CAGE

# Working Paper

828/2026
November 2026

# Who Pays for AI?

Nikhil Datta,
Peter John Lambert,
Patrick Schneider

ISSN: 2978-0276
Grant number: ES/Z504701/1

UNIVERSITY OF WARWICK
Economics

UKRI Economic and Social Research Council

# Who Pays for AI?*

Nikhil Datta Peter John Lambert Patrick Schneider

November 5, 2026

**Abstract**

We document who pays for AI tools using a panel of millions of open banking and cardholder accounts from 2022 to mid-2026, covering billions of transactions. We find that subscriptions to products like Claude and ChatGPT have roughly doubled every 12 months since their introduction, reaching one in twenty UK adults by mid-2026. Adoption is five times higher for under 30s than over 60s, 27 percent higher for men, and grows 1.8 times larger from the bottom to the top quintile of earnings. While ChatGPT held 74 percent of direct subscription spending in late 2025, its share fell 21 percentage points by mid 2026 as Claude surged from 7 to 29 percent of direct spending, driven partly by growth in users who work in technical and professional roles and pay for the most expensive plans. High-frequency banking data are well suited to study the rapidly evolving market for paid consumer AI products.

**Keywords:** artificial intelligence; technology adoption; subscriptions; transaction data; occupational exposure; product competition
**JEL Codes:** O33, D12, L86, J24, C81, E21

*Datta: University of Warwick and CEP, London School of Economics; Lambert: University of Warwick and London School of Economics (corresponding author, email: Peter.Lambert@warwick.ac.uk); Schneider: Imperial Business School. This work was supported by the Centre for Competitive Advantage in the Global Economy (CAGE), funded by the Economic and Social Research Council (ESRC), Grant Number [ES/Z504701/1]. All errors are our own. Up-to-date monthly indices used in this paper can be viewed and downloaded at https://whopaysfor.ai.

# 1 Introduction

Within three years of ChatGPT's public release, between one-third and three-fifths of adults had experimented with generative artificial intelligence (GenAI).[1] This rate of diffusion is unprecedented for a new technology, driven in large part by the costless ability to experiment with free tools like ChatGPT and Claude. In this paper we document the market transactions underpinning this diffusion, identifying users who value these tools enough to pay for them.

Existing data sources are typically not suited to measuring paid adoption at scale. Household surveys rely on self-reported usage that includes users with a one-off curiosity along with those who make a recurring financial commitment. The rich usage data from single AI providers offers a lens into the nature of usage (Chatterji et al., 2025; Handa et al., 2025), but cannot study competing sellers, non-adopters, or the known demographics of users. Bank and card transaction records overcome both limitations by allowing us to observe actual financial commitments across competing providers.[2]

To study the diffusion of paid AI tools, we combine anonymized Open Banking records with credit and debit card transactions covering 14.2 million accounts, with information on account holders' age, gender, and residential location. Our data span multiple payment networks, while the Open Banking component allows us to observe earnings directly and recover employer, industry, and occupational information that is typically absent from card transaction data. This allows us to examine how paid adoption varies with both income and the work individuals perform. We also account for subscriptions concealed within generic app-store charges, which represent approximately two-fifths of estimated active subscriptions.

Our measurement approach identifies monthly subscriptions to AI providers such as OpenAI, Anthropic, Gemini, or Midjourney. In most cases, the transaction description directly identifies the provider. This allows us to record the date a consumer crosses the threshold from free use to a paid commitment. We can also follow subscription spells, cancellations, renewals, and switches across different AI tools in the market.

While this paper's main contribution is descriptive, we also introduce a novel methodological contribution which allows us to recover occupation classifications from finan-

1. Based on representative surveys of US, EU, and UK users (Bick, Blandin, and Deming, 2026; Pew Research Center, 2026; Ofcom, 2026; Bick et al., 2026).

2. The use of administrative banking and card transactions collected in a panel data structure has held real promise for the real-time measurement of household consumption, income dynamics, and other macroeconomic shocks (Ganong and Noel, 2019; Cox et al., 2020; Baker and Kueng, 2022; Chetty et al., 2024).

cial transactions data. The ability to expand our accountholder information to cover the jobs they hold is important for the study of AI adoption given the propensity of the literature to closely link the tasks workers perform with their AI usage (Eloundou et al., 2024). Occupations may therefore be a first-order confounder when assessing the role of demographics such as gender, age, or income. The extent to which a worker's job predicts personal financial commitments also helps us understand the barrier between personal and professional drivers of adoption and whether employers or employees foot the bill.

To measure occupations from financial transactions, we first build a list of transactions which unambiguously reveal an account holder's job, including payments to professional regulators, royal colleges, chartered institutes, and other occupation-specific dues.[3] We use these unambiguous signals to build a supervised classification model which uses high-dimensional representations of account transactions to assign a probabilistic classification of each account's likely occupation.[4] These probabilities are combined with demographics and occupation-distribution priors from official statistics, providing a mapping of millions of accounts into 26 broad occupational groups in a way that matches population moments and cross-tabulation information closely.

Another methodological contribution of this work that extends beyond our specific study is the imputation of transactions intermediated by app-stores. Because the price and billing cycles of the AI tools we study are known, we can screen out most app-store transactions right away. For the remaining potential AI transactions with a generic app-store description, we again leverage high-dimensional representations of all consumer transactions and a supervised classification model. The model in this case is trained to identify masked AI transactions which do not have a generic app-store description. We can further inform this classification by using discontinuities in the subscription behaviour around things like new model releases to fine tune these classifications. We show that this approach is reliable, changes the demographic interpretations very minimally, but does have a sizeable impact on the overall level of estimated consumers of paid AI products.

The analysis in this paper yields three distinct sets of descriptive findings.

First, we document the rapid diffusion of paid AI. The number of active subscribers has roughly doubled every 12 months since early 2024, reaching 5.1 percent of UK

3. Examples include the General Medical Council, Solicitors Regulation Authority, Royal College of Nursing fees.

4. We omit AI transactions at all stages of this classification pipeline, so our main descriptive findings are not contaminated. Those working in other domains of household finance would do well to include these transactions, as they too carry real signal in general about a person's job.

adults (2.8 million people) by June 2026 with an annualized consumer spend exceeding £1.0 billion. A few general-purpose AI tools account for roughly nine in ten active subscriptions, with specialist tools such as Midjourney making up the balance. By mid 2026, more UK adults were signing up for paid AI for the first time than were joining major streaming platforms such as Netflix and Spotify. This growth in overall spending between January 2024 and July 2026 has been predominantly driven along the extensive margin, with 88 percent of the rise coming from more customers, and the remaining 12 percent coming from higher average monthly spending per subscriber.

Second, we find that the adoption of paid AI is highly uneven across age, gender, and earnings. Paid adoption is nearly six times higher among adults aged 20–29 than among those aged 60 or higher, a pattern that persists after controlling for income and demographics. Men are 27 percent more likely to hold an active subscription than women and spend 23 percent more per month when they do, reflecting greater male uptake of expensive "power-user" plans. Across the earnings distribution, cumulative adoption is 78 percent higher in the top fifth of earners than in the bottom fifth. However, adding controls step by step reveals that nearly two-thirds of the raw relationship between salary and paid AI adoption is explained by differences in occupation, industry, location, age, and gender; once these characteristics are accounted for, the relationship between salary and AI adoption is no stronger than for mainstream consumer subscriptions such as Netflix and Spotify. At the same time, early demographic and earnings gaps have narrowed among new subscribers: between early 2023 and mid-2026, the male-to-female ratio among first-time adopters fell from 3.0 to 1.2, and the earnings premium of new adopters relative to the average salaried worker fell by more than half.

What people do at work is a central determinant of whether they pay for AI out of pocket. Across professions identified from professional body membership and registration payments, active direct subscription rates vary twelvefold—from roughly 20 percent among IT professionals and 13–16 percent among university lecturers, finance professionals, project managers, and marketers to under 3 percent among electricians, nurses, and social care workers. Our transaction-based adoption rates align closely with both US survey measures of workplace AI usage ($r = 0.8$ to 0.9; Bick, Blandin, and Deming, 2026) and task-based measures of occupational exposure to AI ($r = 0.8$; Eloundou et al., 2024). Occupational AI exposure remains a strong predictor of personal adoption even after accounting for individual take-home pay, age, gender, and region. Geographically, the rate of paid AI adoption in London (6.8 percent) is nearly twice the rest of Great Britain's (3.6 percent), with pronounced local clusters in university and knowledge-economy centers such as Cambridge, Oxford, and Edinburgh.

Third, while OpenAI's ChatGPT dominated the consumer market through 2024 and most of 2025, competition intensified sharply from late 2025 as Anthropic's Claude captured a rapidly growing share of both subscribers and consumer AI spending. Between December 2025 and July 2026—following the release of Claude Opus 4.5 and the rapid spread of the coding tool Claude Code—Claude's share of direct consumer spending on AI subscriptions quadrupled from 7.4 percent to 29.4 percent, while ChatGPT's direct spending share fell by 21 percentage points from 73.5 percent to 52.5 percent. ChatGPT still leads by a wide margin in active assistant subscribers (71.2 percent versus 27.1 percent in July 2026), and Claude's spending share exceeds its subscriber share because 6.2 percent of its subscribers pay for £80+ monthly power-user plans, compared with just 0.3 percent on ChatGPT—a twenty-to-one gap in power-user plan uptake.

The large shift in market share across AI providers can be further broken down by the kinds of jobs users perform: software and IT, engineering, management consultants, finance professionals, and academic researchers are more likely than average to choose Claude, while school teachers, healthcare workers, and HR staff disproportionately favor ChatGPT. Most subscribers pay for only one general-purpose assistant at a time, while 4.2 percent of assistant subscribers pay for two or more tools in the same month. Claude's rapid growth in 2026 reflects both entry by new subscribers, especially following major model releases, and switching from ChatGPT.

Our study contributes to three main areas of research. First, we build on a long literature measuring technology diffusion (Griliches, 1957; David, 1990; Bresnahan and Trajtenberg, 1995), across both economies (Caselli and Coleman, 2001; Comin and Hobijn, 2004, 2010) and households (Goolsbee and Klenow, 2002; Rogers, 2003). More directly, our study complements studies documenting the rapid diffusion of generative AI (Bick, Blandin, and Deming, 2026; Humlum and Vestergaard, 2025b) and its workplace productivity effects (Noy and Zhang, 2023; Peng et al., 2023; Dell'Acqua et al., 2023; Brynjolfsson, Li, and Raymond, 2025). A large number of these studies have focused predominantly on the labor market effects, from task exposure (Eloundou et al., 2024) to earnings and employment (Humlum and Vestergaard, 2025a; Brynjolfsson, Chandar, and Chen, 2026; Lambert and Schindler, 2026). Our work studies the adoption of paid AI from the consumer side, but we link accounts to employment status, employer, industry, and occupation, suggesting household adoption is strongly tied to the benefits of these tools at work. Similar studies of diffusion have been conducted using the internal logs from the suppliers of such tools (Chatterji et al., 2025; Handa et al., 2025). We complement survey and single-provider studies by observing out-of-pocket subscription payments continuously across all competing sellers, in a representative consumer panel. Our study also provides a direct connection

between consumer spending on AI and workplace task exposure.

Second, we contribute to household finance and empirical macroeconomics using bank and card transaction records (Baker and Kueng, 2022; Chetty et al., 2024). Although transaction datasets have transformed the measurement of household spending (Gathergood et al., 2021), income dynamics (Ganong and Noel, 2019; Hacıoğlu-Hoke, Känzig, and Surico, 2021), and fiscal stimulus (Cox et al., 2020), their application to labor-market questions has historically been limited by the lack of occupation and industry identifiers in banking records, a gap our paper addresses.

Third, our study connects to studies in industrial organization in markets for digital goods (Goldfarb and Prince, 2008; Einav and Levin, 2014; Goldfarb and Tucker, 2019), subscription pricing (Mussa and Rosen, 1978; DellaVigna and Malmendier, 2006; Grubb and Osborne, 2015), and consumer switching (Klemperer, 1995; Chen and Hitt, 2002). Though our work is primarily descriptive, our results speak to multi-homing, switching frictions, and offer some stylised facts about the market structure for paid consumer AI tools.

The remainder of the paper proceeds as follows. Section 2 describes our data and measurement approach; Section 3 examines aggregate growth; Section 4 analyzes how adoption varies across age, gender, income, occupation, and geography; Section 5 studies competition between providers; and Section 6 concludes.

## 2 Data, Measurement, and Validation

This section describes our transaction dataset and how we use it to measure paid AI adoption and individual worker characteristics across the United Kingdom. We first describe the card and bank transaction records that make up the data, and our sample selection approach. Using these data, we show how to identify subscriptions to AI products, via direct payments to providers and indirect ones through app store subscriptions, and to measure labor market characteristics like salary, employer, and occupation. We then detail our weighting procedure so estimates match the UK adult population. Both the AI subscription measure and occupation classification involve inference steps, and we conclude the section with a validation exercise showing that the classifiers we use are very reliable.

## 2.1 Payment Records and Sample Selection

We analyze de-identified UK consumer bank and card records from a financial data aggregator covering 14.2 million accounts between January 2022 and August 2026 (Table 1 and Appendix C). The dataset combines transactions from two distinct types of accounts. The first are 'open-banking' current accounts, capturing debit card purchases, direct debits, bank transfers, and incoming salary deposits.[5] The second are individual credit and debit card accounts.

For each transaction in either account type, we observe an account identifier, account holder age, gender, and residential postcode, as well as the transaction date, amount, currency, payment channel, merchant category code of the seller, and anonymized statement text description. For the open banking subset we additionally observe income data, which allows us to identify salary, employer, and occupation.

We use the combined dataset for our headline population estimates where possible, relying on the open banking subset only where necessary, to use its richer labor market information.

We distinguish between active 'main' accounts and secondary ones so that inactive accounts do not artificially lower per-account adoption rates.[6] We define an active main account as one with an average of at least 20 non-AI transactions per month and median monthly non-AI spending of at least £100 over the preceding twelve months. This filter yields 457,072 active accounts in June 2026 (and 390,404 salaried open-banking accounts spanning 2022–2026; Table 1, Panel C).[7]

## 2.2 Identifying AI Subscriptions

Consumers can subscribe to AI products either directly, by signing up via the websites of providers, or indirectly, channeling the payment through app stores. Our treatment of these different channels necessarily differs.

We identify direct payments by text matching the transaction description to known sellers of AI products, and transaction amounts to known subscription tiers (Table 1,

5. Under UK open-banking regulations, consumers can connect their everyday bank current accounts and credit cards to financial management apps, allowing us to observe anonymized bank statements including incoming salary payments, direct debits, and card purchases; see Appendix C.

6. This is a common approach in existing work dealing with transaction data in which accounts, rather than individuals, are the unit of analysis (e.g. Ganong and Noel, 2019; Cox et al., 2020; Baker and Kueng, 2022; Chetty et al., 2024).

7. The main account screen helps to address the fact that our unit of analysis is the account, rather than person. In our benchmark results we treat main accounts as if they each represented a distinct person. We explore sensitivity to this assumption in Appendix C.

Panel A). The set of potential sellers is in the directory of consumer AI products detailed in Appendix A and Table A1. This directory includes all the main general purpose AI assistants (e.g. ChatGPT, Claude, Gemini, Microsoft Copilot, Perplexity, DeepSeek, Mistral, Grok, and Poe), and excludes pre-existing software packages to which AI features have been added to existing products (e.g. Microsoft 365, Canva, and Adobe).[8]

Consumers also subscribe to AI assistants indirectly, via mobile apps on the iOS and Android platforms. These platforms mask the content of the transaction with generic labels such as `apple.com/bill` and `google play`. To account for indirect billing, we infer AI subscriptions in two steps, detailed in Appendix B. First, we track recurring transactions that match the exact prices for different tiers of major AI plans, and measure the excess of these payments relative to the pre-AI trend.[9] Second, among users paying these recurring fees, we predict the likelihood of their subscribing to AI based on demographics and non-AI spending patterns, using a model trained on users who pay directly. We scale these probabilities so that they sum to the measured excess, and use them as weights in all analyses that include app-store subscriptions. Incorporating these indirect payment channels as well as direct payers is important because younger, female, and lower-salaried users are more likely to pay through these channels than directly.

Having identified individual transactions, we link repeated payments into subscription spells. A spell begins at its first payment, and is treated as canceled when (a) there is no new payment within 65 days of the last one, and (b) the account remains active (Appendix A).

## 2.3 Measuring Salaries, Employers, and Occupations

Our labor market statistics come from the open-banking sample of current accounts from which we use data on incoming payroll deposits and payments to professional bodies to measure income, employer industry and occupation (Table 1, Panel C, and Appendix D). Among the 1.7 million open-banking accounts we observe, 491,228 receive observable salaries. For each of these, we measure take-home salary from the regular amount deposited. We also use the text information in the transaction to

8. When automated privacy filters replace personal names such as "Claude" with `<name>`, or when payment processors truncate `claude.ai` to `claud`, we identify the payment using the remaining web domain and exact subscription price, and drop any ambiguous transactions.

9. We validate that this measure indeed captures actual AI subscriptions by showing that the excess jumps around major release dates and co-moves closely with direct OpenAI subscriptions.

identify the employer. This allows us to identify an employer industry[10] by matching employer names to business records at Companies House for private companies and public-sector directories for organizations such as National Health Service trusts, local authorities and schools. We are able to match 65% of the sample with observable salaries into an industry (Appendix Table A7).

While identifying salary and employer industry is relatively straightforward, finding a worker's occupation is much more difficult because standard banking records do not record it. To overcome this data limitation, we classify a subset of workers into occupations[11] using evidence of their membership in professional organizations. For the remainder, we infer likely occupation using a machine-learning model trained on the characteristics and spending behavior of the people whose occupations we observe (Appendix D).

To build our occupation classification we start by compiling a list of 66,000 organizations linked to specific occupations in the UK. These include 466 statutory professional regulators, royal colleges, chartered institutes, and occupation-specific trade unions. Examples include the General Medical Council for doctors and the Solicitors Regulation Authority for solicitors.

We use outbound payments to these organizations to find their members, and therefore members of the occupations they represent. In all, 46,419 of the salaried open-banking sample accounts have such direct payments, identifying their membership of 65 different occupations. We supplement this occupation measure with a second—when an account receives income from an employer whose workforce is concentrated in a particular occupation (such as airlines or law firms), we count a person as having that occupation. This identifies a further 24,180 salaried accounts for a total of 70,599 in which we directly observe their occupation.

To extend this occupational measure across the remaining observations, we train a machine learning model on this direct sample to classify occupation based on employer industry, salary, age, gender, local area and non-AI spending.[12] We use this model to assign likelihoods for the remaining open-banking accounts having an occupation in 26 broad groups.

10. We classify industries at the 'section' level of the Standard Industrial Classification (SIC 2007); there are 21 possible sections.

11. We classify occupations based on the UK Standard Occupational Classification (SOC 2020).

12. AI spending is excluded from the spending measure so the model never uses information about AI adoption to infer occupation.

## 2.4 Reweighting to the UK Adult Population

All aggregate results in this paper are reweighted to reflect the demographic mix of the adult population of the UK (55.5 million people 18 and over). Specifically, we construct monthly account weights adjusting for age, gender, region, after-tax income, and local neighborhood using public data from the Office for National Statistics (ONS) and HM Revenue and Customs (HMRC). Subsample analysis uses adjusted weights that reflect the sample's characteristics relative to the same aggregate population. More detail is available in Appendix C.

## 2.5 Validation using External Benchmarks

We assess the reliability of our data and inferred measures against external benchmarks and through out-of-sample validation exercises, with results reported in Table 1, Panel B, and Appendix Figure A2. We validate the data by checking measured income moments against the official ONS Annual Survey of Hours and Earnings (ASHE). Our measured net pay percentiles in the middle of the distribution (30th to 90th) closely match the official series. Similarly, we capture the January 2024 cut in National Insurance contributions: take-home pay increases by an average of £31, compared to the £33 implied by the income tax schedule.

We can further validate our predictive models by checking their out of sample fit. To test the indirect measure of subscriptions we hide merchant names for direct subscribers and run the model to separate AI subscribers from those paying the same amounts for other things. The exercise shows the classifier works much better than random, measured with an area-under-the-curve (AUC) metric of 0.9.[13] A similar exercise with the occupation classifier yields an AUC of 0.87 across professions, suggesting it performs far better than random assignment.

# 3 Growth of Paid AI Across the Population

## 3.1 Share of Adults Paying and Total Spending

We begin by examining how rapidly paid consumer AI has spread across the adult population in the United Kingdom.

13. We use the Area Under Curve (AUC) metric to report how well a given model performs. The AUC is (linearly) equivalent to the Wilcoxon-Mann-Whitney U-statistic and represents the probability that a classifier gives a higher rank to a randomly selected positive observation than to a randomly selected negative observation. A value of 0.5 means prediction is as good as random, and a value of 1 implies perfect classification; see Appendix B.

Figure 1 plots active paid AI subscribers per 1,000 UK adults from January 2023 to June 2026. From a negligible base at the start of 2023, paid consumer AI has spread rapidly across the British adult population, roughly doubling every 12 months since early 2024. By June 2026, 50.6 per 1,000 UK adults hold an active paid subscription to an AI product—equivalent to about 2.8 million people. Roughly three-fifths of active subscriptions (29.0 per 1,000 UK adults, or 1.6 million people) are billed directly to bank accounts or cards on provider websites, while the remainder are billed through app stores or other payment processors (Appendix Table A2).[14]

In pound terms, UK consumer spending on AI products has grown from almost nothing, prior to February 2023, into a billion-pound annual market by mid-2026 (Appendix Figure A6a). During 2024, UK consumers spent £3,442 per 1,000 adults on paid AI; in 2025, annual spending more than doubled to £8,387 per 1,000 adults. Growth accelerated further in the first half of 2026: by July 2026, monthly spending reached £1,600 per 1,000 adults, equivalent to an annualized spending rate exceeding £1.0 billion across the UK adult population.

Comparing our 5.1% paid subscription rate with population surveys highlights the sharp difference between costless experimentation and paying out of pocket. UK surveys by Ofcom, Deloitte, and the Digital Regulation Cooperation Forum find that 35% to 62% of UK adults have tried a generative AI tool (Digital Regulation Cooperation Forum, 2025; Ofcom, 2026; Deloitte, 2026), similar to the 45% to 50% trial rates reported in US and European surveys (Bick, Blandin, and Deming, 2026; Gambacorta, Jappelli, and Oliviero, 2025; Pew Research Center, 2026; Bick et al., 2026). Leading providers offer competent free versions of their models, and roughly nine in ten users who have tried AI in any form do not currently pay for a subscription. Dividing our 5.1% adult subscriber share by the 35%–62% who have ever tried AI implies an 8% to 15% subscription rate among ever-users; this is in line with OpenAI's global disclosure of 50 million consumer subscribers out of 900 million weekly active users (5.6%) as well as US bank-statement estimates (OpenAI, 2026; Bank of America Institute, 2026; PNC Financial Services Group, 2026).

## 3.2 New Subscribers and Drivers of Spending Growth

Total consumer spending on AI can grow either because more adults pay in a given month or because active subscribers spend more per person. Between January 2024 and July 2026, the share of UK adults billed for AI in a given month rose nearly sixfold from 1.0% to 5.8%, whereas average monthly AI spending per active subscriber rose

14. Appendix Table A2 and Appendix Figure A1 show that our national estimates are robust across alternative sample definitions and reweighting methods.

much more modestly from £22 to £28 on direct billing, representing roughly 0.4% to 0.5% of active subscribers' total monthly card spending—about the price of a single streaming subscription (Appendix Figure A6c). Separating total spending growth between January 2024 and July 2026 into growth in the number of subscribers versus growth in spending per subscriber shows that nearly nine-tenths (88%) of total market growth comes from more adults paying each month, while the remaining 12% comes from higher average spending per subscriber as some users upgrade to expensive power-user plans or hold multiple subscriptions.[15]

# 4 Who Pays for AI?

Who crosses the threshold from free use to paying out of pocket for AI? In this section we present our main results, showing patterns of AI adoption across Age, Gender, Income, Occupation, and Geography. While there are stark differences across each of these, two main findings emerge. First, gender and earnings gaps reported elsewhere appear to be narrowing over time as a broader cross-section of the population adopts paid AI. Second, the strong income gradient of adoption is weakened a lot by controlling for occupation and industry, suggesting adoption may be more associated with exposure in work settings.

Figures 2 and 3 show how paid adoption varies across age, gender, and net take-home salary, and how the profile of first-time subscribers has broadened over time.

## 4.1 Age

Panel (a) of Figure 2 plots the share of adults who have ever paid for AI since November 2022 across five age bands, comparing raw population-weighted rates (navy bars) with adjusted rates that hold monthly net salary, gender, and UK region constant (orange bars). Paid AI adoption falls nearly sixfold across the adult life cycle, from 17.9% among adults aged 20 to 29 to 3.0% among adults aged 60 and over (a ratio of 5.9 to 1), and remains nearly fivefold higher at ages 20 to 29 than at ages 60 and over after accounting for salary, gender, and region (18.3% versus 3.8%, a ratio of 4.8 to 1). This strong youth pattern mirrors earlier waves of personal computer and internet adoption (Goolsbee and Klenow, 2002; Goldfarb and Prince, 2008; Deming and Noray, 2020).

