (2)$$

where S_0 is the sum of all elements of W . The numerator adds up, over all pairs of neighboring municipalities, the products of their deviations from the mean. Like a correlation coefficient, I is positive when municipalities with high values tend to have neighbors with high values, negative when high values are surrounded by low values, and close to its expectation $-1/(n-1)$ when values are distributed independently of location. Because the sampling distribution of I depends on the arrangement of municipalities, we obtain p-values by permutation: the observed values are randomly reassigned across municipalities 999 times, and the p-value is the share of permutations that produce an I at least as large as the observed one. The local version of the statistic (Anselin, 1995), $I_i = z_i \sum_j w_{ij} z_j$, computed with row-standardized weights and standardized z , identifies for each municipality whether it and its neighbors are jointly high (high-high), jointly low (low-low), or dissimilar (high-low, low-high), with significance again assessed by conditional permutation (199 draws). We compute the global statistic for the empirical-Bayes smoothed adoption rate (Marshall, 1991) of the 862 municipalities with at least five respondents, for first-period adoption, and for the residuals of a regression of the adoption rate on the five municipal covariates used throughout (log population density, the DID share, the share aged 65 or older, the 2022–2026 log change in the working-age population, and log distance to Tokyo). The 50 km band is our benchmark. Because the value of I depends on the weights, Table 6 also reports it for 30 km and 100 km bands and for row-standardized inverse-distance weights within 100 km, the weights used in the spatial regressions below. The last of these answers the question whether the spatial clustering of adoption is explained by the observable characteristics of municipalities.

The second tool is a pair of spatial regression models that relax the independence assumption in two different ways. In the spatial lag model, the adoption rate of a municipality depends on the weighted average of the adoption rates of its neighbors through a parameter ρ , and a positive ρ is consistent with adoption spilling over from one municipality to the next. In the spatial error model, the adoption rate depends only on the municipality's own covariates, but the unobserved determinants are allowed to be correlated across neighbors through a parameter λ , and a positive λ indicates that neighboring municipalities share unobserved influences without any spillover in adoption itself. We estimate both models, and a model with both parameters, by maximum likelihood with inverse-distance weights truncated at 100 km, so that no municipality

is without neighbors. The comparison of ρ and λ indicates which of the two explanations of clustering the data favor.

3.6. Local diffusion and the intensive margin

Following Goolsbee and Klenow (2002), we add to the hazard model the share of respondents in the same secondary medical area and age group who had adopted before period t , excluding the respondent, as a measure of local exposure to earlier adopters. Because the reflection problem (Manski, 1993) prevents a causal interpretation, we report this as an association. Finally, among adopters we regress the frequency of work-related use, an indicator of work use, an indicator of a paid subscription, and the frequency of emotional or conversational use on the municipal covariates, to compare regional gradients on the intensive margin with those on the extensive margin.

4. Results

4.1. Cumulative adoption and convergence

Figure 1 maps current use by secondary medical area, Figure 2 maps cumulative adoption at the end of each of the four periods, and Figure 3 plots cumulative adoption by population-density quintile. Table 2 reports the underlying weighted shares. Adoption is monotonic in population density in every period. In the quintile with the lowest population density, 5.3% of residents had adopted by the end of 2023, 13.8% by the end of 2024, 24.6% by mid-2025, and 35.4% by January 2026. In the quintile with the highest population density, the corresponding shares are 9.6%, 23.0%, 37.2%, and 49.6%. The ratio between the two rose from 0.55 to 0.60, 0.66, and 0.71, while the difference widened from 4.3 to 9.2, 12.6, and 14.2 percentage points. Relative convergence and absolute divergence therefore coexist, as expected in the early phase of a diffusion process in which later cohorts are drawn increasingly from initially lagging places but the leading places continue to add adopters.

4.2. Hazard estimates

Table 3 reports the hazard estimates. Column (1) includes only period dummies and prefecture fixed effects. The coefficients on log population density interacted with the period dummies are positive and statistically significant in every period (0.062, 0.077, 0.035, and 0.048, the first, second, and fourth at the 1% level), and the coefficients on log distance to Tokyo are negative and significant at the 1% level in every period. The distance coefficient declines in magnitude from -0.140 in the first period to -0.078 in the fourth, which is the hazard-model counterpart of the relative convergence in Table 2.

Column (2) adds the individual controls. The population-density coefficients fall to 0.017, 0.035, 0.001, and 0.020, only the second of which remains significant at the 5% level, and the distance coefficients become indistinguishable from zero. Column (3) adds the municipal covariates: the population-density coefficients turn slightly negative (-0.025 , -0.006 , -0.041 , and -0.022 , the third significant at the 10% level), the DID share is positive and significant at the 10% level, and the share aged 65 or older and the change in the working-age population are not significant. A test of equality of the population-density coefficients across the four periods in this specification does not reject ($\chi^2(3) = 2.86$, $p = 0.41$), so the hazard model provides no

evidence of a change over time in the conditional population-density gradient. Column (4) applies the survey weights and yields the same pattern, with the DID share significant at the 5% level. Column (5) replaces population density with UEA type: relative to center cities, suburbs adopted significantly less only in the first period (-0.131 , 1% level), and municipalities outside UEAs adopted significantly less only in 2024 (-0.617 , 1% level). After the individual controls, therefore, the geography of adoption is flat except for a transitory lag outside the urban employment areas in 2024.

4.3. Composition versus place

The decompositions in Table 4 quantify this conclusion. The Oaxaca–Blinder decomposition of the 12.2 percentage point difference in current use between center cities (46.1%) and municipalities outside UEAs (34.0%), both unweighted and excluding lapsed users, attributes 5.8 points to differences in observed characteristics and 6.4 points to differences in returns. Within the explained part, age (1.7 points), education (1.7 points), working from home (2.0 points), and desk work (0.5 points) account for almost everything, and income, firm size, and employment status contribute nothing. The detailed unexplained components in the lower part of the table should not be read individually. For a continuous variable such as age, the unexplained entry equals the difference in the age coefficients between the two groups multiplied by the mean age, and its size depends on where age is measured from. The large positive entry for age (0.144) is offset by the negative constant (-0.086) and by the household-size entry, and a different origin for age would redistribute these entries without changing their sum (Oaxaca and Ransom, 1999; Jann, 2008). Only the total unexplained component (6.4 points) is invariant to such choices, and we interpret only the total.

The Gelbach decomposition tells the same story for the continuous gradient. The coefficient of current use on log population density falls from 0.0270 without controls to 0.0097 with the full set of controls, so that 64% of the raw gradient is explained by observed characteristics. Of the explained 0.0172, age and sex account for 0.0035 (13% of the raw coefficient), education for 0.0050 (19%), employment status and firm size for -0.0002 , income and household size for essentially zero, and working from home and desk work for 0.0089 (33%). The largest single contributor is thus the concentration of desk jobs and remote work in densely populated municipalities, followed by education and age. Whether remote work is a characteristic of residents or a channel through which place matters cannot be settled by the decomposition itself. Section 4.8 and Section 6 return to this question with additional evidence. The decomposition does bound the answer: if working from home and desk work are attributed entirely to place, the compositional share of the raw gradient falls from 64% to 31% (0.0083 of 0.0270), and the corresponding bound for the Oaxaca–Blinder gap is that between 27% (3.3 of 12.2 percentage points) and 47% (5.8 points) of the center–outside gap is compositional.

4.4. Aging

Table 5 reports the aging estimates. With only sex and population density as controls, the coefficient on the municipal share aged 65 or older is -2.47 and significant at the 1% level. With the individual controls it falls to -0.66 and remains significant at the 5% level. Roughly three-quarters of the raw association is thus compositional, but a place effect remains: residents of a given age adopt more slowly in municipalities where the share of older residents is ten percentage points higher, with a hazard about 6% lower. The change in the working-age population between

2022 and 2026 is positive and significant only at the 10% level. When the aging share is interacted with the period dummies, the coefficients are -0.92 (5% level) in the first period and -1.34 (1% level) in 2024, and they are small and insignificant in the two 2025 periods. The place effect of aging is therefore confined to the earliest phase of diffusion.

The municipal-level results in Table 6 are consistent. In a regression of the empirical-Bayes smoothed current-use rate on the municipal covariates across the 862 municipalities with at least five respondents, the coefficient on the aging share is -0.087 (1% level), which implies that a municipality with a ten percentage point higher share of older residents has a current-use rate 0.9 percentage points lower, and the coefficient on log distance to Tokyo is -0.004 (1% level). Population density itself is not significant in this municipal-level regression once the aging share and distance to Tokyo are included.

