The Household Impact of Generative AI: Evidence from Internet Browsing Behavior

Michael Blank* Gregor Schubert† Miao Ben Zhang‡

September 16, 2025

[PRELIMINARY AND INCOMPLETE - DO NOT CITE]

Abstract

How does generative AI affect households in practice? This paper provides the first comprehensive evidence on U.S. households' usage of generative AI using a large panel of household-level internet browsing data from 2021 to 2024. We show that households tend to use ChatGPT in the context of productive non-market activities, such as education or job search. We develop a novel measure of a household's exposure to generative AI based on its pre-ChatGPT internet use patterns and show that a 1 SD higher potential benefit is associated with a 20% higher chance of having tried ChatGPT by December 2023. We also find evidence of a household "generative AI divide:" high-income and younger households adopt generative AI more than low-income and older households, and predicted benefits and differential responsiveness to the benefits can explain 2-10% of the income gap and 12-28% of the age gap in adoption. Finally, we quantify how generative AI changes household internet use: generative AI adopters reduce the amount of their browsing time spent on the productive activities associated with ChatGPT, and increase the share of their browsing time spent on leisure activities.

Keywords: Household Finance, Generative AI, Internet Browsing, Technology Adoption, Inequality.

*Stanford Graduate School of Business, Email: blankm@stanford.edu
†UCLA Anderson School of Management, Email: gregor.schubert@anderson.ucla.edu.
‡USC Marshall School of Business, Email: Miao.Zhang@marshall.usc.edu

Since the release of ChatGPT in November 2022, there have been rapid advances in generative AI technology. The deployment of this technology has led to much discussion about its potential effects, both positive–productivity gains, digital consumption, democratization of digital skills like coding–and negative–labor market displacement, effect on inequality, disruption of learning and early-career skill accumulation. A particular focus has been on the effects of the technology on labor markets, with an active empirical debate over whether and how generative AI is impacting workers.

In this paper, we consider a margin of generative AI's impact that has not received as much attention or academic analysis: the impact on households' online browsing behaviors and outcomes. While there is so far limited empirical evidence on how generative AI affects households as consumers outside of a workplace context, this impact deserves at least as much attention as labor market effects: Empirically, the average American prime-age adult is more likely to use ChatGPT for their own personal browsing than for their job.[1] Even over the longer run, much of generative AI's hypothesized impact lies *outside* of the workplace: it may transform education (Khan, 2024), social media interactions (Ghani et al., 2022), entertainment media content (Zhang et al., 2025), and household production tasks, such as financial planning (Lo and Ross, 2024) and shopping (Lei and Liu, 2025).

We make four key contributions to understanding how generative AI adoption affects private households, focusing on ChatGPT use: (1) We provide evidence of the speed of adoption of generative AI by private households and document an increasing "generative AI divide" with higher-income and younger households adopting the technology at higher rates. (2) We characterize *how* generative AI is used by households and show that private ChatGPT use happens in the context of internet tasks associated with household production tasks (e.g. education, job search, informational research), and is less likely when households are engaging in leisure activities, such as games or social media use. (3) We develop a new methodology for characterizing a household's exposure to generative AI, based on the potential applicability of LLMs for a household's prior online activities and show that this exposure can explain a substantial share of the adoption gaps by income and age. (4) We

[1]For example, in a February 2024 survey, Pew Research found that 23% of all adults had used ChatGPT, 17% of all adults had used it for entertainment, and 17% to learn something new, while only 12% of all adults (20% of those who were employed) had used it for a work task.Similarly, Bick et al. (2024) found in a representative survey fielded in August and November 2024 that 33.7 % of all respondents had used generative AI outside of work, while 26.4% of the subset of workers that were employed had used it at work. An representative survey by Elon University in January 2025 found that, among LLM users, 51% say their main usage is for "personal, informal learning", double the number that use it mainly for work purposes (24%). URL: https://imaginingthedigitalfuture.org/wp-content/uploads/2025/03/ITDF-LLM-User-Report-3-12-25.pdf

quantify how households *change* their online activities after they adopt generative AI.

Studying the impact of generative AI tools on private household behavior has to overcome severe data constraints: Ideally, we would want data for a broad set of households that is informative about a broad range of behaviors *before* adopting generative AI, the timing and nature of generative AI usage, and changes in behavior *after* generative AI adoption. Typical data sources used to evaluate household economic impacts tend to be limited in one or more dimensions as a result of data collection methods and privacy restrictions: Surveys and guided experiments give detailed information on usage, but only at specific moments in time or hard-to-extrapolate contexts—and tend to collect limited information on what people do when not using generative AI tools. Moreover, respondents may underreport their AI use due social desirability bias (Ling and Imas, 2025). Similarly, microdata collected by generative AI companies (e.g. by Anthropic) provide a wealth of information about what people are using the tools to do, but with little information about their broader behaviors and outcomes when not using a chatbot.

In this paper, we overcome these issues by utilizing a novel source of data for studying generative AI–Comscore's web browsing panel that allows us to comprehensively document the impact of generative AI on household online activities. This data allows us to observe all browsing behavior of a large, diverse panel of U.S. households. Three particular features make it uniquely suited to studying generative AI's impacts on households: (1) It is *high-frequency and comprehensive*, containing all of a household's browsing activity on a given machine. (2) It is a *multi-year panel* spanning both pre-adoption and post-adoption household behavior, allowing us to more plausibly estimate underlying drivers and the impact of adoption. (3) It includes large sub-samples of households across the income and age distribution. This allows us to characterize *heterogeneity* in the take-up and effects of generative AI tools.

Exploiting this data's unique features, we establish three main findings regarding generative AI's effect on households. First, we establish that households predominantly utilize ChatGPT in the context of browsing activities that are productive–such as tasks related to building human capital (e.g. education and job search). As with all non-experimental usage data not collected by generative AI companies themselves, our data does not allow us to directly observe what a household is using ChatGPT to do. However, our data allows us to develop a novel technique to estimate which tasks a given household is using ChatGPT for, based on the following intuition: when a person uses ChatGPT to complete a task, they likely are utilizing other websites that complement their work in the chatbot. For example,

a person who asks ChatGPT to help debug their code may also go to Github. Or a person who asks ChatGPT to help think through which jobs to apply to may also visit an online job board. This implies that we can plausibly deduce what activities a person is doing inside ChatGPT by observing their online activities before and after their use of ChatGPT. By comparing activities within narrow time windows around ChatGPT use to the online activities of other households, and of the same household at other times, we can both determine the activity context of ChatGPT use, and to what degree this pattern is due to either selection effects between users and non-users, or due to targeted task-shifting within the user group. We find some evidence of both: ChatGPT users systematically differ from non-users in their online activities, but this difference is even more pronounced for browsing sessions involving ChatGPT.

The distinct online activities associated with ChatGPT use raise the question of whether differences in a household's potential benefits from generative AI *cause* differences in adoption. Our second key finding establishes that potential benefits explain a substantial share of heterogeneity in ChatGPT adoption between different demographic groups. To show this, we first develop a novel measure of the ex ante usefulness of generative AI to a household based on the composition of its online browsing *before* the release of ChatGPT. We use web-scraping and LLMs to systematically classify the overlap between chatbot capabilities and the activities on a website, for essentially the universe of websites visited by U.S. households.

We aggregate this website-level overlap across a household's complete browsing activity *before* the November 2022 release of ChatGPT to arrive at our ex ante measure of potential benefits—the household's generative AI "exposure". A household has a higher exposure if, prior to November 2022, a higher share of their browsing time was spent on sites with a high overlap with ChatGPT capabilities. We find that a household's predicted benefit of using ChatGPT strongly predicts whether they actually adopt ChatGPT: a 1 SD higher exposure predicts a 1.8 pp higher rate of having tried ChatGPT by December 2023–a 20% increase.

While this link between household benefits and adoption of generative AI may not seem suprising, it is critical for understanding the fact–found in previous papers and confirmed in our data–that there are large gaps in generative AI use by income, age, and other demographics. Given the large potential impacts of generative AI, understanding what drives adoption gaps is of paramount importance for both policymakers and technology executives. We decompose observed gaps in adoption into variation explained by differences in ex ante exposure or differential take-up of generative AI in response to exposure, and unexplained

variation. We find that potential benefits can explain 2-10% of the gap in ChatGPT use between high- and low-income households, and 12-28% of the gap in use between young and old households. Moreover, we show that households with programming knowledge have both greater potential benefits from using ChatGPT and respond more strongly to these benefits: our exposure measure explains 21-24% of the greater ChatGPT use among households that have accessed Stack Overflow (as a proxy for programming knowledge) in the past.

Having established *how* private households use generative AI, and that different exposure to potential benefits from the technology can explain adoption gaps, our third key finding measures how ChatGPT use *changes* household behavior and welfare. With our data, we can shed light on which online activities on other websites ChatGPT *substitutes* or *complements*. A key identification concern is that households may select into using ChatGPT at points in time when their online activities make it a particularly useful tool (e.g. when starting a new job search), confounding estimates of short-run changes in behavior. To address this issue, we use an instrumented difference-in-differences approach, using exogenous variation in household exposure to generative AI before the ChatGPT release as an instrument for post-ChatGPT adoption. We find that households overall spend *less* time on the "productive" online activities that we previously found ChatGPT to be associated with, and instead spend more time on activities, such as shopping, that ChatGPT is less likely to be used for. This finding suggests that using ChatGPT for these household production tasks substitutes for time spent on these activities on other websites, and increases the share of time spent on leisure activities online. The additional time spent on leisure activities provides evidence that ChatGPT use has positive welfare effects on households.

These findings have important implications for the societal effects of generative AI: The use of ChatGPT in the context of productive online activities, and the ability to increase leisure time as a result of ChatGPT jointly provide compelling evidence that ChatGPT increases household non-market productivity in addition to the labor market productivity effects found by prior studies. Together with the scale and speed of private household adoption that we document, these household benefits are likely to represent a substantial part of the economic impact of generative AI that needs to be taken into account when designing regulation and quantifying the welfare effects of this new technology.

At the same time, the heterogeneity in adoption that we observe highlights the importance of better understanding the dynamic effects of using generative AI: if use outside of work complements workplace productivity or affects human capital development, inequality in initial adoption can exacerbate disparities in who benefits from generative AI in the

future. Our findings highlight two channels for promoting household adoption by lagging groups and the resulting benefits: (1)Identifying pathways to generative AI from specific online activities and guiding households towards these to increase the perceived benefits from chatbots. (2) Increasing take-up of generative AI for households that have potential benefits but are not availing themselves of tools like ChatGPT. Given the importance of differential take-up and pre-existing technology awareness that we document, targeting technology literacy efforts, training or subsidies at groups with low take-up rates (such as older households) may lead to a greater level and a more even distribution of benefits from the technology.

Related literature. Several studies of the effect on productivity in specific occupations have found that generative AI tools tend to have different effects depending on the initial productivity of workers: In an experiment estimating the effect of giving workers access to ChatGPT to help with professional writing tasks Noy and Zhang (2023) show that lower-performing workers improved more with the help of generative AI. Management consultants (Dell'Acqua et al., 2023) , software developers (Peng et al., 2023), and call center employees (Brynjolfsson et al., 2023) all became more productive, but with previously less productive workers benefiting the most. These effects are not limited to workplace performance: For example, one study by Yu et al. (2024) looked at changes in academic writing proficiency among students at a public university after ChatGPT was released and found significant improvements, particularly for students from backgrounds (e.g. foreign-born) associated with lower English proficiency.

Previous studies also raised concerns around unequal adoption with regard to earlier generations of AI. McElheran et al. (2024) show that, in 2018, AI use in firms was highly concentrated in particular cities and regions, which they call an "AI divide" across geographies. Humlum and Vestergaard (2025) find in a large-scale survey in Denmark that younger and less-experienced workers are more likely to use ChatGPT, while women and workers with lower earnings were less likely to use the tool. We contribute to this literature by showing that a "generative AI divide" also exists in private household adoption, provide new evidence on what causes such adoption gaps, and document how the technology affects household behavior and non-market work productivity.

I. Comscore data: background, sample, and main variables

A. *Background on Comscore data*

Our analysis utilizes microdata for a large panel of U.S. households whose browsing data is tracked by Comscore. Comscore is a U.S. based, publicly-traded media measurement and analytics company founded in 1999. It specializes in characterizing the online behavior and demographics of different websites' user bases. To do so, Comscore recruits a large number of U.S. adults ("panelists") who, in exchange for certain incentives, allow the company to collect comprehensive data on their browsing activities.[2] A panelist installs software on their personal computer (be it a desktop or laptop) that, in a continuous and real-time manner, tracks the URL and granular timestamp (second of the day) of each unique website that the panelist visits.

We have access to the raw microdata. For each day in the panelist's sample period, we observe the set and timing of all website visits that the panelist makes on their personal computer. Critically, our data includes a generic but time-consistent identifier for each panelist, `machine_id`. Although, in accordance with Comscore's data privacy policies, we do not observe the real-life identify of panelists, we can still link an individual panelist's browsing activity over time via `machine_id`. This allows us to associate a panelist's browsing behavior prior to the widespread release of generative AI tools with their behavior after the release. Because of this continuity, and to simplify the exposition, we will interchangeably refer to machines as "households" throughout.

B. *Construction of main browsing variables*

We primarily work with three types of variables that are derived from the raw Comscore browsing data. For each of these types, we make heavy use of fields that Comscore has developed and cleaned, and that are included as part of our data access.

Measuring browsing intensity. The second type of variable is made up of proxies for the quantity and nature of a panelist's browsing activity. In our analysis, both the treatment

[2]These incentives can entail free software, applications, or utilities. Further details are available at:

(eg. overlap of prior browsing activity with ChatGPT) and outcome measures (eg. changes in browsing activity after adopting ChatGPT) require proxies for the intensity of panelists' browsing activity across different types of online tasks.

There are two properties that, collectively, our set of proxies should satisfy. First, they must allow us to compare the amount of browsing that panelists do across different types of sites, including sites with different user interfaces (such as sites with clickable links, chat sessions, videos, and so on). Second, the proxies should be informative about the type and breadth of online tasks that panelists utilize certain sites to perform (such as handling personal finances, staying informed about current events, chatting with friends, and so on). Put simply: we want to measure browsing activity in a way that is not sensitive to the "technology" of how panelists use different sites' services, but that is nonetheless informative about the economic tasks that the different sites help panelists perform.

To motivate why we focus on particular measures of browsing activity, it is informative to first consider a proxy that *does not* satisfy either of the two desired properties. Concretely, suppose that for a particular panelist on a particular day, we want to measure the amount and potential usefulness of the browsing that the panelist does on a certain website. Perhaps the simplest proxy for this would be the *click count*: the number of times throughout the day that the panelist clicks onto the pages of the website. Although seemingly sensible, this proxy would not satisfy either of the two properties described above. To see why this is so for the first property, note that for some sites, it is common to click on many different pages. For example, on a news site like `wsj.com`, a panelist may read several different articles. But other sites, such as video streaming services, often entail few clicks, even for a long intensity of use (eg. watching a movie on `netflix.com`). To permit an apples-to-apples comparison in browsing intensity across sites, the proxy must look through heterogeneity in how many clicks are required to use a site's services. As for the second property, suppose that the day includes several distinct browsing "sessions"–periods of continuous browsing activity in which the panelist completes different online tasks. For example, the panelist may develop and debug a Python script during the morning, and decide on which movie to stream during the night. If the panelist goes to a site–say, in this example, `wikipedia.com`–during both sessions, we would ideally want our browsing intensity proxy to reflect the multiple uses of this site. A proxy that simply sums all clicks onto the site during the day, irrespective of their intra-day timing or context, will not do so.

A simple count of clicks is thus an unsatisfactory proxy in two ways: it does not permit a clean comparison of activity across different types of sites, and it does not capture a more

abstract, task-based notion of browsing intensity. To develop more suitable measures of browsing intensity, we exploit the granular, high-frequency nature of the Comscore data. In particular, Comscore's data team has devised and refined an algorithm that (i) divides a panelist's browsing activity into distinct intra-day sessions (the identifier `session_id`) and (ii) for each session, allocates the total duration of the session to the different websites that the panelist visits (the variable `duration`, with units in seconds). We relegate a more detailed description of the algorithm behind the creation of `session_id` and `duration` to the appendix. But the underlying idea is intuitive. A session (unique value of `session_id`)[3] is a set of adjacent website visits that, based on their timestamps, appear to be part of an essentially continuous browsing session. The duration of browsing for a certain site during the session, `duration`, is the number of seconds that a panelist appears to spend on the site after first going to it (cumulated across all non-continuous visits to the site during the session).

Demographic characteristics Finally, we make extensive use of information that Comscore collects on panelists' demographic characteristics and location within the U.S. The two most important such variables for our analysis are the household income and age of each panelist.[4] For additional controls in some specifications, we also use variables for the city in which the panelist resides, as well as composition of the panelist's household (number of members and whether or not it contains children).

As described below, much of our analysis entails estimating changes in panelists' browsing behavior around the release of ChatGPT in November 2022. To avoid our measures of household demographics being endogenously influenced by whether a panelist adopts ChatGPT, we always use (unless otherwise noted) a *pre-ChatGPT release snapshot* of these measures. In particular, we take the value of the demographic variables that the panelist reported in the month that is closest to–but not after–November 2022. For example, if a panelist updated their reported demographic information in July 2022 and then again in December 2022, we would use the values for July 2022.

[3] More precisely, `session_id` takes generic values that partition the browsing activity of a given panelist. A session is thus defined by the pair `machine_id` × `session_id`.

[4] Precisely, Comscore asks panelists to report the total income of their household and the age of their household's "head" member.

C. *Constructing sample of web-browsing households*

We now describe our baseline sample of panelists from Comscore, first laying out the filters that we apply to the raw Comscore data and then analyzing the representativeness of the set of panelists that pass these filters.

Sample filters. To construct the baseline sample for our analysis, we start with the universe of panelists in Comscore whose browsing data is collected around the November 2022 release of ChatGPT.[5] There are more than 330,000 unique panelists (associated with more than 5.2 million panelist X month pairs) who are in the Comscore panel for at least one month prior to, as well as one month after, November 2022.[6] But not all of this data is suitable for our analysis. We only keep panelists who pass three data filters. Table X in the appendix shows how the sample size changes as each successive filter is applied.

The first filter is that the panelist's browsing data must come from their *own personal computer*. This drops the small minority (less than 5%) of panelists whose data come from online activity on their employer-owned machines. The rationale for this filter is that we want to capture the online tasks that correspond to home production / leisure and in which panelists actively choose to engage.[7]

The second filter is that we drop the approximately 15% of panelists for whom Comscore is not able to collect demographic and locational data.[8] This information–in particular on a panelist's age, location, and income–serves as the basis of our analysis on the correlates and drivers of different individuals' adoption of generative AI tools. Moreover, it allows us to more credibly estimate the causal effect of ChatGPT adoption on panelists' browsing behavior, since we can compare the outcomes of panelists within granular demographic and regional bins.

[5] We discuss why our analysis is centered on the release of ChatGPT in the next subsection

[6] We consider a panelist to be in the panel in a given month if the panelist has browsing activity during that month that has strictly positive duration (according to Comscore's `duration` variable).

[7] Of course, it is possible that some panelists engage in some work tasks on their personal computers. This would seem especially plausible for panelists who work from home and/or are in occupations with blurry lines between work tasks and tasks that build a worker's own human capital (eg. a self-initiated coding project by a software engineer).

[8] Comscore collects this information on an ongoing basis by inducing active panelists to occasionally fill out short surveys. We primarily utilize information that Comscore collects before November 2022. The main reason for this is to ensure that cross-panelist variation in stated age and income is not potentially contaminated by whether the panelist adopts ChatGPT or not. A secondary benefit is that this data is collected closer to when the panelist opts into the Comscore panel, and hence is seemingly not subject to selection bias that could arise from survey attrition.