15. Because total AI spending per adult equals the share of adults paying multiplied by average AI spending per active payer, we decompose the log change in monthly spending per adult into the sum of the log changes in these two components (84% from the share of adults paying on direct web subscriptions alone); see Appendix E.

## 4.2 Gender

Panel (b) of Figure 2 tracks monthly active paid AI subscribers per 1,000 UK adults separately for men and women from February 2023 through July 2026. Men subscribe at higher rates throughout the sample, reaching 62.2 active payers per 1,000 men versus 48.6 per 1,000 women in July 2026—a 27% gender gap (a ratio of 1.3)[16]. Men also spend 23% more per month when they subscribe (£26 versus £21), reflecting higher male uptake of £80+ power-user plans. This gender gap in adoption (see e.g. Aldasoro et al., 2024; Humlum and Vestergaard, 2025b) is, however, narrowing: the male-to-female ratio among first-time monthly adopters fell from 3.0 to 1 in early 2023 to 1.2 to 1 by the first half of 2026.

## 4.3 Income

Figure 3 examines how paid AI adoption varies with monthly net take-home pay across salaried open-banking accounts. Unlike transaction studies that proxy for income using total card spending, which mixes income with differences in card usage and saving behavior across the income distribution (Dynan, Skinner, and Zeldes, 2004), our open-banking records measure net salary deposits directly.[17] Across monthly net take-home salary brackets (Figure 3a), cumulative paid AI adoption is 88% higher in the top bracket than in the bottom bracket before adding controls, rising from 10.0% among workers taking home under £1,500 per month to 18.8% among those taking home over £4,000 per month, and this difference persists after accounting for age, gender, and UK region.

As with gender, the earnings gap between new subscribers and the average worker has narrowed as the market has grown. First-time payers in early 2023 earned 1.7 times the average salaried worker's take-home pay, compared with 1.3 times the average by mid-2026.

Panel (b) of Figure 3 asks how much of this overall relationship with income reflects differences in who works in which jobs and locations versus a direct effect of income, and compares paid AI adoption against everyday consumer streaming subscriptions (Spotify and Netflix). We use a linear probability model to relate adoption to log monthly take-home pay across four specifications that add controls step by step.[18]

16. The gap is larger on direct web payments alone because women purchase a higher share of subscriptions through mobile app stores.

17. On the same open-banking accounts, the estimated relationship between paid AI adoption and log card spending is roughly one-third smaller than the relationship with log net salary, showing that proxying for income with card spending understates the true relationship with earnings.

18. Adoption is a binary outcome measuring whether a person has ever adopted AI.

Before adding controls, a one-unit increase in log monthly take-home pay is associated with a 5.6 percentage point increase in paid AI adoption, which is larger than the corresponding income relationships for both Spotify (4.0 percentage points) and Netflix (3.3 percentage points).

Accounting for age and gender reduces the paid AI slope to 5.0 percentage points, and accounting for local authority district reduces it to 4.4 percentage points. Crucially, accounting further for occupation and employer industry cuts the remaining paid AI income slope by more than half to 2.0 percentage points—showing that nearly two-thirds (64%) of the raw income gap in paid AI adoption is explained by occupation, industry, location, and demographics, and leaving a conditional income response slightly smaller than those of Spotify (3.1 percentage points) and Netflix (2.6 percentage points).[19]

## 4.4 Occupation

Paid AI adoption also varies sharply across occupations. Figure 4 reports active June 2026 paid AI adoption across 34 professions identified directly from registration fees and membership dues paid to professional regulators, royal colleges, and trade unions.[20]

Figure 4 presents rates of AI uptake as of June 2026 by occupation. Adoption spans a twelvefold range across the displayed professions (and more than ninefold across professions with at least 150 active accounts). Adoption is highest among IT professionals (20.3%), lecturers (16.0%), and finance & insurance professionals (15.8%); and lowest for nurses and midwives (2.5%), firefighters (2.5%), and social care workers (1.7%). Including mobile app-store subscriptions raises subscription rates across the board while leaving the ranking of occupations almost unchanged (with a rank-rank correlation of 0.9; Appendix Table A9).

Our measures of AI adoption across occupations also accord with existing measures of occupation-based exposure to, and on-the-job adoption of, AI. Figure 5 compares

19. As shown in Appendix D, this two-thirds reduction is robust both when restricting occupational controls to directly observed professional membership and employer labels and when estimating the regressions on directly billed web subscriptions alone. Within named employers with at least 50 salaried accounts in our dataset (where job seniority and occupation vary across pay grades inside the firm), active direct paid AI adoption rises from 5.9% in the bottom fifth of earners within the firm to 10.0% in the top fifth, whereas Netflix and Spotify adoption are flat across within-firm pay levels.

20. As described in Section 2 and Appendix D, we identify a worker's specific profession when they pay registration fees or membership dues to one of 466 occupation-specific bodies (such as the General Medical Council or Solicitors Regulation Authority). Figure 4 plots the 34 professions meeting our minimum sample-size rules on card accounts where privacy filters do not obscure merchant names; Appendix Tables A8–A10 report full results across all 65 professions and across all payment channels.

12-month active paid assistant adoption across 51 professions against two external benchmarks (Appendix Table A10).

In Panel (a), UK paid AI adoption aligns closely with US survey measures of self-reported workplace AI usage from Bick, Blandin, and Deming (2026), explaining $R^2 = 69.4\%$ of the employment-weighted variation across the 51 professions (and $R^2 = 90.0\%$ across all 26 broad UK occupational groups covering the entire workforce).

In Panel (b), UK paid AI adoption also aligns closely with the task-based GPT-4 exposure index of Eloundou et al. (2024), which measures the share of an occupation's tasks that can be sped up by large language models ($R^2 = 63.1\%$): moving up by one standard deviation in occupational task exposure is associated with a 5.1 percentage point higher rate of paid AI adoption. In individual-level regressions across our open-banking sample (Appendix Table A10, Panel C), both workplace survey usage and GPT-4 task exposure remain strong, statistically precise predictors of personal out-of-pocket AI adoption even after accounting for age, gender, region, and individual take-home pay.

## 4.5 Industry

Grouping salaried workers by the UK Standard Industrial Classification (SIC 2007) industry of their employer reveals a parallel pattern across sectors (Appendix Figure A5a). Among salaried workers linked to an employer industry, the average cumulative adoption rate is 14.2%.[21] Within this group, adoption across major sectors is highest in software and information technology (27.0%), advertising and market research (26.6%), publishing and broadcasting (24.2%), education (18.2%), and management consulting (16.6%), and lowest in human health (12.2%), retail (11.8%), and public administration and defense (11.6%).

These industry-ranked rates of adoption could reflect people buying subscriptions for work in expectation of reimbursement. While we cannot observe reimbursement arrangements directly, we find no evidence of them in transfers from employers. In our open-banking records only 13.8% of consumer AI subscription charges are followed within 30 days by any non-payroll transfer from the account holder's employer, virtually identical to the background rate of employer transfers in random non-AI months on the same accounts (12.6%).

21. Recall that cumulative adoption identifies accounts that have ever paid for AI; the headline 5.1% adoption rate discussed earlier is the current proportion of the population paying in June 2026. The comparable cumulative rate among all salaried workers is 15.6%.

### 4.6 Geography

Occupational and industrial sorting also produces geographic concentration across the United Kingdom. Mapping cumulative paid AI adoption across 2,533 Great Britain postcode districts (Figure 6) or 361 Local Authority Districts (Appendix Figure A7) reveals pronounced clustering in London and major university or technology hubs such as Oxford, Cambridge, Edinburgh, and Reading.

Across Great Britain in June 2026, the active paid AI rate in London (6.8%) is nearly twice as high as outside London (3.6%), with a nearly fourfold gap between the top and bottom tenths of local authorities (3.8 to 1). Across local authorities, cumulative adoption correlates strongly with both the local share of residents in managerial and professional occupations ($r = 0.7$) and local disposable household income per head ($r = 0.7$). Indeed, a single summary index combining local professional employment, household income, and young-adult population share accounts for 64% of the geographic variation across Britain.[22]

## 5 Competition, Pricing Tiers, and Switching Across AI Providers

### 5.1 Market Shares and Provider Differentiation

Because our transaction dataset follows the same individuals across competing sellers over time, we can measure both how spending shares have shifted across leading AI providers and how different groups of workers choose between them.

For nearly three years after launching ChatGPT Plus in February 2023, OpenAI held a near-monopoly in paid consumer AI. As late as December 2025, ChatGPT captured 73.5% of directly billed consumer AI subscription spending (Figure 8) and was held by 88.3% of active general-purpose assistant subscribers (Appendix Table A2), whereas Anthropic's Claude accounted for just 7.4% of standalone subscription spending and 6.5% of active assistant subscribers.

Over the first seven months of 2026, following Anthropic's release of Claude Opus 4.5 in November 2025 and rapid uptake of its coding and desktop tools (Claude Code), Anthropic's share of direct consumer spending quadrupled from 7.4% in December

22. This summary index is the first principal component of the local managerial and professional employment share, ONS Gross Disposable Household Income per head, and the share of residents aged 20–34 across Great Britain local authorities; see Appendix E.

2025 to 29.4% in July 2026, while ChatGPT's spending share fell by 21 percentage points from 73.5% to 52.5% (Figure 8).[23] Measured by active general-purpose assistant subscribers (Appendix Table A2), ChatGPT is held by 71.2% of assistant subscribers in July 2026 versus 27.1% for Claude, while Claude generates substantially higher average revenue per subscriber. These changes in market share happened in the context of explosive growth in the overall market (Figure 1), so both firms had substantial growth in revenues, despite their changing relative slices.

Rather than competing for an identical pool of customers, Claude and ChatGPT attract users from very different occupations and industries. Figure 7 plots each profession's Claude share among direct subscribers of Claude or ChatGPT minus the 22.6% overall average across the 16 professions with at least 10 direct subscribers to each provider.[24] While ChatGPT retains a majority of cumulative direct subscribers in every profession, technical, financial, and analytical professions lean strongly toward Claude relative to the average across all subscribers, led by finance and insurance professionals (+13.3 percentage points), university and college lecturers (+8.1 percentage points), civil servants (+8.1 percentage points), accountants (+7.4 percentage points), marketing specialists (+4.1 percentage points), medical doctors (+3.6 percentage points), and chartered engineers (+2.3 percentage points). Conversely, classroom, clinical, and human-resource professions rely much more heavily on ChatGPT. Relative to the overall benchmark, Claude's share is lower by 9.1 percentage points among school teachers, 8.2 percentage points among allied health professionals, 8.1 percentage points among human-resource specialists, and 6.7 percentage points among nurses and midwives.

Grouping workers by their employer's industry shows a similar split (Appendix Figure A5b): relative to the 15.3% salaried average, employees in software and IT lean toward Claude alongside workers in publishing, advertising, consulting, scientific R&D, and finance. By contrast, employees in public administration and health strongly favor ChatGPT.

Likewise, linking assistant choice to payments for standalone software-development tools (such as Cursor, GitHub Copilot, Replit, and Lovable) highlights strong sorting among software developers, where consumers who pay for standalone coding tools

23. The spending shares in Figure 8 are computed across directly billed AI subscriptions on card accounts where merchant names are not obscured, reweighted to match the UK adult population (so ChatGPT and Claude together account for 81.9% of spending in July 2026, with Gemini and specialist tools accounting for the remainder; see Section 2 and Appendix C). When mobile app-store subscriptions are included, Claude's spending share rises nearly fourfold from 6.1% in December 2025 to 22.8% in July 2026 (Appendix Table A2).

24. As described in Section 2, professions are identified from registration fees and membership dues paid to professional regulators, royal colleges, and trade unions.

are more than twice as likely to choose Claude for their general assistant subscription as other assistant subscribers (46.8% versus 21.6%; Appendix E).

## 5.2 Pricing Tiers, Switching, and Model Releases

Claude's share of subscription spending exceeds its share of active subscribers. Figure 9 examines the sources of this gap and its rapid growth in 2026, distinguishing subscriber growth from uptake of expensive plans and switching from ChatGPT, and assessing how subscription patterns respond to major model releases.

At the standard subscription tier (£20 per month for ChatGPT Plus and £18 to £20 per month for Claude Pro), the average monthly price paid by UK subscribers remained stable at roughly £19 from early 2023 to mid-2026. Rather than competing through price cuts on the £20 tier, OpenAI and Anthropic expanded their product menus at the top and bottom of the price distribution. At the top of the menu, both providers introduced high-capacity plans priced at or above £80 per month, ChatGPT Pro at £160–£200 per month in December 2024, and Claude Max at £90 and £180 per month in April 2025, which offer much higher usage limits and integrated coding features such as Claude Code. Figure 9(a) shows that by mid-2026, 6.2% of active Claude subscribers pay £80 or more per month, compared with just 0.3% of active ChatGPT subscribers.[25] About a quarter of Claude's direct spending-share gain from December 2025 to July 2026 comes from this power-user upgrading, with the remainder coming from new users (Appendix E.2).

At the bottom of the menu, OpenAI introduced an entry-level plan, ChatGPT Go, priced at £7 to £8 per month, with roughly three-fifths (59%) of its UK subscriptions billed through mobile app stores. Figure 9(b) breaks down UK ChatGPT Go buyers by their prior AI payment history. 70.2% had never previously paid for an AI subscription, 12.2% returned to OpenAI after a previously cancelled subscription, 11.9% moved down from an active higher-priced ChatGPT plan, and 5.8% held or returned from a competing provider. The sub-£10 tier thus primarily brought new mobile users into paid AI rather than poaching or cannibalizing existing subscribers.

By contrast, some of Anthropic's growth is driven by poaching subscribers from ChatGPT. Tracking the direct payment history of new Claude subscribers shows that roughly one-fifth were former ChatGPT subscribers who had cancelled and one-seventh held an active ChatGPT subscription when starting Claude.[26] Switching from

25. In account-level transitions in Appendix Table A11, 6.2% of Claude Pro accounts upgrade to Max plans versus 0.7% of ChatGPT Plus accounts upgrading to higher-priced plans.

26. Within general-purpose assistants, most subscribers pay for only one provider at a time: just 4.2% of active assistant subscribers in a given month pay two or more assistant providers concurrently,

cancelled ChatGPT subscriptions rose sharply in spring 2026. Figure 9(c) tracks direct ChatGPT subscribers who cancel in each month and plots the percentage who start a paid Claude subscription within 60 days. Throughout 2025, only 1.4% of cancelling ChatGPT subscribers switched to Claude within 60 days; between March and May 2026, that 60-day switching rate rose to 11.7% (10.5% on direct web charges alone).

Finally, Figure 9(d) shows that consumer sign-ups respond very differently across providers to major model releases. Because AI subscriptions were growing rapidly over this period even in weeks without a product launch, we compare first-time paid sign-ups in the 14 days following major model launches against the preceding four weeks and benchmark those changes against randomly chosen dates (Figure 9d; event list in Appendix E.2). Major Anthropic releases are associated with a 39.5% jump in new sign-ups compared with a 10.0% average two-week increase on random comparison dates, a net surge of nearly 30 percentage points above normal trend growth, whereas OpenAI releases change first-time sign-ups by just 1.9% (versus 3.4% on random dates) and Google releases by 12.2% (versus 4.3% on random dates). For a challenger growing from a smaller customer base, major model releases attract new payers into the market and accelerate market-share gains.

# 6 Concluding Remarks

Combining transaction-level data from consumer bank accounts, credit cards, and debit cards with individual characteristics, this paper provides a comprehensive portrait of who pays for consumer artificial intelligence. By June 2026, an estimated 5.1% of UK adults held an active paid AI subscription, with 88% of total spending growth since early 2024 driven by an expanding share of subscribers rather than higher average spending per subscriber.

Who adopts paid AI varies sharply across the population. Adoption is nearly fivefold higher among adults aged 20–29 than among those aged 60 and over, even after accounting for income, gender, and geographic location, and 27% higher among men than women. Paid adoption is 78% higher in the top than in the bottom fifth of the earnings distribution, though nearly two-thirds of the relationship between salary and adoption is explained by differences in occupation, industry, location, and demographics. Across Great Britain, we find a fourfold difference between the top and bottom decile of adoption by geographic locations.

Following individual subscribers across sellers reveals sizable competition between

whereas users of specialist voice, image, and coding tools frequently pair them with a general assistant (Appendix E.2 and Appendix Figure A8).

OpenAI's ChatGPT and Anthropic's Claude. Between December 2025 and July 2026, Claude's share of direct AI subscription spending rose fourfold from 7% to 29.4% while ChatGPT's spending share declined from 74% to 53%, as Claude's share of active assistant subscribers expanded from 7% to 27%. This shift reflects rapid subscriber growth with strong occupational sorting of technical, financial, and academic workers. It also reflects an eightfold rise in the rate at which cancelling ChatGPT users switched to Claude. Claude also attracts users who are more likely to pay for premium tools.

Aside from facts about who pays for AI, our paper extends prior work using consumer financial transactions in a number of methodological ways. First, we build and validate a method for recovering the occupations of account holders. Second, we explicitly account for masked transactions through app-store intermediaries, using probabilistic classification of masked transactions and validating using numerous structural breaks. Third, we correct for inactive accounts and reweight our sample to address a modest amount of selection towards younger and more urban households. The detailed Appendix documents both our methodological approach, and sensitivity analysis to these decisions.

Over a longer horizon, it will be possible to track the impacts of paid AI adoption on employment spells, promotion rates, and earnings growth, which we leave as future work. We also omit patterns of substitution between AI tools and other consumption goods, another area we feel is worth pursuing. We encourage others to use statistics from this paper, available at http://whopaysfor.ai. These data are updated monthly, allowing researchers and other interested parties to study this rapidly expanding consumer market and assess the causes and consequences of paid AI adoption.

# References


Aldasoro, Iñaki, Olivier Armantier, Sebastian Doerr, Leonardo Gambacorta, and Tommaso Oliviero. 2024. "The Gen AI Gender Gap." *Economics Letters* 241:111814.

Baker, Scott R., and Lorenz Kueng. 2022. "Household Financial Transaction Data." *Annual Review of Economics* 14:47–67.

Bank of America Institute. 2026. *Not Quite mAInstream: A Consumer AI Adoption Profile.* Bank of America Institute, March.

Bick, Alexander, Adam Blandin, and David J. Deming. 2026. "The Rapid Adoption of Generative AI." Articles in Advance, *Management Science.*

Bick, Alexander, Adam Blandin, David J. Deming, Nicola Fuchs-Schündeln, and Jonas Jessen. 2026. "Mind the Gap: AI Adoption in Europe and the U.S." Working Paper 34995. National Bureau of Economic Research.

Bresnahan, Timothy F., and Manuel Trajtenberg. 1995. "General Purpose Technologies 'Engines of Growth'?" *Journal of Econometrics* 65 (1): 83–108.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. 2026. "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence." Working Paper. Stanford Digital Economy Lab.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2025. "Generative AI at Work." *The Quarterly Journal of Economics* 140 (2): 889–942.

Caselli, Francesco, and Wilbur John Coleman. 2001. "Cross-Country Technology Diffusion: The Case of Computers." *American Economic Review* 91 (2): 328–335.

Chatterji, Aaron, Thomas Cunningham, David J. Deming, Zoë Hitzig, Christopher Ong, Carl Yan Shan, and Kevin Wadman. 2025. "How People Use ChatGPT." Working Paper 34255. National Bureau of Economic Research.

Chen, Pei-Yu, and Lorin M. Hitt. 2002. "Measuring Switching Costs and the Determinants of Customer Retention in Internet-Enabled Businesses: A Study of the Online Brokerage Industry." *Information Systems Research* 13 (3): 255–274.

Chetty, Raj, John N. Friedman, Michael Stepner, and The Opportunity Insights Team. 2024. "The Economic Impacts of COVID-19: Evidence from a New Public Database Built Using Private Sector Data." *The Quarterly Journal of Economics* 139 (2): 829–889.

Comin, Diego, and Bart Hobijn. 2004. "Cross-Country Technology Adoption: Making the Theories Face the Facts." *Journal of Monetary Economics* 51 (1): 39–83.

———. 2010. "An Exploration of Technology Diffusion." *American Economic Review* 100 (5): 2031–2059.

Cox, Natalie, Peter Ganong, Pascal Noel, Joseph Vavra, Arlene Wong, Diana Farrell, Fiona Greig, and Erica Deadman. 2020. "Initial Impacts of the Pandemic on Consumer Behavior: Evidence from Linked Income, Spending, and Savings Data." *Brookings Papers on Economic Activity* 2020 (2): 35–82.

David, Paul A. 1990. "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox." *American Economic Review* 80 (2): 355–361.

Dell'Acqua, Fabrizio, Edward McFowland III, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R. Lakhani. 2023. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality." Working Paper 24-013. Harvard Business School.

DellaVigna, Stefano, and Ulrike Malmendier. 2006. "Paying Not to Go to the Gym." *American Economic Review* 96 (3): 694–719.
Deloitte. 2026. *Digital Consumer Trends 2026.* Deloitte LLP.
Deming, David J., and Kadeem Noray. 2020. "Earnings Dynamics, Changing Job Skills, and STEM Careers." *The Quarterly Journal of Economics* 135 (4): 1965–2005.
Digital Regulation Cooperation Forum. 2025. *Understanding Consumer Use of Generative AI.* Digital Regulation Cooperation Forum, April.
Dynan, Karen E., Jonathan Skinner, and Stephen P. Zeldes. 2004. "Do the Rich Save More?" *Journal of Political Economy* 112 (2): 397–444.
Einav, Liran, and Jonathan Levin. 2014. "Economics in the Age of Big Data." *Science* 346 (6210): 1243089.
Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2024. "GPTs are GPTs: Labor Market Impact Potential of LLMs." *Science* 384 (6702): 1306–1308.
Felten, Edward, Manav Raj, and Robert Seamans. 2021. "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses." *Strategic Management Journal* 42 (12): 2195–2217.
Frey, Carl Benedikt, and Michael A. Osborne. 2017. "The Future of Employment: How Susceptible Are Jobs to Computerisation?" *Technological Forecasting and Social Change* 114:254–280.
Gambacorta, Leonardo, Tullio Jappelli, and Tommaso Oliviero. 2025. "Exploring Household Adoption and Usage of Generative AI: New Evidence from Italy." BIS Working Paper 1298. Bank for International Settlements.
Ganong, Peter, and Pascal Noel. 2019. "Consumer Spending during Unemployment: Positive and Normative Implications." *American Economic Review* 109 (7): 2383–2424.
Gathergood, John, Fabian Gunzinger, Benedict Guttman-Kenney, Edika Quispe-Torreblanca, and Neil Stewart. 2021. "Levelling Down and the COVID-19 Lockdowns: Uneven Regional Recovery in UK Consumer Spending." *Covid Economics* 67:24–52.
Goldfarb, Avi, and Jeff Prince. 2008. "Internet Adoption and Usage Patterns are Different: Implications for the Digital Divide." *Information Economics and Policy* 20 (1): 2–15.
Goldfarb, Avi, and Catherine Tucker. 2019. "Digital Economics." *Journal of Economic Literature* 57 (1): 3–43.
Goolsbee, Austan, and Peter J. Klenow. 2002. "Evidence on Learning and Network Externalities in the Diffusion of Home Computers." *Journal of Law and Economics* 45 (2): 317–343.
Griliches, Zvi. 1957. "Hybrid Corn: An Exploration in the Economics of Technological Change." *Econometrica* 25 (4): 501–522.
Grubb, Michael D., and Matthew Osborne. 2015. "Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock." *American Economic Review* 105 (1): 234–271.
Hacıoğlu-Hoke, Sinem, Diego R. Känzig, and Paolo Surico. 2021. "The Distributional Impact of the Pandemic." *European Economic Review* 134:103680.
Handa, Kunal, Alex Tamkin, Miles McCain, Saffron Huang, Esin Durmus, Sarah Heck, Jared Mueller, et al. 2025. "Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations." arXiv: 2503.04761.

Humlum, Anders, and Emilie Vestergaard. 2025a. "Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI." Working Paper 33777. National Bureau of Economic Research.

———. 2025b. "The Unequal Adoption of ChatGPT Exacerbates Existing Inequalities among Workers." *Proceedings of the National Academy of Sciences* 122 (1): e2414972121.

Klemperer, Paul. 1995. "Competition when Consumers have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade." *The Review of Economic Studies* 62 (4): 515–539.

Lambert, Peter John, and Y. Schindler. 2026. "The Broken Ladder: AI, Remote Work, and Early-Career Hiring." Working Paper.

Mussa, Michael, and Sherwin Rosen. 1978. "Monopoly and Product Quality." *Journal of Economic Theory* 18 (2): 301–317.

Noy, Shakked, and Whitney Zhang. 2023. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *Science* 381 (6654): 187–192.