4.5. Labor shortages

Table 7 reports the workplace analysis for employed respondents as average marginal effects from probit models. The workplace job openings-to-applicants ratio is positively associated with work-related adoption: a one-point increase in the ratio is associated with a probability of using generative AI at work that is 0.6 percentage points higher (significant at the 1% level). The ratio is unrelated to purely private adoption (-0.02 percentage points, insignificant). Excluding the two prefectures for which we substituted prefectural ratios leaves the estimates unchanged. The contrast between the two outcomes is what makes a demand-pull reading possible: labor scarcity in the local labor market is associated with generative AI entering work, not with residents using it at home. The size of the association, however, gives the reading little weight, as we show next. Density and aging at the workplace are not associated with either outcome. The association is statistically significant but economically small. The marginal effect of 0.6 percentage points per one-point increase in the ratio compares with a mean work-adoption rate of 36.5% among employed respondents. A one-point increase, from a ratio of 1 to a ratio of 2, is a very large change in labor market conditions: it is more than twice the interquartile range across offices (0.98 to 1.38) and exceeds the difference between the 5th and 95th percentiles of offices (0.74 to 1.79). Moving an office from the 25th to the 75th percentile of tightness is associated with a work-adoption probability only 0.3 percentage points higher. Local labor scarcity, as measured here, therefore explains very little of the variation in who uses generative AI at work. Part of the reason may be measurement: the office-level ratio aggregates over all occupations, whereas job openings-to-applicants ratios differ enormously across occupations, from well below one for clerical work to several for construction and care work, and the occupations in which generative AI is useful are not those in which shortages are most acute. The survey does not record occupation in enough detail to construct occupation-specific ratios, so the estimates should be read as the association with the overall tightness of the local labor market, which is a coarse proxy for the scarcity relevant to any particular job.

Two features of the design limit the scope for reverse causality, that is, for early adoption of generative AI in a local labor market to have reduced vacancies and hence the ratio. First, the ratio is an aggregate over all establishments in the jurisdiction of a public employment security office, whereas adoption is measured for an individual worker. No single respondent's adoption can move the office-level ratio, and work-related use by 23% of respondents (Table 1), most of whom started in 2025, is far below the scale at which generative AI could have altered aggregate vacancies over 2023–2025. Second, the ratio can be measured before most adoption took place.

Columns (5) and (6) of Table A1 replace the 2023–2025 mean by the fiscal 2023 ratio alone, a period in which fewer than one in ten respondents had started to use generative AI at all. The average marginal effect of the fiscal 2023 ratio on work-related adoption is 0.8 percentage points per point and significant at the 1% level, and its effect on purely private adoption is 0.2 percentage points and insignificant, the same pattern as in Table 7. What the design cannot rule out is joint determination by unobserved establishment characteristics: labor markets that are tight because they host growing, technologically advanced firms may also be labor markets in which those firms introduce generative AI. We therefore describe the association as consistent with, rather than as evidence of, a demand-pull mechanism.

4.6. Local diffusion and spatial autocorrelation

Table 8 reports the local diffusion estimates. The lagged share of earlier adopters in the respondent's secondary medical area and age group is positively associated with the adoption hazard in the unweighted specification (0.27, 5% level), but the coefficient falls to 0.10 and loses significance with survey weights and municipal covariates. We regard the evidence for local spillovers as weak.

The spatial statistics in Table 6 and Figure 4 add the following. Panel A of Table 6 reports the global Moran's I of equation (2). For the smoothed current-use rate, I is 0.254 (permutation $p = 0.001$), a clear indication that municipalities with high adoption are surrounded by municipalities with high adoption. For first-period adoption the statistic is much smaller (0.047, $p = 0.001$), and for the residuals of a regression of the current-use rate on the municipal covariates it falls to 0.022 ($p = 0.008$). With the binary distance bands, therefore, most of the spatial clustering of adoption is accounted for by the observable characteristics of municipalities, above all their age structure and their distance from Tokyo. The lower block of panel A shows that the size of the statistic depends on the weights, as it must: a 30 km band gives larger values (0.423 for current use, 0.089 for the residuals), a 100 km band smaller ones (0.111 and 0.006, the latter significant only at the 10% level), and in both cases the covariates absorb most of the clustering. Row-standardized inverse-distance weights within 100 km, which give most weight to the nearest neighbors, tell a somewhat different story: the statistic for current use is 0.320 and that for the residuals 0.231 ($p = 0.005$), so that with these weights a substantial part of the correlation between neighboring municipalities survives the covariates. This residual correlation among close neighbors is what the spatial error model in panel B picks up. Panel B of Table 6 reports the municipal regressions. Column (1) is ordinary least squares. Column (2) adds the spatial lag of the dependent variable, that is, the weighted average of the neighbors' adoption rates. Its coefficient ρ is -0.000 and insignificant, so a municipality's adoption rate is not higher when its neighbors' rates are higher, once its own covariates are held constant. Column (3) instead allows the error terms of neighboring municipalities to be correlated. The spatial error parameter λ is 0.52 and significant at the 1% level, so neighboring municipalities do share unobserved influences on adoption. Column (4) includes both parameters and reproduces this pattern ($\rho = -0.016$, $\lambda = 0.53$). Across the four columns the coefficients on the aging share and on distance to Tokyo are essentially unchanged, which indicates that the spatial correlation of the errors does not bias them. The spatial evidence thus points to clustering of unobserved conditions, such as regional industry structure or infrastructure, rather than to adoption spilling over between municipalities. The local Moran statistics in Figure 4 identify 115 high-high municipalities,

almost all in the Tokyo metropolitan area, and 16 low-low municipalities, mainly in inland Hokkaido. No municipality outside a UEA belongs to a high-high cluster.

4.7. The intensive margin

Table 9 restricts the sample to current users and asks whether the regional covariates are related to how the technology is used. Each column takes a different measure of the intensity of use as the dependent variable: the frequency of work-related use (column 1), an indicator for any use at work (column 2), an indicator for a paid subscription (column 3), and the frequency of conversational or emotional use (column 4). The mean of each dependent variable is shown at the foot of the table. The two frequency measures are estimated by ordinary least squares, and for the two indicators we report average marginal effects from probit models. The rows report the estimates for residence population density and for the municipal aging share from models with the full set of individual controls and prefecture fixed effects. None of the four measures is related to population density. A one log point difference in population density is associated with differences of -0.004 and 0.014 on the five-point frequency scales, and with differences of -0.04 and 0.46 percentage points in the probabilities of working with generative AI and of paying for it, none of them significant. Only the frequency of work-related use is negatively associated with the aging share (-0.510), at the 10% level. In the terminology of Comin and Mestieri (2018), the regional differences we document lie entirely on the extensive margin: where people live is associated with whether they have adopted, not with how they use the technology once they have.

4.8. Lapsed users

Diffusion is not the end of the process: 7.7% of respondents, and 17.5% of those who have ever used generative AI (weighted), report that they used it in the past but do not use it now. Table A2 asks whether lapsing has a geography (panel B). Among ever-users, the weighted lapse rate is 17.0% in UEA centers, 18.3% in suburbs, and 19.0% outside UEAs, and in probit models with prefecture fixed effects none of the regional covariates is associated with lapsing: the coefficients on log population density, distance to Tokyo, the DID share, and UEA type are small and insignificant with and without the individual controls, and the coefficient on the municipal aging share is positive but unstable across specifications (significant at the 10% level without individual controls, insignificant with them, and significant at the 5% level with survey weights). Lapsing is instead strongly individual. The simple weighted lapse rates in panel A of Table A2 rise from 14.5% among ever-users under 30 to 21.4% among those aged 65 or older. They are 13.7% among university graduates against about 20% for other education groups, and 13.1% among those who work from home and 14.8% among desk workers, against 18.5% and 19.1% among their counterparts. Panel C of Table A2 reports the average marginal effects of the individual controls from the model in column (2) of panel B, which hold the regional covariates and the other individual characteristics constant. In the model with individual controls (panel C of Table A2), each additional year of age raises the probability of lapsing by 0.18 percentage points, a university degree lowers it by 4.1 points relative to junior high or high school, working from home by 3.7 points, and desk work by 2.8 points, all significant at the 1% level, whereas sex is unrelated to lapsing. The transition from trial to sustained use, like the intensity of use, therefore depends on the complementary skills and job tasks of the individual rather than on the place of residence. Because lapsed users lack an adoption date, these results also indicate the direction of

the selection that their exclusion introduces into the hazard sample: the excluded group is older and less educated than the retained adopters, but not more rural.

