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

Michael Blank† Gregor Schubert‡ Miao Ben Zhang§

November, 2025

Abstract

This paper studies the impact of generative AI on U.S. households using detailed Internet browsing data from over 200,000 households' home devices during 2021–2024. Our analyses of households' de facto adoption and usage of ChatGPT reveal several new findings. First, based on households' browsing of other websites during ChatGPT usage sessions, we find that households tend to use ChatGPT for productive non-market activities, such as education or job searches, rather than for leisure. Second, based on households' pre-ChatGPT browsing patterns, we show that adopting ChatGPT reduces households' overall productive activities and increases leisure activities on home devices. Together, these findings suggest that generative AI increases households' time spent on leisure by making productive activities more efficient. Finally, we find a substantial "generative AI divide" among households, as high-income and younger households adopt generative AI more than low-income and older households.

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

*We thank seminar participants at the OpenAI.
†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

As generative artificial intelligence (AI) is rapidly diffusing into everyday life, it raises important questions about how its benefits are distributed: Which households adopt generative AI, for which purposes, and, most importantly, with what consequences for their time use and economic activities? While recent efforts using surveys (e.g., Bick et al. (2024)) and laboratory settings (e.g., Kosmyna et al. (2025)) provide important early insights, they are limited by a lack of data that links households' usage of generative AI to their other activities.

In this paper, we address this gap by using detailed internet browsing microdata from over 200,000 U.S. households' home devices from 2021 to 2024 to provide the first large-scale evidence on the adoption and impact of generative AI on households. We focus on the household setting as a major margin of generative AI's impact, as the average American prime-age adult is more likely to use ChatGPT for personal browsing than for their job.[1]

Analyzing the household setting complements existing literature in several ways. First, it delivers revealed-preference evidence on what households actually do, rather than what they report in surveys or how they behave in experimental settings. Second, our focus on generative AI's impact "at home" complements existing empirical work on its impact "at work" (Brynjolfsson et al. (2023)). While these studies show substantial disruption of generative AI on jobs, including unequal productivity improvements and job displacement,[2] households' responses to these shocks—such as rebuilding human capital, or using leisure to release stress—arguably happen to a large degree at home. Third, many of the most transformative hypothesized impacts of generative AI lie *outside* the workplace: it may reshape 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 adoption: First, 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. Second, 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 pre-ChatGPT online

[1] Bick et al. (2024) find 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 employed respondents had used it at work. A Pew Research survey in February 2024 found that 23% of all adults had used ChatGPT, 17% had used it for entertainment, and 17% to learn something new, while only 12% of all adults (20% of those employed) had used it for a work task.

[2] See, for example, Brynjolfsson et al. (2023), Eisfeldt et al. (2023), Eisfeldt and Schubert (2024), Humlum and Vestergaard (2025), Chatterji et al. (2025), among others.

activities, and show that this ex-ante exposure and differential take-up rates in response to predicted benefits can explain part of the adoption gaps by income and age. Third, we characterize *how* generative AI is used by households and show that private ChatGPT adoption 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. Fourth, we quantify how households *change* their online activities in terms of productive and leisure 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 to social desirability bias (Ling and Imas, 2025). Similarly, microdata collected by generative AI companies, such as recent studies by Anthropic (Handa et al., 2025a) and OpenAI (Chatterji et al., 2025) provide a wealth of information about what people are doing inside the chatbots, but with little information about their broader behaviors and outcomes when not using a chatbot.

We overcome these challenges by utilizing a novel source of data for studying generative AI use by households: 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, which, importantly, includes activities before and after interacting with generative AI tools. (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.

Our first key finding establishes that potential benefits can explain part of the 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 webscraping 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 in July 2021-June 2022, i.e. *before* the 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 a higher share of its 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 1.0 pp higher rate of having tried ChatGPT by December 2023–a 12% increase.

While this link between household benefits and adoption of generative AI may not seem surprising, 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 2-8% of the gap in ever using ChatGPT between high- and low-income households, and 3-13% of the gap in use between young and old households. Moreover, we show that households with programming knowledge (proxied by having accessed Stack Overflow) have both greater potential benefits from using ChatGPT and respond more strongly to these benefits: our exposure measure explains 15-24% of the greater ChatGPT use among households with programming knowledge, compared to those without.

Having established which private households use generative AI, and that different exposure to potential benefits from the technology can explain adoption gaps, our main finding quantifies how ChatGPT adoption *changes* household behavior and welfare. Using LLMs to classify browsing activities, we categorize websites into those that are predominantly "leisure" sites and those that are associated with (non-market) "productive" activities, such as tasks related to building human capital (e.g., education and job search). With this data, we can shed light on which types of 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 other behaviors. To address this issue, we use an instrumented difference-in-differences approach, using plausibly exogenous variation in household exposure to generative AI well before the ChatGPT release as an instrument for changes in behavior after the post-ChatGPT adoption.

We find that households overall spend less time on productive online activities and instead spend more time on leisure activities. This finding suggests that using ChatGPT for household production tasks may substitute for time spent on productive activities on other websites, while making households more efficient in completing those tasks. The additional time spent on leisure activities provides evidence that ChatGPT adoption has positive welfare effects on households.

To provide support for this mechanism, we exploit the ability to analyze the browsing context in which ChatGPT is used. First, we establish that households predominantly utilize ChatGPT in the context of browsing activities that are productive. While our data does not allow us to directly observe what a household is doing inside ChatGPT, we can 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 prepare cover letters for their job search may also visit an online job board. This implies that we can plausibly deduce what activities a person is doing on average inside ChatGPT by observing their online activities before and after. By comparing household activities within narrow time windows around ChatGPT use and overall, we find that ChatGPT adopters systematically differ from non-users in their online activities—their online activies are more productive—and that this difference is even more pronounced for browsing sessions involving ChatGPT. This evidence supports the idea that ChatGPT makes productive online tasks more efficient, which allows households to reallocate freed-up browsing time to leisure. In line with this mechanism, we provide a conceptual framework to map our empirical estimates to changes in household welfare.

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: (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. As we document in this paper, ChatGPT impacts household demand for digital services provided on other websites. A number of studies that have found changes in browsing behavior after the release of ChatGPT: Lyu et al. (2025) study effects on different articles within Wikipedia and find that articles more similar to chatbot output experienced a greater decline in views after the release of ChatGPT. Padilla et al. (2025) show that Google searches for information decline after LLM adoption, while navigational searches do not. They also show that online ad exposures and visits to educational websites and Stack Overflow fall, while visits to social media sites and Wikipedia are not affected. The change in visits to Stack Overflow post-ChatGPT is also documented by del Rio-Chanona et al. (2024) and Burtch et al. (2024), who find large declines in engagement, suggesting that generative AI users are resolving their programming questions productively with the help of ChatGPT. Our study complements this literature by focusing on households' changes in their overall browsing behavior and documenting an increase in the leisure share of online browsing.

Our findings on the household non-market productivity effects of generative AI complement several studies that have documented the technology's productivity effects in different labor markets and in business settings: 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. Some studies document productivity effects in other non-workplace setting: 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. Jiang et al. (2025) show that the effects on leisure of using generative AI for *market* work may differ from the effects of ChatGPT in the private household setting that we study: they find that workers with greater generative AI exposure work longer hours and experience less leisure time.

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, and provide new evidence on what predicts such adoption gaps, and document important consequences in the form of differences in non-market work productivity.

Our findings suggest positive changes in household welfare as a result of ChatGPT adoption changing non-market work productivity and allowing households to spend more online time on leisure sites. Our conceptual framework for this household time reallocation in response to a shock to task productivity builds on the model in Aguiar et al. (2021), as well as empirical work by Brynjolfsson et al. (2025), who measured the value to consumers of free digital goods.