Ofcom. 2026. *Adults' Media Use and Attitudes Report 2026.* Ofcom, April.

OpenAI. 2026. "Scaling AI for Everyone," February. Accessed August 28, 2026. https://openai.com/index/scaling-ai-for-everyone/.

Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer. 2023. "The Impact of AI on Developer Productivity: Evidence from GitHub Copilot." arXiv. arXiv: 2302.06590.

Pew Research Center. 2026. *Americans and AI 2026: Chatbots, Smart Devices and Views on Impact.* Pew Research Center, June.

PNC Financial Services Group. 2026. *Consumer Health Check: April 2026—No Stopping for G as Yet.* PNC Financial Services Group, April.

Rogers, Everett M. 2003. *Diffusion of Innovations.* 5th ed. New York: Free Press.

**Figure 1:** Paid Standalone AI Subscriptions Doubled Annually Since 2024, Reaching 5.1% of UK Adults

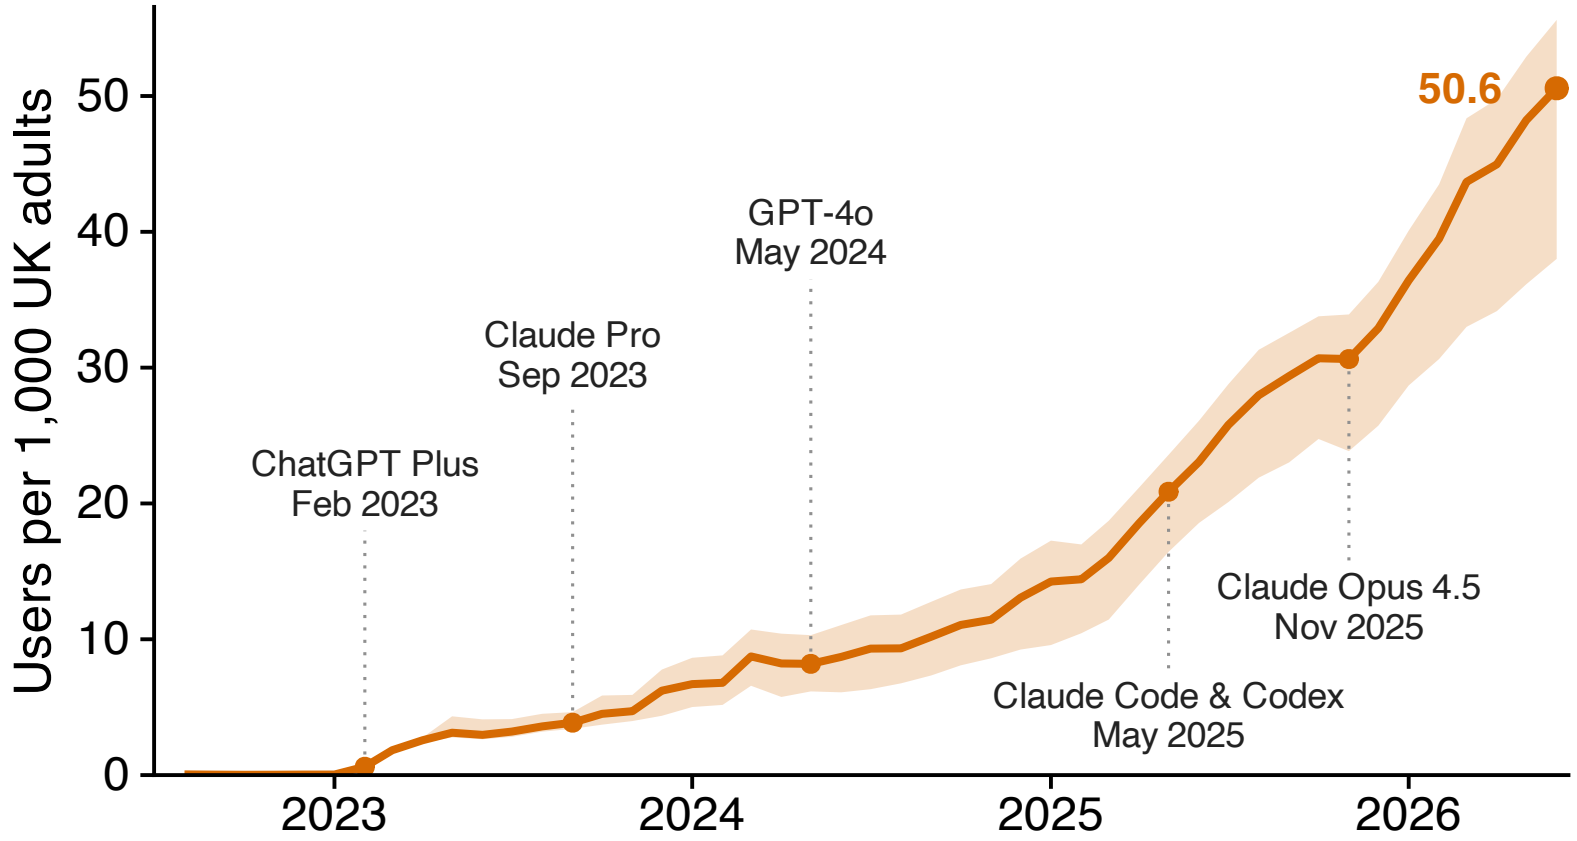


**Notes:** This figure plots active monthly paid standalone AI subscribers per 1,000 UK adults from January 2023 to June 2026. The series combines direct web subscriptions with mobile app-store subscriptions (described in Section 2 and Appendix B), with the shaded band showing 95% bootstrap confidence intervals. Active accounts are reweighted each month using demographic, regional, and income benchmarks to represent the UK adult population (Appendix C). Vertical dotted lines mark major product and model launches.

# **Figure 2:** Young Adults and Men Adopt Paid AI at Substantially Higher Rates

(a) Raw and Adjusted Cumulative Adoption by Age Band

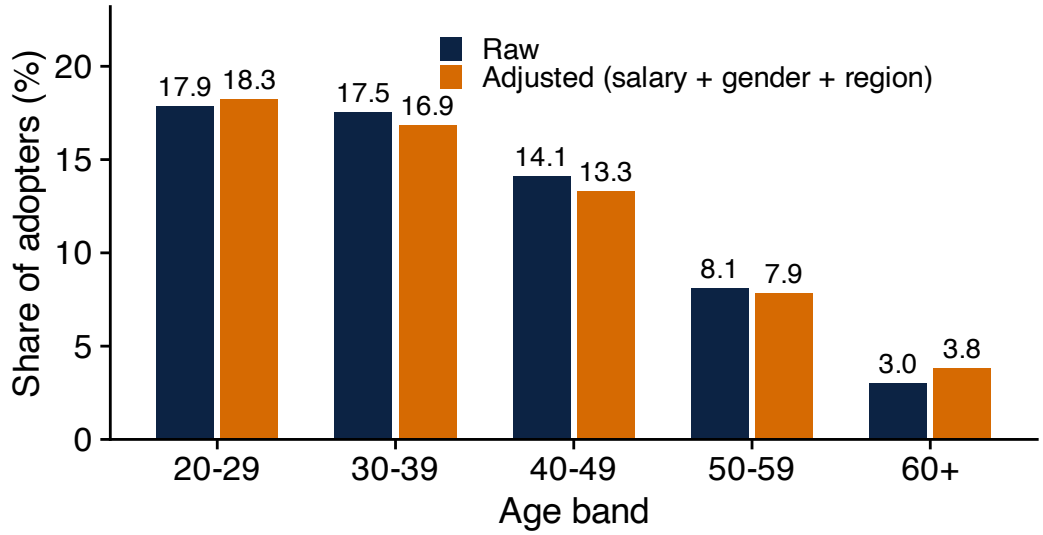


(b) Active Paid AI Subscribers per 1,000 Adults by Gender

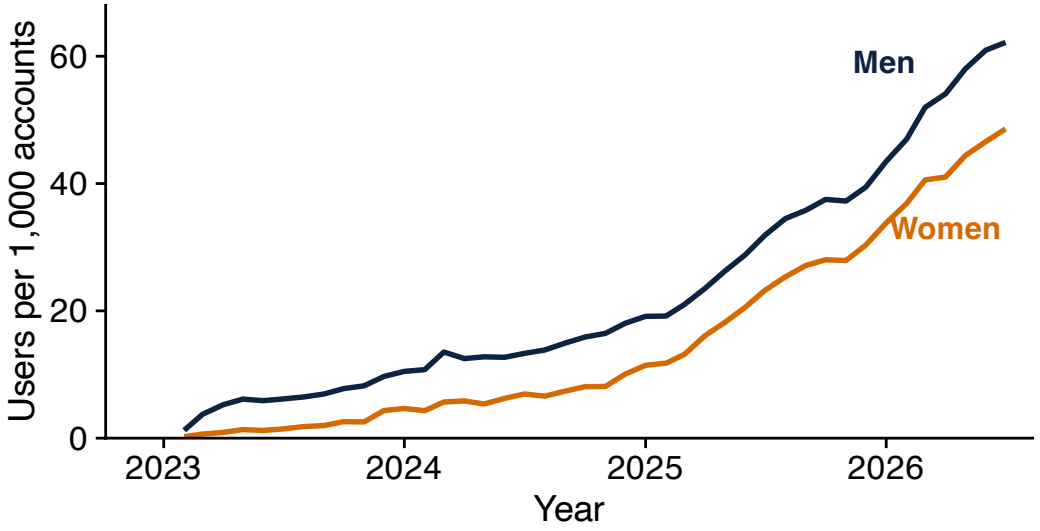


**Notes:** Panel (a) plots the percentage of adults in each age band who have ever held a paid AI subscription since November 2022, shown before adding controls (navy bars) and after accounting for after-tax income, gender, and UK region (orange bars). Panel (b) plots monthly active paid AI subscribers per 1,000 adults separately for men and women from February 2023 through July 2026. Active accounts are reweighted using demographic and regional benchmarks to represent the UK adult population (Appendix C).

# **Figure 3:** Paid AI Adoption Rises With Earnings, Driven Primarily by Occupation and Industry

(a) Raw and Adjusted Adoption by Monthly Net Salary Band

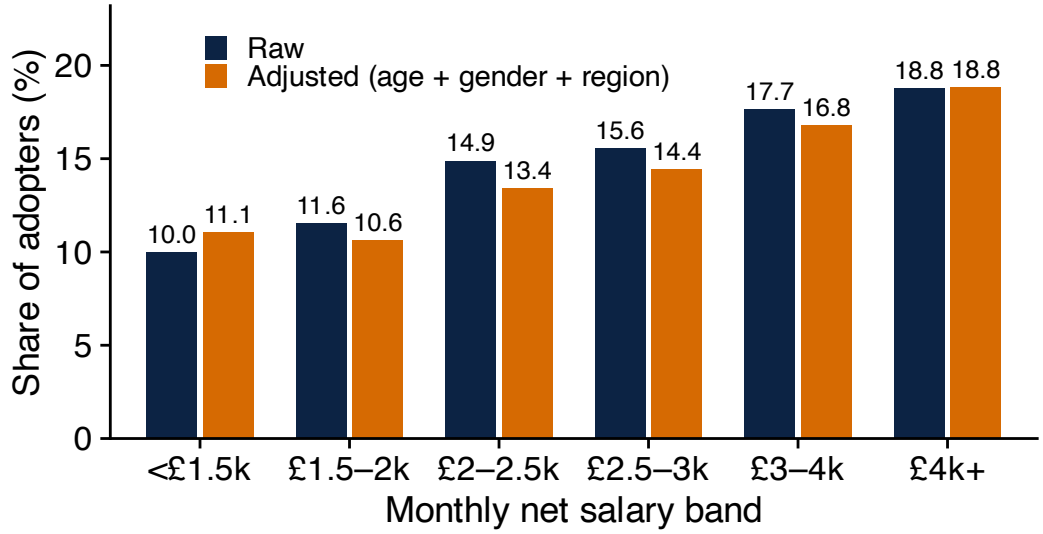


(b) Log-Salary Slopes Across Progressive Control Specifications

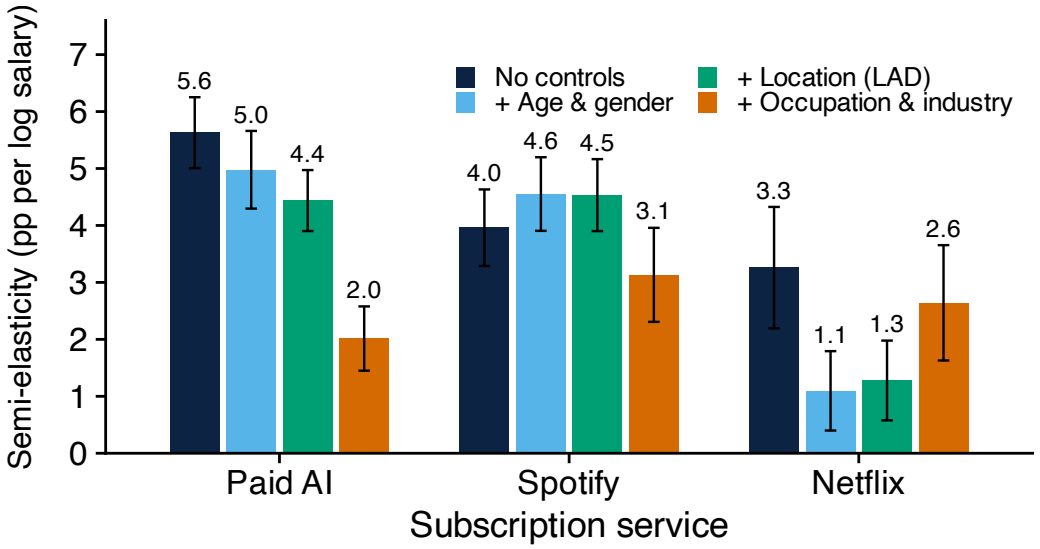


**Notes:** Panel (a) plots the percentage of salaried workers in each monthly net take-home pay band who have ever held a paid AI subscription across 390,404 salaried open-banking accounts, shown before adding controls (navy bars) and after accounting for age, gender, and UK region (orange bars). Panel (b) plots the percentage-point change in cumulative adoption associated with a one-log-point increase in monthly net take-home salary (with 95% confidence intervals clustered by local authority) for Paid AI, Spotify, and Netflix as controls for age and gender, local authority district, and occupation and employer industry are added step by step. Accounts are reweighted to represent the UK adult population; occupational and industry classifications are described in Section 2 and Appendix D.

# **Figure 4:** Paid AI Assistant Adoption Varies Twelvefold Across Occupations

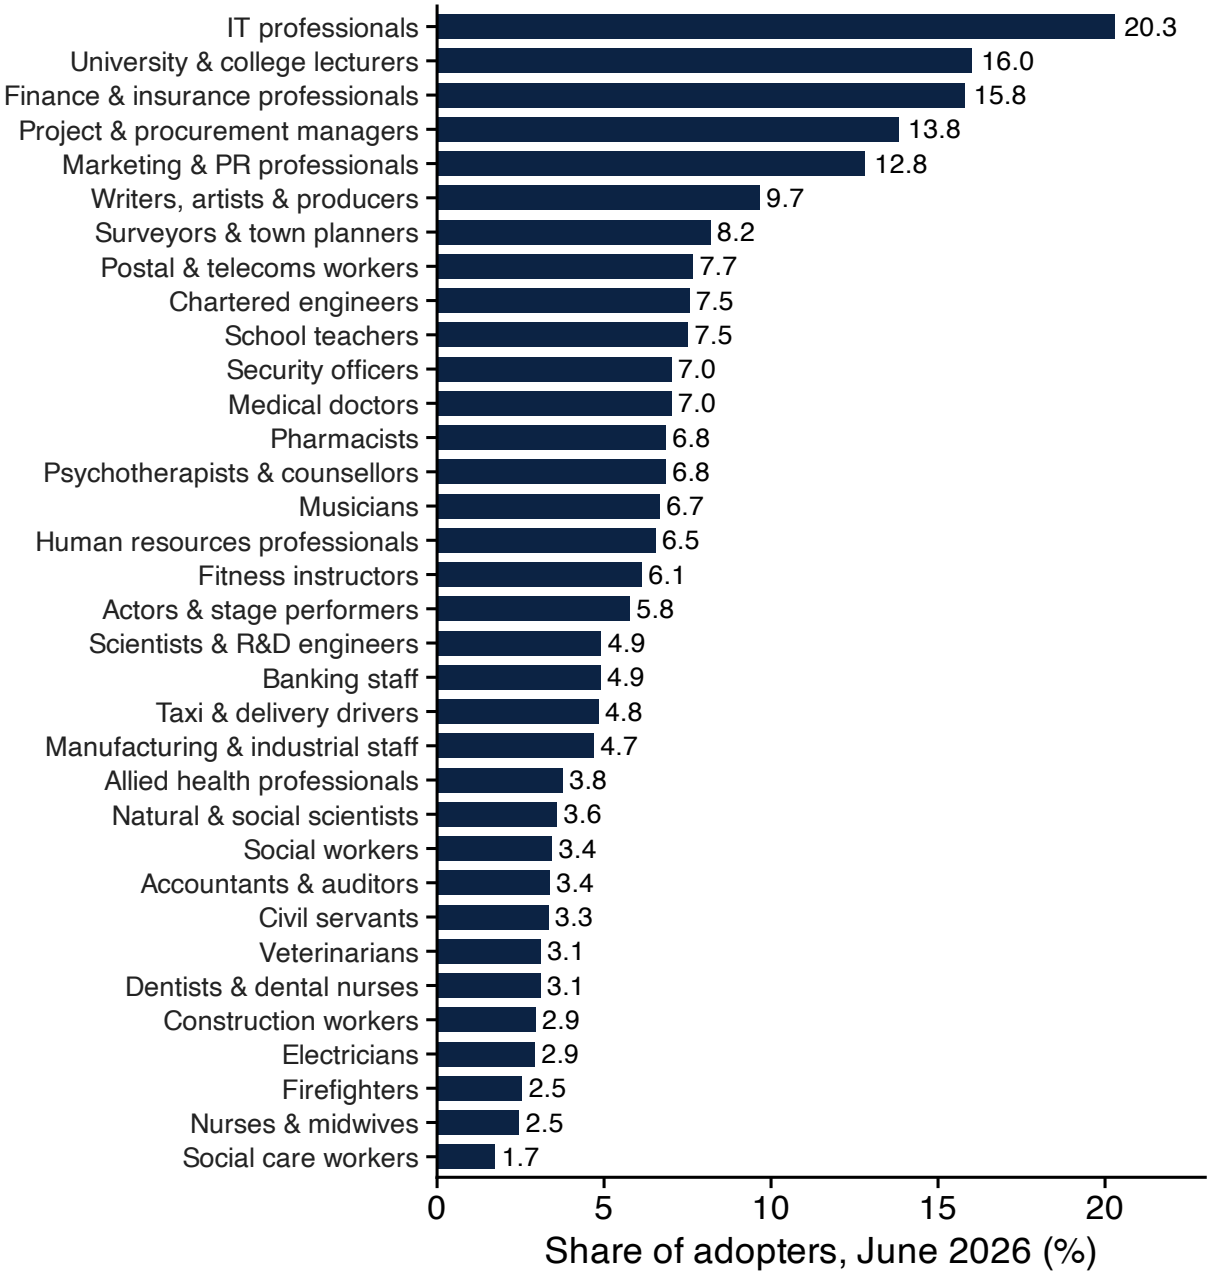


**Notes:** This figure plots the percentage of workers in each profession holding an active direct paid AI subscription in June 2026 across 34 professions with at least 50 active accounts and at least 10 active June subscribers (results across all 65 professions and across all payment channels are reported in Appendix Table A9). Professions are identified directly from registration fees and membership dues paid to professional regulators, royal colleges, and trade unions (Section 2 and Appendix D), evaluated on card accounts where privacy filters do not obscure merchant names, reweighted to match the UK adult population.

# **Figure 5:** Occupational Paid AI Adoption Closely Tracks Workplace Survey Usage and Task Exposure

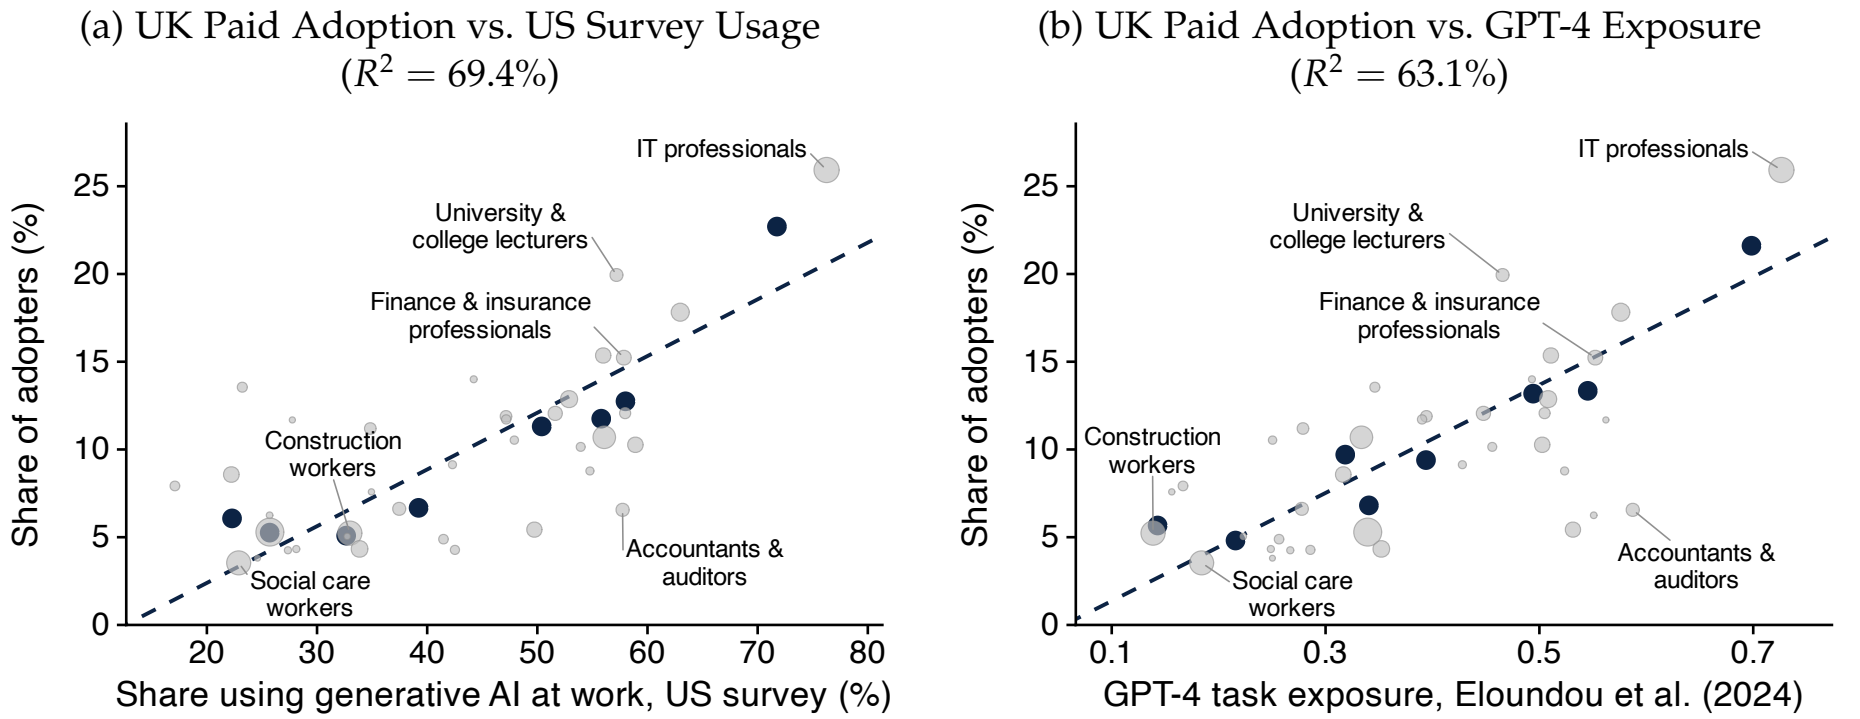


**Notes:** This figure compares 12-month active paid AI assistant adoption across 51 professions in our UK dataset against US survey self-reported workplace generative AI usage from Bick, Blandin, and Deming (2026) in Panel (a) and the task-based GPT-4 exposure index of Eloundou et al. (2024) in Panel (b). Dark navy circles denote eight equal-employment group averages, grey circles denote individual occupations scaled by UK Annual Population Survey employment, and dashed lines show employment-weighted linear fits. Full regression results and occupational classification details are reported in Appendix D.