5. Validity checks

Retrospective adoption dates from a single cross-section are unusual in this literature, and we therefore subject them to a set of checks that we regard as mandatory before the results are interpreted. We examine six questions in turn: (i) whether the reported adoption dates are consistent with external benchmarks and with internal patterns, (ii) whether attaching the current residence to past adoption dates matters, (iii) how the exclusion of lapsed users affects the estimates, (iv) whether the sample represents the regional population and whether selection into the online panel affects the geography of adoption, (v) whether the workplace analysis is sensitive to the way workplaces are located, and (vi) whether the results depend on the length of the periods, on the age of the youngest respondents at the start of the window, or on the number of hypotheses tested. The results are reported in Appendix Tables A1 to A10.

We begin with the validity of recall. The cohort boundaries are anchored to product events named in the question, and two external comparisons are available (Table A3). Nakagomi, Abe, and Tabuchi (2025) report that 21.3% of a JACSIS-based sample had used generative AI in the twelve months to January 2025. In our data, the share who had started by the end of 2024 and were still using is 14.8% (weighted), and adding all lapsed users, whose dates are unknown, gives an upper bound of 22.5%. The external figure lies within these bounds. The two MIC surveys provide benchmarks at two dates. The fiscal 2024 survey (MIC, 2025) found that 26.7% of adults had any experience with generative AI one year before our survey, and the fiscal 2025 survey (MIC, 2026), fielded in January and February 2026, at the same time as our wave, found 52.2% among adults aged 20 or older (58.8% including respondents aged 15 to 19). Our share with any experience is 43.9% with the survey weights and 48.9% without them, and it lies between the two benchmarks. The comparison with the contemporaneous figure is not exact. The MIC sample is drawn from an online monitor panel with equal quotas by age and sex ($N = 1,236$), which gives young respondents a far larger weight than they have in the population, whereas our weighted share reproduces the population age structure. Our unweighted share, whose age composition is closer to that of the MIC sample, is correspondingly closer to the MIC figure. Internally, earlier cohorts use the technology more intensively, as expected if the dates are informative: 71% of first-period adopters use it for work and 23% pay for it, compared with 48% and 10% of the most recent cohort. Finally, if respondents misplace their adoption by one period, merging adjacent periods should absorb the error. Columns (1) and (2) of Table A4 re-estimate the full specification with periods 1 and 2 merged and with periods 3 and 4 merged. The coefficients on population density and on the aging share are within sampling error of the baseline estimates in column (3) of Table 3.

The second check concerns movers. Postal codes refer to the current residence, so that a respondent who moved after adopting is assigned to the wrong municipality. Columns (3) and (4) of Table A4 exclude respondents who have lived at the current address for less than one year (7.3% of the sample) or less than five years (27.5%). The population-density and aging coefficients are unchanged.

The third check concerns the lapsed users who are excluded from the hazard sample because their adoption dates are unknown. Columns (5) to (7) of Table A4 assign the 7.7% of lapsed users to the earliest period, to the latest period, or proportionally to the observed cohort distribution. Every population-density and aging coefficient of the main model remains within one standard error of its baseline value.

The fourth check concerns the representativeness of the sample across regions. Table A5 compares respondents with the 2020 Census by UEA type. The share aged 65 or older among respondents is 21.0% in center cities, 23.8% in suburbs, and 24.2% outside UEAs, against 26.5%, 28.8%, and 38.5% in the Census. The under-representation of older residents is thus most severe outside the urban employment areas, and the weights only partly correct it. Table A6 shows that estimates restricted to municipalities with at least five or at least ten respondents, and estimates with the 2019-based instead of the 2022-based survey weights, give similar results. Municipalities with respondents cover 96% of the population.

The under-representation of older residents outside the urban employment areas is a symptom of a more general selection: only residents who are able and willing to answer an internet survey enter the panel, and this filter is most restrictive among older residents of sparsely populated areas. If the selected older rural residents are more digitally literate than their non-sampled neighbors, the flat conditional geography of Table 3 could understate the place effect in the full population. The data cannot test this directly, but Table A7 estimates the full specification separately by age group. The conditional population-density and distance gradients are flat, or negative, within every age group, including respondents aged 65 or older (population-density coefficients -0.112 , -0.038 , -0.118 , and 0.003 across the four periods, only the third significant at the 5% level). Within the sampled population, therefore, the absence of an urban advantage does not depend on the young, and the selection concern is confined to the difference between sampled and non-sampled residents, which we discuss in Section 6.

The fifth check concerns the workplace data used in Section 4.5. Table A1 excludes, in turn, respondents whose home and workplace postal codes coincide, respondents whose workplace postal code was a large-establishment code, and the two prefectures with substituted ratios. The average marginal effect of the job openings-to-applicants ratio on work-related adoption remains between 0.5 and 0.9 percentage points per point and significant at least at the 5% level. Columns (5) and (6) of the same table, discussed above, use the fiscal 2023 ratio alone.

The final checks concern the specification of the hazard model and multiple testing. Dropping the offset for period length, or excluding respondents who were younger than 15 in November 2022, does not change the results (Table A8). With Benjamini–Hochberg adjustment across the eleven regional coefficients of the full model, no coefficient has an adjusted q-value below 0.30 (Table A9). This is the appropriate summary of column (3) of Table 3, in which the geography of adoption is flat after individual controls.

6. Discussion

Our results speak to three debates. The first concerns the geography of new technologies. The commercial internet was led by large cities in its early phase but its adoption became widespread (Forman, Goldfarb, and Greenstein, 2005). Generative AI in Japan shows the same

pattern in miniature: the quintile of municipalities with the highest population density was almost twice as likely to have adopted by the end of 2023, and the relative gap has narrowed in each subsequent period, while the absolute gap has continued to widen. The urban hierarchy therefore describes the order of adoption rather than a permanent divide, and the flat geography after individual controls implies that a rural resident with the same age, education, and job as an urban resident is about as likely to have adopted.

The second debate concerns the interpretation of regional gaps in the "AI divide" literature. Nakagomi, Abe, and Tabuchi (2025), like studies for other countries, treat urban residence as one of several factors associated with use. Our decompositions show that in Japan a large part of the urban–rural gap is a compositional artifact of the geography of age, education, and work, but how large depends on where the line between composition and place is drawn for working from home and desk work, the group that accounts for the largest share of the explained gradient (33%). These two variables are individual characteristics in the sense that they describe the respondent's own job, but the jobs available in a place, the firms that offer them, and the practices of those firms are attributes of the place. If they are counted as composition, 64% of the population-density gradient and 47% of the center–outside gap are compositional. If they are counted as place, the compositional shares fall to 31% and 27%, and place, understood as the local structure of jobs and working arrangements, accounts for the majority of the gap. Table A10 helps to locate the two variables on this spectrum. Desk work is more common in UEA centers (44%) than outside UEAs (39%), and its population-density gradient survives controls for employment status, firm size, education, age, and sex (panel B of Table A10: a marginal effect of 2.5 percentage points per log point of population density). This is the composition of jobs, which is a feature of place but not one that a policy on generative AI would address. Working from home is a different matter. Among desk workers, 28% in UEA centers and 32% in suburbs work from home at least sometimes, against 14% outside UEAs, and the gap is as large among desk workers in firms with 100 or more employees (28% and 29% against 11%) as in smaller firms. Within the urban employment areas, however, the gradient of remote work with respect to population density is small (about one percentage point per log point among desk workers) and disappears once job characteristics are controlled. The remote-work channel is thus not a smooth function of population density but a discrete disadvantage of the areas outside the urban employment system, which is where the infrastructure and employer practices that make remote work possible are thinnest. Our reading is that the compositional interpretation is appropriate within the urban employment areas, and that outside them a place effect operates through the organization of work. Policies aimed at the geography of adoption would therefore be better targeted at the people and jobs that lag, wherever they are, and at the working arrangements available outside the urban employment areas, than at places as such.