I. Data and Measurement

A. Comscore Internet browsing data

Our analysis uses the microdata for a large panel of U.S. households' internet browsing activity provided by Comscore. Comscore is a U.S. based, publicly-traded media measurement and analytics company that specializes in characterizing the online behavior and demographics of different websites' user bases. To do so, Comscore reaches out to a large number of active internet-using households in the U.S. who, in exchange for certain incentives, allow the company to collect comprehensive data on their browsing activities.[3]

[3] These incentives can entail free software, applications, or utilities, see Comscore Media Metrix Description of Methodology. Each participating household has an application installed on their computer which tracks internet usage. Comscore utilizes these data to produce estimates of online usage patterns as well as for other

The microdata provides detailed internet browsing activities of tens of thousands of households per year, where each household is identified by a unique identifier *machine_id*. To capture active meaningful human visits to a website, Comsore filters out short visits that are less than 3 seconds, times out sessions after 30 minutes of inactivity, and removes non-human visits using a multi-layered approach to detect sophisticated bots and fraud. For each qualified website visit, the data provides the timestamp, the *duration* of the visit in seconds, the URL of the website, a classification code of the URL, and a unique browsing session id that groups the website visit with other adjacent visits in a continuum of browsing period. The data also provides regularly-updated demographic information of the households, including the combined income of the household in 8 bins, the age of the household head in 11 bins, the geographic location in cities, and the number of members in the household.[4]

Sample. Our baseline sample includes household-level browsing data from 2021 to 2024, covering periods before and after the release of ChatGPT on November 30, 2022. We set the twelve months from July 2021 to June 2022 as the *benchmarking period*.[5] To be included in our baseline sample, we require the household to have non-missing information on income and age, have browsing activities for at least six months during the benchmarking period, and have browsing activities for at least one month after the release of ChatGPT. We focus on private household-owned computers by further removing a small set of machines that are labeled as work-owned. Internet Appendix Table IA.1 provides more details on the impact of each filter on the sample size.

Sample representativeness. While Comscore designed the recruiting process to capture a representative sample of internet users in the U.S., we note that maintaining a perfectly representative sample at a particular time, such as during the benchmarking period in our study, can be challenging. Table I compares the composition of our baseline sample with the

studies of online user behavior conducted internally or by their clients.

[4]See data manual at Comscore Web Behavior Database. Compared to the synthesized Comscore web behavior data available on WRDS and used in prior studies, our raw microdata has three advantages: First, the synthesized data includes a curtailed set of about 50,000 households each year, while our raw data includes all households. Second, the synthesized data groups multiple adjacent website visits into a browsing session and provides only a main domain of the browsing session, while our raw data shows all website URLs clicked within the browsing session. Third, our raw data timely reflects any update on households' demographics, such as income, while the the synthesized data does not.

[5]We define the benchmarking period as ending in June 2022 to ensure households' browsing behavior is not contaminated by the test-running of ChatGPT or possible expectation of ChatGPT release. This also prevents short-run household shocks from affecting both the benchmarking period and the post-ChatGPT browsing activity.

U.S. internet-using population from the 2022 American Community Survey (ACS) database. We observe that our sample tends to over-represent households at the low (<$60K) and high ends (>$200K) of the income distribution, and under-represent middle-income ($60K-200K) households. In terms of age, our sample somewhat over-represents the upper middle-age households (45-54 years old) and under-represents younger households (25-44 years old).

To account for these sample discrepancies, we construct a weight factor for each household in our baseline sample so that the joint income and age distribution of our sample households in November 2022 matches that of the 2022 ACS, and we present all statistics using this weight. In addition, To avoid our measures of household demographics being endogenously influenced by household adoption of ChatGPT, we keep households' demographics fixed as of November 2022 in most of our analysis unless otherwise noted.

B. Measuring household browsing behavior

Here, we introduce the main outcome measures of interest that we examine in this study. We begin by introducing two underlying measures of browsing intensity at the website level. Our data includes millions of websites. To draw economic inferences, we aggregate the intensity measures to broader category level in certain analyses, and we further develop our own categorization of websites into productive versus leisure groups for other main analyses.

Browsing intensity measures. Most of our analyses focuses on households' Internet browsing behavior at the monthly frequency. We focus on two measures of a household's browsing intensity of a website in a month. The first is the total number of seconds household j spends on website j during month t, labeled as $duration_{j,i,t}$. This measure captures both intensive and extensive margins of a household's visit to the website. Our second measure of browsing intensity counts the number of times a household visits the website during the month, labeled as $visit_{j,i,t}$. In some illustrations, we aggregate the browsing intensity measures of websites to broad website categories provided by Comscore.

Categorizing productive vs. leisure websites. A key focus of our study is on whether the browsing is for productive tasks or leisure, as motivated by the standard home production framework that we discuss later. To the best of our knowledge there is no existing approach for labeling websites based on the household activity that they enable, so we develop a novel methodology for classifying domains using large language model (LLM). We focused on the top 100K domains by visits between 2022 and 2024, which account for nearly all browsing activity during this period.

The methodology, as illustrated in Figure 1 and detailed in the Internet Appendix A, involves the following three steps: First, for each domain, we obtain a description of each website by using a webscraper to try and access each domain address to retrieve the description and keywords that the website's owner has embedded in "meta tags" in the website's source code. Second, we submit a prompt to an LLM (OpenAI's GPT-4.1 mini model accessed through the Azure API) that asks the LLM to consider the domain, title of the site, and the description from the website's source code and come up with the 5 "main activities" that the website could be used for households.

Third, we submit the website descriptions and list of 5 key activities 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.[6] Our definition of productive and leisure activities closely follows the categorization of general (non-internet) activities in Aguiar et al. (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*. Based on this categorization, we can now measure each household's browsing intensity in productive versus leisure websites within the month, which is the key focus of our study on economic outcomes.

[6] 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.

Figure 1:
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.

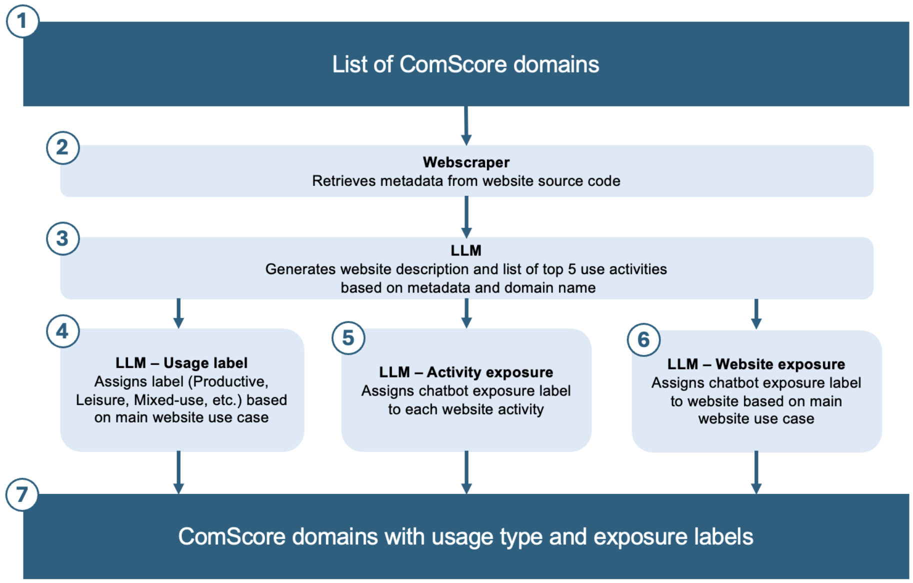

C. *Measuring household exposure to GenAI*

To evaluate the degree to which a household's internet use patterns is likely impacted by the release of ChatGPT, we adapt a methodology from prior work in labor market exposure measures, such as Eisfeldt et al. (2023), to the context of Internet browsing behavior. Specifically, we focus on the overlap between the activities enabled by websites and the functionalities of ChatGPT. First, 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.[7]

we assign each activity a binary exposure label based on whether ChatGPT can perform or complement the website's tasks. We then calculate a continuous exposure score by measuring the share of activities on each website that can be substituted by ChatGPT, with these scores ranging from 0 (no exposure) to 5 (high exposure). Table III lists the websites

[7]We develop this rubric by summarizing the clusters of user activities observed by Handa et al. (2025b) 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. See the announcement at: https://openai.com/index/chatgpt-can-now-see-hear-and-speak/

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.