For our third and final filter, we drop panelists who, in the 12 months leading up to the release of ChatGPT (November 2021 - October 2022), have one or more months without any browsing activity. This is by far the most aggressive filter that we apply: as shown in Table X, of the observations left after the first two filters, only around 25% of unique panelists (35% of panelist - months) pass this last filter. The vast majority of the panelists who are dropped do not pass the filter for a specific reason: they enter the Comscore sample too late (before October 2022 but after November 2021). We still choose to apply this relatively stringent filter, for two reasons. First, and most importantly, it gives us more confidence that for the panelists in our baseline sample, we are able to accurately measure their typical browsing behavior prior to the release of ChatGPT. [9] Second, we do not lose as much data by applying the filter as the unweighted panelist (and panelist - month) counts suggest. In particular, when weighted by browsing duration over November 2021 to October 2022,

Representativeness of sample. We now discuss how the representativeness of our filtered Comscore sample. Given our interest in characterizing the heterogeneity in the adoption of generative AI tools across the income and age distributions, the demographic composition of our sample would ideally be similar to that of the full U.S. population of internet users.[10]

Helpfully, Comscore explicitly designs the process by which it recruits panelists to try to achieve a representative sample. But it is naturally challenging to maintain a perfectly representative sample over time, especially given changes to lifestyles that may interact with home computing behavior (eg. the increased share of works who work from home starting in 2020). We therefore probe the representativeness of our raw (filtered) Comscore sample, and correct for any residual discrepancies between the distribution of demographics

[9] Much of our analysis will center on a panelist's predicted exposure to the release of ChatGPT, as inferred from the composition of their pre-November 2022 browsing behavior. It is thus crucial that the pre-November 2022 data included in the Comscore panel is representative of a panelist's "normal" online activities. There are reasons to be skeptical that this will be the case for panelists who enter the sample after October 2022. This is partly a mechanical result of us not being able to observe a full set of 12 calendar months for these panelists in the pre-ChatGPT period; for certain type of websites, browsing behavior tends to be concentrated in particular months of the year (eg. March and April for tax preparation services), which would lead our exposure measures to be noisy for panelists who enter the sample after November 2021. Moreover, if panelists endogenously select into the Comscore sample–say, for instance, an individual is more likely to opt into the panel if they have recently received a positive shock to their home browsing intensity (eg. being laid off from their job)–then their browsing behavior upon initially entering the panel may be unusually volatile.

[10] There is an alternative goal we could have set: our sample being representative of the full U.S. population, including both internet users and *non-internet users*. We view it as more natural to focus on the population of internet users since, conceptually, it is difficult to think of what a non-user's latent benefit from potentially adopting generative AI tools is.

in our sample relative to the U.S. population of internet users.

The first column of Table I shows the fraction of panelists that, as of November 2022, fall into each income and age bucket contained in the Comscore data. This can be compared to the share of the internet-using U.S. population, as taken from the 2022 American Community Survey (ACS) and shown in the second column. Our Comscore sample tends to overrepresent households at the low (<$60K) and high ends (>$200K) of the income distribution, and accordingly underrepresent middle-income ($60K-200K) households. In terms of age, our sample is more representative of the population, with the largest differences being an overrepresentation of upper middle-age households (45-54 years old) and an underrepresentation of younger households (25-44 years old).

To rectify these sample discrepancies, we weight all statistics and regressions so that the joint income and age distribution of our sample (as of November 2022) matches that of the ACS.[11] As long as Comscore samples a representative set of households within each age X income bin, these weights ensure that our data reflects the behavior of the average internet-using U.S. household. If instead our sample features non-random selection within age X income bins, it would seemingly be towards households who engage in more home browsing and more digitally sophisticated. Such selection would compress variation (relative to the actual amount) in the take-up and associated browsing effects of GenAI tools; this should, if anything, attenuate our estimates.

II. Household Generative AI Adoption Patterns

A. Measuring GenAI use

We focus on the use of the most popular LLM chatbot product–ChatGPT–as a proxy for generative AI adoption by households. There are several reasons for this choice.

First, the initial public release of ChatGPT was a major, largely unexpected technological breakthrough that instantly increased the capabilities and accessibility of generative AI. This is evident in the attention that ChatGPT garnered in news reports and social media following its release on November 30, 2022, as well as the nearly immediate reaction of stock prices to the altered technological landscape (Eisfeldt et al., 2023). The introduction

[11] Precisely, the weight for a given panelist is the ratio of the ACS-implied population share in the panelist's age X income bin to the share of that age X income bin in our (unweighted) Comscore sample.

of ChatGPT thus provides a sharp break in the average household's ability to seek out and utilize sophisticated GenAI tools. Much of our analysis revolves around exploring changes in households' browsing behavior before versus after November 2022, with a focus on whether these changes appear related to the predicted and/or actual adoption of ChatGPT.

Second, as found in surveys such as Bick et al. (2024) and confirmed in our browsing data, the adoption of ChatGPT following its release was remarkably widespread and rapid. This is especially the case outside of the workplace, which is the focus of our paper.

Third and finally, ChatGPT's popularity precipitated the release of many other GenAI-powered tools, especially chatbots such as Anthropic's Claude or Google's Bard (later Gemini) in March 2023. If some panelists in our sample utilize GenAI through one of these alternatives, our ChatGPT-based adoption measure would understate overall GenAI adoption. Although we could work with a measure of GenAI use that explicitly accounts for ChatGPT alternatives, we prefer to work with the simpler ChatGPT-based measure. This is in part because, consistent with survey evidence (Bick et al., 2024), ChatGPT is by far the most widely-used GenAI tool among panelists in our Comscore data. Moreover, for the panelists in our sample who use GenAI tools other than ChatGPT, the vast majority do so only after first trying ChatGPT.

Defining ChatGPT adoption. We now describe how we define the adoption of ChatGPT. It is important to be clear about what we want to measure: the first concept of interest is whether a household *ever* uses ChatGPT. That is, we want to know when a household tries out the technology for the first time, and we define a measure of having ever used ChatGPT that is defined by whether we have observed that household accessing the ChatGPT site at any point in the past. The second concept of interest is if a household found enough value in generative AI to repeatedly access the website at different points in time, i.e. if they have become a "regular user" or "adopter."

This second definition excludes panelists who briefly try ChatGPT but then stop using it. To proxy for this conceptual object in our browsing data, we define our "adoption" measure as a panelist visiting ChatGPT in two consecutive months. We believe that this definition strikes a sensible balance between false positives and false negatives for identifying regular usage. On the one hand, it requires a panelist to exhibit some degree of continuity (consecutive months) in using ChatGPT before being considered an adopter; visiting ChatGPT for a handful of scattered weeks in a year is not sufficient. On the other hand, note that under our definition, once a panelist has visited ChatGPT in two straight months, they are

considered to be an adopter for the remainder of their time in the panel–regardless of if we observe them continuing to regularly use ChatGPT. The rationale is that once a panelist has used ChatGPT for a non-negligible period of time (two months), they may develop other modes of incorporating generative AI into their routine, such as accessing models through APIs rather than the chatbot interface, switching to lesser-known LLM-based chatbots, or using more specialized products incorporating generative AI. Our measure thus allows for some leakage in our ability to observe continued GenAI utilization after a household's initial phase of ongoing engagement.

B. Household use of ChatGPT over time

In this section, we document when and to what extent households adopted generative AI based on their browsing activity. Previous work has documented that a large share of households use ChatGPT outside of work: Bick et al. (2024) find that 35.5% of survey respondents in the second half of 2024 had used generative AI outside of work, and 18% had used ChatGPT in particular in the last week. Pew Research found in March 2023 that 18% of US adults had used it (Pew Research Center, 2025).

Our approach provides a new source of evidence on the speed of adoption based on evidence of ChatGPT use in household browsing behavior. In contrast to survey evidence, our data reflects actual use, permits differentiation between occasional and regular uses, documents the patterns of use over time, and allows us to link use patterns to a household's other online browsing behavior and demographics. However, one disadvantage of our data is that we cannot say anything about generative AI use through channels other than accessing it on a website on a home computer, which means that our evidence represents a lower bound on total use.

Household ChatGPT use over time. We plot the time series of ever using ChatGPT and becoming a regular adopter (as define above) in Figure 1: the share of all households that has ever used ChatGPT rises from zero to 9.2 pp by December 2023 and to 17.4 pp by December 2024, almost doubling within a year. Regular use shows a similarly steep increase but at a lower level: regular adoption rises from 3.6 pp in December 2023 to 8.0 pp by December 2024 (see Table VII for summary statistics).

Comparison to other estimates of adoption. An important validation of our data is whether our time series provides estimates that are comparable to snapshots of household use gleaned from surveys. Our estimated rates of ChatGPT use are comparable to other sources based on surveys: In a February 2024 survey, Pew Research found that 23% of all adults had used ChatGPT, 17% of all adults had used it for entertainment, and 17% to learn something new, while only 12% of all adults (20% of those who were employed) had used it for a work task. This marked an increase from July 2023, when the same survey found 18% of all adults had used the tool, 15% for entertainment and 7% at work (12% of the employed). This is consistent with our measure of use representing both a lower bound in terms of all possible channels through which a household might access ChatGPT, and also focusing on household use outside of work, which is necessarily a subset of all use, as not all households that use ChatGPT at work will also use it privately. The fact that measure of having used ChatGPT shows both similar rapid growth patterns and orders of magnitude of usage is reassuring.

Household ChatGPT use by income and age. Our data also allows us to compare generative AI use across different demographic groups. Figure 2 shows how the share of households ever having used ChatGPT varies across income categories (panel A) and by age of the household head (panel B). Appendix Figure 13 shows the corresponding heterogeneity in becoming regular users. We find that higher income households (> $200K) are substantially more likely to have tried ChatGPT and also to have become a regular adopter at all times since the ChatGPT release. While the relationship is not precisely monotonic, it is generally true that higher income categories tend to have higher rates of use and adoption. This pattern is violated mainly by the lowest-income categories, which is likely due to the fact that many students and recent graduates have high adoption rates and are likely to fall into the lower-income categories.

Panel B of Figure 2 shows that age strongly predicts ChatGPT use patterns: the younger the household head, the more likely is it that they have used ChatGPT, with an adoption gap that increases over time. As summarized in Table VII , by December 2024, 25.4% of households headed by 18-24 year-olds have used ChatGPT, but only 10.3% of households at retirement age (65 years or above).

The generative AI divide. The gap in the use of generative AI documented in Figure 2 recalls concerns over similar patterns of a "digital divide" associated with previous waves

of information and communication technologies between "those people who have access and use of digital media and those who do not" (Van Dijk (2020), p. 1).

Just as access to past technological improvements was mediated by skills, motivation, and socio-economic contexts (Van Dijk, 2020), generative AI access is likely impacted by these characteristics.. An important new insight from our data is therefore that a similar gap exists in the use of generative AI at the household level. Large differences in generative AI adoption among households thus raise the specter of a persistent "Generative AI Divide". This divide would lead to groups that adopt generative AI at higher rates being more likely to benefit from the associated productivity improvements both in the labor market and in their private lives, and to enjoy the utility of the entertainment, companionship and other leisure benefits that chatbots can provide.

In the following sections, we provide additional insights into how this divide arises both from differences in the potential benefits from generative AI and differences in take-up of the technology for a given level of benefits.

III. What Do Households Use ChatGPT For?

A. *Classifying websites by task content*

In order to be able to assess whether households are changing the type of activity that they pursue on the internet after adopting generative AI, we want to classify websites by the types of usage that they enable. In particular, we are interested in distinguishing "productive" from "leisure" uses of the internet, as those can allow us to quantify whether generative AI complements or substitutes for non-work household productivity.

We take the website descriptions an list of 5 key activities that we previously generated and use it as the input for a classification tasks: We apply usage labels to websites by submitting the domain and its description to an LLM with a prompt that asks for two labels: (1) A classification of the domain's main usage into the categories "Productive", "Leisure", and "Ad platform/ CDN" (content delivery network). (2) Whether the domain is "mixed-use" or "single-use" with regard to these usage types.[12] Our definition of productive and leisure activities closely follows the categorization of general (non-internet) activities in Aguiar et al.

[12] In this context, *Mixed-use* domains will include, for instance, email providers and other platforms that are commonly used for both productive and leisure purposes.

(2013): productive uses are those relating to market work, education, childcare, non-market work, civic activities, shopping and personal health care, while leisure includes gaming, social activities, TV, movies, and reading for personal interests.

Importantly, we refined and validated our prompt to ensure that these labels are applied as expected: Two of the authors manually labeled the top 70 domains by page views, and we adjusted the prompt language until the LLM's answers substantially agreed with our own labels on this validation set. Based on this manual labeling, we also adjusted the classification categories to capture ad platforms and content delivery networks as a separate category. While these are not customer-facing sites in the sense that we are interested in, they mechanically generate a large number of page views and we wanted to ensure that these did not get mislabeled as productive or leisure uses of a household's time.

After scaling this process to the full set of domains, we derived a final label for each domain as having a main use that is *Productive*, *Leisure*, or *Ad platforms / CDN*.

B. Inferring ChatGPT uses from high-frequency browsing data

ChatGPT, as a powerful general-purpose AI tool, can potentially be used to fulfill many tasks and needs of households. An intriguing empirical question is what households use ChatGPT for, and whether this purpose differs across households' demographics. In principle, if researchers could observe households' textual inquiry to ChatGPT, they may directly infer the purpose of households' ChatGPT usage. However, such analysis may only be feasible within OpenAI based on their internal data. In addition, the textual inquiry alone may be highly specific and not always reflect the broad task the household is performing on the Internet. For example, an inquiry of "*Can you give me a simple explanation of the Black-Scholes formula?*" to ChatGPT may reflect a student preparing for an exam in finance or a financial analyst making a client presentation or writing a report.

Alternatively, some websites try to publicize their inferred user profile and user interests of ChatGPT based on cross-visitation of ChatGPT and other websites through coarse data, such as on the same day or in the same month, which are usually presented as "what people are also interested in." However, such inference cannot draw a key distinction between whether the cross-visited websites represent the same tasks for using ChatGPT or other tasks households can now do more because of the efficiency gain from ChatGPT. For instance, Similarweb.com finds that chatgpt.com's audience is most interested in "Video

Games Consoles and Accessories." It is unclear whether this means ChatGPT users are likely to enjoy leisure or they use ChatGPT for leisure. As it will be clear, this distinction is an important contribution in our methodology.

In this section, we propose a novel methodology to infer households' purpose of using ChatGPT leveraging on the high-frequency detailed website browsing data. Importantly, we observe the websites that users visit minutes before and after they use ChatGPT ("GPT-window"). This allow us to identify users' purpose of using ChatGPT. We also observe the websites that users visit outside the ChatGPT window ("NonGPT-window"), which allows us to assess how the usage of ChatGPT impacts households' task allocation in general.

Our approach starts with identifying about 3.6 million 30-minute intervals of our sample households' internet visitation from 2022 to 2024. For each interval, we have the duration of the household's visitation of productive, leisure, mixed, and ad websites and a label of whether the household visited openai.com. On average, about XXX% and XXX% of the intervals are labeled as GPT-window in 2023 and 2024, respectively.

C. Tasks associated with ChatGPT use

Figure 3 shows the distribution of the four types of website visitations during the GPT-window. On average, during the 30-minute window when households use ChatGPT, they tasks they focuses on are 83.7% productive, 15% leisure, 1.3% ad. Comparing with the visitation types in the full sample, households' visitation during the GPT-window focuses substantially more on productive tasks, i.e., about 12.3% more, and substantially less on leisure tasks, i.e., about 10.2% less, and also less in ad and mixed.

This simple comparison brings the first glance at what household uses generative AI for. It is perhaps not surprising that ChatGPT, given its superior power in processing text information, substantially facilitate households' performance of productive tasks. Indeed, when we further explore the differences in the categories of internet visitation during the GPT-window and in the full sample, we observe in Figure 4 that the households use GPT much more for education and information searching and much less for entertainment, shopping and finance.

Our findings are in sharp contrast with the crude inference from Similarweb.com which labels chatgpt.com's audience as interested in "Video Games Consoles and Accessories." As we dive deeper later, their inference is not inconsistent with our data, as we observe that

indeed, for households treated by ChatGPT use, they have more bandwidth to enjoy leisure, consistent with the findings in Similarweb.com. Hence, the key advantage of our approach is to use households' internet visitation separately for within the GPT-window and outside the GPT-window to distinguish the "purpose" and "effect" of ChatGPT use.

We end this section by discussing the premises and limitations of our approach and findings. Our approach is based on two premises: First, we view ChatGPT as a *tool* for fulfilling a task the household is working on. Hence, if we observe proxies for the tasks the household is performing on the Internet around her use of ChatGPT, e.g., writing emails or visiting personal relationship therapy websites, we can approximately infer the tasks the household uses ChatGPT for. Our second premise assumes that the household perform one main task within a short time window on the Internet. Together, these premises allow us to approximately infer the tasks that households perform using the Generative AI tool by examining the high-frequency data about the household's website visitations within 15 minute interval of her use of ChatGPT.

Our approach faces two potential sources of noise: First, households may resort to ChatGPT for tasks without visiting other websites. Second, households may be multi-tasking when visiting ChatGPT, making the surrounding websites less informative about the household's purpose of using ChatGPT. While these sources of noise can substantially bias down our analysis, we hope that our billions of high-frequency visitation interviews can still show the average purpose of ChatGPT usage by households as a whole and by demographic categories.

IV. Household Exposure and ChatGPT Use

A. Website-level exposure to GenAI: measurement

We want to construct a measure of the degree to which a household's internet use patterns increase the benefit from using a chatbot. We adapt the approach used in Eisfeldt et al. (2023) for creating firm exposure measures to the browsing behavior setting: First, we determine the overlap between the activities on particular websites and the capabilities of sites like ChatGPT. Then, we aggregate this domain level exposure into a household-level measure using the household's pre-ChatGPT internet usage patterns.

To label the exposure of websites to ChatGPT, we take the website descriptions and lists of activities that we previously generated and submit them individually to an LLM, together with a rubric that defines which types of website activities ChatGPT can be used for. We develop this rubric by summarizing the clusters of user activities observed by Handa et al. (2025) in their analysis of actual chatbot activities observed for users of Anthropic's Claude chatbot. This data provides an empirical starting point for what activities chatbots similar to ChatGPT are actually used for. To account for the additional capabilities of OpenAI's ChatGPT, we additionally incorporate multimodal capabilities that were announced for ChatGPT in September 2023.[13]

We use this rubric in two ways: On the one hand, we provide the domain, its description, and the rubric of activities that ChatGPT could substitute for, and ask the LLM to apply a binary exposure label to the website as a whole. However, this label is necessarily coarse as it does not allow for websites to be partially substituted by ChatGPT usage. Therefore, we separately submit each individual activity on a website (together with the domain and the website description) to an LLM with the same rubric and ask it to determine if this particular activity can be substituted by ChatGPT. The result is an exposure label for each of the 5 activities on each website, which we convert into a continuous exposure score that summarizes the share of the website's activities that are exposed. The full prompt used for this classification and details on the implementation are shown in Appendix A.

Website exposure share. The final result is a binary exposure label and a continuous (activity-based) exposure label for each domain. We show the distribution of these labels in Figure 7: Panel A shows that 39% of all household browsing activities are exposed to ChatGPT (weighting activities by the duration of time spent on their websites)[14] If we aggregate exposure to the website level, we obtain the distribution shown in Panel B: 16% of website browsing duration involves no activities that ChatGPT could be useful for, and only 3% of websites feature *only* exposed activities. The majority of websites by browsing duration (about 55%) are exposed in 40-60% of the activities on the website.

[13] See the announcement at: https://openai.com/index/chatgpt-can-now-see-hear-and-speak/

[14] Implicitly, this assumes that all activities on a website are happening simultaneously, as we cannot allocate browsing duration within a website to different activities.