# **Figure 6:** Paid AI Adoption Clusters in London and Major University and Tech Hubs

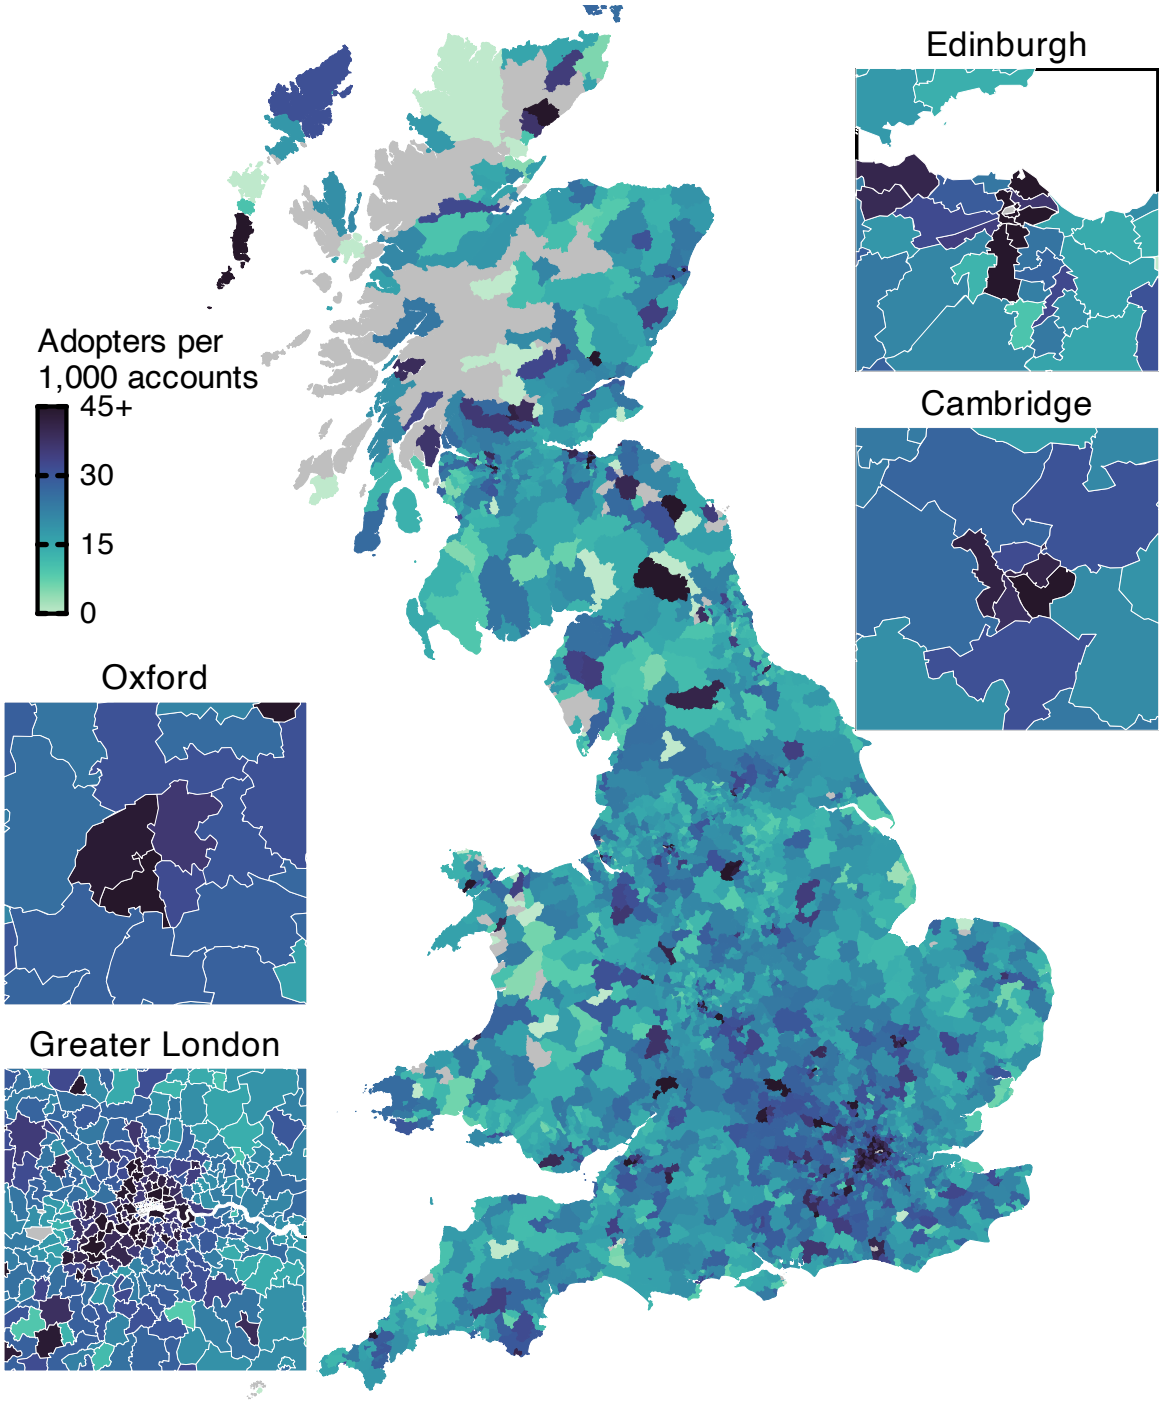


**Notes:** This figure maps cumulative paid AI adopters per 1,000 accounts across 2,533 Great Britain postcode districts (with districts containing fewer than 50 accounts shaded grey). Surrounding the main map are four urban zoom insets on the same colour scale for Edinburgh, Cambridge, Oxford, and Greater London. Population-reweighted comparisons across Local Authority Districts are reported in Section 4 and Appendix E.

**Figure 7:** Technical, Financial, and Academic Professions Lean Toward Claude Over ChatGPT

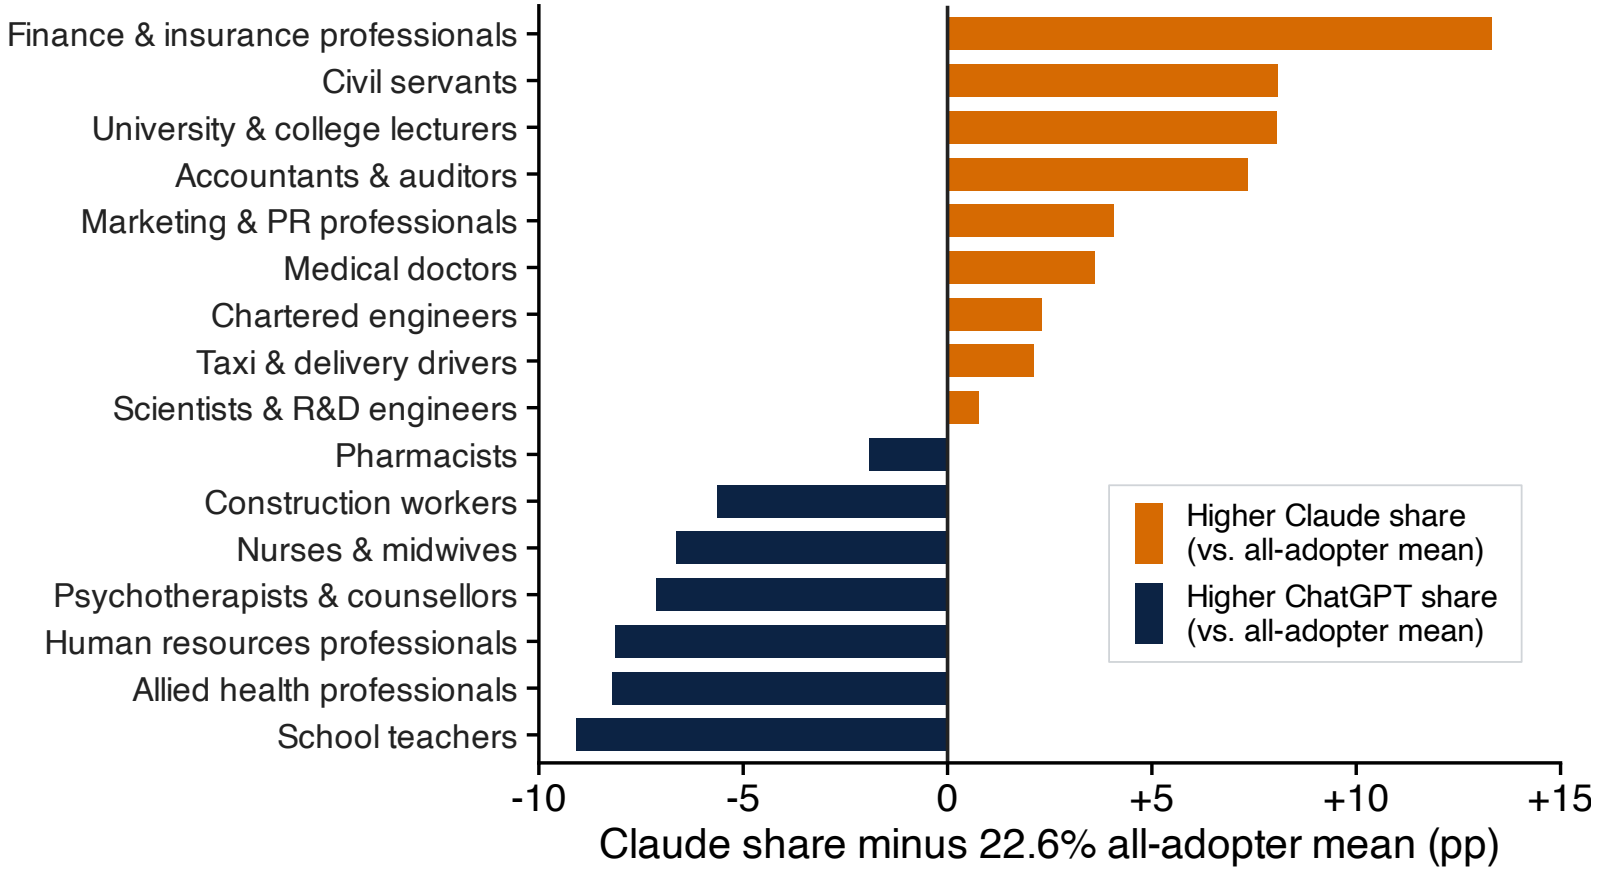


**Notes:** This figure plots provider choice across the 16 professions with at least 10 cumulative direct subscribers to each provider (Claude and ChatGPT), identified from registration fees and membership dues paid to professional regulators, royal colleges, and trade unions. Each horizontal bar reports the percentage of direct Claude or ChatGPT subscribers in that profession who pay for Anthropic's Claude minus the 22.6% average across all subscribers. While ChatGPT holds a majority of cumulative direct subscribers in every profession, positive values (orange bars) indicate professions that lean toward Claude relative to the overall average and negative values (navy bars) indicate professions that lean toward ChatGPT (Appendices C and D).

**Figure 8:** Claude's Share of Consumer AI Spending Quadrupled in Early 2026 as ChatGPT's Fell

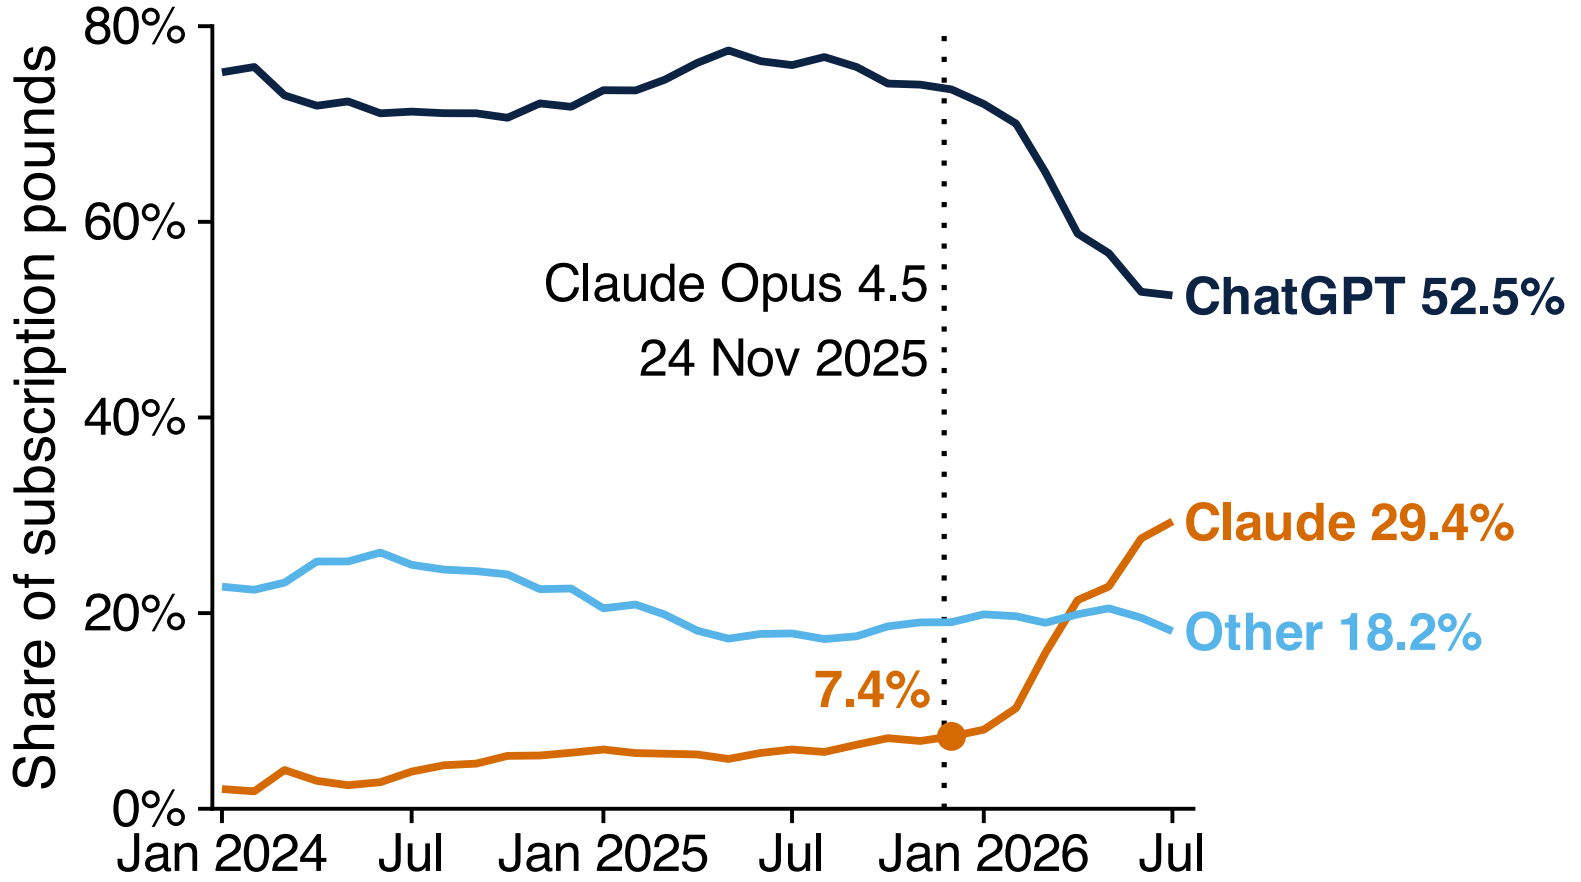


**Notes:** This figure plots monthly provider shares of directly billed standalone consumer AI subscription spending from January 2023 to July 2026 on card accounts where privacy filters do not obscure merchant names, reweighted to match the UK adult population (Section 2 and Appendix C). Active subscriber shares and supplementary competition analyses are reported in Appendix E.2.

# **Figure 9:** Pricing Tiers, Switching, and Responses to Model Releases

(a) Claude Users Take Up £80+ Plans at 20× ChatGPT

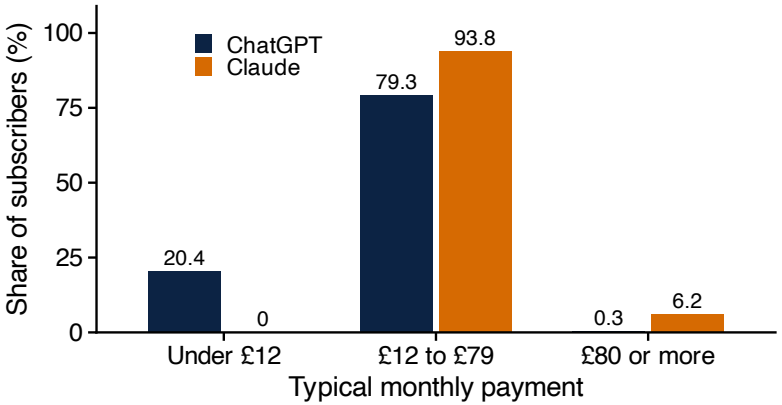


(b) ChatGPT Go Primarily Attracts First-Time Payers

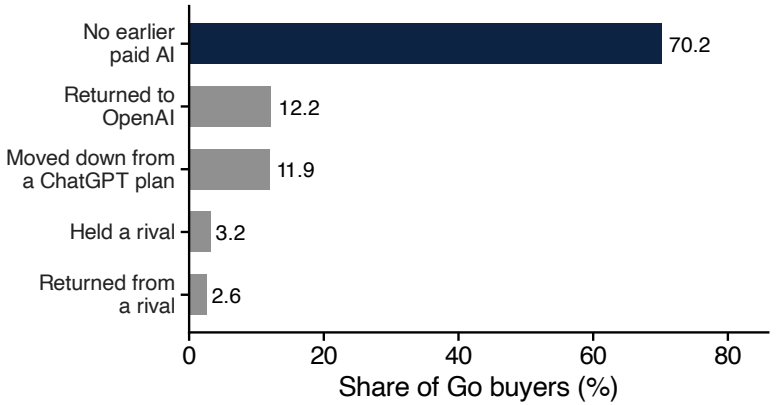


(c) Switching From ChatGPT to Claude Surged

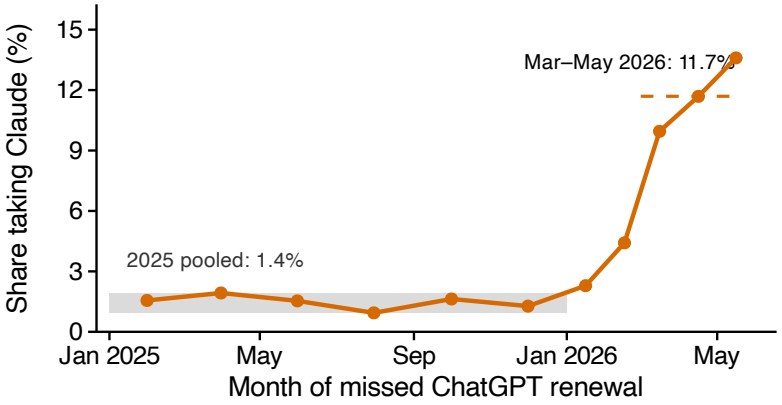


(d) Anthropic Releases Trigger Sharp Sign-Up Surges

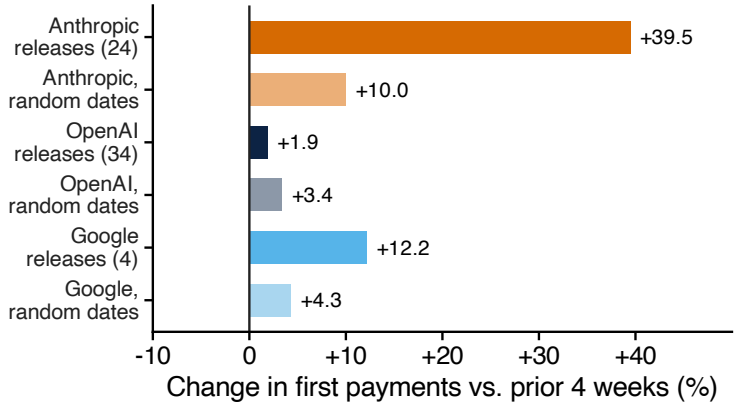


**Notes:** Panel (a) plots the distribution of active ChatGPT and Claude subscribers across monthly price tiers (under £12, £12 to £79, and £80 or more). Panel (b) breaks down buyers of OpenAI's entry-level ChatGPT Go plan (£7 to £8 per month) by their prior AI payment history. Panel (c) plots the percentage of cancelling ChatGPT subscribers in each month who start a paid Claude subscription within 60 days. Panel (d) compares the percentage change in first-time provider sign-ups in the 14 days following major model launches against pre-release levels and 300 random comparison dates. Accounts are reweighted to represent the UK adult population; see Appendix E.2 for further details.

# **Table 1:** Identifying and Measuring Paid AI in Consumer Transaction Records

**Panel A. Examples of Transaction Descriptions and Classification Rules**

| Raw description (as received) | Amount | Month | Channel / MCC | Category | Matching rule |
|---|---|---|---|---|---|
| `openai *chatgpt subscr san francisco ca` | £20.00 | Jul 2026 | Direct / 5734 | Direct AI | Text match `openai`; ChatGPT Plus £20 plan |
| `claude.ai subscription san francisco ca` | £18.00 | May 2026 | Direct / 5734 | Direct AI | Text match `claude.ai`; Claude Pro £18 plan |
| `anthropic* claude sub san francisco ca` | £90.00 | Jul 2026 | Direct / 5734 | Direct AI | Text match `anthropic`; Claude Max £90 plan |
| `www.perplexit` | £15.84 | Jun 2026 | Direct / 5734 | Direct AI | Text match `perplexit`; $20 USD plan |
| `cursor, ai powered ide new york ny` | £14.91 | Oct 2025 | Direct / 5734 | Direct AI | Text match `cursor`; $20 USD plan |
| `apple.com/bill` | £19.99 | Jun 2026 | App store / 5735 | Mobile AI | Monthly iOS subscription at £19.99 Plus price |
| `apple.com/bill` | £8.00 | Jun 2026 | App store / 5735 | Mobile AI | Monthly iOS subscription at £8.00 Go price |
| `google *google play` | £18.00 | Jun 2026 | App store / 5817 | Mobile AI | Monthly Android subscription at £18.00 Pro price |

**Panel B. External Validation Benchmarks**

| Validation check | Benchmark / evaluation sample | Validation result |
|---|---|---|
| Mobile app-store model (merchant names hidden) | Known direct AI vs. non-AI £20 subscribers | 0.9 AUC; sharp jump at ChatGPT Go launch |
| Open-banking net salary distribution (P30–P90) | ONS Annual Survey of Hours and Earnings (ASHE) | Closely matches official net pay percentiles |
| Jan 2024 statutory National Insurance rate cut (12% to 10%) | 390,404 salaried open-banking accounts | +£31/mo observed vs. +£33/mo schedule |
| Consumer digital subscription penetration (Netflix) | BARB Establishment Survey (2022–2026) | Tracks BARB UK Netflix benchmark |
| Out-of-sample accuracy of occupational classification model | 70,599 workers with known occupations (65 professions) | 0.9 AUC (fee payments hidden) |

**Panel C. Analysis Samples and Linked Worker Characteristics (January 2022–August 2026)**

| Sample or linked characteristic | Definition / source | Sample size / coverage |
|---|---|---|
| All consumer payment records | All consumer bank and card accounts | 14.2M accounts (8M weekly transactions) |
| Active main account sample (headline national panel) | $\geq$ 20 non-AI transactions/mo, $\geq$ £100/mo spend | 457,072 active accounts (Jun 2026) |
| Open-banking salaried sample | Regular incoming bank salary deposits | 390,404 active salaried accounts |
| Employer industry (21 UK SIC industry sections) | Companies House & public-sector directories | 19,876 employers (65% of salaried) |
| Individual occupation (65 professions & 26 broad UK groups) | 466 professional bodies/unions + spending model | 70,599 directly identified + 1.7M predicted |

**Notes:** Panel A shows illustrative transaction descriptions from our UK bank and card dataset for direct AI subscriptions and recurring mobile app-store charges (`apple.com/bill` and `google *google play`). Panel B summarizes our external validation checks across mobile app-store subscriptions, open-banking salaries, consumer subscriptions, and worker occupations (Appendices B–D). Panel C summarizes our analysis samples and linked worker characteristics (Appendix C).

# Who Pays for AI?

## Online Appendix

Nikhil Datta  Peter John Lambert  Patrick Schneider

# A Classifying Transactions as AI

This appendix details our direct classification of AI payments from raw bank and card merchant descriptors.

## A.1 AI Product Directory and Classification Rules

We construct a comprehensive directory of standalone AI products across our bank account and card transaction records between January 2022 and August 2026, identifying 249 sellers. Table A1 lists the top 50 sellers, which account for 89.4% of transactions, and 87.8% of spending. For each product, our directory records the billing-descriptor patterns that identify the seller, its subscription plan prices, and its product launch dates.

For dedicated AI providers (such as OpenAI, Anthropic, Midjourney, ElevenLabs, Suno, Perplexity, xAI, and Cursor), a matching merchant description directly identifies an AI purchase. Where a parent technology company also sells non-AI consumer cloud storage or developer software under the same corporate billing stem (such as Google One, Microsoft Copilot Pro, or GitHub), we require either a charge at the exact registered AI subscription price point or an explicit AI signal in the billing string. Pre-existing software suites that later added AI features to existing subscription plans (such as Microsoft 365, Canva, Adobe, Grammarly, Notion, Google Workspace, and Duolingo) are excluded from our main AI measures.