The exposure label for each household is then derived by aggregating the exposure of the domains they visit, weighted by the time spent on each site. We show the distribution of these labels in Figure IA.4: 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)[8] 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.[9]

The key intuition behind this approach is that households whose browsing activities are more aligned with tasks that generative AI can assist with are more likely to benefit from or adopt such technologies. We compute household-level exposure by summing the activity-based exposure labels for each website visited by the household, adjusting for the proportion of total browsing time spent on each site. This method assumes that a higher share of activities on a website that generative AI can perform increases the likelihood that the household will find the technology beneficial. To ensure the accuracy of this exposure measure, we use browsing data from the July 2021 to June 2022 period, before the public release of ChatGPT, to avoid any endogeneity issues from early adoption effects.

Specifically, we define a household's exposure to generative AI as the expected share

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

[9] Panel A of Appendix Figure IA.3 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. Panel B of Appendix Figure IA.3 shows that 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.

share of highly generative AI-exposed websites in its internet browsing. That is, we compute

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

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. Here, we focus on websites with high exposure (4 or 5 out of 5 activities), as switching costs are likely to lead households to require a minimum level of substitutability (and implied productivity benefits) to want to use ChatGPT instead of its previous websites of choice. Empirically, this measure allows us to examine the relationship between pre-existing household exposure to generative AI and subsequent adoption of ChatGPT.

D. Measuring GenAI adoption

We focus on ChatGPT, the most popular LLM chatbot, as a proxy for generative AI adoption by households for several reasons. First, ChatGPT's public release on November 30, 2022, marked a breakthrough that significantly enhanced the accessibility and capabilities of generative AI, as seen in widespread media attention and immediate market reactions (Eisfeldt et al., 2023). This release provided a sharp break in households' ability to access sophisticated GenAI tools, making it a key event in our analysis of changes in browsing behavior before and after November 2022. Second, surveys such as Bick et al. (2024) and our browsing data confirm that ChatGPT adoption spread rapidly, particularly outside the workplace, making it easy to identify its use consistently in the Comscore data. Lastly, while other GenAI tools like Anthropic's Claude and Google's Bard emerged after ChatGPT, we focus on ChatGPT for simplicity, as it remains the most widely used tool among our sample, with most panelists who use alternatives first trying ChatGPT.

We construct two measures of households' adoption of ChatGPT. The first one captures 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 one is the intensity of the household's usage of ChatGPT, which is the number of seconds the household browsed the ChatGPT website within the month.

II. Household Exposure and Generative AI Adoption

In this section, we first establish that ChatGPT adoption patterns over time in our internet browsing data are consistent with evidence from other data sources, and vary substantially with household demographics. Then, we show that a substantial share of this heterogeneity in household adoption can be predicted by differences in the share of substitutable website activities in pre-ChatGPT browsing behavior.

A. Household adoption of ChatGPT over time

In this section, we document the extent and timing of generative AI adoption by households based on their browsing activity. Previous studies show a significant share of households use ChatGPT outside of work: Bick et al. (2024) finds that 35.5% of respondents used generative AI outside work in the second half of 2024, with 18% using ChatGPT in the past week, while Pew Research reported 18% of U.S. adults had used it in March 2023 (Pew Research Center, 2025). Our approach offers new evidence on adoption speed by tracking actual ChatGPT usage in household browsing, distinguishing between occasional and regular use, and linking usage patterns to other browsing behaviors and demographics. However, a limitation of our data is that it only captures use via websites on home computers, likely underestimating total generative AI usage.

Household ChatGPT adoption over time. We plot the time series of ever using ChatGPT and becoming a regular adopter (as define above) in Figure 8: the share of all households that has ever used ChatGPT rises from zero to 7.6 pp by December 2023 and to 14.0 pp by December 2024, almost doubling within a year. Regular use shows a similarly steep increase but at a lower level: it rises from 3.0 pp in December 2023 to 6.3 pp by December 2024 (see Table II 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 adoption 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 adoption by income and age. Our data also allows us to compare generative AI use across different demographic groups. Figure 3 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 IA.1 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 3 shows that age strongly predicts ChatGPT adoption 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 II , by Q4 2024, 19.6% of households headed by 18-24 year-olds have used ChatGPT, but only 9.2% of households at retirement age (65 years or above).

The generative AI divide. The gap in the use of generative AI documented in Figure 3 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.

B. 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 5 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 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 adoption, we estimate regressions of the form

$$\text{ChatGPT Use}_{j,t} = \text{HHGenAIExp}_j + \text{Char}_{j,t} + \text{FE}_t + \text{FE}_j, \tag{2}$$

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 adoptionr (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 Q4 2023, Q2 2024, and Q4 2024 (12, 18, and 24 months after the ChatGPT release).[10] 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 predicted change in the usage measure with regard to a change in the share of a household's browsing activity that is highly exposed to generative AI from zero (no exposure) to 100% (fully exposed). That is, the estimates in column (1) suggest that households with a one SD higher exposure (a 15 pp difference) saw 1.0 pp higher rates of having tried ChatGPT by December 2023, and a 1.4 pp higher rate of ever using it by December 2024. Regular use rates are, by definition, smaller than rates of ever having tried, 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 0.6 pp to 0.9 pp between December 2023 and 2024.

These effects of greater ex ante benefits of generative AI on its use are economically large: a 1 SD change in exposure is associated with an increase in the rate of having tried ChatGPT relative to the average of 12%, 13% and 10% in Dec. 2023, June 2024, and Dec. 2024. The corresponding effect on regular use is an increase of 18%, 17%, and 13% 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 explore the effects on the overall online activity that involves ChatGPT: We find that a one SD increase in exposure is associated, for instance, with a 0.03 pp (18%) increase in total share of browsing time spent on ChatGPT, and 0.1 extra visits (19% more)

[10]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.

per quarter to the ChatGPT site for the average household in Q2 2024. Note that these estimates become less precise by Q4 2024, as the pre-ChatGPT browsing behavior defining exposure becomes more distant in time from the technology use that we are trying to predict.

C. Household exposure and heterogeneous ChatGPT adoption

Can the effect of pre-existing exposure to generative AI benefits help us understand the heterogeneity in take-up of generative AI technologies like ChatGPT across demographic groups?

Adaptability drivers. 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 adoption can be seen from panel A of Table II: high-income households have 1.1 pp (10%) higher exposure to generative AI in their browsing behavior before the release of ChatGPT than low-income households. Similarly, the young households have 1.9 pp (18%) higher exposure than old households. As we found in Table IV that exposure significantly predicts ChatGPT adoption, 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 adoption between groups, we estimate differential take-up rates using specifications of the form

$$\text{ChatGPT Use}_{j,t} = \sum_k \beta^k \, \text{Group}_k \times \text{HH GenAI Exposure}_j + X'_{j,t}\phi + \varepsilon_{j,t}, \tag{3}$$

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

Table IA.2, panels A and B show differences in exposure effects by income group and by age group as of Q2 2024 (see Appendix Table IA.2 for estimates in other periods). Perhaps surprisingly, we find in panel A that there are no significant differences in the effect of exposure to generative AI benefits on ChatGPT adoption 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 IA.2, 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 50% smaller, and effects on regular use are 69% smaller in column (2). This means that–at least for age gaps in ChatGPT adoption–*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 adoption gaps. More formally, we can quantify the drivers of differences in ChatGPT adoption 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 Tables 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, IA.2, and IA.2 and the differences in exposure between groups shown in Table II to quantify the share of the gaps in ChatGPT adoption that is explained by the different components (or the residual). This decomposition is shown in Table VI. 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 adoption, this number will be zero, while 100% would mean that ChatGPT adoption 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 adoption gaps by income & age. In panel A of Table VI, we find that a modest share of 2-8% 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 4-6% of the gap in becoming a regular user. Most of this explanatory power in earlier periods comes from higher-income households having modestly higher take-up rates than lower-income households, while the latent exposure to benefits gains in importance for predicting ever using and regularly using ChatGPT by Q4 2024.