The third debate concerns demographic change and labor scarcity. Acemoglu and Restrepo (2022) show that aging induces the adoption of robots, which substitute for scarce middle-aged manual workers. For generative AI we find the opposite sign for aging and the expected sign for labor scarcity. Municipal aging slows adoption beyond its compositional effect, but only in the first two years, which is consistent with a supply-side constraint, the thinness of the complementary skills and social networks through which a new technology spreads in older communities, that weakens as the technology becomes easier to use. Workplace labor shortages, in contrast, are associated with somewhat more work-related adoption and with no private

adoption. The two findings are compatible: aging regions adopt later as consumers, and tight labor markets adopt slightly more as workplaces. The demand-pull association is, however, small: a one-point increase in the job openings-to-applicants ratio, a very large change in labor market conditions, is associated with a work-adoption probability only about 0.6 percentage points higher, and moving across the interquartile range of local tightness with about 0.3 points. Local labor shortages, as measured by the overall office-level ratio, therefore explain very little of who uses generative AI at work. This is itself informative for the policy debate on whether labor shortages will accelerate the diffusion of AI in Japan: on our evidence, the pull from local scarcity of labor is detectable but weak, at least in the first three years of the technology.

Four limitations qualify these conclusions. First, all associations are correlational. Workplace labor shortages and work-related adoption may be jointly determined by firm characteristics we do not observe, even though the timing check in Table A1 makes reverse causality from adoption to vacancies unlikely, and the flat geography after controls does not identify why the residents of densely populated areas are younger, more educated, and more likely to work at desks.

Second, the sample is an online panel, and the selection this implies is not neutral with respect to our main result. Participation requires the ability and willingness to complete internet surveys, a filter that removes few young urban residents and many older rural ones. Table A5 shows that respondents aged 65 or older are under-represented by 5 percentage points in UEA centers and by 14 points outside UEAs. Inverse probability weights correct the age and regional composition of the sample but not the selection within cells: the older rural residents who are in the panel are, by construction, those for whom the digital barriers of their place are lowest. Two consequences follow. The cumulative adoption levels we report are those of the population that participates in online surveys, and are upper bounds for the population as a whole, most clearly outside the urban employment areas. More importantly for our argument, the selection biases the estimated place effect toward zero. In the full population, a resident of a sparsely populated and aging municipality faces thinner digital infrastructure, fewer adopters among family, colleagues, and neighbors from whom to learn, and fewer local services that use the technology. These are exactly the barriers that keep the non-sampled residents out of an online panel. The flat conditional geography of Table 3, and its flatness within the oldest age group in Table A7, therefore establish that residents with the same characteristics and the same access to the online world adopt at the same rate wherever they live, not that place is irrelevant for those without that access. Our estimates of the place effect should be read as lower bounds, and the compositional interpretation as a statement about the digitally connected population. A test of the place effect in the full population requires a probability sample with offline recruitment, such as the MIC communications usage survey, and the retrospective adoption question could be added to such a survey at low cost.

Third, the adoption question counts trial use, and the resulting dates are subject to measurement error. A respondent who tried ChatGPT once in 2023, stopped, and started to use it in earnest in 2025 may report either date. A respondent who tried once and never returned is a lapsed user and is excluded from the hazard sample. The first type of error, if respondents date adoption at the first trial, moves adoption dates toward the early periods and overstates early adoption. If respondents date it at the start of sustained use, or if the well-documented tendency to recall events as more recent than they were (forward telescoping) dominates, dates move toward

the later periods and understate the first-period hazard. Errors of this kind that are unrelated to place shift the period intercepts α_t but not the place coefficients β_s . Errors that are related to place, for example if early adopters in cities are heavier users who recall their start date more precisely, could bias the early-period gradient in either direction. Three observations limit the concern. The cohort boundaries are anchored to product events named in the question. The intensity of current use declines monotonically across cohorts, as it should if the dates carry information. And merging adjacent periods, which absorbs errors of one period in either direction, leaves the estimates unchanged (Table A4). The first-period estimates are nevertheless best read as describing adoption that persisted to the survey date rather than the first contact with the technology.

Fourth, the retrospective dates are coarse and are attached to the current residence and job, so that the analysis of the intensive margin, of lapsing, and of the workplace refers to the survey date rather than to the adoption date.

7. Conclusion

This study reconstructs the spatial diffusion of generative AI in Japan from the retrospective adoption dates of 27,630 respondents to the 2025 wave of the JACSIS survey, linked through postal codes to municipal statistics and to office-level labor market indicators. Because every respondent could adopt from the same date, the data support a discrete-time hazard analysis of adoption without the left truncation that has complicated diffusion studies of most technologies.

Our empirical analysis yielded four primary findings. First, adoption follows the urban hierarchy, and the gap between high- and low-population density municipalities has narrowed in relative terms while widening in absolute terms. Second, the geographic gradient is largely compositional: 64% of the population-density gradient is explained by the age, education, and job characteristics of residents, and the coefficients on population density and distance to Tokyo are indistinguishable from zero once these are controlled. Third, municipal aging is associated with slower adoption beyond its compositional effect, but only during the first two years of diffusion. Fourth, the tightness of the workplace labor market is associated with work-related adoption and not with private adoption, a pattern consistent with labor scarcity pulling the technology into workplaces, but the association is economically small: a one-point increase in the job openings-to-applicants ratio is associated with a work-adoption probability less than one percentage point higher. Regional differences are confined to the extensive margin: neither the intensity of use among adopters nor the probability that an ever-user stops using the technology, which is 17.5% overall, is related to place, whereas both follow age, education, and job content.

For the international debate on Japan's low adoption rate, the implication is that the national average is a weighted average of a demographic and occupational structure that adopts slowly, rather than the sum of a metropolitan Japan that has adopted and a rural Japan that has not. Future research should use the next JACSIS wave, in which the adoption question will be repeated, to replace retrospective cohorts with observed transitions and to model the hazard of exit from use dynamically as the panel progresses, and should link the workplace analysis to establishment data on the introduction of generative AI tools.

References

- Acemoglu, D. (2010). When does labor scarcity encourage innovation? *Journal of Political Economy*, 118(6), 1037–1078.
- Acemoglu, D., & Restrepo, P. (2022). Demographics and automation. *Review of Economic Studies*, 89(1), 1–44.
- Anselin, L. (1995). Local indicators of spatial association: LISA. *Geographical Analysis*, 27(2), 93–115.
- Anthropic. (2025). Anthropic Economic Index report: Uneven geographic and enterprise AI adoption. September 2025. <https://www.anthropic.com/research/anthropic-economic-index-september-2025-report> (accessed 9 September 2026).
- Cabinet Secretariat of Japan. (2023). Digital Garden City Nation Initiative: Comprehensive Strategy, revised 26 December 2023. https://www.cas.go.jp/jp/seisaku/digital_denen/ (accessed 8 September 2026).
- Bick, A., Blandin, A., & Deming, D. J. (2025). The rapid adoption of generative AI. *Management Science*, forthcoming.
- Chatterji, A., Cunningham, T., Deming, D., Hitzig, Z., Ong, C., Shan, C. Y., & Wadman, K. (2025). How people use ChatGPT. NBER Working Paper No. 34255.
- Clemens, M. A., Lewis, E. G., & Postel, H. M. (2018). Immigration restrictions as active labor market policy: Evidence from the Mexican bracero exclusion. *American Economic Review*, 108(6), 1468–1487.
- Comin, D., & Hobijn, B. (2010). An exploration of technology diffusion. *American Economic Review*, 100(5), 2031–2059.
- Comin, D., & Mestieri, M. (2018). If technology has arrived everywhere, why has income diverged? *American Economic Journal: Macroeconomics*, 10(3), 137–178.
- Forman, C., Goldfarb, A., & Greenstein, S. (2005). How did location affect adoption of the commercial Internet? Global village vs. urban leadership. *Journal of Urban Economics*, 58(3), 389–420.
- Gelbach, J. B. (2016). When do covariates matter? And which ones, and how much? *Journal of Labor Economics*, 34(2), 509–543.
- Goolsbee, A., & Klenow, P. J. (2002). Evidence on learning and network externalities in the diffusion of home computers. *Journal of Law and Economics*, 45(2), 317–343.
- Griliches, Z. (1957). Hybrid corn: An exploration in the economics of technological change. *Econometrica*, 25(4), 501–522.
- Hornbeck, R., & Naidu, S. (2014). When the levee breaks: Black migration and economic development in the American South. *American Economic Review*, 104(3), 963–990.
- Humlum, A., & Vestergaard, E. (2024). The unequal adoption of ChatGPT exacerbates existing inequalities among workers. *Proceedings of the National Academy of Sciences*, 121(50), e2414972121.

Jann, B. (2008). The Blinder–Oaxaca decomposition for linear regression models. *Stata Journal*, 8(4), 453–479.

Jenkins, S. P. (1995). Easy estimation methods for discrete-time duration models. *Oxford Bulletin of Economics and Statistics*, 57(1), 129–138.