B. Website-level exposure to GenAI: mapping to tasks

In order to better understand which types of online behavior are most likely to be substituted by ChatGPT use, we first explore the degree to which exposure concentrates in particular websites. The distribution of exposure scores across website categories provided by Comscore is shown in Appendix Figure 14. Panel A shows the distribution of website types in each exposure category: the majority of unexposed websites are entertainment sites, while the more exposed sites are comprised of education, lifestyle, and general information sites to a disproportionate extent, according to our methodology. Sites related to productive home activity (which include travel booking, shopping, health information, government services, home renovations, and technology information) make up a substantial share of all but the least exposed category.

What is the share of exposure categories within different website categories? As shown in Panel B of Appendix Figure 14, most social media and entertainment websites have low generative AI exposure (score of 0 or 1), while most business, education, job search, or lifestyle sites are categorized as having high exposure (score of 4 or 5). Some categories show large variation in the exposure of websites in that group, with a lot of variation in exposure across sites within the general information, search, personal finance, and productive home activity categories.

Website examples by exposure. What websites do these different exposure categories consist of? Table III shows the websites that make up the highest browsing duration shares of each exposure category. The lowest exposure categories (0 or 1 out of a maximum of 5) is dominated by social media platforms like Facebook, Instagram, TikTok, and Twitch; video streaming services like Netflix, Hulu, and Youtube, and other service sites like AOL, Craigslist, and Weather.com. The other end of the exposure score spectrum (4 or 5 out of 5) is dominated by education and knowledge sites, like Wikipedia, Quizlet, Edmentum, or Desire2Learn. The middle range of exposure (scores of 2 or 3) consists of many sites that often are multi-purpose (e.g. MSN, Yahoo), search engines (e.g. Bing, Google), and platforms for shopping and trading, such as Amazon, Ameritrade, Etsy, and Ebay.

Exposed activities. What types of exposed activities are driving most of the website exposure that we observe? As shown in Figure 8, more than half (54.5%) of the website exposure that we compute comes from information and research-related activities that ChatGPT could be used for. Advisory content constitutes 17.9% of the exposed activities, followed by data analysis (5.1%), content creation (4.9%) and technical support (3.5%) content on web-

sites.

C. *Measuring household-level exposure to GenAI*

While the previous section defined exposure to generative AI at the website level, we want to use this data to measure the degree to which households are likely to find generative AI useful based on their browsing behavior. Here, we are able to exploit the fact that we observe a household's full browsing activity to map from website exposure to the exposure of a household's browsing. The intuition for this mapping is similar to that in labor market-focused exposure measures (Eloundou et al., 2023; Eisfeldt et al., 2023), where the applicability of generative AI to a greater share of a worker's activities is assumed to make it more likely that generative AI impacts that worker. Analogously, we hypothesize that households for whom a greater share of their internet browsing involves activities that generative AI can be used for should be more likely to adopt generative AI in their personal lives. We can also then test this hypothesis empirically by estimating the effect of browsing-based exposure on household adoption of generative AI.

We define a household's exposure to generative AI as the share of its internet browsing that that is on websites with a high share of generative AI-exposed activities. That is, we compute

$$\text{HHGenAIExposure}_i = \sum_j \phi_{ij} \mathbb{1}[E_j \in \{4,5\}],$$

where E_j is the count of activities (out of 5) on website j that are generative AI exposed, and ϕ_{ij} is the share of a household's browsing duration during a pre-period that is spent on website j.

Empirically, we measure this exposure during the July 2021-June 2022 period. We choose a time period *before* the release of ChatGPT to ensure that the effects of generative AI on browsing do not affect our measure of a household's underlying pre-existing potential benefit from the technology.[15] Moreover, we allow for a 6-month gap between this measurement period and the release of ChatGPT (which would have started affecting browsing behavior in Dec. 2022), to avoid short-term mean reversion effects in browsing behavior.

[15] This is similar to the intuition for measuring "exposure" shares in shift-share instruments during a period preceding the estimation sample (Jaeger et al., 2018).

D. *Exposure to GenAI by household income and age*

One key advantage of the Comscore data is the ability to distinguish household usage patterns by income, both for browsing activity in general and ChatGPT usage in particular. While previous studies have found that higher income occupations are more exposed to generative AI (see, e.g. Eloundou et al. (2023); Eisfeldt et al. (2023)) and that higher income workers are more likely to use generative AI at work (Bick et al., 2024), ours is the first study of inequality in *private household* exposure to generative AI benefits.

Figure 9 shows that there is substantial variation in the degree to which the browsing activities of households with different incomes ex ante predict greater benefits from using generative AI: Panel A of the figure plots the average generative AI exposure of pre-ChatGPT internet use for each income group. The figure shows that households with less than $\$60K$ in income use websites that have relatively low exposure to ChatGPT, while higher income households have higher exposure. Moreover, there is a monotonic relationship between higher income and higher exposure for households with income above $\$60K$. The relationship is not monotonic for lower incomes, which may be due to the fact that the lower household income categories will pool households with low lifetime income, young workers at the beginning of their careers with high expected future lifetime income, and students for whom their own income is not reflective of their actual access to resources.

Panel B shows that exposure generative AI benefit in a household's browsing patterns declines monotonically with the age of the household head: exposure is about 25% higher for very young households relative to household with heads above retirement age (>65 years).

Panel C and D show the breakdown of internet use in each income and age category by website exposure level. The graph shows that the greater exposure among high-income households is due to both a monotonic increase in the browsing duration of high-exposure sites (score of 4 or 5) and a monotonic decrease in time spent on sites with no generative AI exposure (score of 0) as incomes increase above $60K.

E. *Household exposure and ChatGPT adoption*

To understand whether households that are more likely to benefit from using ChatGPT are more likely to adopt the technology, we first consider whether greater exposure based on a household's pre-ChatGPT browsing behavior predicts greater adoption.

Time patterns by exposure. Figure 10 plots adoption rates by quintile of exposure of the household: as the Figure shows, households with greater exposure adopt ChatGPT more rapidly, starting right after its release in November 2022, and continue to show higher adoption rates until the end of our data in December 2024. This relationship between exposure and adoption is monotonic, with adoption speed and levels across quintiles aligning with the ordering of the quintiles by exposure. The cumulative gap in having used ChatGPT between the highest and lowest exposure quintile is 24.4 pp in the highest exposure quintile compared to 10.5 pp in the lowest, by December 2024. Regular use–becoming an "adopter"–has a gap of 12.8 pp to 4.4 pp between the same groups.

This empirical pattern validates that the potential benefit to households from adopting generative AI does, in fact, drive differential adoption, which both validates our measure of household exposure and provides, to our knowledge, the first systematic empirical evidence that household characteristics with regard to browsing behavior can explain inequality in adoption of generative AI among private households.

Effect of exposure on GenAI adoption. To estimate the size of the effect of exposure on both the internal and external margin of ChatGPT use, we estimate regressions of the form

$$\text{ChatGPT Use}_{j,t} = \text{HHGenAIExp}_j + \text{Char}_{j,t} + \text{FE}_t + \text{FE}_j, \tag{1}$$

where GenAIExp_j is the household's pre-ChatGPT exposure to generative AI, and $\text{Char}_{j,t}$ is a vector of household characteristics. We include saturated interactions of income (8 categories), age (6 categories), household size (6 categories), and presence of children (yes or no) as fixed effects in all cross-sectional household level regressions to capture the effect of different demographics.

We measure *extensive margin* adoption effects by using a dependent variable that reflects whether the household has ever used ChatGPT or becomes a regular ChatGPT "adopter" (i.e. shows repeated use patterns - see Section II), such that the estimated coefficient on exposure indicates the degree to which exposure leads to adoption. To quantify the *intensive margin* of ChatGPT adoption, we limit the sample to all households that ever use ChatGPT and estimate the effect on measures of the average share of browsing duration spent on ChatGPT per month, and the average number of visits to the ChatGPT site per month (measured over a lagged 3-month period).

We do this estimation at different time horizons, corresponding to December 2023, June

2024, and December 2024 (12 and 18 months after the ChatGPT release).[16] Note that the flexible demographic controls mean that these estimates are holding the observable demographics constant, i.e. they cannot be explained by the differences in exposure patterns across demographic groups that we previously documented.

Extensive margin effects. The results are shown in Table IV, where columns (1) and (2) show the extensive margin effects. The interpretation of the coefficients is that they show the elasticity of the usage measure with regard to the share of a household's browsing activity that is exposed to generative AI. That is, the estimates in column (1) suggest that households with a one SD higher exposure (a 13 pp difference) saw 1.8 pp higher rates of having tried ChatGPT by December 2023, and a 2.8 pp higher rate of ever using it by December 2024. Adoption (regular use) rates are, by definition, smaller than use rates, so the estimated effects in column (2) also tend to be smaller, with the effect of a 1 SD higher exposure on adoption growing from 0.9 pp to 1.7 pp between December 2023 and 2024.

These effects of greater ex ante benefits of generative AI on its use are economically large: in the extensive margin sample, exposure has a standard deviation of 13 pp (and a mean of 11 pp). This means that a 1 SD change in exposure is associated with an increase in the rate of having tried ChatGPT relative to the average of 20%, 16% and 16% in Dec. 2023, June 2024, and Dec. 2024. The corresponding effect on adoption is an increase of 25%, 26%, and 21% relative to the mean in the same periods.

Intensive margin effects. The claim that the exposure measure captures potential benefits from generative AI use implies that once a household tries out the technology, it should also use it more intensively. This is confirmed by the results in columns (3) and (4) of Table IV, where we limit the sample only to households that have already tried ChatGPT: A one SD increase in exposure is associated with 0.4, 0.25, and 0.1 extra visits per month to the ChatGPT site for the average household that has already tried it in Dec. 2023, June 2024, and Dec. 2024, with all but the last of these estimates based on statistically significant coefficients. The fact that these relative effect sizes tend to decline over time is intuitive, as both the underlying average rates of use of ChatGPT (the denominator) increase over time, and as the pre-ChatGPT browsing behavior that defines ex-ante exposure becomes more distant in time from the technology use that we are trying to predict.

[16] While our sample does extend to December 2024, the Comscore sample experiences some attrition over time, so we are trading off sample size against extending the time since the ChatGPT release. To avoid a large reduction in sample size, we record the last status (used / adopted ChatGPT) for households that leave the sample between June and December 2024 and include them in the December 2024 estimation.

F. Heterogeneity by income and age in GenAI take-up

Given the large potential productivity impacts of generative AI adoption, and associated concerns about inequality, we next explore the degree to which adoption patterns vary by income.

ChatGPT use timing by income and age. Figure 13 shows the differences in the speed of adoption for households in different demographic categories: Panel A of the graph shows how the share of households that has ever tried ChatGPT varies across income categories. We find that higher income households (> $200K) are substantially more likely to have tried ChatGPT and also to have become an adopter. While the relationship is not precisely monotonic, it is generally true that higher income categories tend to have higher rates of use and adoption. Panel B shows the pattern of ChatGPT use by age category: there is a monotonic relationship between the age of the household head and the (lack of) adoption of ChatGPT: the youngest households are more than three times as likely to have tried the technology by December 2024 as the retirement age households.

This pattern is also summarized in Table VII, where the income and age groups are combined into coarser buckets corresponding to young (<34 years), middle, and old (>55 years) households, and low (<$60K), middle ($60-100K) and high (>$100K) household income groups.

Decomposing ChatGPT adoption. What drives these differences in adoption of generative AI between different demographics groups? We explore the following two hypotheses: (1) Some groups experience greater benefits from generative AI and therefore seek out the technology more (differences in exposure). (2) Some groups are more likely to use the technology to realize a given benefit from generative AI (technology "take-up"). Greater adaptability could result from better information, greater resources, complementary skills, or other advantages in technology access that are associated with particular groups' characteristics.

The potential for differences in ex-ante exposure to explain differences in ChatGPT use can be seen from panel A of Table VII: high-income households have 2 pp higher exposure to generative AI in their browsing behavior before the release of ChatGPT than low-income households. Similarly, the gap in exposure between young and old households is 2.3 pp. As we found in Table IV that exposure significantly affects ChatGPT use, these exposure gaps can also drive gaps in technology adoption. However, this assumes that the effects of

exposure on generative AI use are the same across groups, which we can test empirically.

Heterogeneity in exposure effects. In order to test the hypothesis that differences in take-up may also be able to explain differences in ChatGPT use between groups, we estimate differential take-up rates using specifications of the form

$$\text{ChatGPT Use}_{j,t} = \sum_k \beta^k \, \text{Group}_k \times \text{HH GenAI Exposure}_j + X'_{j,t}\phi + \varepsilon_{j,t}, \tag{2}$$

using the same extensive margin variables as in Table IV to capture different dimensions of ChatGPT use.

Table V, panels A and B show differences in exposure effects by income group and by age group. Perhaps surprisingly, we find in panel A that there are no significant differences in the effect of exposure to generative AI benefits on ChatGPT use across the different income categories. This means that better adaptability of higher income households does not arise from a stronger link between benefits and use.

Does the effect of exposure to GenAI benefits differ by age? In panel B of Table V, we find that adaptability *does* vary by age: column (1) shows that older households' take-up of ever using generative AI in response to ex ante exposure is 64% smaller (0.14 reduction relative to a 0.22 baseline coefficient for the young group) in Dec. 2023, and regular adoption effects are 69% smaller in column (2), with similarly large differences in effects in December 2024, shown in columns (3) and (4). This means that–at least for age gaps in ChatGPT use–*both* differences in exposure, and differences in take-up rates conditional on exposure may play a role in explaining gaps in technology adoption.

Oaxaca-Blinder decomposition of ChatGPT use gaps. More formally, we can quantify the drivers of differences in ChatGPT use between a "high" (H) and "low" (L) adoption group based on a demographic of interest (X) using a Oaxaca-Blinder decomposition of the following form:

$$\Delta_{H-L}\text{ChatGPT Use} = \underbrace{(\bar{X}^H - \bar{X}^L)\beta}_{\Delta_{H-L}\text{Exposure}} + \underbrace{\bar{X}^L(\beta^H - \beta^L)}_{\Delta_{H-L}\text{Take-up}} + \underbrace{(\bar{X}^H - \bar{X}^L)(\beta^H - \beta^L)}_{\Delta_{H-L}\text{Exposure}\times\Delta_{H-L}\text{Take-up}} + \varepsilon,$$

where β is the average effect coefficient estimate from Table IV, while β^H and β^L are group-specific exposure effect estimates.

We can use this decomposition to combine the coefficient estimates from Tables IV and V and the differences in exposure between groups shown in Table VII to quantify the share

of the gaps in ChatGPT use that is explained by the different components (or the residual). This decomposition is shown in Table VII. Each panel shows first the gap in ever having used ChatGPT (or regular adoption), then columns that show the predicted effects of the difference in household exposure, the difference in take-up conditional on exposure, and the interaction between the two terms (i.e. the effect of higher exposure coinciding with higher take-up rates in some groups). The fifth column in each panel shows the total share of the gap between the demographic groups that is explained by the sum of the different exposure effect components. If groups do not differ in their exposure, or exposure does not matter for ChatGPT use, this number will be zero, while 100% would mean that ChatGPT use gaps can be explained entirely by differences in browsing behavior-based exposure and take-up rates. Note that we are focusing here on only one proxy for the potential benefits of generative AI for a household–the range of websites used in the year before June 2022–so we would expect a substantial role for other factors not captured here.

Explaining ChatGPT use gaps by income & age. In panel A of Table VII, we find that a modest share of 2-10% of the gap in having ever used ChatGPT between low- and high-income households is explained by the potential benefits or take-up rates, and 4-10% of the gap in becoming a regular user. Moreover, most of this explanatory power comes from higher-income households having higher exposure than lower-income households.

Panel B does the same analysis for differences in ChatGPT use by age group. Once we combine differences in adaptability with differences in exposure between young and old households, we can explain 12-28% of the gap in having used ChatGPT and 18-40% the regular adoption gap. Most of this explanatory power comes from the effect in column (8) of greater take-up rates among young households driving higher ChatGPT use than for older households with similar exposure.

G. Programming knowledge and household ChatGPT adoption

One important driver of generative AI benefits in the labor market is the degree to which a worker's job involves programming, as LLMs excel at generating code and assisting in coding-related tasks (Peng et al., 2023; Cui et al., 2025). As a result, workers in software development and similar occupations have generally been at the forefront of adopting generative AI technologies.

Households that are familiar with, or actively involved in, programming-related tasks

should be more likely to benefit from using generative AI. We can test whether this difference in technical knowledge is an important driver of observed differences in ChatGPT use even at the private household level.

As a proxy for whether a household has knowledge of programming, we record whether the household's machine ever accessed Stack Overflow, a popular discussion forum where programmers ask questions and receive answers from other programmers, during the Jan. 2019-June 2022 period.[17] To validate that this is indeed a good proxy for interest in ChatGPT, we plot the rate of having tried ChatGPT in Appendix Figure 16 for households that have used Stack Overflow and those that have not, and find a substantially faster adoption rate for the group with programming knowledge.

Stack Overflow users also have substantially higher ex-ante exposure to generative AI in their browsing patterns, averaging 15.2% exposure relative to 10.8% among non-users. In panel C of Table V, we estimate whether this exposure also affects the tech-savvy households differentially: we find that take-up of ever using ChatGPT in response to household exposure is significantly larger among Stack Overflow users—with a coefficient that is approximately three times as large as that for non-users in both Dec. 2023 and June 2024. This suggests that past experience with technology mediates take-up of generative AI by households: when comparing two households that do tasks on the internet with similar applicability of generative AI, the one with past programming experience is substantially more likely to actually use ChatGPT.

How much of this gap in adoption between these tech-savvy households and the rest do these differences in exposure and take-up rates explain? Panel C of Table VII shows that gaps in ever using ChatGPT are substantial, rising from 22 pp in Dec. 2023 to 33 pp in Dec. 2024, while the adoption gap rises from 13 pp to 23 pp. However, differences in ex-ante exposure based on browsing patterns, and particularly the much higher estimated take-up rates, can together explain 21-24% of the gap in ever using, and 30-42% of the gap in adopting ChatGPT as a regular user. This means that household browsing behavior can be used to derive proxies for technical knowledge that matter for generative AI adoption, and that a large share of the greater adoption among more tech-savvy households is explained by their response to greater benefits from ChatGPT use that are captured in browsing behaviors in earlier years.

[17] We also include other coding-related sites hosted by the same domain (Stack Exchange), including subsites related to Computer Science, Code Review, Unix, Web Apps, , Database Administration, Game development, and Software engineering.

V. Impact of ChatGPT adoption on home browsing behavior

A. *Approach to identification*

In this section we estimate the effect that ChatGPT use had on households' overall browsing behavior: which types of online tasks were substituted by ChatGPT and which types of tasks were crowded in. We want to obtain an estimate that is plausibly causal: among households that use ChatGPT, what is the amount and composition of their *actual* home browsing activity relative to their *counterfactual* activity if they had not adopted ChatGPT? To arrive at such an estimate, we implement an instrumental variables approach that is based on the evidence in the previous section on which types of households utilize ChatGPT following its November 2022 release.

Naive OLS approach To motivate why we must implement an instrumental variables approach in the first place, it is instructive to first consider a "naive" OLS approach that utilizes households' observed decision to use ChatGPT or not. In particular, we consider dynamic quarterly regressions of the form

$$\text{browsing_outcome}_{j,t} = \gamma^{OLS}\ \text{Household GenAI User}_j \times \text{post_gpt}_t + \lambda_j + \text{Char}_{j,t} + \varepsilon_{j,t},$$

where the left-hand side is some browsing outcome of household j in quarter t, Household GenAI User$_j$ is an indicator variable that equals one if the household ever uses ChatGPT following its November 2022 release, and post_gpt$_t$ is a dummy variable that equals one starting in 2023Q1.[18]

The regressions include both household fixed effects as well as quarter - age bucket - income bucket - region fixed effects. As such, the estimate γ^{OLS} is based on comparing the evolution of the browsing behavior of households who have similar demographics but differential take-up of ChatGPT. Intuitively, then, γ^{OLS} identifies the causal effect of using ChatGPT only if households that take up ChatGPT would have exhibited the same changes

[18] For quarterly regressions, we cannot precisely capture the pre- vs. post-period around the November 30, 2022 release. We pick 2023Q1 as the first post-release quarter because, in practice, users of ChatGPT in our sample do not appear to start using ChatGPT, and demonstrate different behavior across other websites, until (and generally after) the first months of 2023. However, all of our results are quantitatively and qualitatively similar if we instead define 2022Q4 as the first post-release quarter.

in browsing behavior after the release of ChatGPT as observably-similar households who did not use ChatGPT.