# **Table A1:** Top 50 Standalone AI Products in Direct UK Bank and Card Billing

| # | Product | Accounts | Transactions | Spend (£k) | # | Product | Accounts | Transactions | Spend (£k) |
|---|---|---|---|---|---|---|---|---|---|
| 1 | OpenAI (ChatGPT) | 58,294 | 329,463 | 6,285 | 26 | Gamma | 924 | 3,821 | 69 |
| 2 | Anthropic (Claude) | 14,928 | 65,728 | 1,783 | 27 | OpenArt | 859 | 3,177 | 74 |
| 3 | Use AI (use.ai) | 6,678 | 18,793 | 211 | 28 | Pixella.ai | 850 | 1,851 | 30 |
| 4 | Midjourney | 6,228 | 35,601 | 555 | 29 | Motion (usemotion) | 842 | 1,948 | 141 |
| 5 | Google (Gemini / One AI) | 4,638 | 19,038 | 353 | 30 | Kling (Kuaishou) | 808 | 3,479 | 56 |
| 6 | QuillBot | 4,162 | 11,801 | 290 | 31 | Base44 (Wix) | 808 | 2,413 | 112 |
| 7 | ElevenLabs | 4,119 | 17,281 | 213 | 32 | OpusClip | 801 | 3,105 | 65 |
| 8 | Suno | 2,649 | 11,034 | 134 | 33 | Runway | 800 | 2,274 | 59 |
| 9 | TeachMate AI | 2,482 | 17,916 | 147 | 34 | Wordtune | 753 | 2,848 | 81 |
| 10 | BitePal | 2,011 | 3,469 | 82 | 35 | Otter.ai | 745 | 4,661 | 110 |
| 11 | JustDone AI | 1,921 | 3,560 | 48 | 36 | OpenRouter | 695 | 3,259 | 39 |
| 12 | GitHub Copilot | 1,668 | 10,252 | 92 | 37 | ChatBot AI / Chatbot App (Paddle) | 651 | 1,802 | 37 |
| 13 | Lovable | 1,660 | 8,177 | 260 | 38 | VEED.io | 562 | 2,035 | 63 |
| 14 | Cursor (Anysphere) | 1,582 | 8,538 | 248 | 39 | HeyGen | 562 | 1,907 | 61 |
| 15 | Remini (Bending Spoons) | 1,500 | 4,266 | 25 | 40 | Reface | 540 | 1,710 | 34 |
| 16 | Looka (AI logo maker) | 1,342 | 2,661 | 109 | 41 | Lovi (AI skin scanner) | 539 | 772 | 22 |
| 17 | Perplexity | 1,310 | 4,254 | 89 | 42 | GPTZero | 538 | 1,494 | 32 |
| 18 | Speechify | 1,297 | 2,343 | 115 | 43 | Descript | 528 | 2,703 | 68 |
| 19 | Higgsfield | 1,172 | 4,105 | 161 | 44 | AI-Pro.org | 518 | 1,710 | 25 |
| 20 | Replit | 1,159 | 7,974 | 300 | 45 | ScriptGPT | 502 | 578 | 18 |
| 21 | InVideo | 1,105 | 3,442 | 110 | 46 | scite.ai | 481 | 1,627 | 24 |
| 22 | Candy.ai | 1,060 | 2,909 | 38 | 47 | Aigeeny | 478 | 3,160 | 238 |
| 23 | xAI (Grok / SuperGrok) | 962 | 3,050 | 56 | 48 | Google Colab Pro | 466 | 2,163 | 28 |
| 24 | Undetectable AI | 953 | 2,765 | 53 | 49 | Crayo | 423 | 901 | 17 |
| 25 | Zing Coach | 940 | 1,941 | 61 | 50 | Replika (Luka) | 422 | 881 | 16 |
| **Top 50 standalone AI products** | | **110,974** | **656,640** | **13,338** | *Share of all standalone AI* | | *90.1%* | *89.4%* | *87.8%* |
| Other 199 standalone AI products | | 20,619 | 77,828 | 1,859 | **All 249 standalone AI products** | | **123,223** | **734,468** | **15,197** |

**Notes:** This table lists the top 50 standalone AI products ranked by distinct paying accounts in our UK consumer and business bank and card records between January 2022 and August 2026, alongside summary totals for the remaining 199 standalone AI products with observed payments and all 249 standalone AI products combined. Because some accounts pay more than one standalone AI product over the sample period, distinct accounts across the top 50 products (110,974) and the remaining 199 products (20,619) sum to more than the 123,223 unique accounts across all 249 standalone AI products. General software suites that later added AI features to existing plans (such as Microsoft 365, Canva, Adobe, Grammarly, and Notion) are excluded from standalone AI measures. Classification rules are detailed in Appendix A.

# B App-Store Subscription Imputation

Consumers who subscribe to AI assistants through the iOS ChatGPT, Claude, Gemini, Perplexity, or Grok apps, or via Android devices, are billed by Apple or Google rather than directly by the AI provider. On UK bank and card statements, these transactions appear under generic aggregator descriptors, principally `apple.com/bill` (10.9 million charges across 540,735 accounts in 2025) and `google *google play` (3.5 million charges across 190,595 accounts). Ignoring in-app billing undercounts consumer AI adoption. This appendix details our two-track methodology for recovering app-store AI subscriptions.

## B.1 Time Series Imputation using Excess-Mass Decomposition

At the aggregate level, app-store AI assistant subscriptions generate sharp, dated structural breaks in the frequency of recurring `apple.com/bill` and `google *google play` charges at the exact price points for AI subscriptions. For example, Apple and Google bill UK consumers at fixed price tiers: £19.99 and £20.00 per month for ChatGPT Plus and Claude Pro iOS subscriptions, £18.99 per month for Google One AI Premium, £18.00 per month for Android Claude Pro, £6.99 to £8.00 per month for ChatGPT Go, and £199.99 per month for high-capacity power-user plans.

To construct our aggregate population best estimate (Table A2 and Figure A1, Panel (a)), we first isolate recurring monthly app-store sequences defined as runs of two or more charges to Apple or Google at the exact same sterling amount separated by 25 to 35 days. For each target AI price tier $p$ active in month $t$, we normalize active spell counts by the monthly active account base $A_t$ ($r_{p,t} = N_{p,t}/A_t$) and scale the twelve-month pre-May-2023 baseline rate $r_{p,\text{pre}}$ by a non-AI placebo price-tier growth index $\Gamma_t$ estimated across adjacent non-AI app-store price tiers, yielding the time-varying counterfactual $\hat{N}_{p,\text{pre},t} = A_t \cdot r_{p,\text{pre}} \cdot \Gamma_t$ and the post-launch excess mass:

$$E_{p,t} = \max\{0,\, N_{p,t} - \hat{N}_{p,\text{pre},t}\}, \qquad \pi_{p,t} = \frac{E_{p,t}}{N_{p,t}}, \tag{B.1}$$

with empirical Bayes shrinkage applied across narrow composite tiers where an account's pre-existing iCloud or Google One storage plan is billed on the same invoice date.

Adding these excess-mass estimates across Apple App Store and Google Play recurring runs, together with the PayPal layer reported in Table A2, increases June 2026 active paid standalone AI penetration on the primary sample from the direct lower bound of 29.0 per 1,000 UK adults to 50.6 per 1,000 UK adults with a 95% bootstrap confidence interval of [38.0, 55.6] across 500 draws. Figure A1(a) shows the impact on our time-series estimates.

Across alternative counterfactual baselines for $\hat{N}_{p,\text{pre},t}$, the June 2026 app-store increment varies between +21.8 per 1,000 under a static pre-launch rate ($\Gamma_t \equiv 1$), +20.5 per 1,000 under our baseline placebo-tier growth adjustment, +18.9 per 1,000 under

linear tier-neighbor drift, and +16.8 per 1,000 if composite storage-plus-AI invoices are excluded entirely.

## B.2 Demographic Imputations and Sensitivity Analysis

To incorporate app-store subscribers into individual-level analyses of age, gender, income, industry, and tier choice, we assign each account $i$ in the candidate recurring app-store pool $\mathcal{C}_{p,t}$ at price tier $p$ in month $t$ an imputed AI subscription probability $\hat{p}_{i,t}^{\text{app}}$. To ensure that excess payment runs reconcile one-for-one with unique incremental payers in our composite indicator $Y_{i,t} = D_{i,t}^{\text{direct}} + (1 - D_{i,t}^{\text{direct}})\hat{p}_{i,t}^{\text{app}}$, the candidate pool $\mathcal{C}_{p,t}$ restricts to unique accounts without a concurrent direct standalone AI subscription ($D_{i,t}^{\text{direct}} = 0$, deduplicating the $< 0.4\%$ of accounts with multiple app-store runs in the same month by retaining the highest-tier run, and netting out the empirical cross-channel overlap rate when mapping run-level excess $E_{p,t}$ to unique-account excess mass $E_{p,t}^{w} = 1{,}000 \times E_{p,t}^{\text{acct},w} / A_t^{w}$ evaluated under the specification's weights $w_{i,t}$ so that $\sum_i w_{i,t} = 1{,}000$, yielding a population-weighted increment of +20.5 per 1,000 adults in June 2026 versus +26.3 per 1,000 unweighted in Table A2). We evaluate three candidate allocation rules:

1. *Rule 1 (Homogeneous excess-mass prior only):* Assigns every post-May-2023 candidate app-store spell at price tier $p$ in month $t$ the uniform cohort excess-mass share $\pi_{p,t} = E_{p,t}^{w} / W_{p,t}$, where $W_{p,t} = \sum_{i \in \mathcal{C}_{p,t}} w_{i,t}$.

2. *Rule 2 (Account historical baseline subtraction):* Excludes accounts whose total monthly Apple or Google spend in 2021–2022 already averaged at least £15 per month, attributing only net new app-store spending above the account's pre-AI baseline.

3. *Rule 3 (Preferred spell-cadence and supervised probability model):* Restricts candidate spells to post-launch recurring runs with inter-charge intervals strictly between 25 and 35 days at exact AI assistant tiers on accounts with zero pre-May-2023 recurring charges at that tier, and combines the cohort excess-mass target $E_{p,t}^{w}$ with out-of-fold predictions $\hat{g}(X_i) \in [\epsilon, 1 - \epsilon]$ ($\epsilon = 10^{-4}$) from a gradient-boosted decision tree trained on direct web AI subscribers using exclusively non-AI account features $X_i$ (age band, gender, ITL1 region, log non-AI card spending, count of non-AI digital subscriptions, ONS Output Area Classification supergroup, and urban-rural postal class) via Bayes odds scaling:

$$\hat{p}_{i,t}^{\text{app}} = \frac{\kappa_{p,t}\,\hat{g}(X_i)}{1 - \hat{g}(X_i) + \kappa_{p,t}\,\hat{g}(X_i)}, \qquad \text{where } \kappa_{p,t} \text{ solves } \sum_{i \in \mathcal{C}_{p,t}} w_{i,t}\,\hat{p}_{i,t}^{\text{app}} = E_{p,t}^{w}, \quad \text{(B.2)}$$

with boundary conventions $\hat{p}_{i,t}^{\text{app}} = 0$ if $E_{p,t}^{w} = 0$ (or $\mathcal{C}_{p,t} = \emptyset$) and $\hat{p}_{i,t}^{\text{app}} = 1$ if $E_{p,t}^{w} \geq W_{p,t}$ (in practice $E_{p,t}^{w} / W_{p,t} \in [0.18, 0.74]$ is strictly interior in every post-launch month, guaranteeing a unique finite $\kappa_{p,t} > 0$). An account's composite adoption indicator is $Y_{i,t} = D_{i,t}^{\text{direct}} + (1 - D_{i,t}^{\text{direct}})\hat{p}_{i,t}^{\text{app}}$, so $\sum_i w_{i,t} Y_{i,t} = \sum_i w_{i,t} D_{i,t}^{\text{direct}} + \sum_p E_{p,t}^{w}$.

We validate this imputation approach in three ways. First, imputed spells expand steadily after the late-May 2023 UK rollout of the ChatGPT iOS app and co-move with direct OpenAI subscriptions ($R^2 = 0.84$). Second, on pre-launch placebo cohorts, the excess-mass estimator achieves 95.1% coverage with a six-month error of only 2.0%. Third, to verify that demographic gradients do not depend on transporting web-trained classifier weights $\hat{g}(X_i)$ to app-store users, we compare subgroup ratios across our different imputation rules and against direct non-parametric price-tier excess mass computed separately within each demographic cell. We find little sensitivity, for example the age 20–29 to 60+ adoption ratio ranges tightly from 5.7 to 5.9, the male-to-female active ratio ranges from 1.24 to 1.29, and the top-to-bottom salary-band ratio ranges from 1.84 to 1.90.

# **Table A2:** Decomposition of Active June 2026 Paid Standalone AI Subscribers and Cross-Design Sensitivity of Headline Paper Moments

**Panel A. Active June 2026 Paid Standalone AI Subscribers by Billing Channel, Sample, and Feed Type (per 1,000)**

| Population and weighting specification | Accounts | Direct benchmark [95% CI] | + App store | + PayPal | Best estimate [95% CI] | Best / Direct |
|---|---|---|---|---|---|---|
| *Population-reweighted UK adult estimates* | | | | | | |
| Rolling primary, pooled feeds (headline) | 443,728 | 29.0 [28.4, 29.6] | +20.5 | +1.06 | 50.6 [38.0, 55.6] | 1.74 |
| Rolling primary, 99th-percentile weight cap | 443,728 | 29.8 [29.2, 30.4] | +21.0 | +1.08 | 51.9 [39.0, 57.1] | 1.74 |
| Rolling primary, open-banking feeds only | 273,418 | 34.6 [33.8, 35.4] | +20.7 | +1.08 | 56.4 [42.4, 62.0] | 1.63 |
| Rolling primary, card-only feeds only | 170,310 | 18.1 [17.3, 18.9] | +19.1 | +1.00 | 38.2 [28.7, 42.1] | 2.11 |
| Fixed balanced primary, pooled feeds | 172,988 | 31.1 [30.0, 32.2] | +20.8 | +1.32 | 53.2 [40.3, 58.5] | 1.71 |
| *Unweighted consumer account panels* | | | | | | |
| Rolling primary accounts, unweighted | 457,072 | 34.2 [33.7, 34.7] | +26.3 | +1.42 | 61.9 [45.8, 68.3] | 1.81 |
| Fixed balanced primary accounts, unweighted | 181,939 | 32.5 [31.7, 33.3] | +23.8 | +1.57 | 57.9 [43.7, 63.7] | 1.78 |
| All transacting consumer records, unweighted | 1,437,225 | 15.3 [15.1, 15.5] | +12.4 | +0.58 | 28.3 [21.8, 31.0] | 1.85 |

**Panel B. Cross-Design Sensitivity of Headline Diffusion, Heterogeneity, and Competition Moments**

| Empirical moment | Unweighted Direct | Weighted Direct | Reweighted Direct (Unmasked) | All-Channel (Direct + Mobile) | Primary exhibit |
|---|---|---|---|---|---|
| Age gradient: 20–29 vs. 60+ raw ratio (adjusted ratio) | 5.70 (3.27) | 6.08 (4.66) | 6.18 (4.66) | 5.90 (4.76) | Figure 2a |
| Gender gradient: Men vs. women active Jul 2026 per 1,000 (ratio) | 44.9 vs. 28.4 (1.58) | 39.7 vs. 24.1 (1.65) | 38.7 vs. 23.5 (1.65) | 62.2 vs. 48.6 (1.28) | Figure 2b |
| Salary gradient: $>$ £4,000 vs. $<$ £1,500 ever paid % (ratio) | 14.2% vs. 5.0% (2.81) | 11.8% vs. 4.6% (2.55) | 10.7% vs. 4.1% (2.60) | 18.8% vs. 10.0% (1.88) | Figure 3a |
| Occupation spread: Top vs. bottom class active Jun 2026 % (ratio) | 15.9% vs. 1.8% (8.89) | 16.0% vs. 1.8% (8.94) | 16.0% vs. 1.7% (9.21) | 18.9% vs. 4.3% (4.44) | Figure 4 |
| External validation $r$: US survey usage / GPT-4 exposure ($\beta$) | 0.89 / 0.79 | 0.88 / 0.77 | 0.89 / 0.76 | 0.85 / 0.75 | Figure 5 |
| Industry adoption: Software & IT vs. all salaried ever paid (%) | 12.5% vs. 5.2% | 19.7% vs. 7.3% | 18.6% vs. 6.8% | 27.0% vs. 14.2% | Figure A5a |
| Claude share of subscription pounds: Dec 2025 $\rightarrow$ Jul 2026 (%) | 7.6% $\rightarrow$ 30.9% | 7.9% $\rightarrow$ 29.7% | 7.4% $\rightarrow$ 29.3% | 5.0% $\rightarrow$ 22.5% | Figure 8 |
| Claude share of active subscribers: Dec 2025 $\rightarrow$ Jul 2026 (%) | 6.5% $\rightarrow$ 27.8% | 6.7% $\rightarrow$ 27.4% | 6.5% $\rightarrow$ 27.0% | 4.6% $\rightarrow$ 20.3% | Section 5.1 |
| Power-user tier share ($\geq$ £80/mo): Claude vs. ChatGPT (%) | 8.74% vs. 0.51% | 7.80% vs. 0.46% | 7.32% vs. 0.44% | 6.21% vs. 0.31% | Figure 9a |
| 60-day switching to Claude after ChatGPT exit: 2025 $\rightarrow$ 2026 (%) | 1.4% $\rightarrow$ 9.3% | 1.3% $\rightarrow$ 10.6% | 1.3% $\rightarrow$ 10.5% | 1.4% $\rightarrow$ 11.7% | Figure 9c |

**Notes:** Panel A decomposes active June 2026 paid standalone AI subscribers per 1,000 across account samples, weighting specifications, and billing channels, with 95% bootstrap confidence intervals in brackets. Panel B compares headline paper moments across four estimation specifications: unweighted direct payments, population-weighted direct payments, reweighted direct payments on accounts with unmasked billing descriptions, and all-channel estimates incorporating app-store subscriptions. In Panel B, the occupation spread row reports the outcome-blind comparison across professions with $N \geq 150$ primary accounts, external validation $r$ reports the unweighted correlation across the 51 professions, and the all-channel Claude pound share row includes the unallocated PayPal layer in the standalone AI spending denominator (corresponding to 6.1% $\rightarrow$ 22.8% across provider-attributed direct and app-store spending).

# **Figure A1:** App-Store Subscription Imputation: Billing-Layer Decomposition and Launch Timing

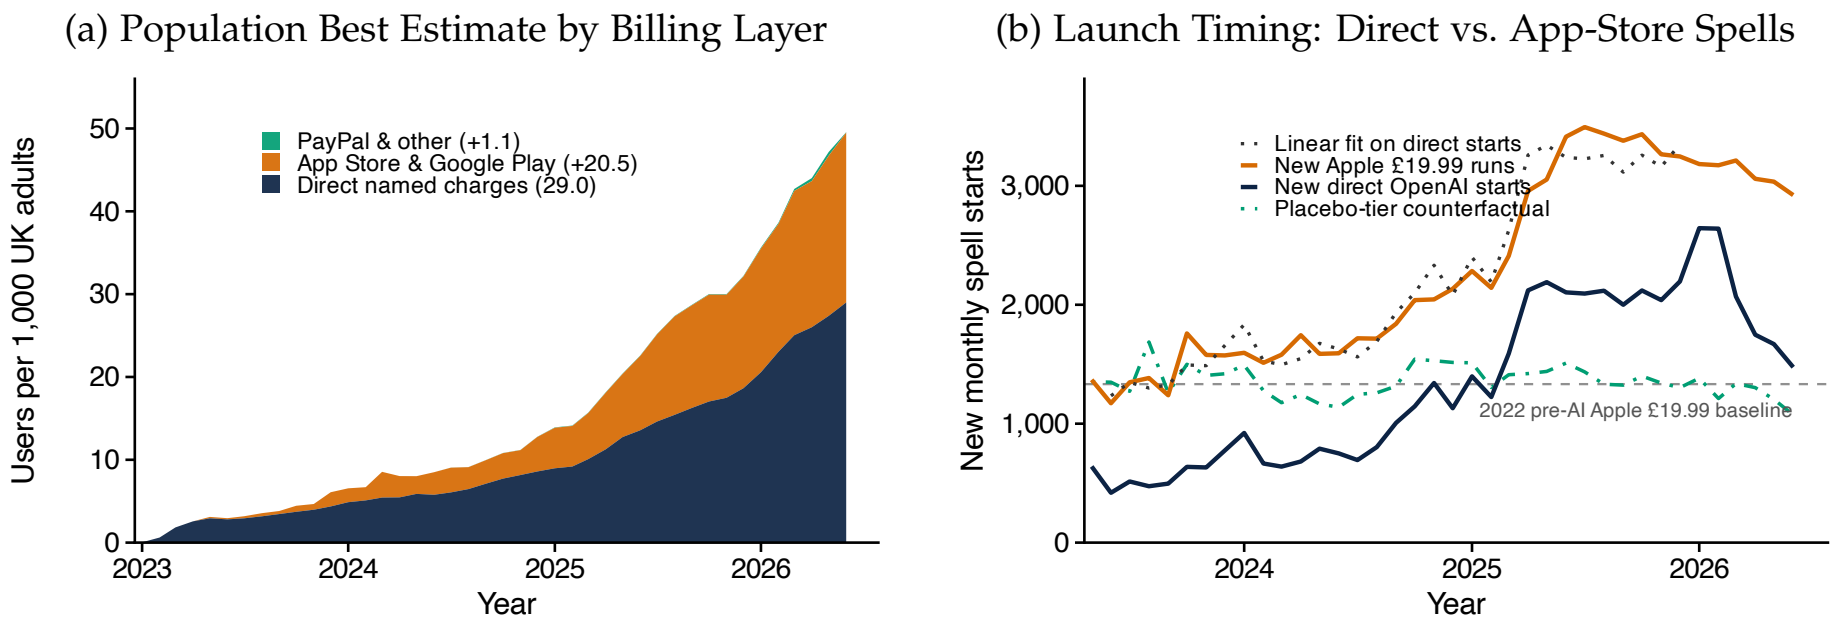


**Notes:** Panel (a) decomposes active paid standalone AI subscribers per 1,000 UK adults in June 2026 into directly classified payments, imputed Apple App Store and Google Play subscriptions, and PayPal-routed subscriptions. Panel (b) compares monthly initiation counts of direct web OpenAI Plus subscriptions and new recurring £19.99 Apple App Store spells from May 2023 through June 2026 against the 2022 pre-AI baseline (dashed line) and the non-AI placebo-tier counterfactual (dot-dash line).

# C Sample Construction, Population Weighting, and External Validation

Raw administrative bank and card panels contain secondary credit cards and dormant accounts, and their demographic composition skews younger and more urban than the national adult population as they are biased towards fintech providers. This appendix details our sample-screening ladder, our two-stage reweighting to the adult population, our regression adjustment specifications, and our supplementary external validation checks.

## C.1 The Account Sample Ladder and Descriptive Statistics

Table A3 shows the count and paid AI subscriptions across progressively stricter sample selection criteria. In the unrestricted universe of 1,424,855 consumer records with at least one transaction in June 2026, direct standalone AI penetration is 15.3 per 1,000 accounts, falling predominantly as the denominator grows for low utilisation cards. Table A4 compares adoption rates across alternative definitions of an 'active' account, showing for example that screening cards which have at least six active months since 2022 raises adoption rates to 17.0 per 1,000.

Our analysis focuses on two primary-account samples that isolate consumers whose main day-to-day payment instrument is observed:

1. *Balanced primary panel ($N = 178,109$ active in June 2026; 201,038 ever-active across 2022H1–2026H1):* Consumer accounts that transact in every half-year from 2022H1 through 2026H1, maintain median monthly non-AI expenditure of at least £100, and average at least 20 non-AI transactions per month. Direct June 2026 penetration on this sample is 33.9 per 1,000 accounts. We use this balanced panel when measuring baseline 2022 account characteristics (Tables A5 and A6).

2. *Rolling primary sample ($N = 457,072$ active in June 2026):* Accounts meeting the same primary activity screen ($\geq$ 20 non-AI transactions per month and $\geq$ £100 monthly spend) over a rolling twelve-month window. Because the rolling screen retains younger adults and newly opened fintech accounts that entered the panel after 2022, it captures 71.4% of all active direct AI subscribers in the raw data while yielding a similar direct adoption rate (34.2 per 1,000). Within our open-banking feeds, 390,404 primary salaried accounts have verified recurring payroll deposits for our earnings and industry analyses.

Because de-identified aggregator records do not link accounts held by the same person across different banking institutions, our headline national estimate translates calibrated primary accounts to UK adults one-for-one. Paying for an AI subscription on an unlinked secondary card would cause per-account rates to understate individual adoption, whereas holding multiple accounts that simultaneously clear our monthly activity threshold would expand the account denominator relative to individuals;

allowing the net account-to-person conversion factor to vary by ±15% around one-to-one places individual UK adult payers between 2.4 million (4.4%) and 3.3 million (6.0%).