Kanemoto, Y., & Tokuoka, K. (2002). Proposal for the standards of metropolitan areas of Japan. *Journal of Applied Regional Science*, 7, 1–15. (In Japanese.)

Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. *Review of Economic Studies*, 60(3), 531–542.

Marshall, R. J. (1991). Mapping disease and mortality rates using empirical Bayes estimators. *Journal of the Royal Statistical Society: Series C*, 40(2), 283–294.

Moran, P. A. P. (1950). Notes on continuous stochastic phenomena. *Biometrika*, 37(1/2), 17–23.

Ministry of Internal Affairs and Communications (MIC). (2025). White Paper on Information and Communications in Japan 2025 (令和 7 年版情報通信白書), Section on individual use of AI. Tokyo.
<https://www.soumu.go.jp/johotsusintokei/whitepaper/ja/r07/html/nd112210.html> (accessed 9 September 2026).

Ministry of Internal Affairs and Communications (MIC). (2026). White Paper on Information and Communications in Japan 2026 (令和 8 年版情報通信白書), Part I, Chapter 1, Section 1: Individual use of AI. Tokyo.
<https://www.soumu.go.jp/johotsusintokei/whitepaper/ja/r08/pdf/n1110000.pdf> (accessed 12 September 2026).

Nakagomi, A., Abe, N., & Tabuchi, T. (2025). Emerging generative AI divide: Personal, positional, and resource-based factors associated with use and reasons for non-use. *Telematics and Informatics*, 102360.

Oaxaca, R. L., & Ransom, M. R. (1999). Identification in detailed wage decompositions. *Review of Economics and Statistics*, 81(1), 154–157.

Yamada, H., Kursim, L. A., Matsushima, M., Nakagomi, A., & Tabuchi, T. (2026). Anticipating displacement: AI job replacement anxiety, work motivation, and mental well-being in Japan. Keio-IES Discussion Paper Series, DP2026-015, Institute for Economic Studies, Keio University.

Tables

Table 1. Adoption and use of generative AI by urban employment area (UEA) type

Variable	UEA center	UEA suburb	Outside UEA
Current user	0.372	0.358	0.254
Ever used	0.448	0.438	0.314
Started 2022/11–2023/12 and still using	0.065	0.052	0.031
Used in the past, not currently	0.076	0.080	0.060
Adopted by December 2023	0.070	0.056	0.033
Adopted by December 2024	0.168	0.159	0.070
Adopted by June 2025	0.292	0.273	0.165
Adopted by January 2026	0.403	0.389	0.271
Uses generative AI for work	0.233	0.221	0.167
Paid subscription	0.051	0.043	0.021
Mean age	48.0	48.2	50.5
Weighted N	16,612	9,818	1,169

Notes: Inverse probability weighted (2022 weights) means for the 27,086 respondents with a residence municipality. "Adopted by" rows give the cumulative share of respondents who started using generative AI by the date and were still using it at the survey, and lapsed users are excluded from the numerator. UEA types follow the 2020 urban employment area definitions of Kanemoto and Tokuoka (2002).

Table 2. Cumulative adoption of generative AI by quintile of residence population density

Quintile of municipal population density	Adopted by 2023-12	Adopted by 2024-12	Adopted by 2025-06	Adopted by 2026-01	Current user	Weighted N
Q1 (lowest population density)	0.053	0.138	0.246	0.354	0.326	9,647
Q2	0.062	0.151	0.272	0.383	0.353	8,089
Q3	0.064	0.169	0.292	0.407	0.376	4,654
Q4	0.077	0.195	0.324	0.444	0.411	3,015
Q5 (highest population density)	0.096	0.230	0.372	0.496	0.460	2,192
Ratio Q1/Q5	0.55	0.60	0.66	0.71		
Difference Q5 – Q1 (percentage points)	4.4	9.2	12.6	14.2		

Notes: Inverse probability weighted shares of respondents who had adopted by each date (lapsed users excluded). Quintiles of 2020 population density are defined over municipalities, so the weighted number of respondents differs across quintiles.

Table 3. Discrete-time hazard of generative AI adoption: population density, distance to Tokyo, and period

	(1)	(2)	(3)	(4)	(5)
	Prefecture FE only	Plus individual controls	Plus municipal covariates	Column 3 with survey weights	UEA type instead of density and distance
Log population density x 2022/11-2023/12	0.062*** (0.019)	0.017 (0.018)	-0.025 (0.024)	-0.001 (0.034)	
Log population density x 2024	0.077*** (0.017)	0.035** (0.017)	-0.006 (0.022)	0.005 (0.032)	
Log population density x 2025H1	0.035** (0.016)	0.001 (0.015)	-0.041* (0.021)	-0.009 (0.032)	
Log population density x 2025H2-2026/01	0.048*** (0.016)	0.020 (0.016)	-0.022 (0.022)	-0.011 (0.030)	
Log km to Tokyo x 2022/11-2023/12	-0.140*** (0.028)	-0.022 (0.033)	-0.010 (0.033)	-0.009 (0.050)	
Log km to Tokyo x 2024	-0.117*** (0.029)	-0.004 (0.035)	0.009 (0.035)	-0.002 (0.052)	
Log km to Tokyo x 2025H1	-0.109*** (0.025)	-0.001 (0.032)	0.011 (0.032)	0.015 (0.049)	
Log km to Tokyo x 2025H2-2026/01	-0.078*** (0.027)	0.026 (0.032)	0.039 (0.032)	0.038 (0.048)	
DID population share, residence (2020)			0.137* (0.072)	0.226** (0.103)	
Share aged 65+, residence (2022-01)			-0.469 (0.344)	-0.527 (0.502)	-0.308 (0.325)
Log change in pop. aged 15-64, 2022-26, residence			0.427 (0.411)	0.127 (0.617)	0.232 (0.401)
UEA suburb x 2022/11-2023/12					-0.131*** (0.048)
UEA suburb x 2024					0.044 (0.039)
UEA suburb x 2025H1					-0.002 (0.039)
UEA suburb x 2025H2-2026/01					-0.004 (0.039)
Outside UEA x 2022/11-2023/12					-0.289 (0.183)
Outside UEA x 2024					-0.617*** (0.174)
Outside UEA x 2025H1					-0.068 (0.123)
Outside UEA x 2025H2-2026/01					0.070 (0.109)
Observations	86,629	86,629	86,629	86,629	86,629

Notes: Complementary log-log estimates of equation (1) on the person-period file (86,629 person-periods and 11,385 adoption events, lapsed users excluded). Coefficients are log hazard ratios. All columns include period dummies, prefecture fixed effects, and an offset for the log length of the period in months. Column (2) adds the individual controls (age and its square, sex, education, employment status, firm size, household income, household size, working from home, desk work). Column (3) adds the municipal covariates. Column (4) re-estimates column (3) with inverse probability weights. Column (5) replaces population density and distance with UEA type interacted with period (reference: UEA center). Cluster-robust standard errors (municipality of residence) in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4. Composition versus place: Oaxaca–Blinder and Gelbach decompositions

Panel A. Oaxaca–Blinder decomposition of current use, UEA center vs. outside UEA	Estimate	SE
Mean, UEA center	0.4612	0.0040
Mean, outside UEA	0.3396	0.0187
Difference	0.1217	0.0191
Explained (total)	0.0577	0.0075
Unexplained (total)	0.0640	0.0175
Explained, by covariate group		
Age	0.0175	0.0057
Sex	-0.0013	0.0007
Education	0.0165	0.0025
Employment status	0.0002	0.0005
Firm size	-0.0008	0.0011
Household income	0.0004	0.0006
Household size	0.0008	0.0008
Working from home	0.0195	0.0024
Desk work	0.0049	0.0015
Unexplained, by covariate group (not separately identified, see notes)		
Age	0.1444	0.0570
Sex	-0.0233	0.0523
Education	-0.0085	0.0416
Employment status	-0.0499	0.0464
Firm size	0.0096	0.0219
Household income	-0.0004	0.0265
Household size	0.0780	0.0386
Working from home	0.0113	0.0056
Desk work	-0.0115	0.0138
Constant	-0.0857	0.1199
Panel B. Gelbach decomposition of the coefficient on log population density	Estimate	SE
Coefficient on log population density, no controls	0.0270	0.0020
Coefficient on log population density, full controls	0.0097	0.0019
Age and sex (13% of raw coefficient)	0.0035	0.0006
Education (19% of raw coefficient)	0.0050	0.0004
Employment status and firm size (-1% of raw coefficient)	-0.0002	0.0002
Household income and size (0% of raw coefficient)	-0.0000	0.0002
Working from home and desk work (33% of raw coefficient)	0.0089	0.0005
Total explained (64% of raw coefficient)	0.0172	0.0009

Notes: Unweighted linear probability models of current use, lapsed users excluded. Panel A: pooled-coefficient Oaxaca–Blinder decomposition (Jann, 2008) of the difference between residents of UEA centers (N = 15,246) and residents of municipalities outside UEAs (N = 642), with robust standard errors. The detailed unexplained components are reported for completeness only: for continuous variables they depend on the origin of measurement (for age, the entry equals the difference in the age coefficients times mean age and is offset by the constant), and for categorical variables on the choice of base category (Oaxaca and Ransom, 1999). Only the total unexplained component is invariant. Panel B: Gelbach (2016) decomposition of the change in the coefficient on log residence population density between a regression with no controls and one with all individual controls (N = 25,465). The entries sum to the total change in the coefficient ($0.0270 - 0.0097 = 0.0172$).