A priori, we believe that this OLS identification condition is quite unlikely to hold. This is for two reasons. First, as we've shown and explored in the previous sections, a household's decision to utilize ChatGPT seems highly endogenous to sociodemographic characteristics (e.g., income and age) and home browsing needs (e.g., execution of coding tasks). Even after controlling for observable characteristics as of the end of 2022, such as income, age, and region, it is likely that variation across households in actual use of ChatGPT is driven, at least in part, by residual variation in unobserved determinants of households' exposure to GenAI technologies.[19] Second, even if we could perfectly control for households' *ex ante* characteristics, it is likely that *ex post* shocks to households' home browsing capabilities and/or needs affect their actual use of ChatGPT. For example, suppose that some households in our sample get laid off at some point during 2023 or 2024. These households may be expected to increase their overall time spent on home browsing, as well as tilt their home browsing towards productive uses such as searching for a new position. Both of these effects could make the household more likely to start using ChatGPT. Under such an example, even if the use of ChatGPT itself had no causal effect on the household's browsing behavior (or any other outcome), γ^{OLS} would still indicate a potentially large effect of ChatGPT use.

These concerns with the naive OLS approach appear to be reflected in the actual estimates shown in Appendix Figure 18. Panel (B) shows that observed users of ChatGPT increase their share of browsing activity on websites with a high degree of overlap to ChatGPT. At first glance, it may seem surprising that on average, users of ChatGPT *increase* their browsing of even websites for which ChatGPT serves as a potential substitute. But this would make complete sense if the naive OLS estimates are driven, in part, by households who adopt ChatGPT due to unobserved shocks of the sort described in the previous paragraph.[20] Such households have a greater need and/or desire to complete browsing tasks for which ChatGPT or related, non-GenAI sites are helpful, and do so through some combination (averaged across households) of ChatGPT and non-ChatGPT websites. This interpretation is further reinforced by the increased level of overall browsing activity evident in Panel (C) of Figure 18.

[19] For example, among households with the same income and age, those who use ChatGPT for home browsing may be more likely to have been introduced to the technology at work. To the extent that GenAI-exposed jobs are in occupations that employ workers with different skills and/or productivity levels (Eisfeldt and Papanikolaou, 2013), this would lead to bias in our estimate of γ^{OLS}.

[20] The pre-trends evident in Figure 18 suggest that in addition, ChatGPT use is still correlated with ex ante household characteristics that affect future browsing behavior, even after controlling for observable demographic characteristics as we do.

Instrumental variables approach: idea Given the endogeneity of the take up of ChatGPT, it is thus necessary to estimate its impacts via an instrumental variables approach. A valid IV would be one that leads to significant differences in the actual use of ChatGPT due to differences in *ex ante* ChatGPT adoption incentives/costs that are *orthogonal* to the household's post-GPT release behavior (including any ex post shocks).

Our analysis in the previous section suggests a potential instrument: the exposure of a household to ChatGPT's release given their ex ante browsing behavior. This measure is, as documented in the previous section, strongly related to households' actual take up of ChatGPT. Moreover, it plausibly satisfies the above-stated orthogonality condition. Once we condition on observable sociodemographic characteristics and broad differences in pre-GPT release browsing behavior (eg. use of productive- vs. leisure-oriented websites), variation in exposure is the result of differences among otherwise similar households in the precise tasks done through online browsing.

Instrumental variables approach: implementation We implement this IV approach be again defining household-level GenAI exposure to be the share of ex ante browsing activity in sites with a high degree of overlap–defined as 4/5 or 5/5 of activities begin potentially replaceable with ChatGPT,

$$\text{HHGenAIExposure}_i = \sum_j \phi_{ij} \mathbb{1}[E_j \in \{4,5\}],$$

again computed over July 2021 to June 2022. With this instrument in hand, we then estimate the two-stage system in which the reduced form equation is given by

$$\text{browsing_outcome}_{j,t} = \gamma^{RF}\, \text{HHGenAIExposure}_j \times \text{post_gpt}_t + \lambda_j + \text{Char}_{j,t} + \varepsilon_{j,t},$$

and the first-stage equation is given by

$$\text{adopt}_{j,t} = \gamma^{FS}\, \text{Household GenAI Exposure}_j \times \text{post_gpt}_t + \lambda_j + \text{Char}_{j,t} + \eta_{j,t},$$

B. Instrumental variables estimates

Overall browsing effects Figure 12 shows the dynamic IV estimates for the reduced form equation. Panel (B) shows that once instrumented with our exposure measure, the

use of ChatGPT *reduces* households' share of browsing time spent on ChatGPT-overlapping websites. Panel (B) shows that while ChatGPT use does not seem to lead to a significantly different amount of home browsing (if anything, the effect is slightly negative), Panel (C) shows that it shifts the composition of browsing *towards* leisure-oriented websites. This is especially the case starting in 2024.

Table VI presents our 2SLS regressions when the estimates are pooled over all quarters. As indicated by the dynamic estimates in Figure 12, while the effect of instrumented take-up of ChatGPT on the overall amount of browsing is not statistically significant, the amount of browsing activity deveoted to leisure-oriented sites goes up in an economically and statistically significant way in 2024.

Tasks crowded in vs. crowded out by ChatGPT Our results so far indicate that households that use ChatGPT tilt their browsing behavior away from productive sites whose tasks can plausibly be replaced with ChatGPT and towards leisure-oritented sites. Table IX provides insight into exactly which browsing activities are crowded in vs. crowded out by ChatGPT use. In particular, it shows pooled 2SLS estimates (using the exact same design as for the main browsing results) for dependent variables equal to the share of duration in different online activities (utilizing Comscore's classification of the activities done on websites).

Looking across the different regressions, we see that ChatGPT use crowds out the activities that–based on the high-frequency statistics of Figure 4–households utilize ChatGPT to conduct (such as education) and crowds in leisure-related activities such as online communications (eg. Facebook messenger and Discord) and, to a lesser extent, entertainment sites.

VI. Potential implications of home browsing effects

A. *Back-of-the-envelope welfare calculation*

[INCOMPLETE]

B. Outcomes beyond home browsing

[INCOMPLETE]

C. Potential interaction with labor demand shifts

[INCOMPLETE]

VII. Conclusion

[INCOMPLETE]

REFERENCES

Aguiar, Mark, Erik Hurst, and Loukas Karabarbounis, 2013, Time use during the great recession, *American Economic Review* 103, 1664–1696.

Bick, Alexander, Adam Blandin, and David J Deming, 2024, The rapid adoption of generative ai, Technical report, National Bureau of Economic Research.

Brynjolfsson, Erik, Danielle Li, and Lindsey R Raymond, 2023, Generative ai at work, Technical report, National Bureau of Economic Research.

Cui, Zheyuan Kevin, Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz, 2025, The effects of generative ai on high-skilled work: Evidence from three field experiments with software developers, *Available at SSRN 4945566* .

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, *Harvard Business School Technology & Operations Mgt. Unit Working Paper* .

Eisfeldt, Andrea L., and Dimitris Papanikolaou, 2013, Organization capital and the cross-section of expected returns, *The Journal of Finance* 68, 1365 – 1406.

Eisfeldt, Andrea L, Gregor Schubert, Miao Ben Zhang, and Bledi Taska, 2023, The labor impact of generative ai on firm values, *Available at SSRN 4436627* .

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, 2023, Gpts are gpts: An early look at the labor market impact potential of large language models, *arXiv preprint arXiv:2303.10130* .

Ghani, Miharaini Md, Wan Azani Wan Mustafa, Mohd Ekram Alhafis Hashim, Hafizul Fahri Hanafi, and Durratul Laquesha Shaiful Bakhtiar, 2022, Impact of generative ai on communication patterns in social media, *Journal of Advanced Research in Computing and Applications* 26, 22–34.

Handa, Kunal, Alex Tamkin, Miles McCain, Saffron Huang, Esin Durmus, Sarah Heck, Jared Mueller, Jerry Hong, Stuart Ritchie, Tim Belonax, et al., 2025, Which economic tasks are performed with ai? evidence from millions of claude conversations, *arXiv preprint arXiv:2503.04761* .

Humlum, Anders, and Emilie Vestergaard, 2025, The unequal adoption of chatgpt exacerbates existing inequalities among workers, *Proceedings of the National Academy of Sciences* 122, e2414972121.

Jaeger, David A, Joakim Ruist, and Jan Stuhler, 2018, Shift-share instruments and the impact of immigration, Technical report, National Bureau of Economic Research.

Khan, Salman, 2024, *Brave new words: How AI will revolutionize education (and why that's a good thing)* (Penguin).

Lei, Ke, and Yixuan Liu, 2025, When ai becomes a shopping advisor: A study on the impact of generative ai review on consumer purchase decision, *SAGE Open* 15, 21582440251357671.

Ling, Yier, and Alex Imas, 2025, Underreporting of ai use: The role of social desirability bias, *Available at SSRN* .

Lo, Andrew W, and Jillian Ross, 2024, Can chatgpt plan your retirement?: Generative ai and financial advice, *Harvard Data Science Review* .

McElheran, Kristina, J Frank Li, Erik Brynjolfsson, Zachary Kroff, Emin Dinlersoz, Lucia Foster, and Nikolas Zolas, 2024, Ai adoption in america: Who, what, and where, *Journal of Economics & Management Strategy* 33, 375–415.

Noy, Shakked, and Whitney Zhang, 2023, Experimental evidence on the productivity effects of generative artificial intelligence, *Available at SSRN 4375283* .

Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer, 2023, The impact of ai on developer productivity: Evidence from github copilot, *arXiv preprint arXiv:2302.06590* .

Pew Research Center, 2025, 34% of u.s. adults have used ChatGPT, about double the share in 2023, Survey methodology and topline results available.

Van Dijk, Jan, 2020, *The digital divide* (John Wiley & Sons).

Yu, Renzhe, Zhen Xu, Sky CH-Wang, and Richard Arum, 2024, Whose chatgpt? unveiling real-world educational inequalities introduced by large language models, *arXiv preprint arXiv:2410.22282* .

Zhang, Ruihan, Borou Yu, Jiajian Min, Yetong Xin, Zheng Wei, Juncheng Nemo Shi, Mingzhen Huang, Xianghao Kong, Nix Liu Xin, Shanshan Jiang, et al., 2025, Generative ai for film creation: A survey of recent advances, in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 6267–6279.

Main tables

Table I:
Distribution by income and age, Comscore sample vs. ACS

This table shows the distribution with respect to household income and age of our main Comscore sample (first column) and the ACS (second column). For the Comscore sample, income and age are from the most recent month prior to November 2022 that the panelist updated their demographic information. The ACS distribution is among all internet-using households in the 2022 survey.

	Comscore sample	**ACS**
HH Income (% of Panelists)		
<$25K	16.52	13.94
$25K - $40K	15.13	10.45
$40K - $60K	15.92	14.33
$60K - $75K	5.84	9.89
$75K - $100K	9.41	13.38
$100K - $150K	14.34	17.50
$150K - $200K	2.92	9.19
>$200K	19.91	11.33
Head of HH's Age (% of Panelists)		
18 - 24	7.28	4.09
25 - 34	9.89	16.14
35 - 44	13.29	18.50
45 - 54	30.52	17.69
55 - 64	18.67	18.64
65 and over	20.34	24.93

Table II:
ChatGPT exposure and use over time: by age & income

The table shows average rates of ChatGPT use for all households, and by age and income groups. See text for definitions of income and age categories. The sample for this analysis is the same as the regression sample in column (1) of Table .

		By Income			By Age		
	All	Low	Middle	High	Young	Middle	Old
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: % Household High Exposure Browsing Share							
Pre-ChatGPT (12 mths)	11.1	10.0	10.8	12.0	12.3	11.8	10.0
Panel B: % Tried ChatGPT							
December 2023	9.2	7.6	7.6	11.8	12.9	11.3	5.7
June 2024	11.9	9.7	10.1	15.2	17.2	14.3	7.5
December 2024	17.4	14.2	14.6	22.3	25.4	21.5	10.3
Panel C: % ChatGPT Adopters							
December 2023	3.6	3.0	3.1	4.6	5.4	4.6	2.0
June 2024	5.5	4.5	4.5	7.1	8.4	7.0	2.9
December 2024	8.0	6.6	6.5	10.3	12.1	10.2	4.3

Table III:
Top 10 websites in aggregate duration share over pre-GPT release period, split by fraction of activities that overlap with ChatGPT

0/5 Overlap		1/5 Overlap		2/5 Overlap	
Website	**Share**	**Website**	**Share**	**Website**	**Share**
facebook.com	24.65	youtube.com	34.94	msn.com	31.64
weather.com	4.06	aol.com	7.12	yahoo.com	15.97
pluto.tv	3.13	instagram.com	3.10	amazon.com	8.55
tubitv.com	2.64	veryfast.io	1.78	bing.com	7.38
netflix.com	1.59	craigslist.org	1.71	espn.com	0.94
jigsaw-world.com	1.38	walmart.com	1.63	robinhood.com	0.85
hulu.com	1.37	wellsfargo.com	1.25	ameritrade.com	0.83
trontv.com	1.27	tiktok.com	1.03	centurygames.com	0.77
twitch.tv	1.26	discord.com	0.97	etsy.com	0.76
sadpandastudios.com	0.93	microsoftonline.com	0.89	blackboard.com	0.71

3/5 Overlap		4/5 Overlap		5/5 Overlap	
Website	**Share**	**Website**	**Share**	**Website**	**Share**
google.com	43.14	intuit.com	6.85	wikipedia.org	9.51
live.com	5.59	indeed.com	3.94	gingersoftware.com	5.97
instructure.com	4.52	newsandpromotions.com	2.86	wwnorton.com	3.33
ebay.com	3.69	reddit.com	2.82	quizlet.com	2.30
pushflow.org	3.47	oculus.com	2.66	aarp.org	1.71
sunbiz.org	2.98	office.com	2.56	state.gov	1.58
microsoft.com	1.32	fandom.com	2.54	edmentum.com	1.43
twitter.com	1.27	edgenuity.com	2.37	ny.gov	1.42
bankofamerica.com	0.85	cengage.com	2.06	desire2learn.com	1.16
buildinglink.com	0.78	fidelity.com	1.71	atitesting.com	1.15

Table IV:
Exposure effects on the extensive and intensive margin of ChatGPT use

This table shows estimates of machine-level cross-sectional specifications of the form

$$\text{ChatGPT Use}_{j,t} = \gamma\, \text{Household GenAI Exposure}_j + X'_{j,t}\phi + \varepsilon_{j,t},$$

where household generative AI exposure is measured based on the ChatGPT substitutability share of browsing activity in the 12 months preceding the release of ChatGPT. Specifically, *HH GenAI Exposure* is the browsing duration share of websites that are in the high exposure (4 or 5 out of 5) category based on the website activities. The dependent variable in columns 1 and 2 is the extensive margin of ChatGPT use: whether a machine has ever used ChatGPT, and whether it ever became a ChatGPT adopter (repeated user in sequential months). In columns 3 and 4, the dependent variable is the intensive margin of ChatGPT use, after conditioning on whether the machine has ever used ChatGPT. Column 3 shows the average monthly share of browsing time spent on the ChatGPT site during the previous 3 months (including the reference month), and column 4 shows the average monthly visits to the site over the same period. All regressions include saturated fixed effects that capture the interaction between income categories, age categories, household size, and an indicator for the presence of children. Heteroskedasticity-robust standard errors clustered at the company level in parentheses: * p<0.10, ** p<0.05, *** p<0.01.

	Extensive margin		Intensive margin	
Dependent variable:	Tried ChatGPT	Adopted ChatGPT	Share of Browsing Time	Visits per Month
	(1)	(2)	(3)	(4)
Panel A: December 2023				
HH GenAI Exposure	0.14***	0.07***	0.04***	3.14***
	(0.02)	(0.02)	(0.01)	(0.97)
Observations	20,817	20,817	1,709	1,892
Dep. var. mean	0.094	0.037	0.011	1.262
Indep. var. SD	0.128	0.128	0.136	0.138
Panel B: June 2024				
HH GenAI Exposure	0.15***	0.11***	0.01**	1.94**
	(0.02)	(0.02)	(0.01)	(0.89)
Observations	20,817	20,817	2,127	2,422
Dep. var. mean	0.120	0.055	0.009	1.550
Indep. var. SD	0.128	0.128	0.131	0.133
Panel C: December 2024				
HH GenAI Exposure	0.22***	0.13***	0.02	0.75
	(0.03)	(0.02)	(0.01)	(1.12)
Observations	19,990	19,990	2,723	3,397
Dep. var. mean	0.173	0.080	0.023	2.255
Indep. var. SD	0.128	0.128	0.130	0.132
Sample	All	All	Used ChatGPT	Used ChatGPT
Income x Age x HH Size x Children FEs	X	X	X	X

Table V:
Exposure effects on the extensive margin of ChatGPT use: by age & income

This table shows estimates of machine-level cross-sectional specifications of the form

$$\text{ChatGPT Use}_{j,t} = \beta\,\text{HH GenAI Exposure}_j + \textstyle\sum_k \gamma_k\,\text{Group}_k \times \text{HH GenAI Exposure}_j + X'_{j,t}\phi + \varepsilon_{j,t},$$

where household generative AI exposure is measured based on the ChatGPT substitutability share of browsing activity in the 12 months preceding the release of ChatGPT. Specifically, *HH GenAI Exposure* is the browsing duration share of websites that are in the high exposure (4 or 5 out of 5) category based on the website activities. The dependent variable is the extensive margin of ChatGPT use: whether a machine has ever used ChatGPT (columns 1 and 3), and whether it ever became a ChatGPT adopter (repeated user in sequential months), shown in columns 2 and 4. The omitted reference category for the interaction terms is the"low income" group in panel A, the "young" group in panel B, and the non-tech group in panel C. All regressions include saturated fixed effects that capture the interaction between income categories, age categories, household size, and an indicator for the presence of children. Heteroskedasticity-robust standard errors clustered at the household level in parentheses: * p<0.10, ** p<0.05, *** p<0.01.