Table A5 compares the baseline 2022 characteristics of future paid AI adopters ($N = 14,798$) against the balanced primary panel as a whole ($N = 201,038$). Prior to the launch of ChatGPT Plus, future adopters were 3.6 years younger on average, had 23% higher baseline monthly card spending, held more non-AI digital subscriptions, and were more likely to be observed on an open-banking current-account feed.

To verify that demographic and earnings gradients in paid AI adoption are not driven by differences in account feed type or baseline spending intensity, Table A6 estimates linear probability and Poisson regressions of cumulative paid AI adoption on age band, gender, income band, account feed type, baseline 2022 spending tercile, and UK region simultaneously ($N = 201,038$). Holding all other recorded characteristics fixed, the adoption gaps by age, gender, and income remain large, retaining 68% to 90% of their unconditional magnitudes.

Similarly, for the regression-adjusted bars in Figures 2(a) and 3(a), we estimate population-weighted linear probability models on the rolling primary sample incorporating both direct and imputed app-store subscriptions:

$$Y_i = \sum_b \beta_b \mathbf{1}[\text{Group}_i = b] + Z_i'\gamma + \varepsilon_i, \tag{C.1}$$

where $Z_i$ controls for after-tax income band, gender, and UK region when estimating age profiles (Figure 2a), and for age band, gender, and UK region when estimating monthly net salary profiles (Figure 3a), with adjusted bars $\hat{\beta}_b + \bar{Z}'\hat{\gamma}$ evaluated at the weighted sample mean $\bar{Z}$.

## C.2 Two-Stage Population Weighting and External Validation

To scale our primary-account samples to the 55.5 million adults aged 18 and over in the United Kingdom, we construct monthly calibration weights $w_{i,t}$ in two stages:

1. *Stage 1 (Within-aggregator selection weight):* Among accounts with at least six active months in the wider aggregator panel, we estimate a logistic propensity model for selection into the rolling primary sample as a function of age band, region, local area classification, account feed type, and baseline non-AI activity deciles, trimming the resulting inverse-propensity weights at the 1st and 99th percentiles.

2. *Stage 2 (Iterative proportional fitting to UK adult benchmarks):* In each month $t$, we rake the trimmed Stage-1 weights to match five official national benchmark distributions from the Office for National Statistics (ONS) and HM Revenue and Customs (HMRC): the joint distribution of Age Band × ITL1 Region, Gender, five net after-tax income bands, ONS 2021 Census Output Area Classification supergroups, and urban-rural postal classes.

Figures A2 and A3 summarize our external validation and weighting diagnostics. In Figure A2, we benchmark our transaction records against external series, documenting the sharp jump in sub-£10 Apple App Store subscriptions upon the UK launch of ChatGPT Go (Panel a), the close alignment between open-banking net salary deposits and net pay percentiles constructed from ONS ASHE data (Panel b), the pass-through of the January 2024 employee National Insurance cut into observed take-home pay (Panel c), and primary-account Netflix adoption relative to BARB household survey benchmarks (Panel d). Figure A3 shows that our aggregate adoption trajectory is stable across weighting schemes with modest design effects (Panels a and b), compares launch-aligned paid AI diffusion against major video streaming services (Panel c), and shows temporary increases in total card spending around both AI and Disney+ subscription starts (Panel d).

# **Table A3:** The Account Sample Ladder: Counts and Direct Standalone AI Penetration in June 2026

| Sample layer | Screening rule | Active accounts | | Direct AI subscribers | | |
|---|---|---|---|---|---|---|
| | | Count | Share (%) | Count | Share (%) | Per 1,000 |
| *Panel A. Broad transacting bases* | | | | | | |
| All transacting records | $\geq$ 1 transaction in June 2026 | 1,424,855 | 100.0 | 21,837 | 100.0 | 15.3 |
| Wider panel | $\geq$ 6 active months since 2022 | 1,211,609 | 85.0 | 20,571 | 94.2 | 17.0 |
| *Panel B. Balanced panel screens (2022H1–2026H1)* | | | | | | |
| Fixed panel (no spend floor) | Active every half-year 2022H1–2026H1 | 419,652 | 29.5 | 7,717 | 35.3 | 18.4 |
| Regular accounts | Active every half-year, median spend $\geq$ £100/mo | 357,037 | 25.1 | 7,490 | 34.3 | 21.0 |
| **Balanced primary panel** | Regular accounts, $\geq$ 20 non-AI txns/mo | 178,109 | 12.5 | 6,029 | 27.6 | 33.9 |
| *Panel C. Rolling 12-month primary screens* | | | | | | |
| **Rolling primary (threshold 20)** | Mean $\geq$ 20 non-AI txns/mo over prior 12 months | 446,429 | 31.3 | 15,581 | 71.4 | 34.9 |
| Rolling threshold 15 | Mean $\geq$ 15 non-AI txns/mo over prior 12 months | 517,714 | 36.3 | 16,655 | 76.3 | 32.2 |
| Rolling threshold 25 | Mean $\geq$ 25 non-AI txns/mo over prior 12 months | 396,766 | 27.8 | 14,662 | 67.1 | 37.0 |
| Rolling 9-of-12 screen | $\geq$ 20 non-AI txns in $\geq$ 9 of prior 12 months | 321,168 | 22.5 | 11,999 | 54.9 | 37.4 |
| In-month activity screen | $\geq$ 20 non-AI txns in June 2026 | 503,831 | 35.4 | 16,216 | 74.3 | 32.2 |
| Spend, debits, and tags screen | $\geq$ £250 spend, $\geq$ 10 debits, $\geq$ 5 tags (9/12 mos) | 349,239 | 24.5 | 12,574 | 57.6 | 36.0 |
| *Panel D. Persistence and subscription-active benchmarks* | | | | | | |
| Continuously active (12 months) | $\geq$ 1 transaction in each of prior 12 months | 549,199 | 38.5 | 13,657 | 62.5 | 24.9 |
| Recurring-subscription payers | $\geq$ 1 non-AI recurring subscription in June 2026 | 631,614 | 44.3 | 17,748 | 81.3 | 28.1 |
| *Panel E. Breakdown by data feed* | | | | | | |
| Wider panel: open banking | Current-account open-banking feeds | 371,212 | 26.1 | 13,646 | 62.5 | 36.8 |
| Wider panel: card-only | Credit and debit card feeds | 840,397 | 59.0 | 6,925 | 31.7 | 8.2 |
| Balanced primary: open banking | Primary open-banking current accounts | 104,136 | 7.3 | 4,674 | 21.4 | 44.9 |
| Balanced primary: card-only | Primary credit and debit card accounts | 73,973 | 5.2 | 1,355 | 6.2 | 18.3 |

**Notes:** Each row reports the number of consumer accounts active in June 2026, the number holding an active direct paid standalone AI subscription, the share of all transacting accounts and active direct subscribers captured by that sample, and the resulting direct subscription rate per 1,000 accounts. Sample screening criteria are described in Appendix C.

# **Table A4:** Direct Paid AI Adoption Across Sample Definitions and Time Windows

| | Definition | Accounts (Jan–Jun 2026) | Share of the stated population (%) | | | |
|---|---|---|---|---|---|---|
| | | | June 2026 | January–June 2026 | July 2025–June 2026 | Jul 2025–Jun 2026, age-reweighted |
| *This paper: paid subscriptions to standalone AI products* | | | | | | |
| All transacting records | every UK consumer record | 2,043,430 | 1.5 | 1.8 | 1.9 | — |
| Wider panel | $\geq$6 active months since 2022 | 1,703,286 | 1.7 | 2.0 | 2.1 | 1.6 |
| Regular accounts | active every half-year 2022H1–2026H1; median monthly spend $\geq$ £100 | 439,616 | 2.0 | 2.4 | 2.8 | 2.2 |
| **Primary accounts** | regular accounts averaging $\geq$20 non-AI transactions a month | 201,038 | 3.3 | 4.1 | 4.7 | 3.6 |
| primary accounts fed by open banking | the whole current account is observed | 119,968 | 4.3 | 5.3 | 6.1 | 4.4 |

**Notes:** Share of transacting consumer accounts holding an active direct paid standalone AI subscription during the stated window (June 2026 single month, 2026H1 six-month window, and 12-month window ending June 2026) across progressively stricter activity screens, before and after post-stratification reweighting to the ONS adult age distribution.

# **Table A5:** Baseline 2022 Descriptive Statistics: Paid AI Adopters vs. Balanced Primary Sample

| | Group mean or share | | Contrast | |
|---|---|---|---|---|
| | Primary accounts | Paid-AI adopters | Ratio | Std. diff. |
| Age (band midpoint, years) | 39.9 | 36.3 | 0.91 | −0.29 |
| Total spend 2022 (GBP) | 20,624.0 | 25,283.8 | 1.23 | 0.29 |
| Spend percentile within card type | 50.0 | 54.8 | 1.10 | 0.17 |
| Online share of spend (%) | 46.3 | 51.0 | 1.10 | 0.23 |
| Transactions in 2022 | 759.4 | 883.6 | 1.16 | 0.27 |
| Active recurring subscriptions (monthly mean) | 9.3 | 11.1 | 1.19 | 0.26 |
| Streaming subscriptions (monthly mean) | 0.5 | 0.7 | 1.29 | 0.20 |
| *Base* | *192,939–201,038* | *14,595–14,798* | | |
| *Age band (% of column)* | | | | |
| <30 | 22.5 | 27.3 | 1.21 | 0.11 |
| 30–39 | 34.2 | 41.7 | 1.22 | 0.16 |
| 40–49 | 22.4 | 22.2 | 0.99 | −0.00 |
| 50–59 | 14.2 | 7.4 | 0.52 | −0.19 |
| 60–69 | 5.3 | 1.2 | 0.23 | −0.18 |
| ≥70 | 1.3 | 0.2 | 0.14 | −0.10 |
| *Base* | *192,939* | *14,595* | | |
| *Card type (% of column)* | | | | |
| Open-banking feed | 59.7 | 74.5 | 1.25 | 0.30 |
| Debit card | 21.4 | 7.4 | 0.35 | −0.34 |
| Credit card | 18.9 | 18.1 | 0.96 | −0.02 |
| *Base* | *201,038* | *14,798* | | |
| *Region (% of column)* | | | | |
| South East | 14.0 | 14.7 | 1.05 | 0.02 |
| London | 12.1 | 18.4 | 1.53 | 0.20 |
| North West | 11.8 | 10.7 | 0.91 | −0.03 |
| Scotland | 10.7 | 9.1 | 0.84 | −0.05 |
| East of England | 8.8 | 8.6 | 0.98 | −0.01 |
| South West | 8.7 | 8.4 | 0.97 | −0.01 |
| Yorkshire and The Humber | 8.1 | 7.0 | 0.86 | −0.04 |
| West Midlands | 8.0 | 7.7 | 0.96 | −0.01 |
| East Midlands | 7.2 | 6.4 | 0.90 | −0.03 |
| Wales | 5.0 | 4.4 | 0.88 | −0.03 |
| North East | 3.8 | 3.0 | 0.78 | −0.04 |
| Northern Ireland | 1.9 | 1.6 | 0.87 | −0.02 |
| *Base* | *199,006* | *14,647* | | |
| *ONS supergroup (% of column)* | | | | |
| Multicultural Metropolitans | 19.4 | 20.8 | 1.07 | 0.04 |
| Urbanites | 16.9 | 17.4 | 1.03 | 0.01 |
| Hard-Pressed Living | 15.6 | 10.1 | 0.65 | −0.15 |
| Suburbanites | 14.0 | 14.0 | 1.00 | −0.00 |
| Ethnicity Central | 10.6 | 14.3 | 1.35 | 0.12 |
| Constrained City Dwellers | 8.2 | 6.5 | 0.79 | −0.06 |
| Rural Residents | 8.2 | 6.0 | 0.73 | −0.08 |
| Cosmopolitans | 7.0 | 10.9 | 1.55 | 0.15 |
| *Base* | *45,265* | *2,807* | | |
| All primary accounts | 201,038 | 14,798 | | |

**Notes:** Balanced primary panel ($N = 201,038$ ever-active accounts, including 14,798 cumulative paid AI adopters). All spending, transaction-count, online-share, and recurring-subscription variables are measured in calendar year 2022 prior to the diffusion of consumer generative AI. Standardized differences divide the difference in group means by the full-sample standard deviation.

# **Table A6:** Joint Multivariate Linear Probability and Poisson Models of Paid AI Adoption

| | | | Gap from base, pp | | Rate ratio to base | |
|---|---|---|---|---|---|---|
| | Share of accounts, % | Adoption rate, % | Alone | Joint | Alone | Joint |
| *Age band* | | | | | | |
| Under 30 | 21.6 | 9.17 | 7.59 (0.18) | 6.08 (0.19) | 5.80 | 4.43 |
| 30–39 | 32.8 | 9.21 | 7.63 (0.16) | 5.30 (0.17) | 5.82 | 3.98 |
| 40–49 | 21.5 | 7.50 | 5.91 (0.17) | 3.82 (0.17) | 4.74 | 3.38 |
| 50–59 | 13.6 | 3.95 | 2.37 (0.16) | 1.29 (0.16) | 2.50 | 2.08 |
| 60 and over (base) | 6.4 | 1.58 | | | | |
| Not recorded | 4.0 | 2.51 | 0.92 (0.21) | 1.86 (0.30) | 1.58 | 2.34 |
| *Gender* | | | | | | |
| Women (base) | 41.9 | 6.93 | | | | |
| Men | 32.0 | 10.58 | 3.64 (0.15) | 3.28 (0.16) | 1.53 | 1.45 |
| Not recorded | 26.0 | 4.09 | −2.85 (0.12) | 2.83 (0.33) | 0.59 | 1.36 |
| *Monthly income band* | | | | | | |
| Under £1,500 (base) | 35.9 | 5.04 | | | | |
| £1,500–2,000 | 20.2 | 6.72 | 1.67 (0.15) | 0.19 (0.15) | 1.33 | 1.07 |
| £2,000–2,500 | 11.4 | 9.79 | 4.75 (0.21) | 2.60 (0.22) | 1.94 | 1.43 |
| £2,500–3,000 | 6.7 | 11.13 | 6.09 (0.28) | 3.62 (0.29) | 2.21 | 1.55 |
| £3,000–4,000 | 5.9 | 12.97 | 7.92 (0.32) | 5.05 (0.33) | 2.57 | 1.71 |
| Over £4,000 | 5.3 | 14.17 | 9.12 (0.35) | 6.21 (0.36) | 2.81 | 1.86 |
| Not recorded | 14.6 | 5.59 | 0.54 (0.16) | 1.13 (0.24) | 1.11 | 1.20 |
| *Record type* | | | | | | |
| Open banking (base) | 59.7 | 9.18 | | | | |
| Credit card | 18.9 | 7.03 | −2.15 (0.16) | −2.34 (0.19) | 0.77 | 0.77 |
| Debit card | 21.4 | 2.56 | −6.62 (0.11) | −5.05 (0.34) | 0.28 | 0.39 |
| *Spending tercile* | | | | | | |
| Bottom (base) | 33.3 | 6.26 | | | | |
| Middle | 33.3 | 6.80 | 0.54 (0.13) | 0.53 (0.13) | 1.09 | 1.07 |
| Top | 33.3 | 9.02 | 2.75 (0.14) | 1.93 (0.15) | 1.44 | 1.30 |
| *Region* | | | | | | |
| London (base) | 11.9 | 11.24 | | | | |
| East Midlands | 7.1 | 6.60 | −4.64 (0.29) | −3.93 (0.29) | 0.59 | 0.65 |
| East of England | 8.7 | 7.20 | −4.05 (0.28) | −3.61 (0.28) | 0.64 | 0.68 |
| North East | 3.8 | 5.72 | −5.53 (0.33) | −4.72 (0.33) | 0.51 | 0.57 |
| North West | 11.7 | 6.67 | −4.57 (0.26) | −3.72 (0.26) | 0.59 | 0.67 |
| Northern Ireland | 1.8 | 6.43 | −4.82 (0.45) | −4.12 (0.45) | 0.57 | 0.63 |
| Scotland | 10.6 | 6.21 | −5.03 (0.26) | −4.48 (0.26) | 0.55 | 0.60 |
| South East | 13.9 | 7.72 | −3.53 (0.26) | −3.29 (0.26) | 0.69 | 0.72 |
| South West | 8.6 | 7.18 | −4.07 (0.28) | −3.46 (0.28) | 0.64 | 0.70 |
| Wales | 4.9 | 6.51 | −4.73 (0.32) | −3.99 (0.32) | 0.58 | 0.64 |
| West Midlands | 7.9 | 7.07 | −4.18 (0.29) | −3.33 (0.29) | 0.63 | 0.71 |
| Yorkshire and The Humber | 8.0 | 6.34 | −4.90 (0.28) | −4.13 (0.28) | 0.56 | 0.63 |
| Not recorded | 1.0 | 7.43 | −3.81 (0.62) | −3.52 (0.61) | 0.66 | 0.68 |
| Accounts, every column | 201,038 | | | | | |

**Notes:** Balanced primary panel ($N = 201,038$). The dependent variable is an indicator for having held a direct paid AI subscription at a standalone seller by July 2026 (scaled in percentage points for LPM). "Alone" columns report bivariate differences (LPM) or rate ratios (Poisson) relative to each characteristic's reference category; "Joint" columns include indicators for all six characteristics simultaneously. Heteroskedasticity-robust standard errors are in parentheses.

# **Figure A2:** Validating Transaction Records Against External Benchmarks

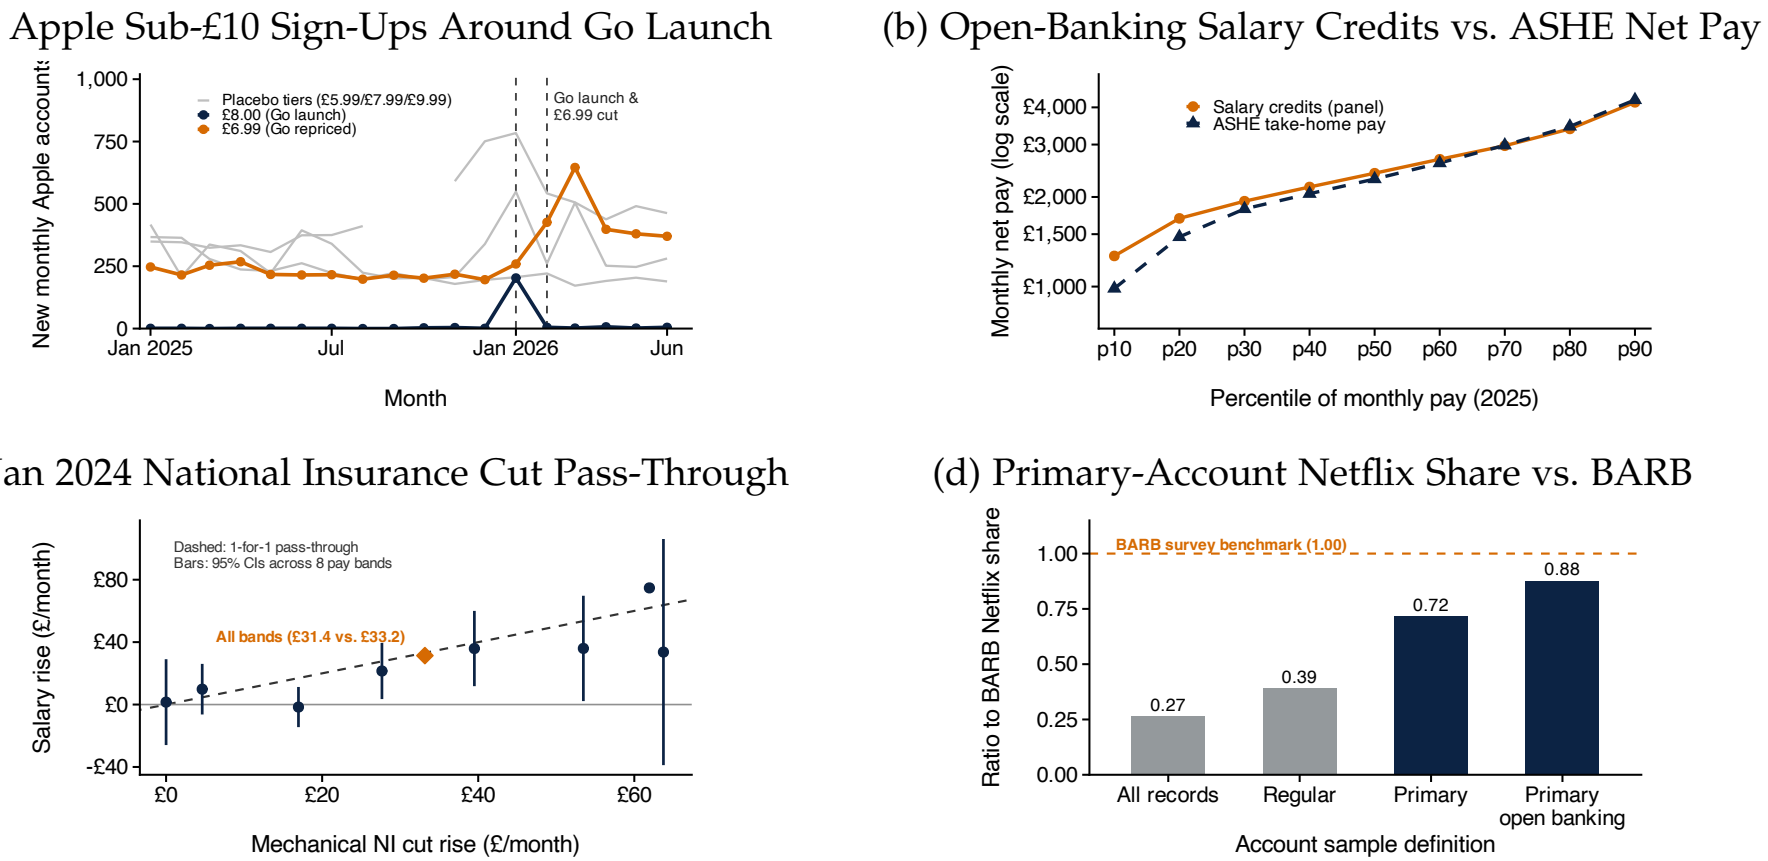


**Notes:** Panel (a) plots monthly new recurring Apple App Store subscription runs at OpenAI's entry-level ChatGPT Go price tiers (£8.00 at its January 2026 UK launch, jumping from ∼49 pre-launch to ∼200, and £6.99 after the February 2026 repricing, reaching 658 total runs per month or 410 net of the ∼248 pre-launch £6.99 non-AI baseline). Panel (b) compares percentiles of monthly net take-home salary credits in our open-banking panel against net take-home pay percentiles constructed from the official UK Annual Survey of Hours and Earnings (ASHE). Panel (c) plots the event-time change in monthly net salary deposits around the January 2024 two-percentage-point cut in UK employee National Insurance contributions. Panel (d) compares the share of rolling primary accounts paying Netflix against the implied per-adult primary-payer benchmark derived from BARB's quarterly UK household Netflix survey (∼59% household penetration scaled across 1.9 adults per UK household to 31%–38% implied per-adult primary payers).