Table 5. Municipal population aging and the hazard of adoption

	(1)	(2)	(3)	(4)
	Aging share, sex and density only	Aging share, full individual controls	Working-age population change, full controls	Aging share by period, full controls
Share aged 65+, residence (2022-01)	-2.468*** (0.295)	-0.662** (0.311)		
Log population density, residence (2020)	0.002 (0.013)	-0.001 (0.013)	0.003 (0.012)	-0.001 (0.013)
Log change in pop. aged 15-64, 2022-26, residence			0.689* (0.366)	
Share 65+ x 2022/11-2023/12				-0.922** (0.466)
Share 65+ x 2024				-1.342*** (0.456)
Share 65+ x 2025H1				-0.137 (0.409)
Share 65+ x 2025H2-2026/01				-0.449 (0.415)
Observations	86,629	86,629	86,629	86,629

Notes: Complementary log-log estimates on the person-period file, with period dummies, prefecture fixed effects, and the log period-length offset in all columns. Column (1) enters the municipal share aged 65 or older with sex and log population density as the only other controls, so that its coefficient includes the compositional effect of older residents adopting less. Column (2) adds the full set of individual controls, including age, so that the remaining coefficient is the place effect of living in an older municipality. Column (3) replaces the aging share with the 2022–2026 log change in the population aged 15–64, with full controls. Column (4) interacts the aging share with the four period dummies, with full controls. Cluster-robust standard errors (municipality of residence) in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6. Spatial autocorrelation of municipal adoption: Moran's I and spatial regression models

Panel A. Global Moran's I (equation 2)	Moran's I	Expected value under no autocorrelation	Permutation p-value (999 draws)	Municipalities
50 km band (benchmark): Current-use rate (EB smoothed)	0.254	-0.0012	0.001	862
50 km band (benchmark): First-period adoption rate	0.047	-0.0012	0.001	862
50 km band (benchmark): Residual of current-use rate on municipal covariates	0.022	-0.0012	0.008	862
30 km band: Current-use rate (EB smoothed)	0.423	-0.0012	0.001	862
30 km band: First-period adoption rate	0.079	-0.0012	0.001	862
30 km band: Residual of current-use rate on municipal covariates	0.089	-0.0012	0.001	862
100 km band: Current-use rate (EB smoothed)	0.111	-0.0012	0.001	862
100 km band: First-period adoption rate	0.020	-0.0012	0.003	862
100 km band: Residual of current-use rate on municipal covariates	0.006	-0.0012	0.054	862
Inverse distance within 100 km (row-standardized): Current-use rate (EB smoothed)	0.320	-0.0012	0.001	862
Inverse distance within 100 km (row-standardized): First-period adoption rate	0.182	-0.0012	0.276	862
Inverse distance within 100 km (row-standardized): Residual of current-use rate on municipal covariates	0.231	-0.0012	0.005	862

Panel B. Municipal regressions of the smoothed current-use rate

	(1) OLS	(2) Spatial lag model	(3) Spatial error model	(4) Spatial lag and error model
Log population density (2020)	0.001 (0.001)	0.001 (0.001)	0.000 (0.001)	0.000 (0.001)
DID population share (2020)	-0.005* (0.003)	-0.005 (0.004)	-0.004 (0.004)	-0.004 (0.004)
Share aged 65+ (2022-01)	-0.087*** (0.019)	-0.087*** (0.020)	-0.086*** (0.021)	-0.085*** (0.021)
Log change in pop. aged 15-64, 2022-26	0.029 (0.028)	0.029 (0.028)	0.027 (0.029)	0.027 (0.029)
Log km to Tokyo Station	-0.004*** (0.001)	-0.004*** (0.001)	-0.006*** (0.001)	-0.006*** (0.001)
Constant	0.460*** (0.009)	0.460*** (0.021)	0.469*** (0.011)	0.476*** (0.022)
Spatial parameters				
Spatial lag of the dependent variable (ρ)		-0.000 (0.046)		-0.016 (0.046)
Spatial error parameter (λ)			0.520*** (0.121)	0.529*** (0.124)
Observations	862	862	862	862

Notes: Sample: the 862 municipalities with at least five respondents. Panel A: Moran's I of equation (2) for the empirical-Bayes smoothed current-use rate, for the share who adopted in the first period, and for the residuals of an OLS regression of the smoothed current-use rate on the five municipal covariates of panel B (log population density, the DID population share, the share aged 65 or older, the 2022–2026 log change in the population aged 15–64, and log distance to Tokyo Station). The benchmark weights are a binary 50 km band between municipality centroids ($w_{ij} = 1$ within 50 km). The lower block uses 30 km and 100 km bands and row-standardized inverse-distance weights within 100 km. P-values from 999 random permutations of the values across municipalities. Panel B: dependent variable is the empirical-Bayes smoothed share of current users. Column (1) is OLS. Column (2) adds the spatial lag of the dependent variable (weighted average of neighbors' rates, parameter ρ). Column (3) allows spatially correlated errors (parameter λ). Column (4) includes both. Spatial weights are inverse-distance weights between municipality centroids truncated at 100 km, row-standardized, estimated by maximum likelihood. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7. Workplace labor market tightness and work-related versus private adoption: average marginal effects from probit models (employed respondents)

	(1) Work-related use	(2) Work-related use	(3) Private use only	(4) Private use only
--	-------------------------	-------------------------	-------------------------	-------------------------

	All	Excluding Gunma and Yamanashi	All	Excluding Gunma and Yamanashi
Job openings-to-applicants ratio, workplace (2023-25)	0.0063*** (0.0024)	0.0063*** (0.0024)	-0.0002 (0.0015)	-0.0002 (0.0015)
Log population density, workplace (2020)	-0.0052 (0.0039)	-0.0048 (0.0039)	0.0039 (0.0029)	0.0035 (0.0029)
Share aged 65+, workplace (2022-01)	-0.1318 (0.0977)	-0.1319 (0.0987)	-0.1101 (0.0767)	-0.1211 (0.0774)
Mean of dependent variable	0.365	0.366	0.155	0.155
Observations	17,814	17,524	17,817	17,527

Notes: Average marginal effects from probit models, that is, the change in the probability of the outcome per unit increase in the regressor, averaged over the sample. The dependent variable in columns (1)–(2) is an indicator for using generative AI for work. In columns (3)–(4) it is an indicator for using generative AI for private purposes only. The job openings-to-applicants ratio is the fiscal 2023–2025 mean for the public employment security office with jurisdiction over the workplace municipality. For Gunma and Yamanashi, which do not publish office-level ratios, the prefectural ratio is used in columns (1) and (3). All columns include the individual controls and prefecture fixed effects. Delta-method standard errors, clustered by workplace municipality, in parentheses. The underlying probit coefficients are in the replication package (T5b_jobrati_o_coef). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8. Lagged local adoption and the hazard of adoption (periods 2–4)

	(1) Unweighted, individual controls	(2) Survey weights, individual and municipal controls
Lagged adoption share, same secondary medical area and age group (leave-one-out)	0.274** (0.108)	0.098 (0.151)
Log population density x 2024	0.037** (0.017)	0.043 (0.028)
Log population density x 2025H1	-0.002 (0.015)	0.021 (0.025)
Log population density x 2025H2-2026/01	0.004 (0.015)	0.014 (0.026)
Observations	60,859	60,859

Notes: Complementary log-log estimates on periods 2–4 of the person-period file. The lagged adoption share is the leave-one-out share of respondents in the same secondary medical area (one of the 335 areas, each a group of municipalities, that prefectures define under the Medical Care Act for planning hospital care) and the same age group who had adopted before period t . Both columns include period dummies and prefecture fixed effects. Standard errors clustered by secondary medical area $\times$ age group cell. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9. The intensive margin among current users: regional covariates and the intensity of use