Panel B does the same analysis for differences in ChatGPT adoption by age group. Once we combine differences in adaptability with differences in exposure between young and old households, we can explain 3-13% of the gap in ever having used ChatGPT and 15-24% of the regular usage gap. The predictive power for the regular use gaps in particular is driven mainly by the much greater take-up rates among young households than for older households with similar exposure.

III. The Household Impact of Generative AI Adoption

A. Objects of interest

The previous sections established that internet browsing behavior can help us understand the time pattern of generative AI use across different demographic groups, and that differences in the potential benefit from generative AI based on pre-ChatGPT browsing patterns can drive differences in ChatGPT adoption post-2022. In this section, we estimate the *effect* that ChatGPT adoption has on households' overall browsing behavior: which types of online tasks were substituted by ChatGPT and which types of tasks were crowded in.

We are particularly interested in estimating the impact that adopting ChatGPT has on the quantity of digital tasks whose productivity ChatGPT plausibly increases, such as tasks related to acquiring information or learning new skills, relative to the impact on the quantity of digital consumption, such as streaming videos for entertainment. In theory, adopting ChatGPT could lead households to tilt the composition of their digital activities either away from or towards the productive websites that ChatGPT displaces. Intuitively–and as we formally model in the following section–if the output that households attain via these websites are "necessary goods," then adopting a tool that allows them to be completed more efficiently may lead a household to spend *less* time on them. To the extent that websites that provide households with leisure are "luxury goods," adopting ChatGPT may lead a household to spend more of their home computing time consuming these goods.

Generic identification challenges 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.

B. *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

$$\text{Browsing outcome}_{j,t} = \gamma^{OLS}\ \text{Household GenAI Adoption}_j \times \text{Post-ChatGPT}_t + \lambda_j + \text{Char}_{j,t} + \varepsilon_{j,t},$$

where the left-hand side is some browsing outcome of household j in quarter t, *Household GenAI Adoption*$_j$ is an indicator variable that equals one if the household ever uses ChatGPT following its November 2022 release, and Post-ChatGPT$_t$ is a dummy variable that equals one starting in 2023Q1.[11]

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 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

[11]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.

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.[12] 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 adoption.

These concerns with the naive OLS approach appear to be reflected in the actual estimates shown in Appendix Figure IA.8. 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.[13] 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 IA.8.

C. *Instrumental variables estimates*

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

[12] 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}.

[13] The pre-trends evident in Figure IA.8 suggest that in addition, ChatGPT adoption is still correlated with ex ante household characteristics that affect future browsing behavior, even after controlling for observable demographic characteristics as we do.

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-ChatGPT 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 by using the household-level GenAI exposure measure defined in equation (1) computed over July 2021 to June 2022. With this instrument, we then estimate the two-stage system in which the second-stage equation is given by

$$\text{Browsing outcome}_{j,t} = \gamma^{SS}\,\widehat{\text{Adoption}}_{j,t} \times \text{Post-ChatGPT}_t + \lambda_j + \text{Char}_{j,t} + \varepsilon_{j,t},$$

and the first-stage equation is given by

$$\text{Adoption}_{j,t} = \gamma^{FS}\,\text{HHGenAIExposure}_j \times \text{Post-ChatGPT}_t + \lambda_j + \text{Char}_{j,t} + \eta_{j,t}.$$

The reduced form regression is given by:

$$\text{Browsing outcome}_{j,t} = \gamma^{RF}\,\text{HHGenAIExposure} \times \text{Post-ChatGPT}_t + \lambda_j + \text{Char}_{j,t} + \varepsilon_{j,t}.$$

Main finding: browsing effects of generative AI exposure Figure 6 shows the dynamic IV estimates for the reduced form equation. Panel (A) shows that households substantially shift their browsing duration *towards* leisure-oriented websites. This is especially the case starting in 2024, which is consistent with the J-curve of technology adoption. In sharp comparison, Panel (B) shows that exposure to generative AI does not lead to similar increase in productive-oriented websites.

Table VII presents our 2SLS regressions when the estimates are pooled over all quarters. As indicated by the dynamic estimates in Figure 6, while the effect of instrumented take-up of ChatGPT on the overall amount of browsing is not statistically significant, the amount of browsing activity devoted 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-oriented sites. Table IX provides insight into exactly which browsing activities are crowded in vs. crowded out by ChatGPT adoption. 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 adoption crowds out the activities that–based on the high-frequency statistics of Figure IA.9–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.

Taken together, these results show that adopting ChatGPT causes households to reallocate their online time away from productive websites and toward leisure websites, with little change in overall browsing time. This pattern is suggestive of a household-production mechanism in which generative AI raises the efficiency of completing necessary 'productive' digital tasks, freeing up time for leisure consumption. To interpret our coefficients in terms of changes in preferences and welfare, the next section develops a simple two-activity model of household time allocation at home.

IV. Conceptual Framework

In this section, we develop a simple conceptual framework that adapts the model of Aguiar et al. (2021) to understand household choices of online activity. This framework allows us to map empirical measures of changes in households' time spent on productive and leisure activities to the impact of generative AI as a productivity-enhancing tool.

A. Setup

Household utility and time allocation problem We consider a household i who, at each time t, spends a fixed amount of time (normalized to one) engaged in digital activities at home. The household allocates an amount of time ℓ_{it} to *digital leisure consumption* (e.g., streaming movies, scrolling social media) and z_{it} to *digital home-production* tasks (paying bills, searching for information, planning, and other chores that could be sped up by ChatGPT), so the period-by-period time budget is

$$z_{it} + \ell_{it} \leq 1 \tag{4}$$

The household optimally allocates time between these two digital activities to maximize flow utility, given by the power utility function

$$v_{it} = \frac{(\theta^{\ell}_{it}\xi^{\ell}_{it}\ell_{it})^{1-(1/\eta^{\ell})}}{1-(1/\eta^{\ell})} + \frac{(\theta^{z}_{it}\xi^{z}_{it}z_{it})^{1-(1/\eta^{z})}}{1-(1/\eta^{z})}, \tag{5}$$

where θ^{ℓ}_{it} and θ^{z}_{it} are preference weights that represent a household's "taste" for different digital activities; ξ^{ℓ}_{it} and ξ^{z}_{it} capture the productivity of that activity, i.e. how much "utility output" time spent on that activity generates; and η^{ℓ} and η^{z} are curvature parameters that govern the diminishing marginal utility with regard to time spent on the different activities: when $\eta^{a} > 0$, activity a behaves like a "time luxury" (analogous to Aguiar et al. (2021) "leisure luxuries") in that the household spends more time on the activity as it gets more productive, while $\eta^{a} > 0$ suggests a "time necessity" which, once it can be done more productively, leads the household to shift time to other activities. Our assumption of power utility is not necessary for the insights discussed in this section and is adopted mainly for expositional convenience.