Period	December 2023		June 2024	
Dependent variable:	Used ChatGPT	Adopted ChatGPT	Used ChatGPT	Adopted ChatGPT
	(1)	(2)	(3)	(4)
Panel A: By Income Group				
HH GenAI Exposure	0.11***	0.07***	0.12***	0.09***
	(0.03)	(0.02)	(0.03)	(0.02)
HH GenAI Exposure x Middle Income	0.08	0.02	0.07	0.04
	(0.05)	(0.03)	(0.05)	(0.04)
HH GenAI Exposure x High Income	-0.00	-0.01	0.02	0.01
	(0.05)	(0.03)	(0.05)	(0.04)
Observations	20,817	20,817	20,817	20,817
Panel B: By Age Group				
HH GenAI Exposure	0.22***	0.13***	0.22***	0.22***
	(0.06)	(0.04)	(0.07)	(0.05)
HH GenAI Exposure x Middle Age	-0.07	-0.05	-0.05	-0.12*
	(0.07)	(0.05)	(0.08)	(0.06)
HH GenAI Exposure x Old Age	-0.14**	-0.09**	-0.12	-0.16***
	(0.07)	(0.04)	(0.08)	(0.06)
Observations	20,817	20,817	20,817	20,817
Panel C: By Tech Knowledge (proxied by Stack Overflow use)				
HH GenAI Exposure	0.09***	0.05***	0.10***	0.07***
	(0.02)	(0.01)	(0.02)	(0.02)
HH GenAI Exposure x Stack Overflow User	0.24*	0.14	0.31**	0.27**
	(0.13)	(0.09)	(0.13)	(0.11)
Observations	20,817	20,817	20,817	20,817
Sample	All	All	All	All
Income x Age x HH Size x Children FEs	X	X	X	X

Table VI:
2SLS regressions of effect of instrumented ChatGPT adoption on exposed browsing activity (share of duration)

This table shows estimates of monthly panelist-level 2SLS regressions in which the reduced-form equation is

$$\text{exposed_browsing}_{j,t} = \gamma^{RF}\ \text{Household GenAI Exposure}_j \times \text{post_gpt}_t + \lambda_j + \text{Char}_{j,t} + \varepsilon_{j,t},$$

and the first-stage equation is

$$\text{adopt}_{j,t} = \gamma^{FS}\ \text{Household GenAI Exposure}_j \times \text{post_gpt}_t + \lambda_j + \text{Char}_{j,t} + \eta_{j,t},$$

The endogenous variable is a dummy for whether household j has, as of month t, adopted ChatGPT. The instrument is the panelist's predicted exposure to ChatGPT, interacted with the dummy post_gpt_t that equals one starting in December 2022. The dependent variable in each regression is a measure of the household's browsing activity on sites exposed to ChatGPT. In Columns (1) and (2), it is the share of the panelist's monthly browsing duration that is in sites for which 4/5 or 5/5 of activities overlap with ChatGPT; in Columns (3) and (4) it is the duration share in sites that have any overlap with ChatGPT. In Columns (2) and (4), the instrument is interacted with two separate post-GPT release dummies: one that is one for December 2022-December 2023 and another that is one for January 2024-December 2024. All regressions include machine fixed effects as well as machine effects for each unique month X income bin X age bin X MSA X household size bin X presence of children indicator. Heteroskedasticity-robust standard errors clustered at the panelist level are in parentheses: * p<0.10, ** p<0.05, *** p<0.01.

	Exposed (high overlap) duration share		Exposed site (any overlap) duration share	
	(1)	(2)	(3)	(4)
Post GPT Release X GPT Adoption	-2.836***		-0.926***	
	(0.451)		(0.175)	
Post Short Run (2023) X Adoption		-2.932***		-0.949***
		(0.512)		(0.198)
Post Medium Run (2024) X Adoption		-2.721***		-0.899***
		(0.451)		(0.186)
Machine FEs	Yes	Yes	Yes	Yes
Quarter X Demographics X Region FEs	Yes	Yes	Yes	Yes
N	371,668	371,668	370,436	370,436
IV F-stat	42.839	18.277	42.936	18.321

Table VII:
Decomposition of age & income gaps in ChatGPT adoption

This table quantifies the drivers of differences in ChatGPT use between a "high" (H) and "low" (L) group based on a demographic of interest (X) using a Oaxaca-Blinder decomposition of the following form:

$$\Delta_{H-L}\text{ChatGPT Use} = \underbrace{(\bar{X}^H - \bar{X}^L)\beta}_{\Delta_{H-L}\text{Exposure}} + \underbrace{\bar{X}^L(\beta^H - \beta^L)}_{\Delta_{H-L}\text{Take-up}} + \underbrace{(\bar{X}^H - \bar{X}^L)(\beta^H - \beta^L)}_{\Delta_{H-L}\text{Exposure}\times\Delta_{H-L}\text{Take-up}} + \varepsilon.$$

The columns for each demographic grouping (income, age, tech knowledge) show first the gap between the highest and lowest adoption group (high minus low income; young minus old age; tech user minus non-tech user). Then, each of the next three columns correspond to the 3 terms in the decomposition shown above. The last column in bold shows the ratio between the sum of the three decomposition terms (the part of the gap explained by the decomposition) and the total gap.

	(A) By Income					(B) By Age					(C) By Tech Knowledge				
	High-Low Gap	Δ E	Δ T	Δ E × Δ T	% **Explained**	Young-Old Gap	Δ E	Δ T	Δ E × Δ T	% **Explained**	High-Low Gap	Δ E	Δ T	Δ E × Δ T	% **Explained**
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
Panel A: % Used ChatGPT															
December 2023	4.13	0.27	-0.04	-0.01	**5.53**	7.23	0.31	1.43	0.33	**28.66**	22.43	0.65	2.86	1.26	**21.29**
June 2024	5.53	0.31	0.21	0.04	**10.07**	9.75	0.35	1.16	0.27	**18.26**	24.52	0.73	3.24	1.42	**21.99**
December 2024	8.05	0.46	-0.22	-0.05	**2.39**	15.06	0.52	1.06	0.25	**12.14**	32.62	1.09	4.72	2.13	**24.34**
Panel B: % ChatGPT Adopters															
December 2023	1.55	0.15	-0.08	-0.02	**3.87**	3.43	0.17	0.93	0.21	**38.42**	13.49	0.36	2.61	1.15	**30.49**
June 2024	2.69	0.21	0.05	0.01	**10.20**	5.52	0.24	1.63	0.37	**40.62**	16.26	0.51	4.44	1.95	**42.42**
December 2024	3.67	0.27	-0.02	-0.00	**6.80**	7.80	0.30	0.91	0.22	**18.29**	23.10	0.64	4.47	2.02	**30.83**

Table VIII:
2SLS regressions of effect of instrumented ChatGPT adoption on amount and leisure composition of browsing activity

This table shows estimates of monthly panelist-level 2SLS regressions in which the reduced-form equation is

$$\text{browsing_outcome}_{j,t} = \gamma^{RF}\ \text{Household GenAI Exposure}_j \times \text{post_gpt}_t + \lambda_j + \text{Char}_{j,t} + \varepsilon_{j,t},$$

and the first-stage equation is

$$\text{adopt}_{j,t} = \gamma^{FS}\ \text{Household GenAI Exposure}_j \times \text{post_gpt}_t + \lambda_j + \text{Char}_{j,t} + \eta_{j,t},$$

The endogenous variable is a dummy for whether household j has, as of month t, adopted ChatGPT. The instrument is the panelist's predicted exposure to ChatGPT, interacted with the dummy post_gpt_t that equals one starting in December 2022. The dependent variable in each regression is a measure of the household's browsing activity. In Columns (1) and (2), it is the monthly log of browsing duration across all sites; in Columns (3) and (4) it is the monthly log of browsing duration in leisure sites; and in Columns (5) and (6), it is the share of the month's total browsing that is in leisure sites. In Columns (2), (4), and (6), the instrument is interacted with two separate post-GPT release dummies: one that is one for December 2022-December 2023 and another that is one for January 2024-December 2024. All regressions include machine fixed effects as well as machine effects for each unique month X income bin X age bin X MSA X household size bin X presence of children indicator. Heteroskedasticity-robust standard errors clustered at the panelist level are in parentheses: * p<0.10, ** p<0.05, *** p<0.01.

	Log browsing duration, all sites		Log browsing duration, leisure sites		Leisure duration share	
	(1)	(2)	(3)	(4)	(5)	(6)
Post GPT Release X GPT Adoption	-1.067		0.391		0.261***	
	(0.725)		(0.642)		(0.101)	
Post Short Run (2023) X Adoption		-1.266		-1.063		0.211*
		(0.811)		(0.750)		(0.109)
Post Medium Run (2024) X Adoption		-0.779		2.763***		0.330***
		(0.880)		(0.967)		(0.118)
Machine FEs	Yes	Yes	Yes	Yes	Yes	Yes
Quarter X Demographics X Region FEs	Yes	Yes	Yes	Yes	Yes	Yes
N	371,680	371,680	345,438	345,438	371,680	371,680
IV F-stat	42.841	18.276	46.736	21.272	42.841	18.276

Table IX:
2SLS regressions of effect of instrumented ChatGPT adoption on amount and leisure composition of browsing activity

This table shows estimates of monthly panelist-level 2SLS regressions in which the reduced-form equation is

$$\text{browsing_outcome}_{j,t} = \gamma^{RF}\ \text{Household GenAI Exposure}_j \times \text{post_gpt}_t + \lambda_j + \text{Char}_{j,t} + \varepsilon_{j,t},$$

and the first-stage equation is

$$\text{adopt}_{j,t} = \gamma^{FS}\ \text{Household GenAI Exposure}_j \times \text{post_gpt}_t + \lambda_j + \text{Char}_{j,t} + \eta_{j,t},$$

The endogenous variable is a dummy for whether household j has, as of month t, adopted ChatGPT. The instrument is the panelist's predicted exposure to ChatGPT, interacted with the dummy post_gpt_t that equals one starting in December 2022. The dependent variable in each regression is a measure of the household's browsing activity. In Columns (1) and (2), it is the monthly log of browsing duration across all sites; in Columns (3) and (4) it is the monthly log of browsing duration in leisure sites; and in Columns (5) and (6), it is the share of the month's total browsing that is in leisure sites. In Columns (2), (4), and (6), the instrument is interacted with two separate post-GPT release dummies: one that is one for December 2022-December 2023 and another that is one for January 2024-December 2024. All regressions include machine fixed effects as well as machine effects for each unique month X income bin X age bin X MSA X household size bin X presence of children indicator. Heteroskedasticity-robust standard errors clustered at the panelist level are in parentheses: * p<0.10, ** p<0.05, *** p<0.01.

	(1) Multi-use	(2) Search	(3) Education	(4) Communications	(5) Social Media	(6) Promotional	(7) Shopping	(8) Technology	(9) News	(10) Entertainment	(11) Information	(12) Finances	(13) Job search	(14) Games
Post Short Run (2023) X Adoption	-0.107	0.119*	-0.378***	0.484***	-0.016	-0.078	0.194***	0.005	0.080	0.042	0.003	-0.091	-0.074***	-0.095*
	(0.093)	(0.062)	(0.095)	(0.098)	(0.035)	(0.055)	(0.069)	(0.052)	(0.052)	(0.054)	(0.026)	(0.065)	(0.025)	(0.049)
Post Medium Run (2024) X Adoption	-0.145	0.092	-0.341***	0.535***	0.085*	-0.142**	0.274***	-0.000	-0.022	0.056	0.027	-0.115	-0.069***	-0.014
	(0.099)	(0.059)	(0.082)	(0.105)	(0.046)	(0.058)	(0.070)	(0.060)	(0.047)	(0.061)	(0.024)	(0.072)	(0.025)	(0.052)
Machine FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Quarter X Demographics X Region FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	371,680	371,680	371,680	371,680	371,680	371,680	371,680	371,680	371,680	371,680	371,680	371,680	371,680	371,680
IV F-stat	18.276	18.276	18.276	18.276	18.276	18.276	18.276	18.276	18.276	18.276	18.276	18.276	18.276	18.276

Main figures

Figure 1: Household ChatGPT use over time

This figure shows the share of households that has ever tried ChatGPT or become a regular ChatGPT adopter.

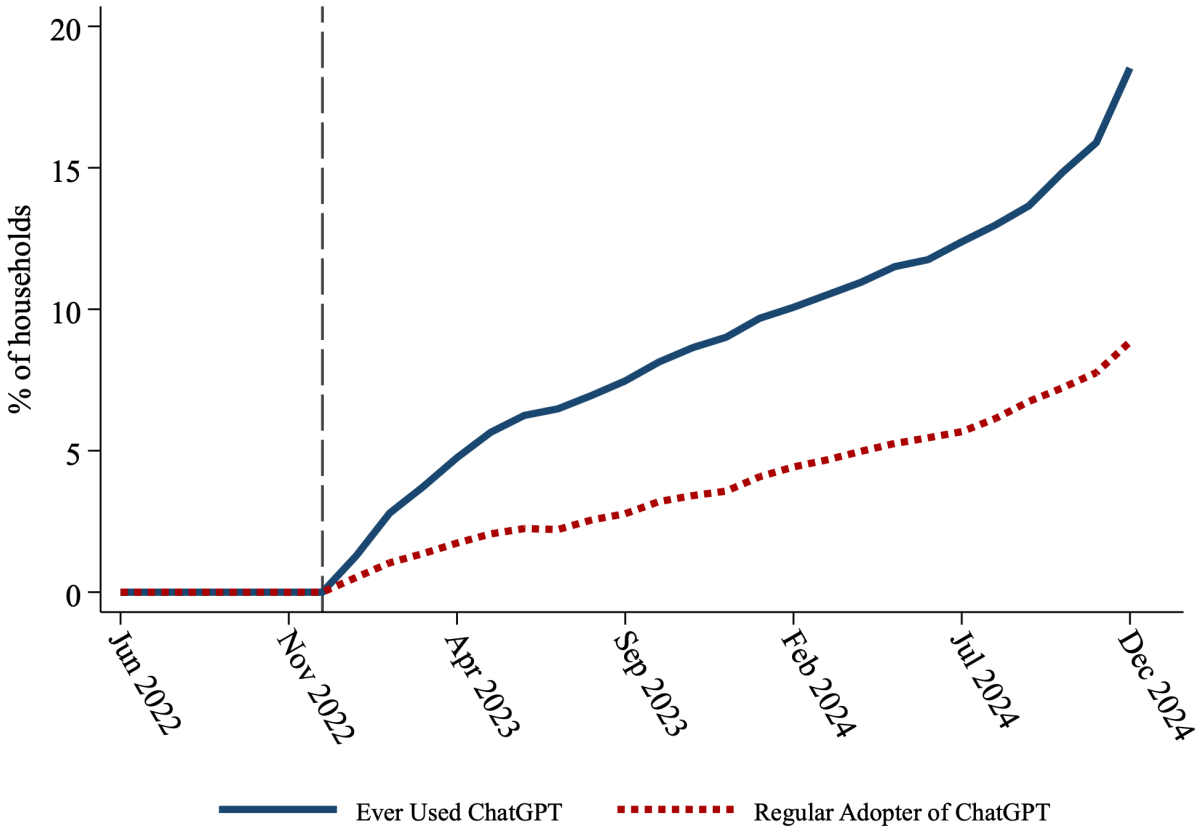
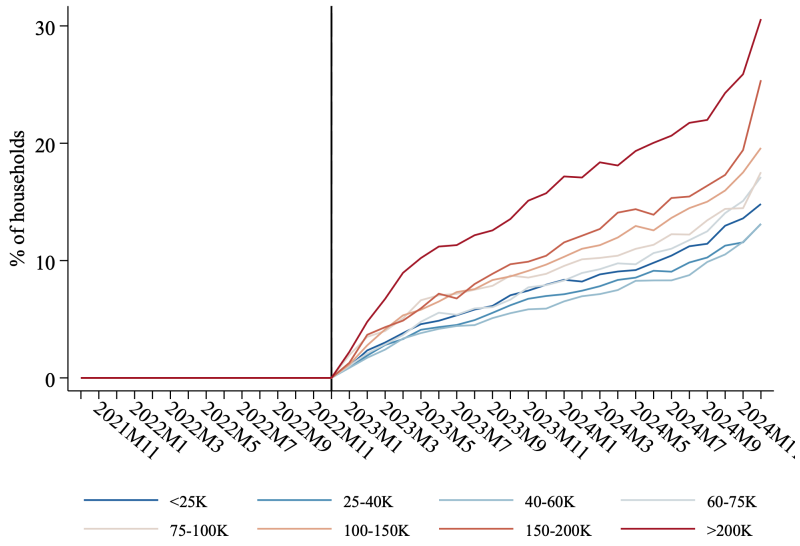
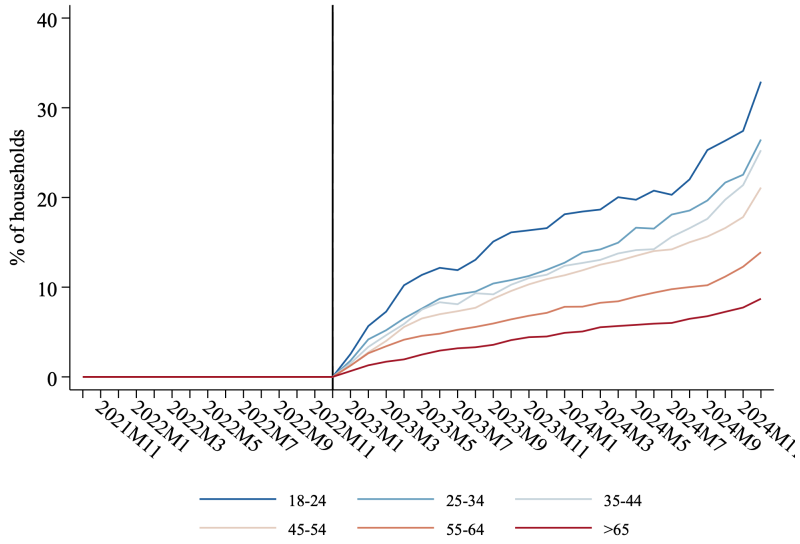


Figure 2: ChatGPT use by household income and age

This figure shows the share of households in each demographic group that has ever tried ChatGPT.

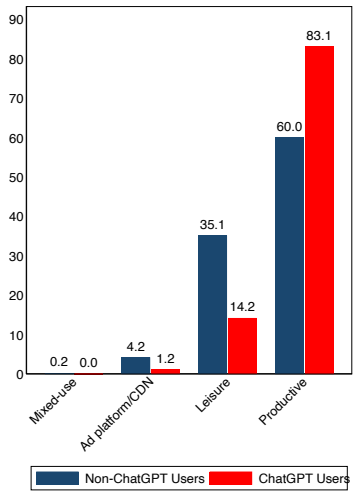


(A) ChatGPT use by income bucket

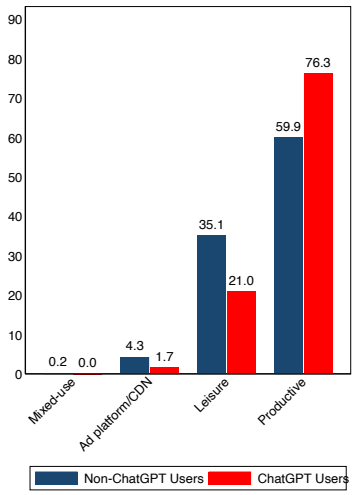


(B) ChatGPT use by age bucket

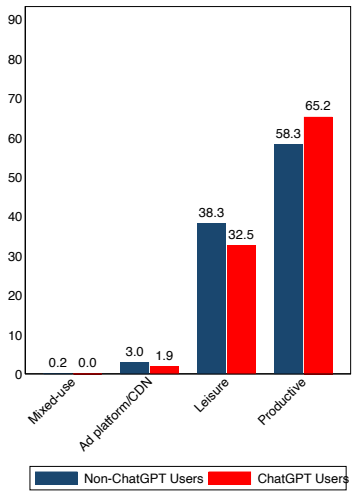
Figure 3: Usage of leisure vs. productive sites in vs. outside ChatGPT use windows



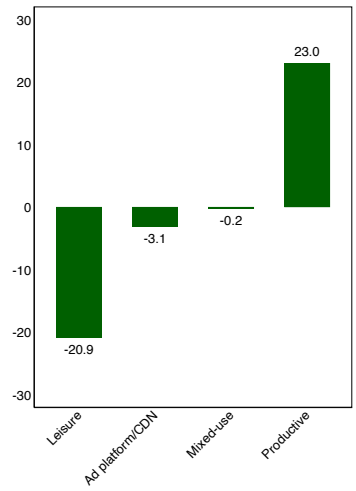
(A) GPT window: browsing shares



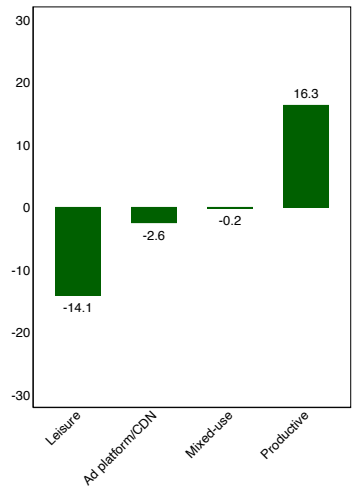
(B) GPT day, outside window: browsing shares