	(1)	(2)	(3)	(4)
Dependent variable:	Frequency of work-related use	Any use at work	Paid subscription	Frequency of conversational or emotional use
Log population density, residence (2020)	-0.0036 (0.0112)	-0.0004 (0.0052)	0.0046 (0.0044)	0.0144 (0.0145)
Share aged 65+, residence (2022-01)	-0.5102* (0.2659)	0.0272 (0.1165)	-0.1377 (0.0967)	0.5888 (0.3594)
Estimator	OLS	Probit, average marginal effects	Probit, average marginal effects	OLS
Mean of dependent variable	2.557	0.593	0.150	1.861
Observations	11,365	11,362	11,362	11,365

Notes: Sample: current users. In columns (1) and (4) the dependent variable is the mean of the survey's frequency items (1 = never to 5 = daily) for work-related uses and for conversational or emotional uses, and the entries are OLS coefficients. In columns (2)–(3) the dependent variable is an indicator, and the entries are average marginal effects from probit models, that is, the change in the probability of the outcome per unit increase in the regressor, averaged over the sample. All columns include the individual controls and prefecture fixed effects. Only the estimates for the two regional covariates are shown. The underlying probit coefficients are in the replication package (T6_intensity_coef). Cluster-robust standard errors (municipality of residence) in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix tables

Table A1. Workplace analysis: sample restrictions and timing of the labor market measure (average marginal effects from probit models)

	(1)	(2)	(3)	(4)	(5)	(6)
	All employed respondents	Excluding home postal code equal to workplace postal code	Excluding workplace postal codes resolved through the establishment file	Excluding Gunma and Yamanashi (prefectural ratios)	FY2023 ratio only, work-related use	FY2023 ratio only, private use only
Job openings-to-applicants ratio, workplace (2023-25)	0.0077*** (0.0021)	0.0055** (0.0025)	0.0089*** (0.0027)	0.0077*** (0.0021)		
Log population density, workplace (2020)	-0.0017 (0.0031)	-0.0050 (0.0036)	-0.0024 (0.0032)	-0.0013 (0.0031)	-0.0008 (0.0032)	0.0064*** (0.0024)
Job openings-to-applicants ratio, workplace (FY2023 only)					0.0079*** (0.0022)	0.0018 (0.0011)
Observations	17,814	13,380	16,832	17,524	17,103	17,097

Notes: Average marginal effects from probit models among employed respondents, with individual controls and prefecture fixed effects. The dependent variable is work-related adoption in columns (1)–(5) and adoption for private purposes only in column (6). Columns (5)–(6) replace the fiscal 2023–2025 mean by the fiscal 2023 ratio alone, measured before most work-related adoption took place. Standard errors clustered by workplace municipality. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A2. Lapsed use among respondents who have ever used generative AI

Panel A. Weighted lapse rates by group

Group	Category	Lapse rate	Weighted N
UEA type	UEA center	0.170	7,450
UEA type	UEA suburb	0.182	4,302
UEA type	Outside UEA	0.190	367
Age group	Under 30	0.144	3,490
Age group	30–44	0.172	3,878
Age group	45–64	0.192	3,312
Age group	65 and over	0.214	1,455
Education	Junior high or high school	0.208	4,001
Education	Vocational or two-year college	0.201	2,610
Education	University or above	0.137	5,294
Education	Other	0.182	230
Works from home	No	0.185	9,766
Works from home	Yes	0.131	2,369
Desk work	No	0.191	7,528
Desk work	Yes	0.148	4,607
	All ever-users	0.175	12,135

Panel B. Probit coefficients on the regional covariates

	(1) Regional covariates only	(2) Plus individual controls	(3) Column 2 with survey weights	(4) UEA type instead of population density and distance
Log population density, residence (2020)	0.023 (0.024)	0.023 (0.024)	0.031 (0.035)	
Log km to Tokyo Station, residence	0.019 (0.040)	-0.013 (0.042)	0.001 (0.057)	
DID population share, residence (2020)	-0.038 (0.103)	-0.025 (0.105)	-0.083 (0.139)	
Share aged 65+, residence (2022-01)	0.866* (0.478)	0.432 (0.482)	1.693** (0.706)	0.254 (0.443)
Log change in pop. aged 15-64, 2022-26, residence	-0.693 (0.581)	-0.565 (0.577)	0.986 (0.929)	-0.375 (0.546)
UEA center				0.000 (.)
UEA suburb				0.001 (0.032)
Outside UEA				-0.037 (0.101)
Observations	13,487	13,486	13,486	13,486

Panel C. Average marginal effects of the individual controls, model of column (2) of panel B

	Average marginal effect (change in probability of lapsing)
Age (years)	0.002*** (0.000)
Female	-0.005 (0.007)
Vocational or two-year college	-0.009 (0.010)
University or above	-0.041*** (0.009)
Other education	-0.041 (0.041)
Works from home	-0.037*** (0.008)
Desk work	-0.028*** (0.007)
Observations	13,486

Notes: Sample: respondents who have ever used generative AI (N = 13,486). The dependent variable equals one for those who no longer use it. Panel A: inverse probability weighted share of ever-users who no longer use generative AI, by group. Panel B: column (1) includes the regional covariates and prefecture fixed effects only. Columns (2)–(4) add the individual controls. Column (3) uses inverse probability weights. Column (4) replaces population density and distance with UEA type (reference: UEA center). Panel C: average marginal effects from the model of column (2) of panel B. Reference categories are male, junior high or high school, not working from home, and not desk work. Cluster-robust standard errors (municipality of residence) in parentheses. * p < 0.10, ** p < 0.05, *** p < 0.01.

Table A3. Recall validity: adoption shares compared with external benchmarks

Quantity	Unweighted	Survey weights	External benchmark
Adopted by December 2024 and still using (lower bound)	0.177	0.148	0.213
Adopted by December 2024, all lapsed users included (upper bound)	0.254	0.225	0.213
Ever used	0.489	0.439	0.267 / 0.522
Current user	0.412	0.362	

Notes: External benchmarks: Nakagomi, Abe, and Tabuchi (2025) report that 21.3% of a JACSIS-based sample had used generative AI in the twelve months to January 2025 (rows 1–2). The MIC (2025) White Paper reports that 26.7% of individuals had experience with generative AI in fiscal 2024, one year before our survey, and the MIC (2026) White Paper reports 52.2% of adults aged 20 or older in the fiscal 2025 survey, fielded in January and February 2026 (row 3, the two figures separated by a slash).

Table A4. Period merging, movers, and bounds for lapsed users

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Periods 1 and 2 merged	Periods 3 and 4 merged	Excluding residents of under 1 year	Excluding residents of under 5 years	Lapsed users assigned to period 1	Lapsed users assigned to period 4	Lapsed users assigned proportionally
Log population density x P1-2 merged	-0.008 (0.019)						
Log population density x P3	-0.039** (0.020)						
Log population density x P4	-0.034* (0.020)						
Share aged 65+, residence (2022-01)	-0.431 (0.342)	-0.437 (0.343)	-0.401 (0.364)	-0.576 (0.455)	-0.341 (0.321)	-0.400 (0.326)	-0.353 (0.328)
Log change in pop. aged 15-64, 2022-26, residence	0.397 (0.404)	0.406 (0.408)	0.245 (0.423)	-0.068 (0.519)	0.217 (0.380)	0.298 (0.397)	0.299 (0.389)
Log population density x P1		-0.011 (0.022)					
Log population density x P2		-0.004 (0.021)					
Log population density x P3-4 merged		-0.039** (0.018)					
Log population density x 2022/11-2023/12			-0.019 (0.025)	-0.019 (0.030)	-0.038** (0.019)	-0.012 (0.023)	-0.019 (0.022)
Log population density x 2024			-0.006 (0.023)	0.008 (0.028)	0.013 (0.021)	0.006 (0.021)	0.006 (0.021)
Log population density x 2025H1			-0.043* (0.022)	-0.046* (0.027)	-0.022 (0.020)	-0.031 (0.020)	-0.028 (0.019)
Log population density x 2025H2-2026/01			-0.020 (0.022)	-0.012 (0.027)	-0.003 (0.020)	-0.024 (0.018)	-0.018 (0.019)
Observations	86,629	86,629	80,837	64,627	88,751	95,117	92,358

Notes: All columns estimate the full specification of Table 3, column (3) (period dummies, population density and distance to Tokyo interacted with period, municipal covariates, individual controls, prefecture fixed effects, log period-length offset). Columns (1)–(2) merge adjacent periods. Columns (3)–(4) exclude respondents who have lived at the current address for less than one year (7.3% of respondents) or less than five years (27.5%). Columns (5)–(7) assign lapsed users to the earliest period, the latest period, or proportionally to the observed cohort distribution. Cluster-robust standard errors (municipality of residence) in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A5. Respondent composition by UEA type compared with the 2020 Population Census

UEA type	Mean age	Female share	Share aged 65+	University share	Share aged 65+ (Census 2020)	N
UEA center	47.5	0.524	0.210	0.550	0.265	16,510
UEA suburb	48.8	0.510	0.238	0.520	0.288	10,390
Outside UEA	49.3	0.485	0.242	0.402	0.385	687

Notes: Unweighted respondent characteristics. The Census column is the population-weighted share aged 65 or older across municipalities of each UEA type.