First order conditions This optimization problem yields the first-order condition for z_{it} (after taking logs of both sides) of

$$\ln(z_{it}) = (\eta^{z} - 1) \times \left(\ln(\theta^{z}_{it}) + \ln(\xi^{z}_{it})\right) - \eta^{z}\ln(\omega_{it}) \tag{6}$$

where ω_{it} is the shadow price of the time constraint. Similarly, the first-order condition for ℓ_{it} is

$$\ln(\ell_{it}) = (\eta^{\ell} - 1) \times \left(\ln(\theta^{\ell}_{it}) + \ln(\xi^{\ell}_{it})\right) - \eta^{\ell}\ln(\omega_{it}) \tag{7}$$

These equations show that households allocate more time to an activity when its taste/productivity parameter is high and less when the shadow cost of time is high.

These two conditions can also be combined to write the difference in time allocated between the two activities as

$$\frac{\ln(z_{it})}{\eta^{z}} - \frac{\ln(\ell_{it})}{\eta^{\ell}} = \frac{\eta^{z} - 1}{\eta^{z}} \times \left(\ln(\theta^{z}_{it}) + \ln(\xi^{z}_{it})\right) - \frac{\eta^{\ell} - 1}{\eta^{\ell}} \times \left(\ln(\theta^{\ell}_{it}) + \ln(\xi^{\ell}_{it})\right) \tag{8}$$

Here, for example, if $\eta_z < 1$, productive digital tasks are a "time necessity" and their relative share of digital time declines when the effective value of z_{it} rises as a result of productivity increases.

Release of ChatGPT We now consider how, within this basic framework, the release of ChatGPT affects the household's time allocation constraints and decisions. We model ChatGPT's release as being a task-biased technological shock to the household's digital home activities. In particular, we analyze the impact of a potential shock to the efficiency of the household's completion of digital production tasks, shifting ξ^z while leaving ξ^ℓ, θ^z and θ^ℓ unchanged.

For simplicity, assume that there are just two periods, $t = pre$ and $t = post$, and that ChatGPT is released in between (i.e., after $t = pre$ and before $t = post$). At the start of $t = post$, the household chooses whether or not to pay some cost to adopt ChatGPT. Adopting ChatGPT increases the efficiency of the household's digital home production, as captured by the shifter ξ^z_{it}. In particular, assume that this object's time pre value equals

$$\xi^z_{i,pre} = \underline{\xi} \tag{9}$$

while its time $post$ value is

$$\xi^z_{i,post} = \underline{\xi} \times \left(1 + \delta \cdot \text{Adopt}_{i,post}\right) \tag{10}$$

for $\text{Adopt}_{i,post} \in \{0,1\}$ an indicator for whether the household chooses to adopt ChatGPT and for $\delta > 0$ the magnitude by which ChatGPT improves efficiency on the productive digital tasks. Note that if the household does not adopt ChatGPT, then their efficiency on these tasks does not change across the two periods.

The household chooses to adopt ChatGPT if

$$v^{\text{Adopt}}_{i,post} - v^{\text{No adopt}}_{i,post} \geq c \tag{11}$$

for $v^{\text{Adopt}}_{i,post}$ and $v^{\text{No adopt}}_{i,post}$ the indirect utility functions if, respectively, the household adopts or does not adopt. We use superscripts 'Adopt' and 'No adopt' to denote the optimal choices under the two technology regimes, holding fixed all other shocks in period $t = \text{post}$. Here c is measured in units of period utility, so the left-hand side is the utility gain from faster productive tasks and the right-hand side is the utility cost of learning and adopting ChatGPT.

B. *Objects of interest: intuition in stylized model*

Objects of interest I: determinants of ChatGPT adoption For a sufficiently small percentage impact of ChatGPT adoption on efficiency δ, and a small (utility) cost of adoption c, the increase in the household's utility from adopting ChatGPT can be approximated as

$$\delta \cdot \left(\theta^{z}_{i,post} \times \underline{\xi} \times z^{\text{No adopt}}_{i,post}\right) \cdot \left(\theta^{z}_{i,post} \times \underline{\xi} \times z^{\text{No adopt}}_{i,post}\right)^{\frac{-1}{\eta^{z}}} \tag{12}$$

meaning that the households adopt ChatGPT if and only if

$$\delta \cdot \left(\theta^{z}_{i,post} \times \underline{\xi} \times z^{\text{No adopt}}_{i,post}\right)^{\frac{\eta^{z}-1}{\eta^{z}}} \geq c \tag{13}$$

Note that, given the FOC (6), the second term on the left-hand side equals $z^{\text{No adopt}}_{i,post} \cdot \omega^{\text{No adopt}}_{i,post}$, where $\omega^{\text{No adopt}}_{i,post}$ is the household's shadow price of home computing time absent ChatGPT adoption. Thus, according to equation (13), a household adopts ChatGPT as long as the total benefit from doing so (the left-hand side)–which is proportional to the overall value of the productive digital tasks that the household would have done even absent ChatGPT–is larger than the cost.

We can express this condition in a more immediately intuitive way by assuming the cost c arises from a required amount of time that a household must spend to successfully adopt the new technology, which we will denote c^{time}.[14] In utility terms, the total cost that household incurs is then $c^{\text{time}} \cdot \omega^{\text{No adopt}}_{i,post}$. Here, multiplying by the shadow value of time converts this time cost into utility units, making it directly comparable to the benefits of adoption. Plugging this into equation (13), and rewriting the left-hand side as $\delta \cdot z^{\text{No adopt}}_{i,post} \cdot \omega^{\text{No adopt}}_{i,post}$ based on the above discussion, a household adopts ChatGPT if and only if

$$\delta \cdot z^{\text{No adopt}}_{i,post} \geq c^{\text{time}} \tag{14}$$

This implies that the household's adoption decision, to first order, boils down to comparing how much time they would have spent on productive tasks absent ChatGPT versus how much time it will take to learn how to use ChatGPT for these tasks. Intuitively, the benefit of increased productive task efficiency, as well as the cost of foregone time on useful home

[14] This specification for the cost can be easily accommodated into the stylized model of this section by augmenting the budget constraint over time as follows:

$$z_{i,post} + \ell_{i,post} + c^{\text{time}} \cdot \text{Adopt}_{i,post} \leq 1$$

computing tasks, scale with the marginal value that the household places on additional time for home computing. As a result, this marginal value cancels out from the two sides of the condition, leaving the adoption decision as a function of a single household-specific object: the intensity of productive digital activities (high in which the household plans to engage. This means that households with a high baseline level of productive digital tasks (high $z_{i,post}^{\text{No adopt}}$) are more likely to adopt ChatGPT, which maps onto our empirical measure of household exposure based on a high share of pre-existing digital activities that are made more productive by generative AI.

Objects of interest II: causal effects of ChatGPT adoption Denote the treatment effect of adopting ChatGPT on the household's time spent on productive digital tasks as $\beta_z^{GPT} \equiv \ln(z_{i,post}^{\text{Adopt}}) - \ln(z_{i,post}^{\text{No adopt}})$, where $z_{i,post}^{\text{Adopt}}$ and $z_{i,post}^{\text{No adopt}}$ are the household's optimal time allocated to productive digital tasks under, respectively, the decision to adopt or not adopt ChatGPT. Denote $\beta_\ell^{GPT} \equiv \ln(\ell_{i,post}^{\text{Adopt}}) - \ln(\ell_{i,post}^{\text{No adopt}})$ as the analogous object for time spent on leisure tasks.