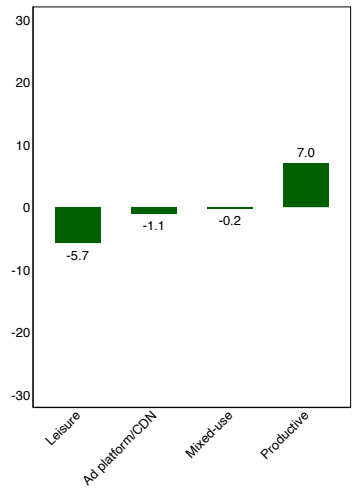
(C) Pre-GPT use: browsing shares



(D) GPT window: differences in shares

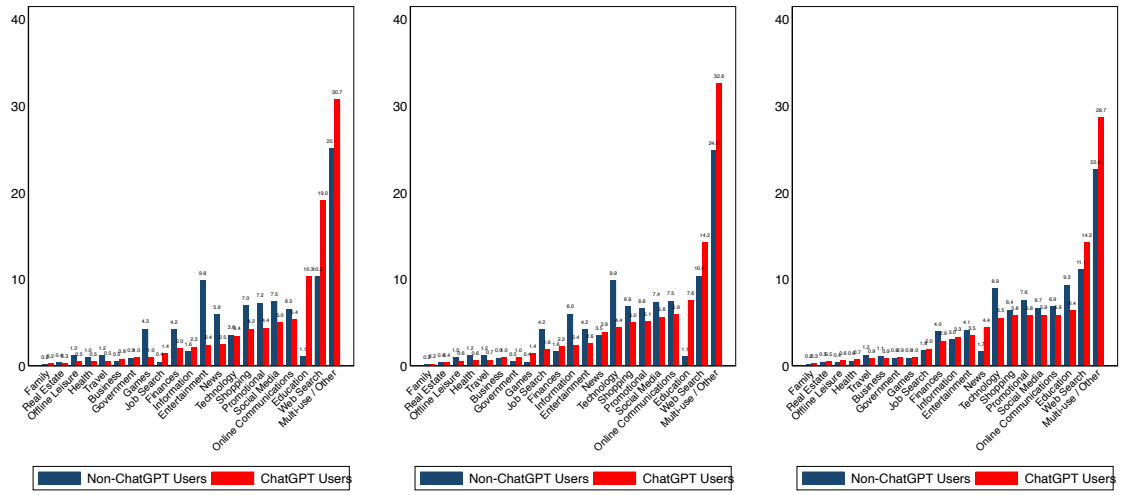


(E) GPT day, outside window: differences in shares



(F) Pre-GPT use: differences in shares

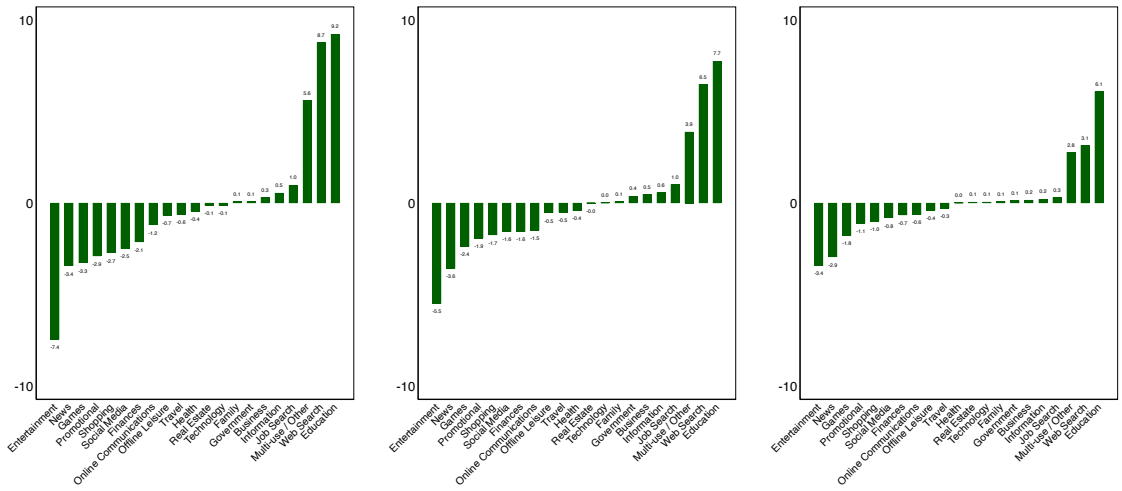
Figure 4: Usage of sites associated with different activities in vs. outside ChatGPT use windows



(A) GPT window: browsing shares

(B) GPT day, outside window: browsing shares

(C) Pre-GPT use: browsing shares

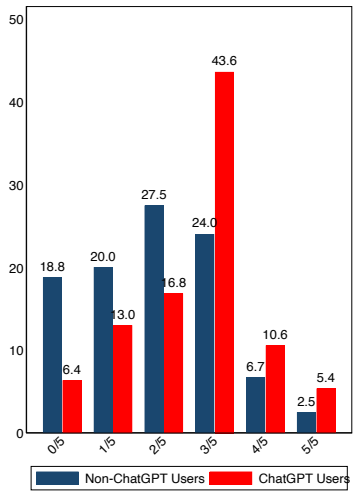


(D) GPT window: differences in shares

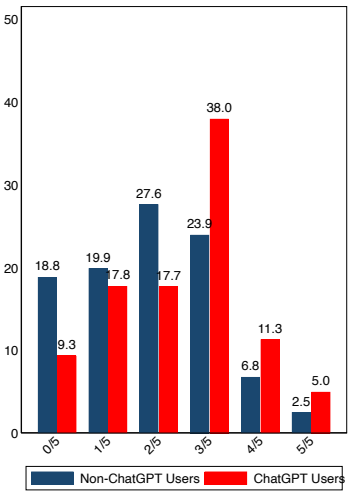
(E) GPT day, outside window: differences in shares

(F) Pre-GPT use: differences in shares

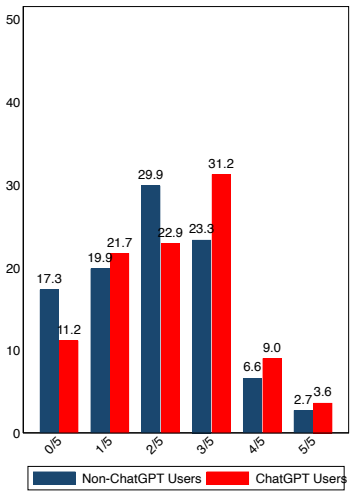
Figure 5: Usage of sites with different GenAI overlap in vs. outside ChatGPT use windows



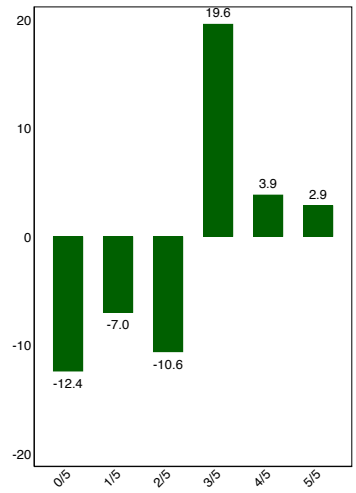
(A) GPT window: browsing shares



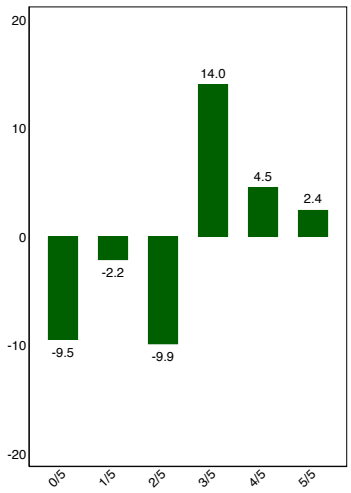
(B) GPT day, outside window: browsing shares



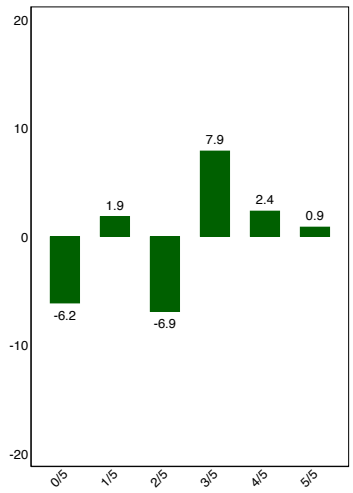
(C) Pre-GPT use: browsing shares



(D) GPT window: differences in shares



(E) GPT day, outside window: differences in shares



(F) Pre-GPT use: differences in shares

Figure 6:
Website classification methodology

This figure illustrates the data processing pipeline used for generating labels that classify websites into having activities exposed to generative AI and into mainly having *Productive* or *Leisure* uses, with separate labels for *Ad platforms & Content distribution networks*, as well as *Mixed-use* sites that cannot be labeled definitively.

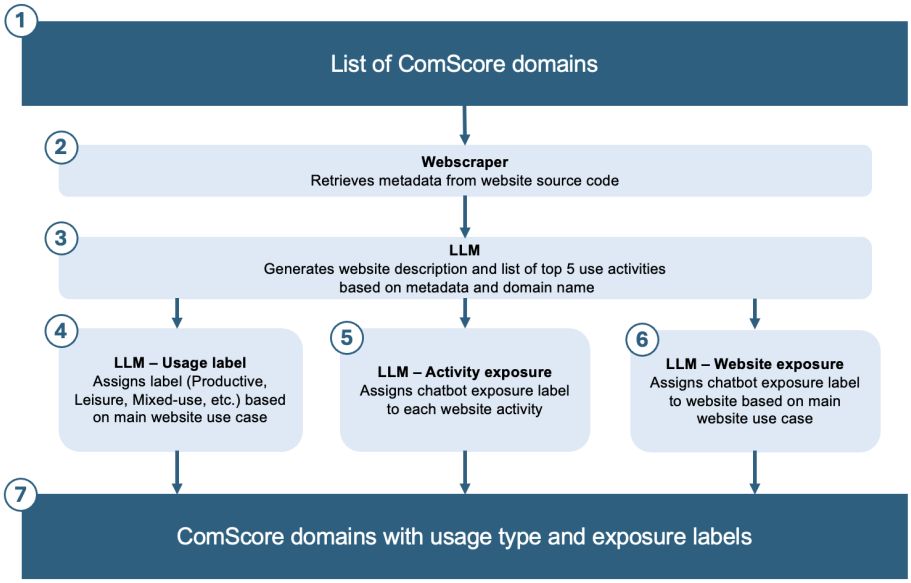
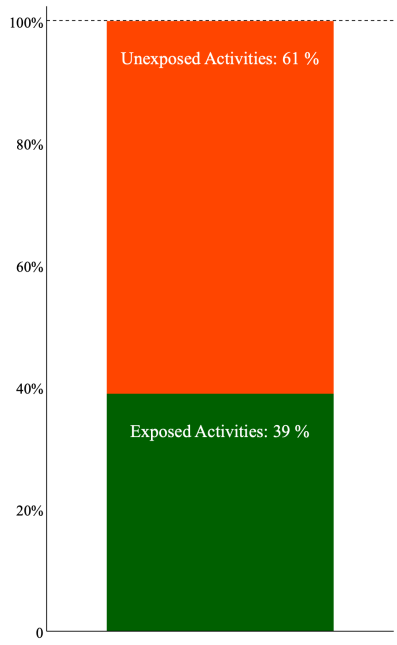
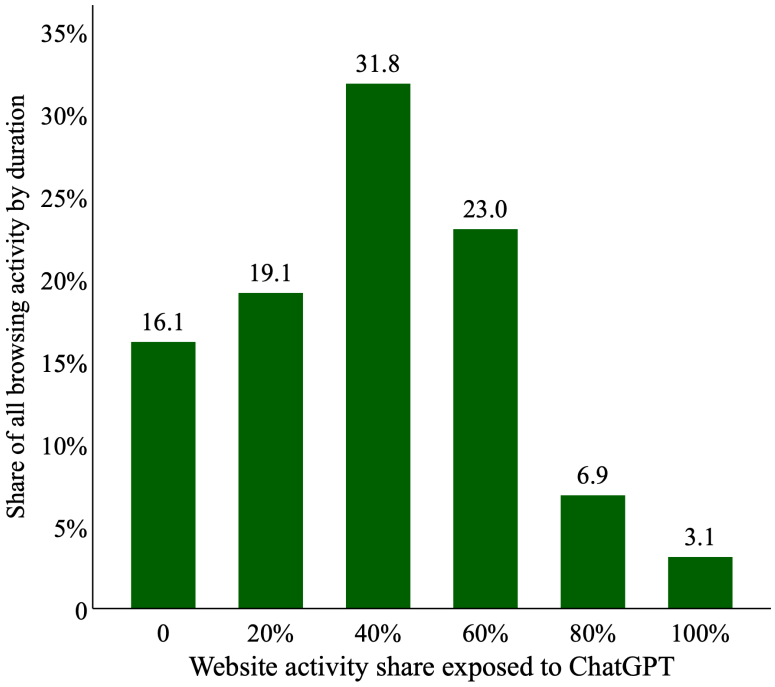


Figure 7:
Exposure of website activity to ChatGPT substitution.

This figure shows the share of online household activity duration that consists of website activities that ChatGPT can be useful for. Panel A shows the share of all activity that is labeled as exposed, while Panel B shows the distribution of browsing duratoin over websites with different levels of exposure, proxied by the share of activities associated with the website that are labeled as exposed.



(A) Duration share of exposed activities



(B) Duration distribution over website exposure levels

Figure 8:
Share of exposed website activity by type

This figure shows the share by duration of different website activity categories in the online household activity that ChatGPT can be useful for.

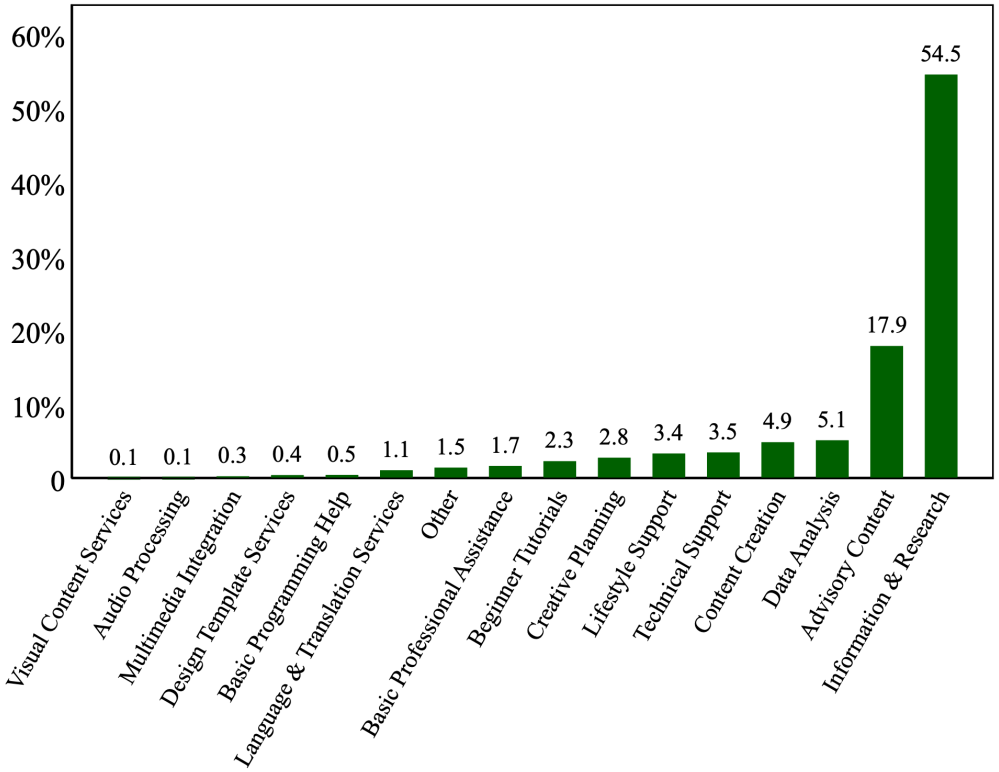
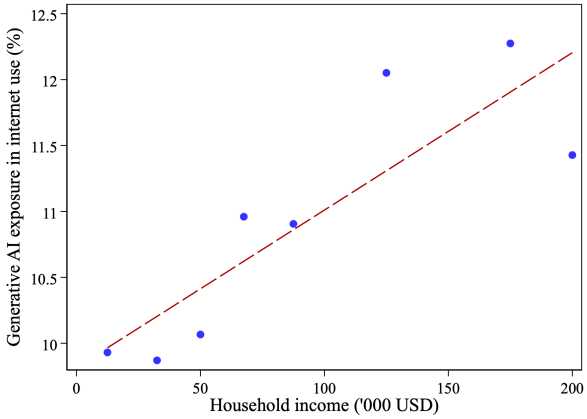
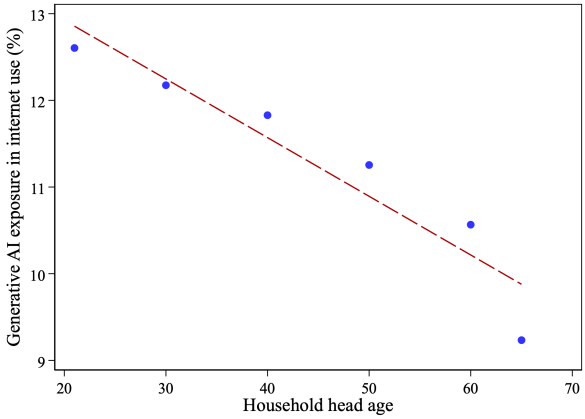


Figure 9: Internet use exposure to generative AI by household income & age.

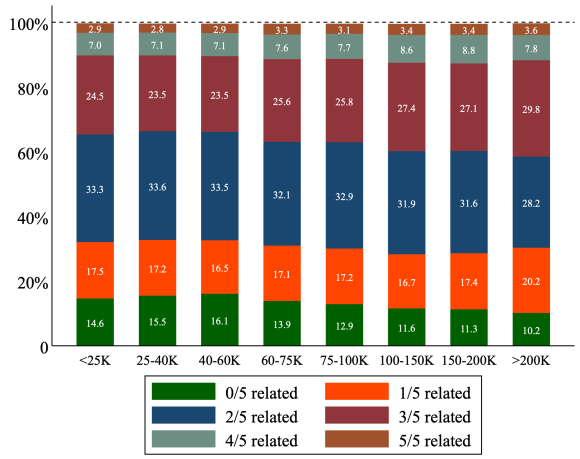
The figures below shows the average exposure to generative AI of internet use of households in different income and age buckets during 12 months preceding the release of ChatGPT by 6 months (July 2021 - June 2022). Within households, website activity is weighted by the duration of use. When aggregating across households, they are weighted by the share of their age x income group in ACS data from 2022. Panel A shows the average exposure in each income category. Income and age categories are mapped to the point on the horizontal axis that corresponds to the middle of the income and age ranges, except the highest category, which is plotted at the value of its lower bound. The red dashed line shows an unweighted linear fit across income categories. Panel B shows the distribution of internet use duration within each income category across website exposure levels.