# **Figure A3:** Population Weighting Trajectories, Design Effects, Launch Diffusion, and Event-Time Spending

(a) Adoption Trajectories Across Weighting Schemes

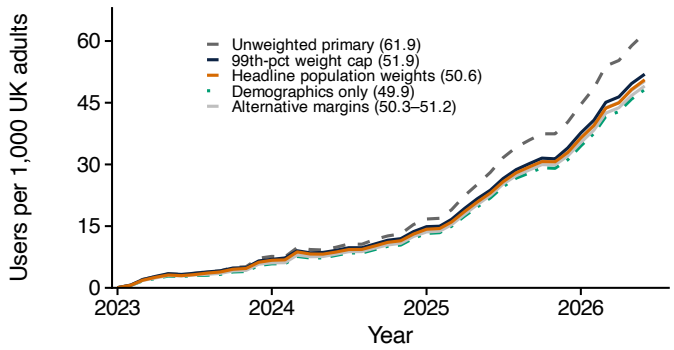


(b) Weighting Design Effect (DEFF) Over Time

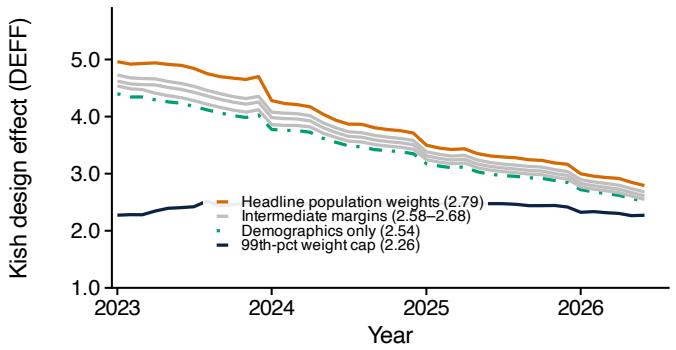


(c) Launch-Aligned Diffusion vs. Streaming Services

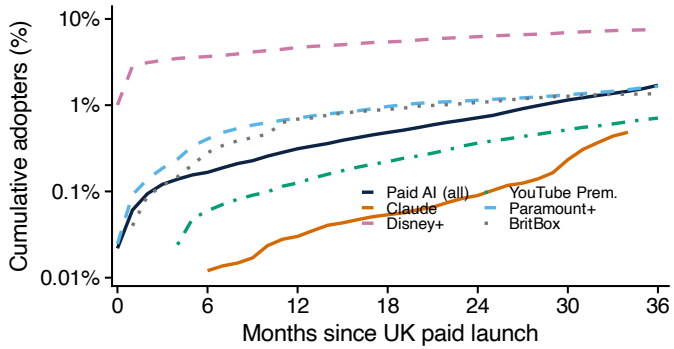


(d) Event-Time Card Spending Around Adoption

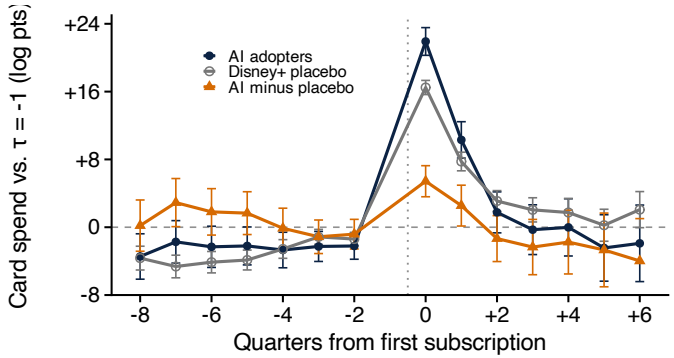


**Notes:** Panel (a) compares monthly active paid standalone AI subscribers per 1,000 across weighting specifications on the rolling primary sample. Panel (b) plots the monthly Kish design effect (DEFF $= 1 + \mathrm{CV}(w)^2$) of the population calibration weights across specifications. Panel (c) compares launch-aligned cumulative direct card adoption trajectories over the first 36 months since UK launch (or panel start for YouTube Premium) for paid AI assistants against major video streaming services, estimated on the unweighted extended 2019–2026 historical card subpanel. Panel (d) plots event-time log total card expenditure around first subscription start for paid AI adopters and a Disney+ adoption placebo.

# D Individual Occupation and Employer Industry Classification

Standard bank and card datasets record demographics and geographic location but lack occupation or employer industry. To overcome this limitation of our data, we exploit two administrative signals in open-banking records. The first is the incoming salary credits from named employers, while the second is payments to professional regulators, royal colleges, learned societies, and trade unions. In this appendix section we provide the details of our classification and validation approach.

## D.1 Hierarchical Occupation Classification and Validation

To classify individual accounts into occupations we use a four-stage approach.

1. *Stage 1: Direct professional fees and statutory registers.* Many UK professions routinely entail an annual registration fee or dues to statutory regulators (such as the General Medical Council, Nursing and Midwifery Council, or Solicitors Regulation Authority), royal colleges, chartered institutes, or other job-specific bodies. Matching outbound payments against our curated directory of 466 job-specific fee-taking entities directly identifies members of these professions.

2. *Stage 2: Single-occupation employers.* We next match workers to employers based on incoming payroll strings. Combined with the salary and demographics of each worker, we can back out a signal of each account holder's occupation. This is a clear signal in many industries, such as airlines, accounting partnerships, law firms, dental practices, veterinary clinics, and haulage fleets. Combining Stage 1 fee payers and Stage 2 employers yields 70,599 salaried accounts with a sharp occupational prior, used for model training and validation (Table A8).

3. *Stage 3: Industry $\times$ age $\times$ net-salary cells.* Within multi-occupation employers (such as hospital trusts, universities, local authorities, retail chains, and banks), we combine resolved employer SIC industry with age-by-gender net pay percentiles constructed from ONS ASHE data to form structured occupational priors, helping separate support staff, technical staff, and managerial roles which co-occur in the same sector.

4. *Stage 4: Transaction classifier.* For the remaining accounts we train and test an out-of-fold gradient-boosted decision tree model, which uses the information from the above three steps to identify the most likely occupation for all accounts. We exclude all AI transactions from this model, though for the best possible predicted occupation this would be a valuable feature in general.

Across the salaried panel, this four-stage classification approach produces 94.2% of accounts with 88.4% concordance against a blinded human audit of 29,436 stratified records. When we mask Stage 1 fee payments and predict occupation solely from employer and non-AI transaction features, the out-of-fold classifier achieves a median AUC of 0.874 across anchor professions (Figure A4a–b).

Table A9 reports paid AI adoption across the 34 fee-verified professions meeting our main-text sample-size thresholds, comparing direct web payments against all-channel estimates incorporating app-store subscriptions. The more than ninefold spread in direct-payment adoption between technical or academic professions and hands-on care or manual roles holds across sample thresholds.

Finally, Table A10 and Figure A4(c)–(d) validate our UK occupational adoption rates against external workplace AI usage surveys (Bick, Blandin, and Deming, 2026) and task-exposure indices (Eloundou et al., 2024; Felten, Raj, and Seamans, 2021). Paid AI adoption correlates strongly with US survey usage across all 26 broad SOC sub-major groups ($r = 0.95$, Panel A) and across our 51 fee-verified professions, aligns closely with generative-AI task-exposure measures while correlating negatively with pre-LLM routine computerization (Frey and Osborne, 2017), and is strongly associated with survey usage and GPT-4 exposure in individual-level regressions controlling for age, gender, region, and take-home pay (Panel C). Likewise, in our earnings regressions (Figure 3b), accounting for occupation, industry, location, and demographics explains roughly two-thirds of the raw income gradient whether we use the full occupational cascade or restrict to directly observed professional membership and employer labels.

# **Table A7:** Employer SIC 2007 Industry Resolution Coverage, Institutional Sectors, and Audit Precision

**Panel A. Employer SIC Division and Named-Organization Resolution Coverage**

| Sample / unit | $N$ | Any SIC division | Confidence $\geq 0.60$ | Confidence $\geq 0.80$ | Named organization | Generic payroll processor |
|---|---|---|---|---|---|---|
| All salary spells | 533,958 | 63.8% | 62.3% | 59.7% | 62.3% | 20.0% |
| Named salary spells (excl. generic BACS) | 427,439 | 79.7% | 77.8% | 74.6% | 77.8% | 0.0% |
| Salaried accounts (longest-duration employer) | 390,404 | 65.0% | 63.4% | 61.0% | 63.5% | 19.9% |
| Salaried open-banking outcome accounts | 374,902 | 64.0% | 62.5% | 60.2% | 62.6% | 20.4% |
| Salaried accounts (any resolved employer) | 390,404 | 68.4% | 66.7% | 64.2% | 66.8% | — |

**Panel B. Resolved Named Employers by Institutional Sector**

| Employer sector / organization type | Distinct organizations | Account–employer links | Share of links (%) |
|---|---|---|---|
| Private-sector company (Companies House) | 15,134 | 175,624 | 58.2% |
| NHS hospital trust or health board | 232 | 30,505 | 10.1% |
| Local authority or municipal council | 373 | 17,908 | 5.9% |
| Registered charity or non-profit | 996 | 12,320 | 4.1% |
| State school or multi-academy trust | 1,576 | 11,124 | 3.7% |
| University or further-education college | 410 | 9,097 | 3.0% |
| Central government or public body | 264 | 6,606 | 2.2% |
| General medical practice | 187 | 568 | 0.2% |
| Other organization types | 741 | 38,237 | 12.7% |
| All distinct named organizations | 19,876 | 301,989 | 100.0% |

**Panel C. Independent Verification Audit of 300 Stratified Resolved Payroll Strings**

| Audit weighting scheme | Audited strings | Organization confirmed | Organization indeterminate | Organization refuted | SIC division confirmed |
|---|---|---|---|---|---|
| Stratified by spell-stratum size | 300 | 87.0% | 11.0% | 2.0% | 80.6% |
| Weighted by accounts $\times$ stratum size | 300 | 85.5% | 13.7% | 0.8% | 78.2% |

**Notes:** This table summarizes employer SIC 2007 industry resolution across open-banking salaried accounts ($N = 319,298$ SIC-resolved accounts out of 491,228 salaried accounts, or 65.0%, including 301,989 accounts linked to 19,876 named organizations and 17,309 accounts resolved directly from industry-specific payroll descriptions). Panel A reports coverage across confidence thresholds and payroll processors; Panel B reports resolved named employers and linked accounts by institutional sector; Panel C reports human-audit verification on 300 stratified payroll strings (80.6% confirmed, 16.7% indeterminate from the narration string alone, and 2.7% refuted, yielding 96.8% precision among determinable strings).

# **Table A8:** Hierarchical Occupation Classification Cascade and Professional Fee Anchor Groups

**Panel A. Occupational Registry Entities by Discovery Channel and Organization Type**

| Discovery channel | Entities | Organization type | | | Label-producing | On $\geq 10$ accounts | Matched in both channels |
|---|---|---|---|---|---|---|---|
| | | Statutory registers | Professional bodies | Trade / employer orgs | | | |
| Curated statutory registers and unions | 205 | 37 | 109 | 59 | 146 | 183 | 182 |
| Top-down SOC2020 institutional enumeration | 55,370 | 212 | 2,657 | 52,501 | 181 | 3,465 | 0 |
| Bottom-up merchant descriptor vocabulary | 10,778 | 52 | 353 | 10,373 | 136 | 65 | 0 |
| Residual high-frequency stem scan | 39 | 0 | 3 | 36 | 3 | 10 | 0 |
| All occupational entities | 66,392 | 301 | 3,122 | 62,969 | 466 | 3,723 | 182 |

**Panel B. Gold-Standard Fee-Paying Accounts Across Major Statutory and Union Families**

| Occupation family | Full panel | Salaried OB universe | Other OB accounts | Card-only panel | New entrants | Lapsed / one-off | Professional body / union |
|---|---|---|---|---|---|---|---|
| Nursing and midwifery (NMC / RCN) | 23,535 | 17,341 | [s] | <10 | 2,044 | 1,019 | 6,211 |
| School teachers (NEU / NASUWT / GTCS) | 15,338 | 11,191 | 4,136 | 11 | 1,449 | 2,245 | 5,622 |
| Allied health professionals (HCPC) | 8,492 | 6,756 | 1,736 | 0 | 2,015 | 505 | 3,640 |
| Social care registers (SSSC / SCW / NISCC) | 7,310 | 5,228 | 2,047 | 35 | 260 | 1,460 | 0 |
| Medical doctors (GMC / BMA / Royal Colleges) | 4,243 | 3,293 | 940 | [s] | 288 | 336 | 3,131 |
| Social workers (SWE) | 3,796 | 2,474 | 795 | 527 | 0 | 659 | 469 |
| Dental professionals (GDC / BDA) | 2,205 | 1,417 | 592 | 196 | 147 | 777 | 336 |
| Pharmacists and technicians (GPhC / PDA) | 1,500 | 1,151 | 349 | 0 | 25 | 239 | 277 |
| Veterinary surgeons and nurses (RCVS) | 936 | 590 | 163 | 183 | 22 | 290 | 276 |
| Security operatives (SIA) | 897 | 676 | 208 | 13 | 0 | 72 | 0 |
| Rail and transport staff (RMT / ASLEF) | 611 | 468 | 143 | 0 | 56 | 37 | 3,239 |
| Registered childminders (Ofsted / PACEY) | 510 | 381 | 129 | 0 | 0 | 159 | 273 |
| Commercial pilots (CAA / BALPA) | 276 | 191 | [s] | <10 | 25 | 23 | 26 |
| Architects (ARB / RIBA) | 263 | 213 | 34 | 16 | <10 | 117 | 127 |
| Solicitors and barristers (SRA / Law Society) | 119 | 97 | 22 | 0 | 0 | 40 | 91 |
| Smaller statutory families (11 groups) | 568 | 329 | 95 | 144 | <10 | 225 | 792 |
| All statutory and union families | 70,599 | 51,796 | 17,665 | 1,138 | 6,341 | 8,203 | 24,510 |

**Notes:** Panel A summarizes occupational registry entities (66,392 total entities, including 466 professional regulators, royal colleges, and trade unions) by discovery channel and organization type, alongside the 70,599 accounts paying statutory, royal-college, or trade-union fees across all feeds (51,796 salaried open-banking accounts, 17,665 other open-banking accounts, and 1,138 credit-card-only accounts). Panel B summarizes ground-truth accounts across statutory professional regulators, royal colleges, chartered institutes, trade unions, and single-occupation anchor employers (where deduplicating the 51,796 salaried fee-payer rows across overlapping professional bodies yields 46,419 unique salaried Stage 1 fee payers, which combine with 24,180 Stage 2 single-occupation employer anchors to form the 70,599 salaried ground-truth anchors used in supervised training).

# **Table A9:** Ground-Truth Occupational Classes, Source Types, and Paid AI Adoption Across the 34 Primary-Panel Professions

| Profession / fee class | Source type | Rolling primary base calibrated to UK adults | | | | All gold-standard members, unweighted | | |
|---|---|---|---|---|---|---|---|---|
| | | Active June 2026 direct [90% CI] | + App-store in-app | Ever paid since 2022 | $N$ | Active June 2026 direct | Ever paid since 2022 | $N$ |
| IT professionals (BCS) | Prof. body / union | 20.3 [10.9, 32.2] | 22.0 | 40.0 | 53 | 6.7 | 22.7 | 163 |
| HE and FE lecturers (UCU) | Prof. body / union | 16.0 [12.8, 19.1] | 17.7 | 23.0 | 684 | 5.9 | 13.1 | 2,019 |
| Insurance and finance (CII) | Prof. body / union | 15.8 [10.6, 22.2] | 18.9 | 20.8 | 290 | 3.0 | 7.9 | 1,171 |
| Project and procurement (APM/CIPS) | Prof. body / union | 13.8 [8.7, 22.1] | 18.3 | 22.4 | 248 | 4.4 | 11.2 | 735 |
| Marketing and PR (CIM/CIPR) | Prof. body / union | 12.8 [8.8, 17.2] | 18.3 | 20.4 | 192 | 4.2 | 11.6 | 662 |
| Artistic and literary occupations (SOC 341) | Prof. body / union | 9.7 [4.6, 14.2] | 11.0 | 16.2 | 109 | 3.8 | 9.2 | 261 |
| Surveyors and planners (RICS/RTPI) | Prof. body / union | 8.2 [4.3, 12.9] | 11.3 | 13.0 | 134 | 4.1 | 7.9 | 340 |
| Postal and telecoms (CWU) | Prof. body / union | 7.7 [4.0, 11.4] | 10.0 | 14.1 | 150 | 2.8 | 9.7 | 507 |
| Chartered engineers | Prof. body / union | 7.5 [5.6, 9.4] | 10.1 | 12.8 | 552 | 3.6 | 9.1 | 1,713 |
| School teachers (NEU/NASUWT/GTCS) | Statutory register | 7.5 [6.9, 8.2] | 9.9 | 19.3 | 5,560 | 3.1 | 12.5 | 15,067 |
| Security operatives (SIA) | Statutory register | 7.0 [4.0, 9.7] | 10.5 | 12.0 | 292 | 2.0 | 6.3 | 889 |
| Medical doctors (GMC/BMA) | Statutory register | 7.0 [5.9, 8.4] | 11.0 | 14.9 | 1,160 | 2.5 | 7.9 | 4,175 |
| Pharmacists and technicians (GPhC) | Statutory register | 6.8 [5.1, 8.8] | 10.6 | 13.0 | 631 | 3.1 | 6.7 | 1,486 |
| Psychotherapy and counselling (BACP) | Prof. body / union | 6.8 [5.1, 9.6] | 9.9 | 12.5 | 798 | 2.7 | 7.4 | 2,394 |
| Musicians (Musicians' Union) | Prof. body / union | 6.7 [4.2, 10.2] | 10.8 | 11.1 | 203 | 2.9 | 7.0 | 647 |
| HR professionals (CIPD) | Prof. body / union | 6.5 [5.1, 7.8] | 10.9 | 13.6 | 1,038 | 2.5 | 7.0 | 3,070 |
| Fitness instructors | Prof. body / union | 6.1 [3.9, 8.7] | 9.4 | 10.5 | 255 | 2.8 | 8.3 | 791 |
| Performing arts (Equity) | Prof. body / union | 5.8 [4.2, 7.3] | 9.8 | 11.3 | 646 | 2.0 | 6.5 | 2,271 |
| Science and engineering union (Prospect) | Prof. body / union | 4.9 [3.6, 5.9] | 8.3 | 11.5 | 1,104 | 2.2 | 7.5 | 3,024 |
| Banking staff | Prof. body / union | 4.9 [2.4, 7.2] | 8.3 | 8.2 | 243 | 1.3 | 3.2 | 950 |
| Taxi and courier drivers | Prof. body / union | 4.8 [4.1, 5.6] | 7.6 | 11.6 | 2,513 | 1.1 | 4.1 | 16,813 |
| Industrial union (Community) | Prof. body / union | 4.7 [2.4, 6.9] | 6.8 | 9.3 | 205 | 2.5 | 6.0 | 635 |
| Allied health professionals (HCPC) | Statutory register | 3.8 [3.3, 4.4] | 6.3 | 8.6 | 3,334 | 1.7 | 5.1 | 8,387 |
| Natural and social scientists | Prof. body / union | 3.6 [1.9, 5.9] | 4.3 | 10.9 | 173 | 2.7 | 10.4 | 594 |
| Social workers (SWE) | Statutory register | 3.4 [2.2, 5.1] | 7.0 | 7.9 | 1,234 | 1.0 | 3.5 | 3,754 |
| Accountancy (ACCA/ICAEW/CIMA) | Prof. body / union | 3.4 [2.2, 4.6] | 6.0 | 8.0 | 844 | 1.0 | 4.7 | 3,449 |
| Civil service unions (PCS/FDA) | Prof. body / union | 3.3 [2.5, 4.1] | 6.3 | 7.5 | 1,722 | 1.2 | 4.3 | 5,129 |
| Veterinary professionals (RCVS) | Statutory register | 3.1 [1.4, 4.7] | 5.1 | 6.8 | 327 | 1.3 | 3.9 | 928 |
| Dental professionals (GDC) | Statutory register | 3.1 [1.8, 4.4] | 6.4 | 6.6 | 759 | 1.1 | 3.0 | 2,180 |
| Construction site operatives (CSCS/CITB) | Prof. body / union | 2.9 [2.2, 3.7] | 6.3 | 7.0 | 1,489 | 1.4 | 4.6 | 3,869 |
| Electricians | Prof. body / union | 2.9 [2.2, 5.1] | 7.4 | 6.5 | 469 | 1.1 | 4.0 | 1,696 |
| Firefighters (FBU) | Prof. body / union | 2.5 [1.3, 3.8] | 5.6 | 4.9 | 434 | 1.0 | 3.1 | 1,197 |
| Nursing and midwifery (NMC/RCN) | Statutory register | 2.5 [2.1, 2.8] | 5.1 | 5.7 | 8,627 | 1.0 | 3.3 | 23,140 |
| Social care registers (SSSC/SCW/NISCC) | Statutory register | 1.7 [1.3, 2.3] | 4.8 | 5.2 | 2,302 | 0.6 | 2.8 | 7,239 |

**Notes:** This table reports paid AI adoption across the 34 fee-verified professions with at least 50 primary accounts and at least 10 active June 2026 adopters (drawn from the 51 primary-sample professions with $N \geq 30$). Columns report organization source type, active June 2026 direct (matching Figure 4) and all-channel (including app-store) adoption rates on the rolling primary sample reweighted to UK adults, cumulative ever-adoption rates, and unweighted direct adoption rates across all ground-truth accounts. Across the outcome-blind set of professions with at least 150 primary accounts regardless of adopter count, direct adoption spans 16.0% (university lecturers) to 1.7% (social care workers), a 9.21-fold ratio (18.9% to 4.3%, or 4.44-fold, on the all-channel series).

# **Table A10:** Occupational Paid AI Adoption Regressions: Survey Usage, Task Exposure, and Account Controls

**Panel A. Occupation-Level Regressions on US Survey AI Usage (Bick, Blandin, and Deming, 2026) (1.74M Open-Banking Universe, Mean 2.1%)**

| Occupation partition | Specification | No fixed effects | SOC1 major group FE | Worker controls (pay, age, gender) | SOC1 FE + worker controls |
|---|---|---|---|---|---|
| SOC1 major groups (9 groups) | Level-level (pp per SD) | 0.627 (0.060) $R^2 = 0.96$ | — | — | — |
| | Log-log (elasticity) | 0.765 (0.121) $R^2 = 0.93$ | — | — | — |
| | Rank-rank slope | 0.968 (0.121) $R^2 = 0.94$ | — | — | — |
| SOC2 sub-major groups (26 groups) | Level-level (pp per SD) | 0.734 (0.034) $R^2 = 0.90$ | 0.473 (0.130) $R^2 = 0.95$ | 0.568 (0.113) $R^2 = 0.91$ | 0.419 (0.212) $R^2 = 0.95$ |
| | Log-log (elasticity) | 1.000 (0.054) $R^2 = 0.87$ | 0.509 (0.171) $R^2 = 0.96$ | 0.611 (0.199) $R^2 = 0.91$ | 0.384 (0.229) $R^2 = 0.97$ |
| | Rank-rank slope | 0.930 (0.073) $R^2 = 0.86$ | 0.593 (0.189) $R^2 = 0.95$ | 0.640 (0.186) $R^2 = 0.91$ | 0.516 (0.303) $R^2 = 0.95$ |

**Panel B. Occupation-Level Regressions on GPT-4 Task Exposure (Eloundou et al., 2024) (Rolling Primary Base, APS-Weighted Mean 10.4%)**

| Occupation partition | Specification | No fixed effects | SOC1 major group FE | Worker controls (pay, age, gender) | SOC1 FE + worker controls |
|---|---|---|---|---|---|
| Fee-verified professions (51 classes) | Level-level (pp per SD) | 5.122 (1.612) $R^2 = 0.63$ | 5.812 (2.578) $R^2 = 0.76$ | 4.422 (1.654) $R^2 = 0.69$ | 5.021 (2.553) $R^2 = 0.78$ |
| | Log-log (elasticity) | 0.931 (0.279) $R^2 = 0.52$ | 1.016 (0.396) $R^2 = 0.70$ | 0.780 (0.165) $R^2 = 0.65$ | 0.661 (0.355) $R^2 = 0.74$ |
| | Rank-rank slope | 0.734 (0.134) $R^2 = 0.54$ | 0.630 (0.229) $R^2 = 0.68$ | 0.620 (0.164) $R^2 = 0.62$ | 0.359 (0.281) $R^2 = 0.72$ |

**Panel C. Individual Account-Level Regressions on Rolling Primary Base (Level-Level Slope, pp per SD)**

| Regressor and occupational assignment ($N$ accounts; clusters) | No fixed effects | Age × gender × region FE | + Individual income band FE | + Occupation mean net pay |
|---|---|---|---|---|
| US survey AI usage across 26 SOC2 groups (191,220 accounts; 26 occupational clusters; mean 6.3%) | 3.21 (0.37) $p_{\text{wcb}} < 0.001$ | 2.71 (0.25) $p_{\text{wcb}} < 0.001$ | 2.34 (0.26) $p_{\text{wcb}} < 0.001$ | 3.24 (0.57) $p_{\text{wcb}} < 0.001$ |
| US survey AI usage across 9 SOC1 major groups (176,643 accounts; 9 occupational clusters; mean 7.4%) | 2.55 (0.25) $p_{\text{wcb}} = 0.010$ | 2.23 (0.25) $p_{\text{wcb}} = 0.006$ | 1.71 (0.27) $p_{\text{wcb}} = 0.006$ | 3.57 (0.85) $p_{\text{wcb}} = 0.028$ |
| GPT-4 task exposure ($\beta$) across 51 gold-standard professions (27,403 accounts; 51 occupational clusters; mean 7.7%) | 1.71 (0.58) $p_{\text{wcb}} = 0.003$ | 1.66 (0.57) $p_{\text{wcb}} = 0.003$ | 1.57 (0.57) $p_{\text{wcb}} = 0.005$ | 1.45 (0.63) $p_{\text{wcb}} = 0.010$ |
| Task exposure first principal component (16 indices, 51 professions) (27,403 accounts; 51 occupational clusters; mean 7.7%) | 1.81 (0.60) $p_{\text{wcb}} = 0.002$ | 1.96 (0.56) $p_{\text{wcb}} < 0.001$ | 1.87 (0.55) $p_{\text{wcb}} < 0.001$ | 1.78 (0.61) $p_{\text{wcb}} < 0.001$ |

**Notes:** All panels use 12-month active paid AI assistant adoption (September 2025–August 2026) as the dependent variable. Panel A regresses UK occupational adoption across SOC 2020 major and sub-major groups on US survey workplace AI usage (Bick, Blandin, and Deming, 2026). Panel B regresses UK occupational adoption across the 51 fee-verified professions on GPT-4 task exposure (Eloundou et al., 2024). Panel C reports individual account-level linear probability regressions on the population-weighted rolling primary panel against the first principal component across all 16 external occupational benchmarks (10 task-exposure indices and 6 observed-usage benchmarks; restricting the principal component strictly to the 10 task-exposure indices yields +1.72 pp per SD, $p_{\text{wcb}} = 0.002$), controlling progressively for demographics, net income band, and occupational mean pay. Heteroskedasticity-robust standard errors (Panels A–B) and occupational wild-cluster-bootstrap $p$-values ($p_{\text{wcb}}$, Panel C) are reported.