By plugging the time constraint (4) into the FOC (8) for the household's allocation of time across the two types of digital activities, we can characterize how the household's relative time allocation changes due to the adoption of ChatGPT. To first order (i.e., for δ close to zero), the relative impact on time allocated to productive tasks versus leisure tasks is

$$\frac{\beta_z^{GPT}}{\eta^z} - \frac{\beta_\ell^{GPT}}{\eta^\ell} \approx \frac{\eta^z - 1}{\eta^z} \times \delta \tag{15}$$

while the impact on time allocated to productive tasks is

$$\beta_z^{GPT} \approx \frac{(\eta^z - 1)}{1 + \left(\eta^z/\eta^\ell\right) \cdot \left(z_{i,post}^{\text{No adopt}}/\ell_{i,post}^{\text{No adopt}}\right)} \times \delta \tag{16}$$

Equation (16) shows that the effect of adoption on the amount of time spent on productive digital activities–i.e., the tasks that are directly made more efficient by ChatGPT–can be either positive or negative. In particular, ChatGPT adoption leads to more (less) time spent on productive tasks, $\beta_z^{GPT} > 0$ ($\beta_z^{GPT} < 0$) if the elasticity of utility with respect to productive browsing, η^z, is greater than (less than) one. Intuitively, if $\eta^z < 1$, then the positive substitution effect from adopting ChatGPT–productive tasks becoming cheaper (in terms of foregone time) to do–is dominated by a negative income effect–the marginal productive task being less valuable the more time that the household has to devote to home computing. The income effect is negative when $\eta^z < 1$, and gets more negative the lower η^z is, due to the decreasing returns that the household gets from completing additional productive tasks.

The time allocated to leisure tasks takes a form that is similar but in which η^z enters with the opposite sign:

$$\beta_{\ell}^{GPT} \approx \frac{(1-\eta^z)\cdot\frac{\eta^\ell}{\eta^z}}{1+\left(\eta^\ell/\eta^z\right)\cdot\left(\ell_{i,post}^{\text{No adopt}}/z_{i,post}^{\text{No adopt}}\right)} \times \delta \tag{17}$$

Adopting ChatGPT leads a household to spend more (less) time on leisure tasks if $\eta^z < 1$ ($\eta^z > 1$). The intuition for why is the same as with productive tasks, but in the opposite direction: when $\eta^z < 1$, a boost in the efficiency of productive tasks means that the household runs into decreasing returns in the value of such tasks, leading them to tilt their time allocation towards leisure tasks. Note that although this positive effect on leisure time is attenuated as η^ℓ itself goes down, the effect remains positive regardless of whether leisure is a luxury ($\eta^\ell > 1$) or inferior / necessity good ($\eta^\ell < 1$).

The key implication for our empirical results is straightforward. When productive digital tasks behave as a 'necessity' ($\eta^z < 1$), the model predicts that a productivity shock to productive tasks reduces time spent on productive browsing and increases time spent on leisure. This is exactly the pattern we estimate: we find that ChatGPT adoption leads to a statistically significant rise in the leisure share of browsing and little to no increase in total browsing duration (see Table VII), while time spent on sites whose activities overlap strongly with ChatGPT falls (Table IX).

C. *Quantifying the Productivity and Welfare Effects of Household GenAI Adoption*

Our conceptual framework shows how a task-biased productivity shock from generative AI maps into changes in households' allocation of time across productive and leisure digital activities. We can combine that framework with our IV estimates to quantify the implied magnitude of the productivity shock generated by ChatGPT and to provide back-of-the-envelope measures of the resulting welfare effects from private household use of the technology.[15]

From time reallocation to an implied productivity shock. In our framework, adoption of ChatGPT raises the efficiency of productive digital tasks by a factor of $1+\delta$ while

[15]This exercise builds on the way Aguiar et al. (2021) back out leisure technology shocks from changes in time allocation, and also on the methodology in Brynjolfsson et al. (2023) for valuing new and free digital goods.

leaving the technology for leisure tasks unchanged. This assumption is conservative in the sense that it will likely lead us to undervalue the total change in welfare due to the generative AI technology shock (see further discussion below). Given values for the curvature parameters η_z and η_ℓ, empirical estimates of pre-adoption time shares of tasks, and our empirical estimates of changes in leisure browsing time β_ℓ^{GPT}, we can invert equation 17 to recover the implied productivity shock δ.

Effective time savings for households. The condition in equation 13 for adopting ChatGPT provides an intuition for the magnitude of δ, as $\delta z_{i,post}^{\text{No adopt}}$ can be interpreted as the additional units of home-computing time that ChatGPT effectively adds to the household's time budget each period, holding the shadow price of time fixed. To then translate these time-savings into a money-metric welfare measure, we can follow Aguiar et al. (2021) in valuing time at the household's shadow wage.

Aggregate gains. Finally, to scale these household-level welfare gains to the aggregate economy, we can combine it with assumptions about the eventual diffusion of ChatGPT and the corresponding number of regular ChatGPT users to estimate the contribution to real income growth. As Brynjolfsson et al. (2025) suggest, this household benefit of generative AI would likely be missed by conventional GDP accounting, which misses the substantial consumer surplus of new free digital goods.

While such calculations would be intentionally simple and rely on several strong assumptions, they would allow us to make some progress on estimating the welfare gains corresponding to the household-level productivity improvements due to generative AI. These welfare gains are likely to be economically meaningful both for adopters and in the aggregate, but are mostly missing from the discourse around the economic impact of generative AI.

Note: We will add explicit calculations based on the methodology laid out in this section in the next versions of this working paper.

V. Mechanism: What Do Households Use ChatGPT For?

In this section, we provide further evidence of the mechanism for *why* ChatGPT adoption would increase the leisure share of household browsing. We show that ChatGPT adoption is mainly associated with "productive" non-market work browsing sessions , which suggests that it is likely to increase productivity mainly of productive internet use.

A. *Inferring ChatGPT adoption context 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 adoptionrs 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 0.2% and 1.1% of the intervals are labeled as GPT-window in 2023 and 2024, respectively.

B. Tasks associated with ChatGPT adoption

Figure 7 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, the 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 provides a first glance at what household uses generative AI for. Given the particular strength of generative AI models in processing text information, it is intuitive that they can support 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 IA.9 that the households use GPT much more for education and information searching and much less for entertainment, shopping and finance.

Our findings go beyond some existing data sources that characterize generative AI users as a whole based on correlations in website visits[16]: The key advantage of our approach is that we can use the full pattern of households' internet use to separate activity within the ChatGPT-window, which suggests ChatGPT's "purpose", and activity outside the ChatGPT-window. By being able to distinguish what ChatGPT is used for from the effect of its use on overall household browsing behavior, we can infer to what degree ChatGPT increases household productivity.

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 Chat-

[16]See, for instance, Similarweb.com which labels chatgpt.com's audience as interested in "Video Games Consoles and Accessories."

GPT 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.

VI. Conclusion

This paper provides the first comprehensive evidence on the adoption and impact of generative AI on U.S. households, utilizing detailed browsing data to explore the timing, usage patterns, and behavioral changes associated with ChatGPT. Our findings show that households predominantly use ChatGPT for productive, non-market activities, such as education and job search, and that generative AI adoption reduces productive online activities while increasing leisure time. This suggests that ChatGPT enhances household welfare by making certain tasks more efficient, thus freeing up time for leisure. Moreover, we document significant demographic gaps in adoption, with younger and higher-income households adopting generative AI at higher rates. These gaps can be partially explained by differences in the potential benefits of generative AI, which are greater for households engaged in activities that align well with ChatGPT's capabilities.

Our analysis also sheds light on how generative AI adoption alters household behavior. By quantifying changes in online activity before and after ChatGPT adoption, we find that households tend to shift time from productive online tasks to leisure activities, demonstrating both a substitution effect and positive welfare implications. These results suggest that while generative AI can increase household productivity, it may also contribute to enhanced well-being by providing more time for leisure. The broader economic impact of these shifts is likely to be substantial, especially as generative AI continues to diffuse into households across the income and age spectrum.