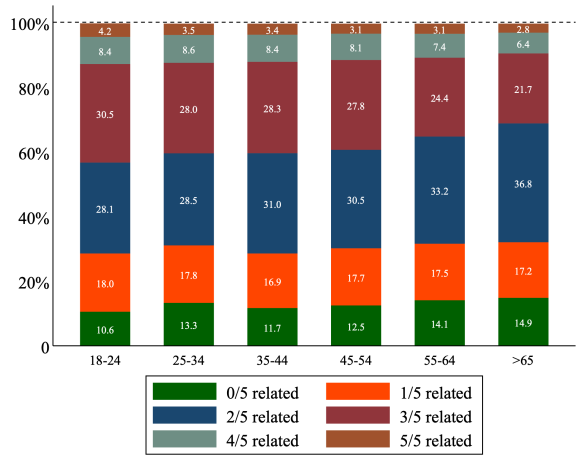
(A) Generative AI exposure by income bucket



(B) Generative AI exposure by age bucket



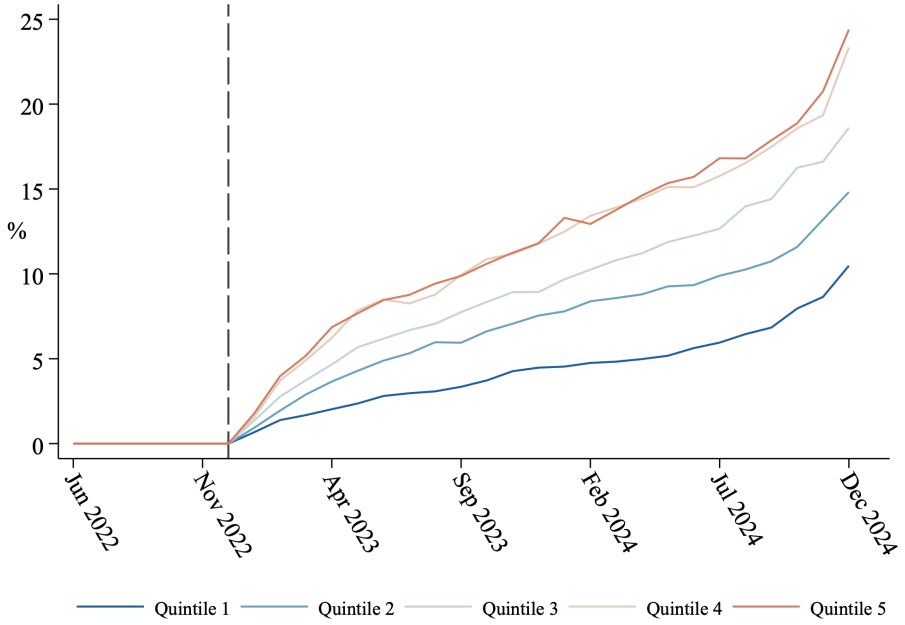
(C) Detailed generative AI exposure categories by income bucket



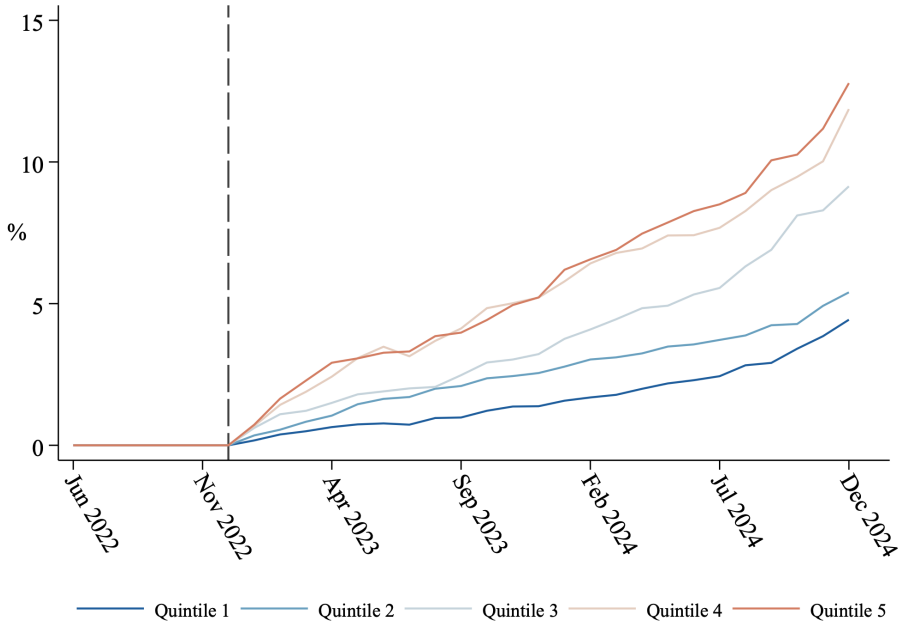
(D) Detailed generative AI exposure categories by age bucket

Figure 10: ChatGPT use by generative AI exposure

This figure shows the share of households in each generative AI exposure quintile that has ever tried ChatGPT (Panel A), and that has ever become a regular user ("adopter") of ChatGPT (Panel B).



(A) Share that has tried ChatGPT



(B) Share of ChatGPT "adopters"

Figure 11: Pre-ChatGPT release duration shares for different site categories, by household income bucket

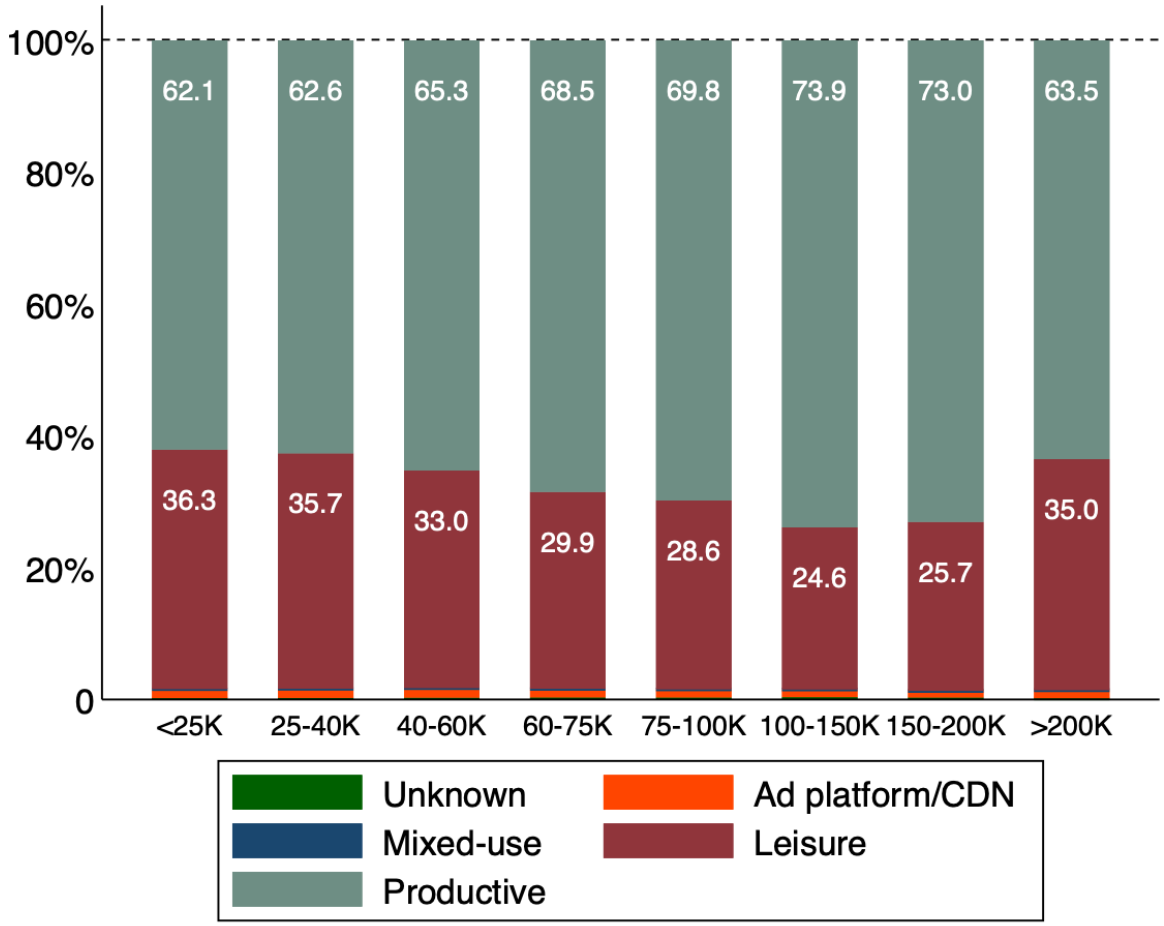
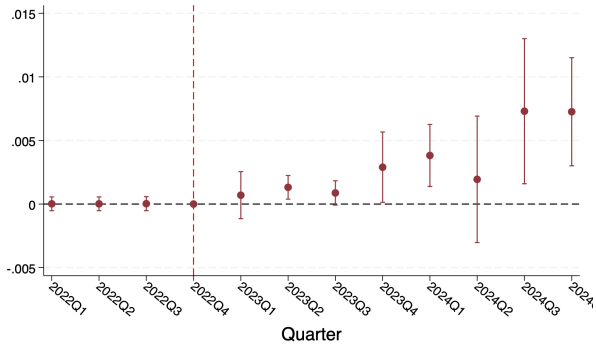
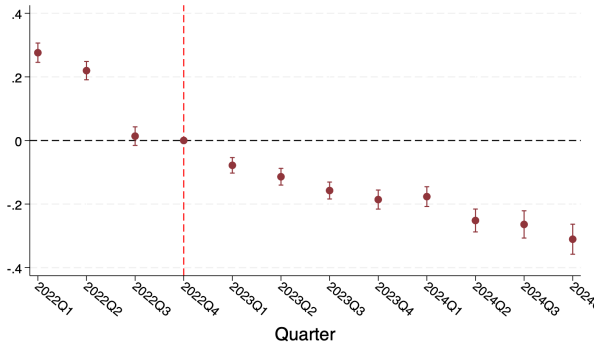


Figure 12: Reduced-form IV estimates of effect of ChatGPT use on amount and composition of browsing activity.

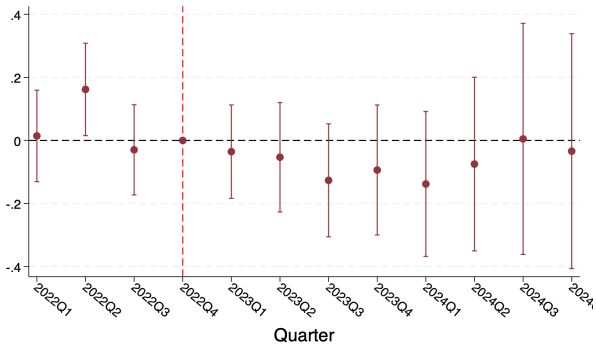
.



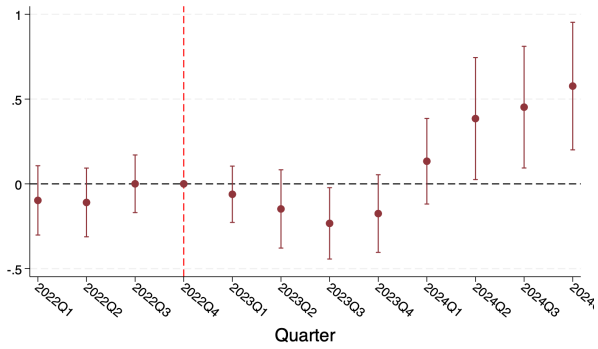
(A) Share of browsing duration in ChatGPT



(B) Share of browsing duration in ChatGPT-overlapping sites



(C) Log browsing duration, all sites



(D) Log browsing duration, leisure sites

Additional results

Appendix A. Methodology Appendix

Appendix A. Website classification

We sort websites in the public use version of CompScore data by total page views in 2023 and focus our labeling on the top 100K websites, which correspond to > 95% of all internet use.

The website classification pipeline then processes this domain-level data through multiple stages to create a dataset of websites with their usage classification—whether the website use is most likely "Productive", "Leisure", or "Mixed-use":

1. **Domain Selection**: We start with a ranked list of domains. Duplicates are removed to ensure each domain appears only once in the dataset.
2. **Metadata Extraction**: For each selected domain, we use the Beautiful Soup package in Python to access the domain and extract metadata from the website's source code, focusing on information contained within 'title', 'description' and 'keyword' tags provided by the website owner.
3. **Website description and activity extraction**: For each domain, we use the GPT-4.1 Mini model from OpenAI (accessed in the May 28, 2025 version of the model via Azure OpenAI) to generate:
    - A concise description (2-3 sentences) of the website's nature and primary function
    - A list of the top 5 meaningful user activities typically performed on the website
    - The full prompt used is the following:

```
I need a detailed analysis of the activities one would expect on the website with domain "{domain}".

Here is the available metadata about this website:
{available_data if metadata else "No additional metadata available."}

Based on this information, please:
1. Provide a concise description (2-3 sentences) of the nature of the website and its primary function.
2. List the top 5 meaningful activities that users would likely perform on this website.

Omit auxiliary activities, such as "set up an account with the website", "provide their payment details to pay for the website" etc.
that are not core to the purpose of the website.

Return your response in this exact JSON format:
{{
```

```
    "website_description": "Description of the website's nature and primary function",
    "top_activities": [
        "Activity 1",
        "Activity 2",
        "Activity 3",
        "Activity 4",
        "Activity 5"
    ]
}}

If the domain name alone doesn't provide enough information for a confident assessment, make reasonable inferences based on similar websites.
```

- Note that this approach relies on the website's metadata if it's available as grounding for the model's response, but if there is no further description also allows the model to rely on the knowledge obtained from its training data (which contains large amounts of internet text) in judging what the likely use of a website is.
- Websites triggering content management policy violation errors from the LLM (because of adult content) are classified as "adult content, likely of a sexual nature" if the LLM refuses to respond.

Usage Classification: productive vs. leisure use websites

Each website is classified along two dimensions using GPT-4.1 mini:

- **Usage Type**: Categorized as "Productive" (market work, education, childcare, etc.), "Leisure" (gaming, social activity, TV, etc.), or "Ad platform/CDN" (advertising networks, content delivery networks)
- **Usage Diversity**: Labeled as either "Single-use" (primarily productive OR leisure) or "Mixed-use" (has important productive AND leisure uses)

The prompt used in the classification is shown below:

```
I need to classify a website based on its domain name and description.

Website Domain: {domain}
Website Description: {website_description}

Please classify this website in two ways:

1. Usage Type: Determine if the website is primarily:
"Productive" (Market Work, Other Income-Generating Work, Job Search, Childcare, Non-market work, Education, Civic Activities, Own Medical Care) or
"Leisure" (TV, Social Activity, Sleep, Eating and Personal Care, Gaming, Other Leisure) or
"Ad platform/CDN" (Advertising networks, ad servers, analytics platforms, tracking services, content delivery networks etc. that primarily serve ads or video for other si

2. Usage Diversity: Determine if the website is
"Single-use" (only has important Productive OR Leisure uses) or
"Mixed-use" (has important Productive uses AND Leisure uses).

Consider:
- The website's primary purpose
- The typical user activities and motivations
- Whether users would engage in both leisure and productive uses on the website
- For ad platforms/CDNs, look for domains related to ad serving, tracking, analytics, content delivery infrastructure, or marketing infrastructure

Return your response in this exact JSON format:
{{
    "usage_type_reasoning": "Brief explanation of why this is classified as productive, leisure, or ad platform/CDN",
    "usage_type": "Productive" OR "Leisure" OR "Ad platform/CDN",
    "usage_diversity_reasoning": "Brief explanation of why this is single-use or mixed-use",
    "usage_diversity": "Single-use" OR "Mixed-use"
}}
```

To validate the robustness of this overall procedure, we also used a separate methodology that first converts the website activities resulting from the first LLM call into semantic embeddings. Then, each website activity is matched to similar activities from the American Time Use Survey (ATUS) that could conceivably be done on a computer, by also embedding the ATUS activities and computing cosine similarities. Next, an LLM is asked to determine which ATUS activity is the main one for this website by choosing one from the list of matches. Finally, this ATUS activity is mapped to the categories of activities that reflect productive or leisure uses based on on the categories of activities in Aguiar et al. (2013). This procedure resulted in labels of productive or leisure uses of websites that were very similar to the more direct procedure described above, but substantially reduced the coverage of domains because for many website activities no good ATUS activity matches could be found.

Appendix B. Website exposure classification

sec:exposure$_prompt$

Here is the prompt used for labeling website activities:

```
I need to classify a website based on its domain name and description to determine if it's likely to be exposed to competition from ChatGPT chatbot use.

Website Domain: {domain}
Website Description: {website_description}

Please classify this website based on whether it provides services or content that users might partially be able to substitute with ChatGPT use.

Here are the website capabilities that ChatGPT might replace:

### 1. Language & Translation Sites
- **Translation Platforms:** Websites offering language conversion services
- **Grammar & Writing Tools:** Sites providing spelling, grammar, and style checking
- **Language Learning Resources:** Educational platforms teaching languages
- **Multilingual Customer Service:** Sites whose primary value is language bridging

### 2. Content Creation Websites
- **Creative Writing Platforms:** Sites offering story, poem, or script generation
- **Business Document Services:** Websites specializing in professional document templates
- **Marketing Copy Generators:** Sites creating advertising text and slogans
- **Email Writing Assistants:** Platforms helping draft professional communications
- **Content Mills:** Websites producing basic informational articles

### 3. Information & Research Sites
- **General Knowledge Platforms:** Encyclopedia and fact-aggregation websites
- **Simple Q&A Forums:** Sites answering straightforward informational questions
- **Educational Summaries:** Websites offering simplified explanations of concepts
- **Product Comparison Sites:** Basic review aggregators without detailed testing
- **Academic Resource Centers:** Sites providing standard explanations of subjects

### 4. Basic Programming Help
- **Code Snippet Libraries:** Websites offering common programming solutions
- **Beginner Programming Tutorials:** Sites explaining fundamental coding concepts
- **Simple Debugging Platforms:** Services that identify basic code errors
- **Standard Documentation Sites:** Websites with common technical explanations

### 5. Data Analysis Services
- **Basic Visualization Tools:** Sites generating simple charts and graphs
- **Data Interpretation Services:** Platforms explaining statistical information
- **Simple Reporting Tools:** Websites creating standard data reports
- **Algorithm Explanation Sites:** Services describing common computational approaches

### 6. Technical Support Forums
- **Common Troubleshooting Platforms:** Sites offering solutions to frequent issues
- **Software Usage Guides:** Websites explaining how to use popular programs
- **IT Knowledge Bases:** Repositories of standard technical solutions
- **Command Reference Sites:** Platforms listing command structures and usage

### 7. Basic Professional Assistance
- **Resume Builders:** Websites generating standard resume formats
- **Cover Letter Templates:** Platforms offering professional letter frameworks
- **Business Plan Generators:** Sites creating basic business documentation
- **Simple Legal Document Services:** Platforms for standard legal forms

### 8. Advisory Content Sites
- **General Financial Advice:** Websites offering basic investment information
- **Career Guidance Platforms:** Sites providing common professional development tips
```

```
- **Business Strategy Blogs:** Publications sharing standard business insights
- **Health Information Resources:** General wellness and nutrition advice sites

### 9. Lifestyle Support Services
- **Relationship Advice Columns:** Sites offering general interpersonal guidance
- **Basic Trip Planning Tools:** Platforms generating standard travel itineraries
- **Fitness Program Generators:** Websites creating common exercise routines
- **Study Aid Resources:** Services offering standard learning assistance

### 10. Creative Planning Tools
- **Event Planning Templates:** Sites offering standard organizational frameworks
- **Project Management Guides:** Platforms providing basic structural approaches
- **Brainstorming Tools:** Websites facilitating ideation and creative processes
- **Organizational Systems:** Services creating standard management frameworks

### 11. Visual Content Services
- **Image Description Sites:** Platforms that caption or explain visual content
- **Basic Visual Analysis Tools:** Services identifying elements in images
- **Chart Interpretation Services:** Websites explaining graphs and visualizations
- **Document Reading Services:** Platforms extracting text from images

### 12. Audio Processing Sites
- **Transcription Services:** Websites converting speech to text
- **Basic Audio Analysis:** Platforms identifying elements in audio content
- **Podcast Summarization:** Services creating text summaries of audio content
- **Voice-to-Text Applications:** Sites offering speech recognition functionality

### 13. Multimedia Integration Platforms
- **Basic Media Conversion:** Services translating between different content formats
- **Simple Virtual Assistants:** Websites offering basic multimodal interactions
- **Content Accessibility Tools:** Platforms making content available across formats
- **Media Description Services:** Sites generating text descriptions of visual/audio content

### 14. Beginner Tutorial Sites
- **Introductory Learning Platforms:** Websites offering fundamental knowledge
- **How-To Guides:** Services providing step-by-step instructions
- **Basic Skill Development:** Sites teaching entry-level capabilities
- **Simplified Explanations:** Platforms reducing complex topics to basics

### 15. Design Template Services
- **Basic UI/UX Resources:** Websites offering standard design patterns
- **Template Libraries:** Platforms providing pre-made design frameworks
- **Simple Logo Generators:** Services creating basic brand elements
- **Design Guidelines:** Sites explaining fundamental aesthetic principles

Determine if ANY activities on the given website is exposed to ChatGPT competition based on the above categories. Note that exposure may come from only some activities be

Return your response in this exact JSON format:
{{
    "exposure_reasoning": "Brief explanation of why this website is or is not exposed to ChatGPT competition",
    "substitution_examples": "If exposed, provide 1-2 specific examples of how users might substitute this website's services with ChatGPT",
    "is_exposed": "Yes" OR "No",
    "exposure_level": "High" OR "Medium" OR "Low" OR "None",
}}
```