# **Figure A4:** Occupational Classification Calibration, Out-of-Sample AUC, and External Validation

(a) Classifier Probability Calibration vs. Human Audit

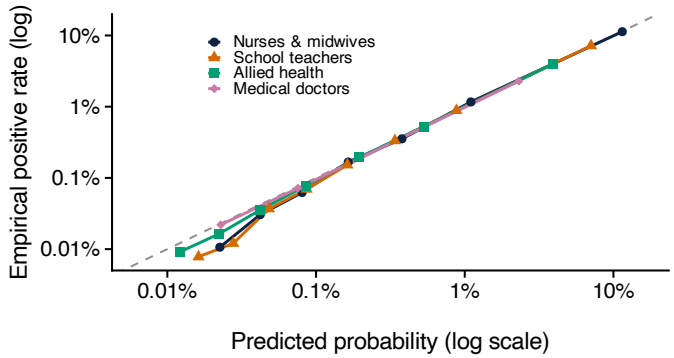


(b) Out-of-Sample AUC Across Anchor Professions

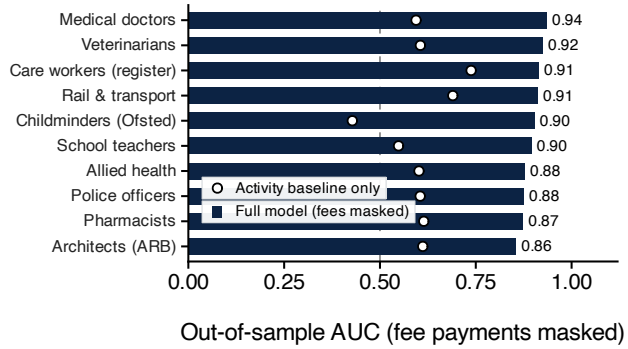


(c) Sub-Major SOC 2020 Adoption vs. Survey Usage

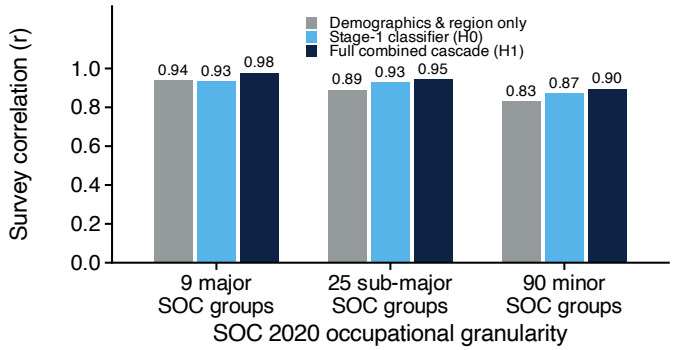


(d) 51 Professions vs. External Exposure Indices

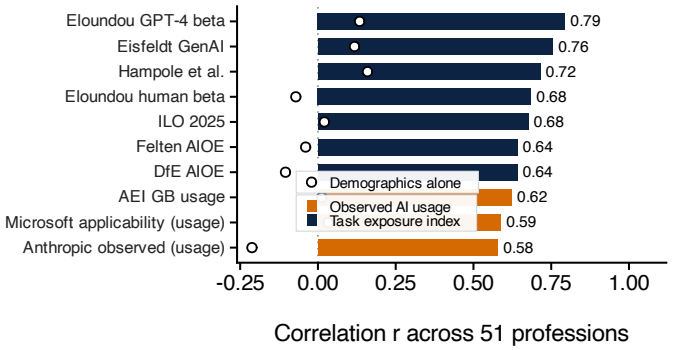


**Notes:** Panel (a) plots human-audit accuracy across deciles of predicted occupational assignment probability. Panel (b) reports out-of-sample AUCs when holding out Stage 1 fee labels and predicting occupational membership from employer and non-AI transaction features for the top 10 fee-masked classes with at least 50 held-out members, alongside an activity-visibility baseline (0.562 AUC from account tenure, transaction count, and income). Panel (c) compares occupational adoption across SOC 2020 major, sub-major, and minor groups against US survey usage across alternative classifier specifications. Panel (d) plots employment-weighted correlations across the 51 fee-verified professions against the top 10 of our 13 generative-AI task-exposure and observed-usage benchmarks (out of 16 total occupational benchmarks; navy dots) alongside demography-implied rates (hollow dots).

# **Figure A5:** Paid AI Adoption and Provider Lean Across Major Employer Industries

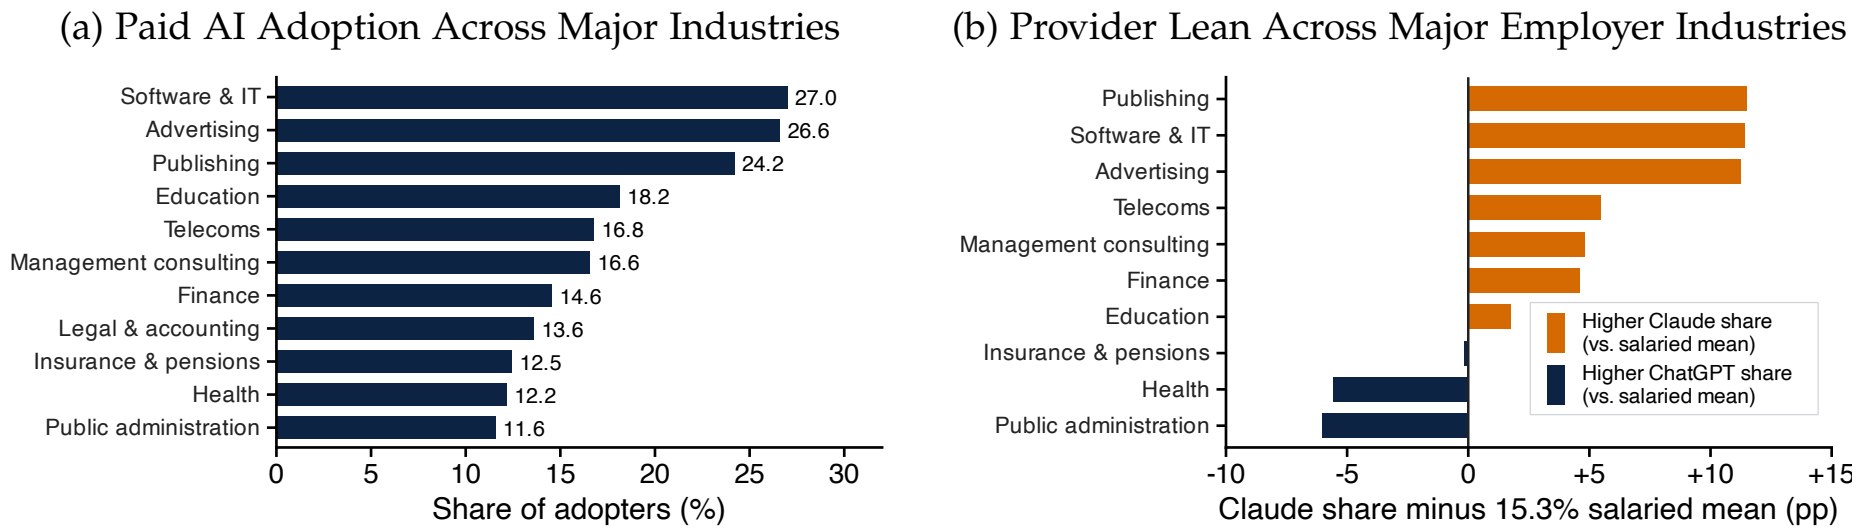


**Notes:** Panels (a) and (b) group salaried open-banking accounts across major UK Standard Industrial Classification (SIC 2007) industry groupings (each with at least 2,500 resolved primary salaried accounts). Panel (a) plots cumulative paid AI adoption by industry. Panel (b) plots each industry's direct Claude share among cumulative salaried payers of Claude or ChatGPT relative to the salaried mean.

# E Supplementary Diffusion, Geography, and Competition

This final appendix presents supplementary empirical exhibits supporting Sections 3–5.

## E.1 Extensive vs. Intensive Spending Margins and Geographic Concentration

**Extensive vs. intensive margins.** Figure A6 decomposes aggregate UK consumer AI expenditure into the share of adults paying in a given month (extensive margin) and average monthly AI spending per active payer (intensive margin), using the identity AI Spend per Adult$_t$ = Active Payer Share$_t$ × AI Spend per Active Payer$_t$. Between January 2024 and July 2026, the monthly active payer share rose nearly sixfold (from 1.0% to 5.8% of UK adults), while average monthly spending per active payer rose by roughly one-quarter (from £21.80 to £27.60, or from 0.37% to 0.47% of active payers' total monthly card spending in Panel c), reflecting uptake of power-user tiers and multi-homing. In log terms, expansion in the active payer share accounts for 88% of aggregate spending growth (84% on directly observed payments). Restricting the first-time subscriber inflow series in Figure A6(b) to accounts continuously observed for at least 12 months prior to January 2026 confirms that new sign-ups reflect genuine adoption rather than newly linked accounts.

**Geographic concentration across Britain.** Figure A7 complements the postcode-district map in the main text (Figure 6) by examining adoption across local authority districts. In June 2026, London's active paid AI assistant adoption rate (6.8%) is nearly double that of Great Britain outside London (3.6%), and the top decile of local authorities records an adoption rate 3.8 times that of the bottom decile. Across Great Britain local authorities, cumulative adoption rises steeply with the Census 2021 share of residents in managerial and professional occupations (Panel a, population-weighted $r = 0.74$) and ONS Gross Disposable Household Income per head (Panel b, $r = 0.68$). A single principal component combining local professional employment share, household income per head, and the share of residents aged 20–34 accounts for 64% of the cross-area variation.

## E.2 Multi-Homing, Cross-Provider Switching, Tier Upgrades, and Event Studies

In this section we provide supplementary evidence on consumer multi-homing, cross-provider switching, vertical plan transitions, and model release event studies, discussed only briefly in Section 5.

**Multi-homing across product scopes.** Within general-purpose AI assistants, 4.2% of active AI subscribers pay two or more subscriptions in the same calendar month, with strong asymmetry by provider. For the full duration of a subscription spell, 27% of Claude subscribers concurrently pay a rival assistant compared with 7.9% of ChatGPT

subscribers (Figure A8a). Figure A8(b) shows the destination providers for switchers who cancel their previous subscription.

**Vertical tier upgrades.** Table A11 tabulates account-level vertical transitions across subscription price tiers for OpenAI and Anthropic. Among Claude Pro subscribers, 6.2% eventually upgrade to higher-priced Claude Max tiers (£90–£180 per month), whereas only 0.74% of ChatGPT Plus subscribers transition to higher-priced Pro or Team plans. In counterfactual accounting on the July 2026 direct-billing cross-section, replacing every £80+ power-user payment with the provider's standard £20 monthly plan reduces Claude's average monthly spending per active subscriber from £28.90 to £20.15 (eliminating nearly all of its premium over ChatGPT's £20.40) and lowers Claude's direct spending share from 29.4% to 23.6%. Applying the same £20 cap between December 2025 and July 2026 reduces Claude's spending-share gain by one-quarter (24%), leaving roughly three-quarters (76%) under the cap.

**Frontier model releases and secular-trend adjustment.** For the model release event study in Figure 9(d), we compile major flagship model and consumer product launches between January 2023 and July 2026 across Anthropic (24 releases), OpenAI (34 releases), and Google (4 Gemini releases). For each event, we compare first-time provider sign-ups in the 14 days following the announcement against the average of the two preceding 14-day windows. To account for rapid baseline growth between announcements, we benchmark these post-release changes against 300 randomly drawn non-release dates over the same sample period. Anthropic releases are associated with a +39.5% raw 14-day jump in sign-ups versus +10.0% on random placebo dates—a net surge of +29.5 percentage points ($p < 0.001$)—compared with net changes of $+7.9$ percentage points for Google and $-1.5$ percentage points for OpenAI.

# **Table A11:** Account-Level Subscription Tier Upgrades to Power-User Plans

| Plan upgrade transition | Accounts on origin tier | Upgraded accounts | Share of origin (%) | Active at or above tier in Jul 2026 |
|---|---|---|---|---|
| Claude Pro → Max 5× (£18/£20 → £90) | 2,242 | 124 | 5.53 | 82 |
| Claude Pro → Max 20× (£18/£20 → £180) | 2,242 | 15 | 0.67 | <10 |
| Claude Max 5× → Max 20× (£90 → £180) | 246 | 31 | 12.60 | 23 |
| ChatGPT Plus → Pro (£20 → £160–£200) | 8,625 | 30 | 0.35 | <10 |
| ChatGPT Plus → Pro 5× (£20 → £80–£100) | 8,625 | 21 | 0.24 | 19 |
| ChatGPT Plus → Team (£20 → £25+/seat) | 8,625 | 13 | 0.15 | <10 |
| ChatGPT Go → Plus (£6.99/£8 → £20) | 1,020 | 49 | 4.80 | 41 |

**Notes:** This table reports account-level vertical upgrade transitions across active subscription price tiers for OpenAI (where 0.59% of ChatGPT Plus subscribers upgrade to £80+ individual Pro plans and 0.15% transition to multi-seat Team plans, totaling 0.74%) and Anthropic (where 6.20% of Claude Pro subscribers upgrade to £90–£180 Claude Max plans). Cells with fewer than 10 underlying accounts are suppressed as $< 10$.

# **Figure A6:** Aggregate Consumer AI Spending, First-Time Adopter Inflows, and Extensive vs. Intensive Margins

(a) Monthly AI Spending per 1,000 UK Adults

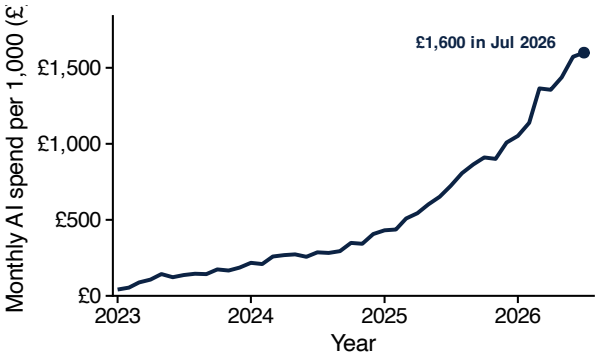


(b) First-Time Adopters vs. Streaming (2026)

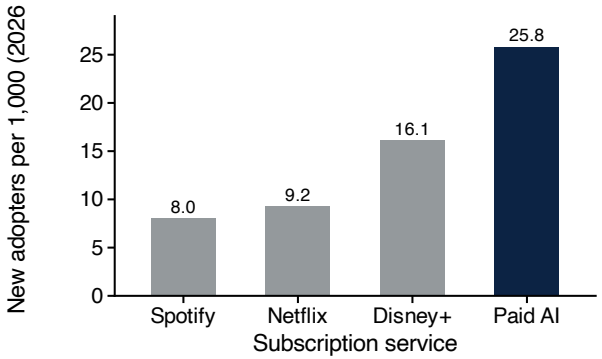


(c) Extensive Payer Share vs. Intensive Budget Share

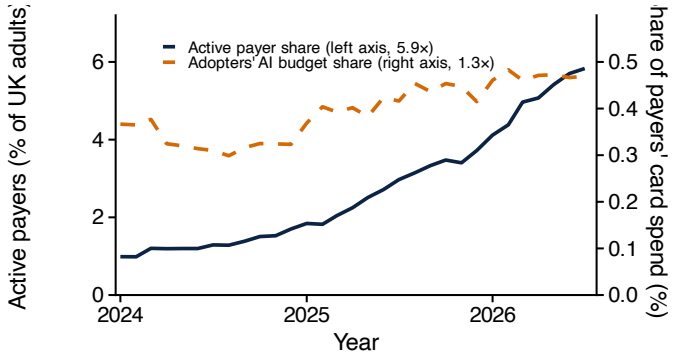


**Notes:** Panel (a) plots monthly consumer expenditure on standalone AI products in pounds per 1,000 UK adults from January 2023 to July 2026. Panel (b) compares cumulative first-observed paid AI adopters per 1,000 UK adults over January–July 2026 against directly observed first-time panel entrants to Disney+, Netflix, and Spotify among accounts continuously observed for at least 12 months prior to January 2026. Panel (c) plots the monthly percentage of UK adults billed for AI (left axis) alongside monthly AI spending per active payer expressed as a percentage of active payers' fixed window-average total monthly card spending (£5,890, right axis). Accounts are reweighted to represent UK adults.

# **Figure A7:** Local Authority Paid AI Adoption vs. Occupational Structure and Household Income

(a) Local Authority Adoption vs. Professional Share

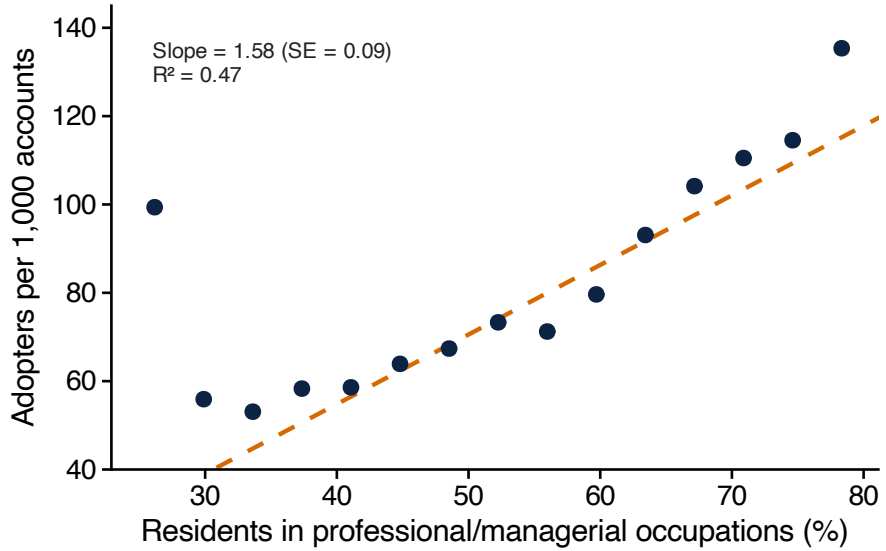


(b) Local Authority Adoption vs. Household Income

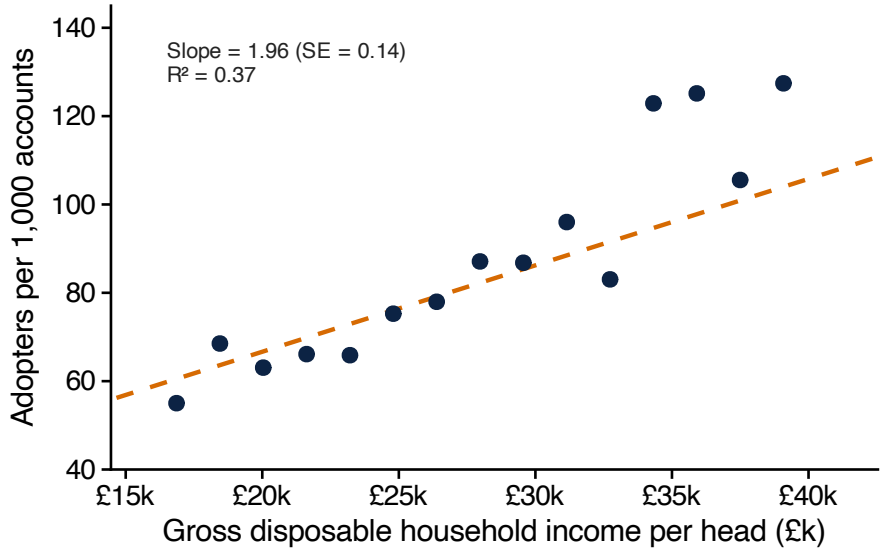


**Notes:** This figure plots cross-area binscatters of cumulative paid AI adoption per 1,000 accounts across Great Britain local authority districts against the Census 2021 share of residents in managerial and professional occupations (Panel a) and ONS Gross Disposable Household Income per head (Panel b).

# **Figure A8:** Spell-Level Rival-Assistant Multi-Homing and 60-Day Cross-Provider Switching

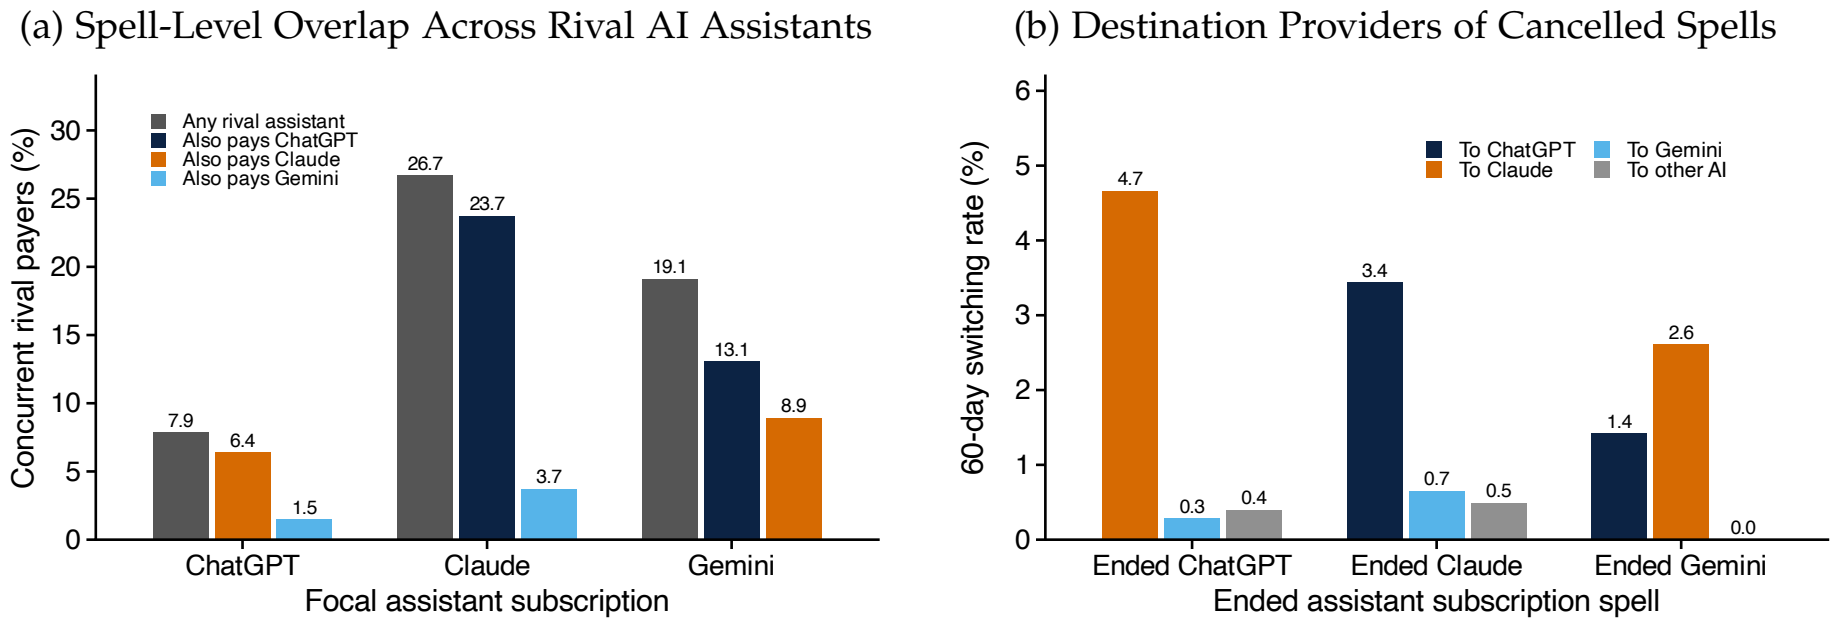


**Notes:** Panel (a) reports the percentage of general-purpose assistant subscribers (ChatGPT, Claude, and Gemini on the balanced primary panel) who concurrently hold an active paid subscription with a rival general-purpose assistant during their subscription spell (corresponding to the 6.0% balanced-panel assistant multi-homing rate vs. 4.2% monthly rolling-panel multi-homing rate; cross-product monthly overlap across all standalone AI sellers averages 35.1% at peak and reaches 44.3% at ElevenLabs and 37.9% at Midjourney). Panel (b) tracks ended assistant subscription spells where the account continues transacting and reports the percentage that initiate a subscription with a competing AI provider within 60 days.