Finally, our findings emphasize the need for policies that address the unequal adoption of generative AI, particularly among older and lower-income households. By improving access to tools and increasing technology literacy in these groups, policymakers can help ensure that the benefits of generative AI are more evenly distributed. As generative AI continues to evolve, understanding its impact on household productivity and welfare will be critical for designing inclusive strategies that maximize its societal benefits while mitigating potential inequalities.

REFERENCES

Aguiar, Mark, Mark Bils, Kerwin Kofi Charles, and Erik Hurst, 2021, Leisure luxuries and the labor supply of young men, *Journal of Political Economy* 129, 337–382.

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, Avinash Collis, W Erwin Diewert, Felix Eggers, and Kevin J Fox, 2025, Gdp-b: Accounting for the value of new and free goods, *American Economic Journal: Macroeconomics* 17, 312–344.

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

Burtch, Gordon, Dokyun Lee, and Zhichen Chen, 2024, The consequences of generative ai for online knowledge communities, *Scientific Reports* 14, 10413.

Chatterji, Aaron, Thomas Cunningham, David J. Deming, Zoe Hitzig, Christopher Ong, Carl Yan Shan, and Kevin Wadman, 2025, How People Use ChatGPT.

del Rio-Chanona, R Maria, Nadzeya Laurentsyeva, and Johannes Wachs, 2024, Large language models reduce public knowledge sharing on online q&a platforms, *PNAS nexus* 3, pgae400.

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, and Gregor Schubert, 2024, Ai and finance, Technical report, National Bureau of Economic Research.

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* .

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., 2025a, Which economic tasks are performed with ai? evidence from millions of claude conversations, *arXiv preprint arXiv:2503.04761* .

Handa, Kunal, Alex Tamkin, Miles McCain, Saffron Huang, Esin Durmus, Sarah Heck, Jared Mueller, Jerry Hong, Stuart Ritchie, Tim Belonax, et al., 2025b, 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.

Jiang, Wei, Junyoung Park, Rachel Jiqiu Xiao, and Shen Zhang, 2025, Ai and the extended workday: Productivity, contracting efficiency, and distribution of rents, 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).

Kosmyna, Nataliya, Eugene Hauptmann, Ye Tong Yuan, Jessica Situ, Xian-Hao Liao, Ashly Vivian Beresnitzky, Iris Braunstein, and Pattie Maes, 2025, Your brain on chatgpt: Accumulation of cognitive debt when using an ai assistant for essay writing task, *arXiv preprint arXiv:2506.08872* .

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* .

Lyu, Liang, James Siderius, Hannah Li, Daron Acemoglu, Daniel Huttenlocher, and Asuman Ozdaglar, 2025, Wikipedia contributions in the wake of chatgpt, in *Companion Proceedings of the ACM on Web Conference 2025*, 1176–1179.

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* .

Padilla, Nicolas, H Tai Lam, Anja Lambrecht, and Brett Hollenbeck, 2025, The impact of llm adoption on online user behavior, *Available at SSRN 5393256* .

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.

Figure 2: Household ChatGPT adoption over time

This figure shows the share of households that has ever tried ChatGPT based on internet browsing data.