Output Data

The final dataset includes the following key fields for each domain:

- Domain name
- Website description
- Top user activities (up to 5)
- Usage type (Productive/Leisure/Ad platform)
- Usage type reasoning

- Usage diversity (Single-use/Mixed-use)
- Usage diversity reasoning

This processed dataset enables analysis of internet usage patterns, classification of time spent online, and economic analysis of productivity and leisure trade-offs in digital environments. “‘

Appendix B. Additional Tables

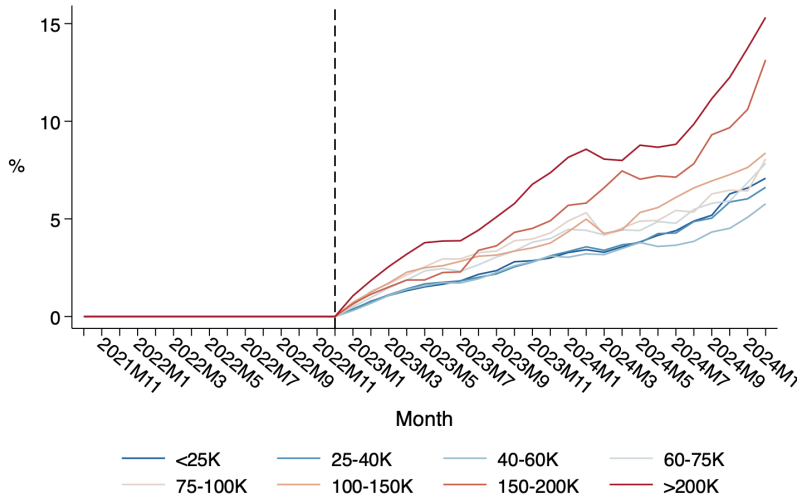
Table X: Sample Restriction Criteria

Exclusion Criteria	Number of Machines		Number of Machine - Months	
	Number	Percent	Number	Percent
None	332,260	100	5,254,024	100
- Machines for work	319,509	96	5,006,439	95
- Machines w/o demographic data	270,741	81	4,257,721	81
- Machines w/ < 12 months of data from 7/2021 to 6/2022	X	X	X	X

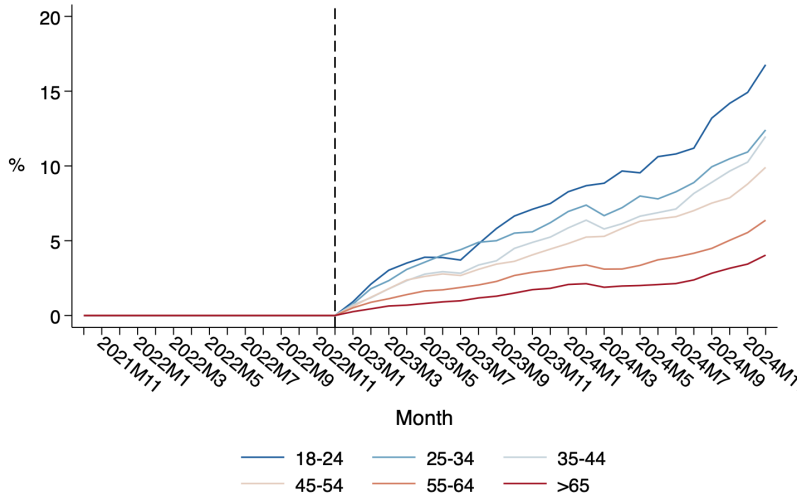
Appendix C. Additional Figures

Figure 13: ChatGPT regular adoption by household income and age

This figure shows the share of households in each demographic group that have become ChatGPT "adopter", i.e. have used ChatGPT in consecutive months.

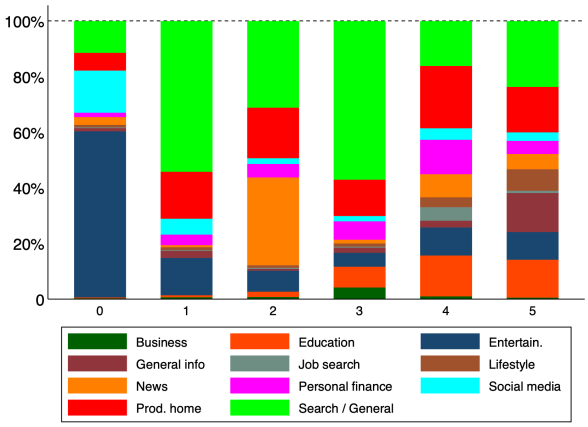


(A) ChatGPT adoption by income bucket

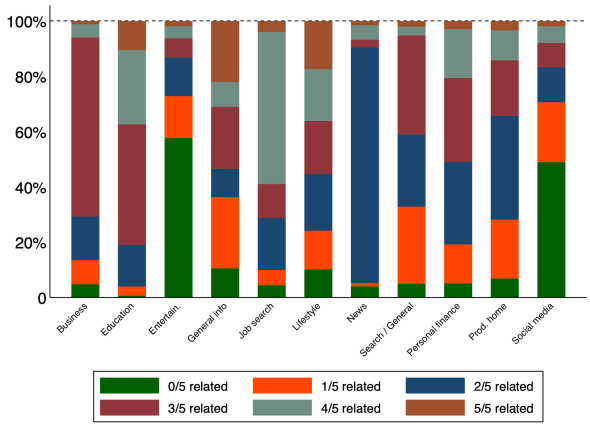


(B) ChatGPT adoption by age bucket

Figure 14: Duration-weighted double sorts of website overlap fractions and Comscore activity categories



(A) Duration shares of overlap fractions



(B) Duration shares of site activity categories

Figure 15: Estimated effect of placebo pre-ChatGPT adoption on GenAI-exposed browsing intensity

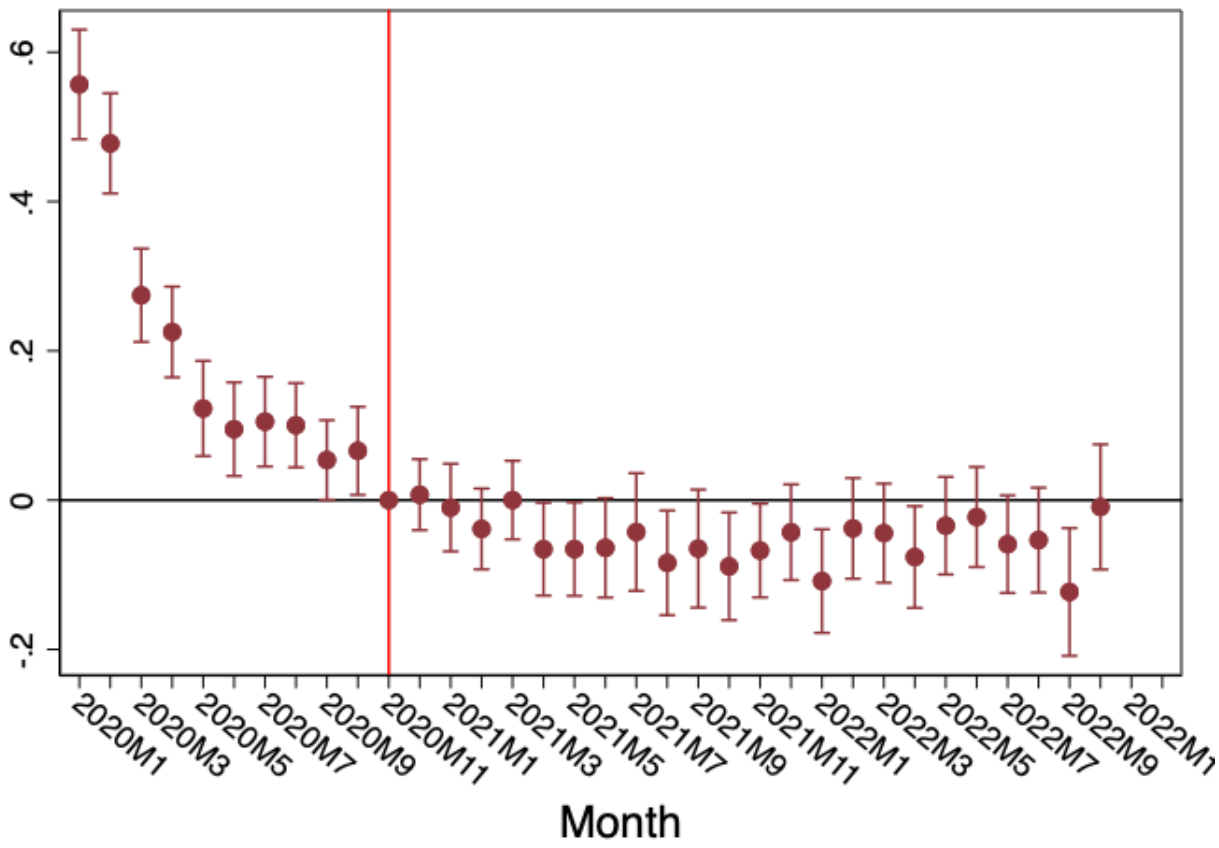


Figure 16: Share of households in the sample that have ever used ChatGPT, by whether they used Stack Overflow (or other coding-related subsites of Stack Exchange) or not.

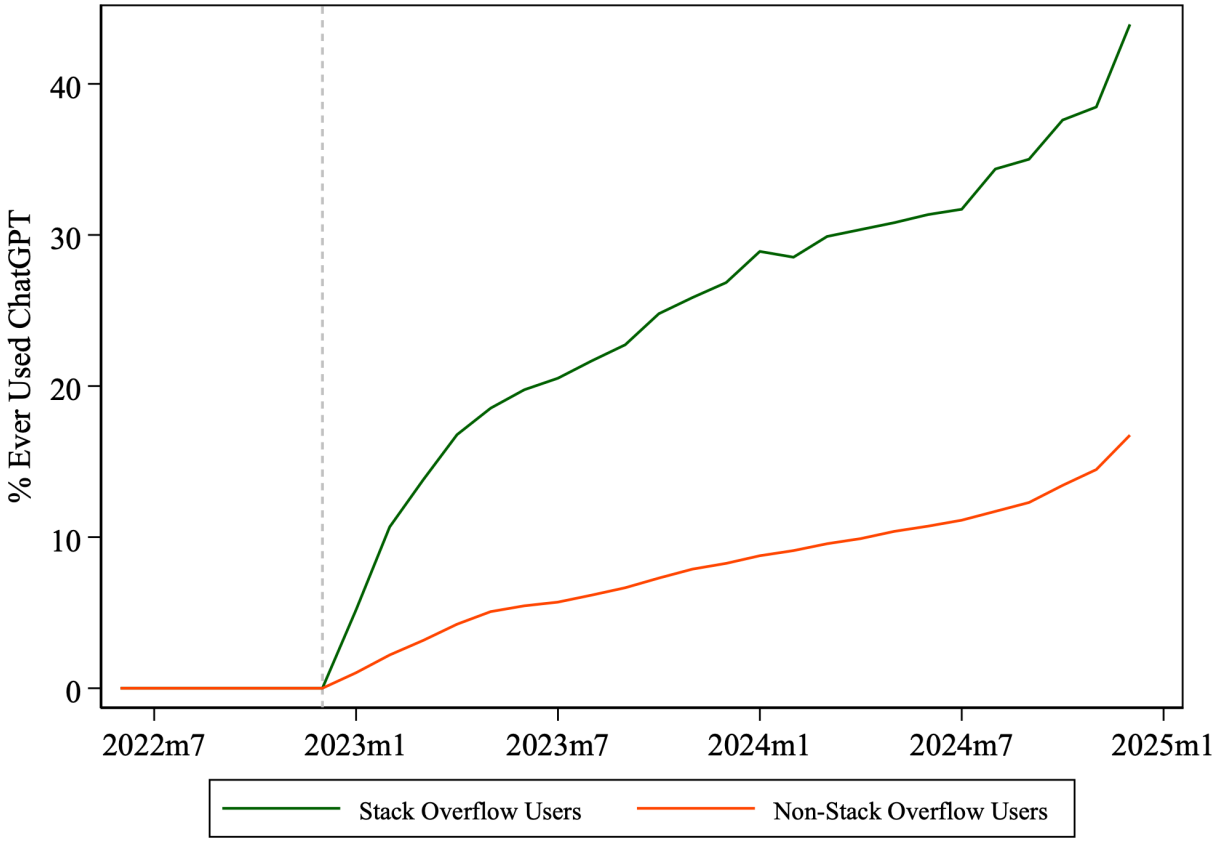
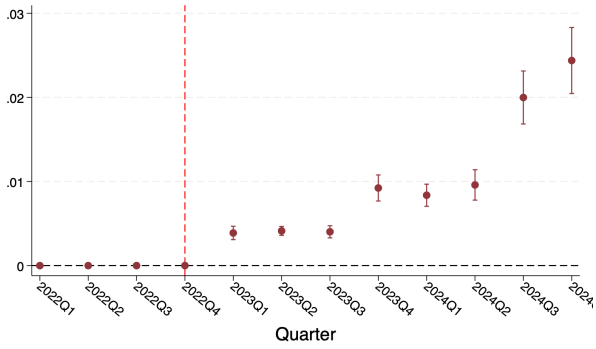


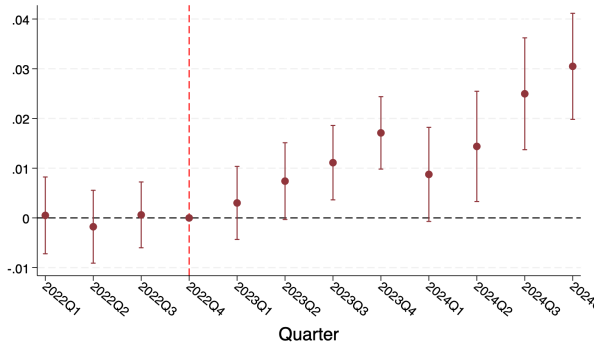
Figure 17

Figure 18: OLS estimates of effect of ChatGPT use on amount and composition of browsing activity.

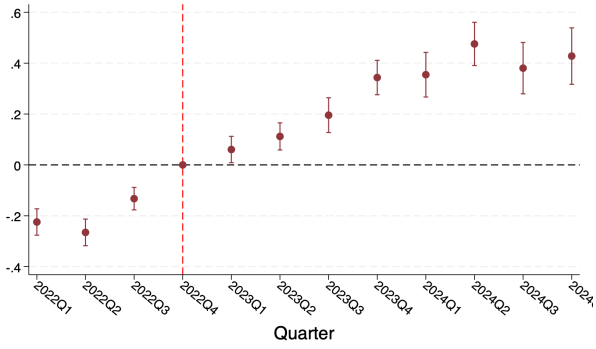
.



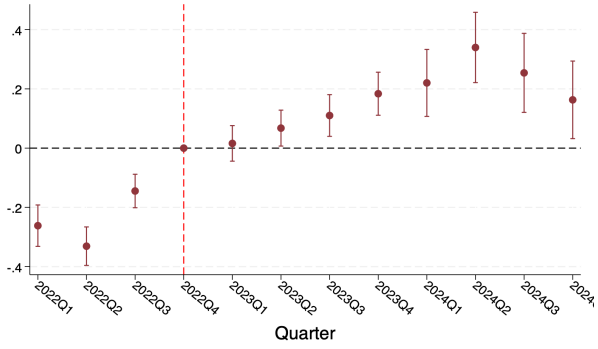
(A) Share of browsing duration in ChatGPT



(B) Share of browsing duration in ChatGPT-overlapping sites



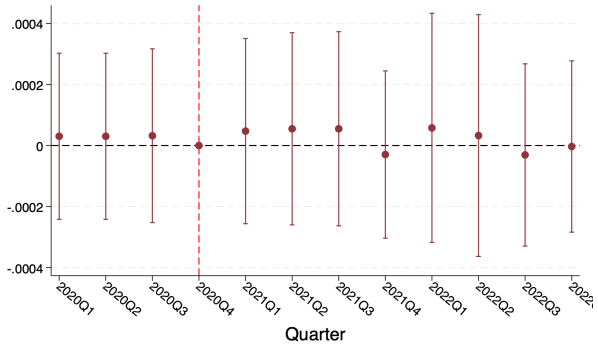
(C) Log browsing duration, all sites



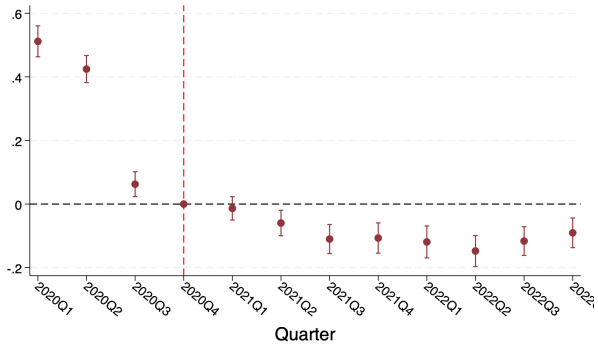
(D) Log browsing duration, leisure sites

Figure 19: OLS estimates of effect of ChatGPT use on amount and composition of browsing activity.

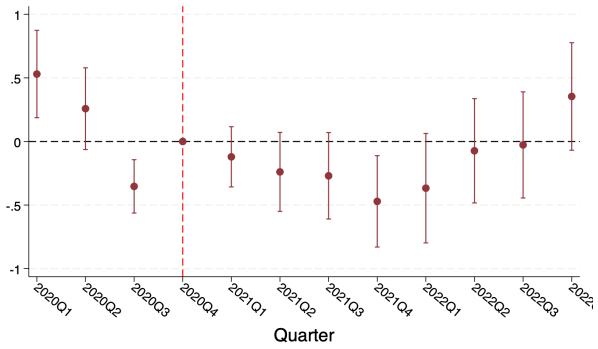
.



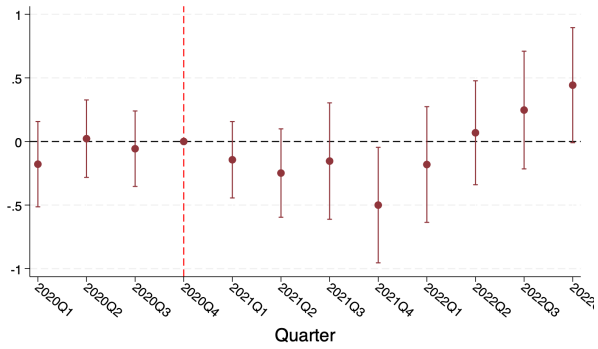
(A) Share of browsing duration in ChatGPT



(B) Share of browsing duration in ChatGPT-overlapping sites



(C) Log browsing duration, all sites



(D) Log browsing duration, leisure sites