Table A6. Representativeness: municipality cell size and alternative survey weights

	(1) Municipalities with 5 or more respondents	(2) Municipalities with 10 or more respondents	(3) Survey weights, 2019 base	(4) Survey weights, 2022 base
Log population density x 2022/11-2023/12	-0.027 (0.026)	-0.018 (0.028)	0.002 (0.039)	-0.001 (0.034)
Log population density x 2024	0.001 (0.024)	0.007 (0.026)	0.010 (0.037)	0.005 (0.032)
Log population density x 2025H1	-0.051** (0.023)	-0.036 (0.025)	-0.008 (0.037)	-0.009 (0.032)
Log population density x 2025H2-2026/01	-0.003 (0.023)	0.014 (0.025)	-0.011 (0.035)	-0.011 (0.030)
Share aged 65+, residence (2022-01)	-0.343 (0.359)	-0.043 (0.383)	-0.622 (0.573)	-0.527 (0.502)
Log change in pop. aged 15-64, 2022-26, residence	0.518 (0.457)	0.759 (0.467)	0.212 (0.719)	0.127 (0.617)
Observations	81,914	76,467	86,629	86,629

Notes: All columns estimate the full specification of Table 3, column (3) (period dummies, population density and distance to Tokyo interacted with period, municipal covariates, individual controls, prefecture fixed effects, log period-length offset). Columns (1)–(2) restrict the sample to respondents living in municipalities with at least five or at least ten respondents. Columns (3)–(4) apply the inverse probability weights based on the 2019 and the 2022 Comprehensive Survey of Living Conditions (the main text uses the 2022 weights). Cluster-robust standard errors (municipality of residence) in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A7. Conditional population-density and distance gradients by age group

	(1)	(2)	(3)	(4)
	Age under 30	Age 30-44	Age 45-64	Age 65 and over
Log population density x 2022/11-2023/12	-0.040 (0.042)	-0.022 (0.043)	0.014 (0.047)	-0.112 (0.072)
Log population density x 2024	-0.032 (0.041)	-0.007 (0.038)	0.008 (0.043)	-0.038 (0.061)
Log population density x 2025H1	-0.037 (0.040)	-0.063* (0.037)	-0.023 (0.040)	-0.118** (0.060)
Log population density x 2025H2-2026/01	-0.078** (0.038)	-0.032 (0.038)	-0.017 (0.040)	0.003 (0.060)
Log km to Tokyo x 2022/11-2023/12	0.017 (0.053)	-0.107* (0.056)	0.030 (0.062)	-0.022 (0.095)
Log km to Tokyo x 2024	0.039 (0.050)	-0.102* (0.058)	0.025 (0.061)	0.093 (0.096)
Log km to Tokyo x 2025H1	0.103** (0.051)	-0.118** (0.055)	0.025 (0.058)	0.024 (0.087)
Log km to Tokyo x 2025H2-2026/01	0.071 (0.050)	-0.044 (0.056)	0.029 (0.058)	0.141 (0.092)
Share aged 65+, residence (2022-01)	-1.420** (0.652)	0.309 (0.613)	-0.448 (0.640)	-1.405 (1.034)
Observations	13,953	23,379	28,090	21,191

Notes: Each column estimates the full specification of Table 3, column (3) (period dummies, population density and distance to Tokyo interacted with period, municipal covariates, individual controls, prefecture fixed effects, log period-length offset) on the respondents of one age group. Cluster-robust standard errors (municipality of residence) in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A8. Age at risk and the period-length offset

	(1)	(2)
	Excluding respondents under 15 in November 2022	Without the period-length offset
Log population density x 2022/11-2023/12	-0.024 (0.024)	-0.025 (0.024)
Log population density x 2024	-0.005 (0.023)	-0.006 (0.022)
Log population density x 2025H1	-0.040* (0.021)	-0.041* (0.021)
Log population density x 2025H2-2026/01	-0.020 (0.022)	-0.022 (0.022)
Share aged 65+, residence (2022-01)	-0.453 (0.345)	-0.469 (0.344)
Log change in pop. aged 15-64, 2022-26, residence	0.417 (0.412)	0.427 (0.411)
Observations	86,338	86,629

Notes: Both columns estimate the full specification of Table 3, column (3) (period dummies, population density and distance to Tokyo interacted with period, municipal covariates, individual controls, prefecture fixed effects, log period-length offset), except as indicated in the column heading. Cluster-robust standard errors (municipality of residence) in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A9. Benjamini–Hochberg adjustment for the regional coefficients of the main specification

Coefficient	p-value	Rank	Benjamini–Hochberg q-value
Log population density × 2025H1	0.051	1	0.310
DID population share	0.056	2	0.310
Share aged 65+	0.173	3	0.499
Log km to Tokyo × 2025H2–2026/01	0.222	4	0.499
Log change in pop. aged 15–64	0.299	5	0.499
Log population density × 2022/11–2023/12	0.300	6	0.499
Log population density × 2025H2–2026/01	0.317	7	0.499
Log km to Tokyo × 2025H1	0.718	8	0.805
Log km to Tokyo × 2022/11–2023/12	0.776	9	0.805
Log population density × 2024	0.779	10	0.805
Log km to Tokyo × 2024	0.805	11	0.805

Notes: The eleven regional coefficients of Table 3, column (3), ranked by unadjusted p-value. The q-values control the false discovery rate at the stated level.

Table A10. Desk work and working from home by UEA type, and their population-density gradients**Panel A. Shares among employed respondents (survey weights)**

UEA type	Desk work	Works from home	Works from home, desk workers	Works from home, desk workers in firms under 100	Works from home, desk workers in firms of 100+	Firm of 100+	Weighted N
UEA center	0.442	0.188	0.280	0.281	0.281	0.485	11,605
UEA suburb	0.398	0.188	0.316	0.349	0.293	0.456	6,703
Outside UEA	0.388	0.089	0.145	0.166	0.114	0.402	774


Panel B. Average marginal effects of log residence population density from probit models

	(1) Works from home, desk workers	(2) Works from home, desk workers with job controls	(3) Desk work, all employed with job controls
Log population density, residence (2020)	0.0096* (0.0054)	0.0084 (0.0052)	0.0246*** (0.0039)
Observations	9,454	9,454	19,615

Notes: Panel A: inverse probability weighted shares among employed respondents. Panel B: average marginal effects of log residence population density from probit models with prefecture fixed effects. Column (1) is estimated among desk workers, column (2) among desk workers adding employment status, firm size, education, age, and sex, and column (3) models desk work among all employed respondents with the same controls. Cluster-robust standard errors (municipality of residence) in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.


Figures

Figure 1. Current use of generative AI by secondary medical area




Notes: Inverse probability weighted share of current users, empirical-Bayes smoothed, for the 333 secondary medical areas with respondents.

Figure 2. Cumulative adoption by secondary medical area at the end of each period




Notes: Panels show the share of residents who had adopted by December 2023, December 2024, June 2025, and January 2026 (inverse probability weighted, empirical-Bayes smoothed).

Figure 3. Cumulative adoption curves by quintile of residence population density



Notes: Inverse probability weighted cumulative adoption at the end of each period, lapsed users excluded. Quintiles of municipal population density are defined over municipalities.

Figure 4. Local Moran clusters of the municipal current-use rate



Notes: Local indicators of spatial association (Section 3.5) for the empirical-Bayes smoothed current-use rate across 862 municipalities with at least five respondents, using a 50 km distance band, 199 conditional permutations, and the 5% significance level. The map classifies 115 municipalities as high-high, 16 as low-low, 14 as high-low, and 69 as low-high.