Panel A: Share that has tried ChatGPT

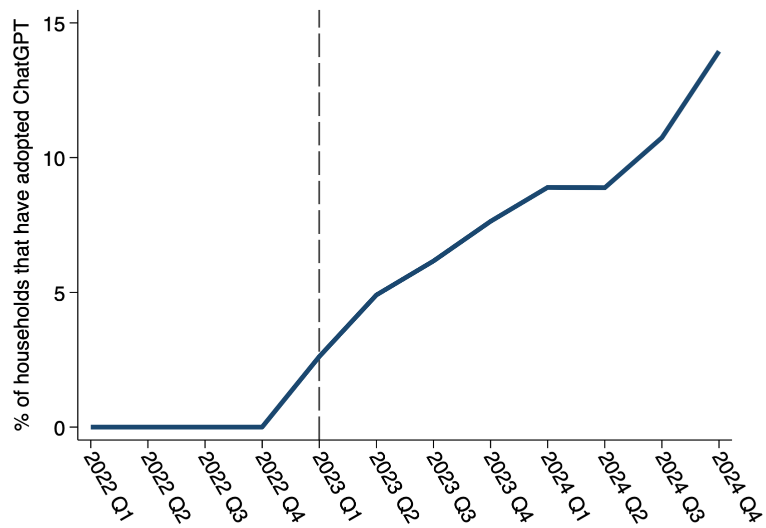

Panel B: Share of browsing duration on ChatGPT

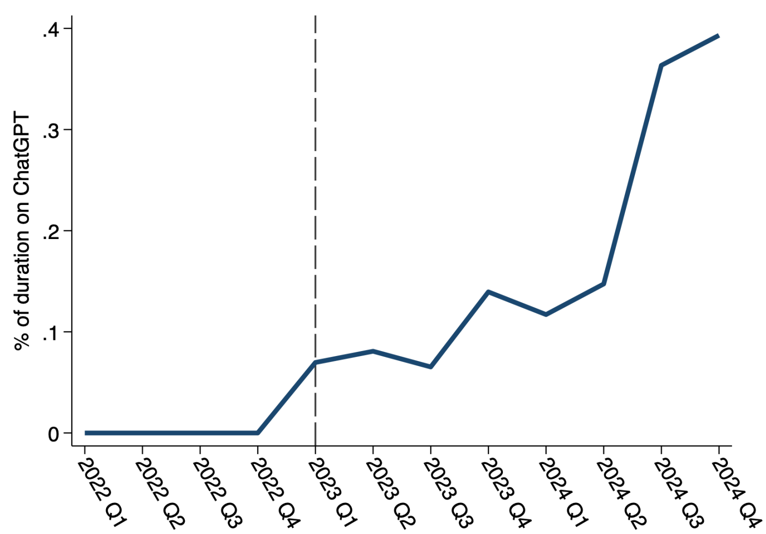

Figure 3: ChatGPT adoption by household income and age

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

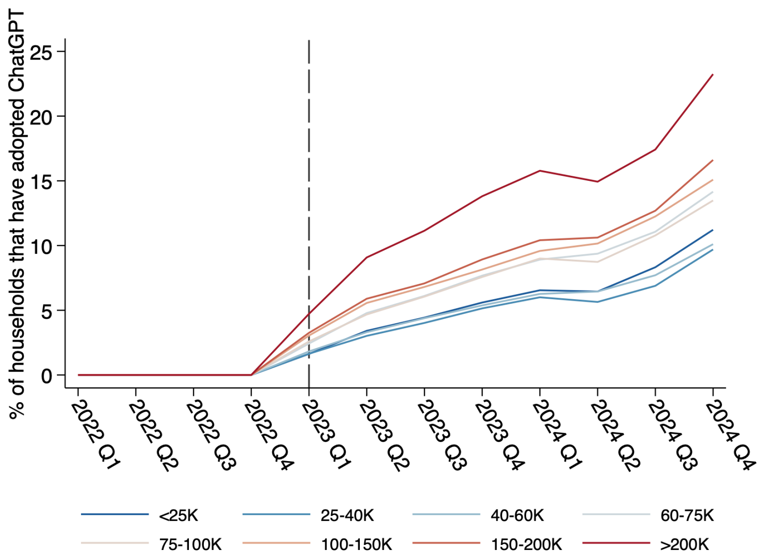

(A) ChatGPT adoption by income bucket

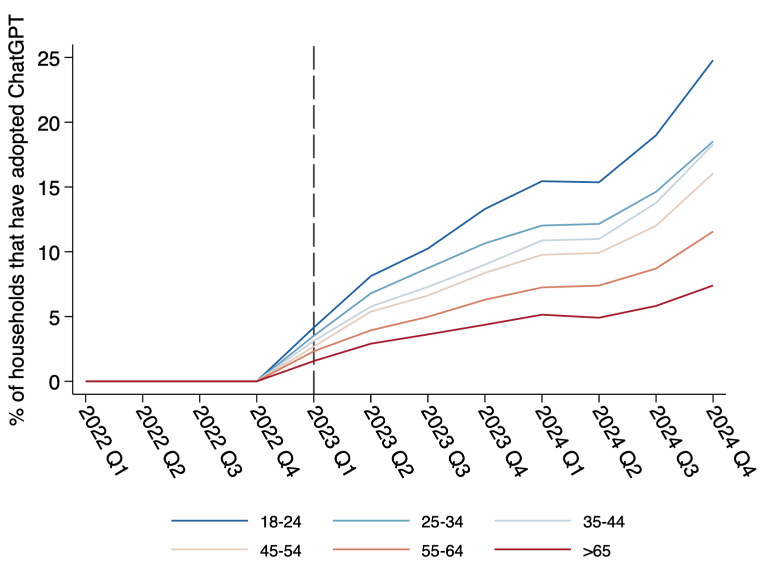

(B) ChatGPT adoption by age bucket

Figure 4: Share of GenAI-exposure websites 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.

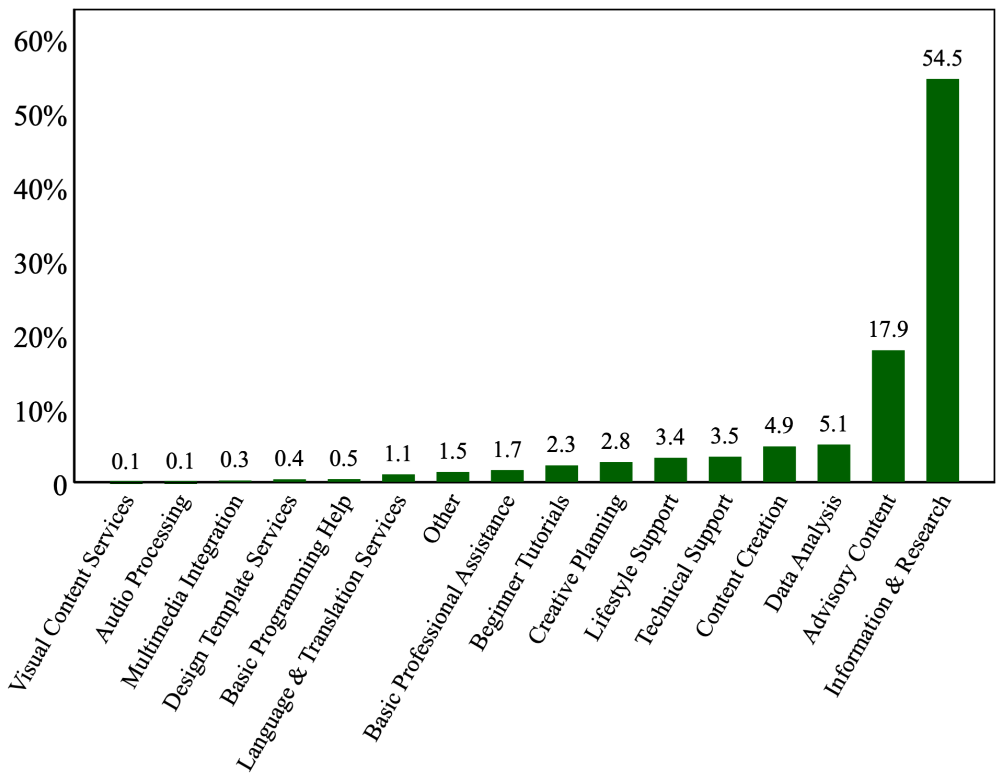

Figure 5: ChatGPT adoption by GenAI exposure

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).

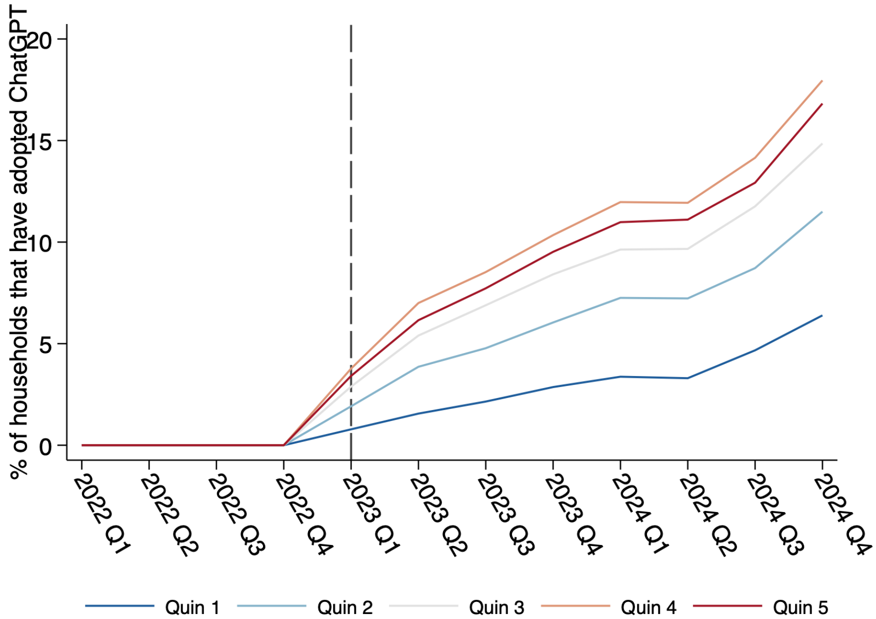

(A) Share that has tried ChatGPT

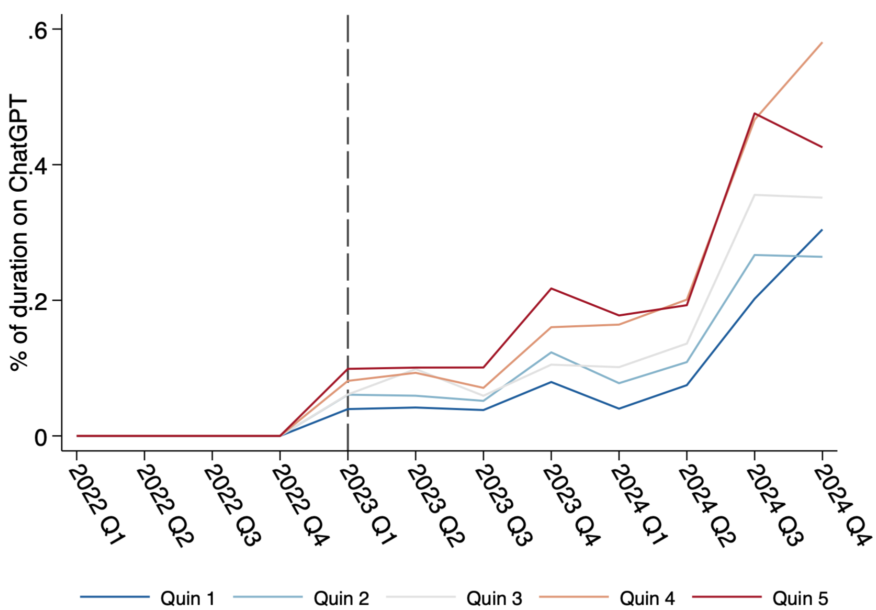

(B) Share of browsing duration on ChatGPT

Figure 6: Impact of GenAI on household productive vs. leisure activities.

This figure plots the results of regressing households' leisure and productive browsing duration on their HHGenAIExposure, defined in equation (1) in each quarter before and after the release of ChatGPT in November 30, 2022. See Section III for more details.

Panel A: Response of log browsing duration of leisure websites

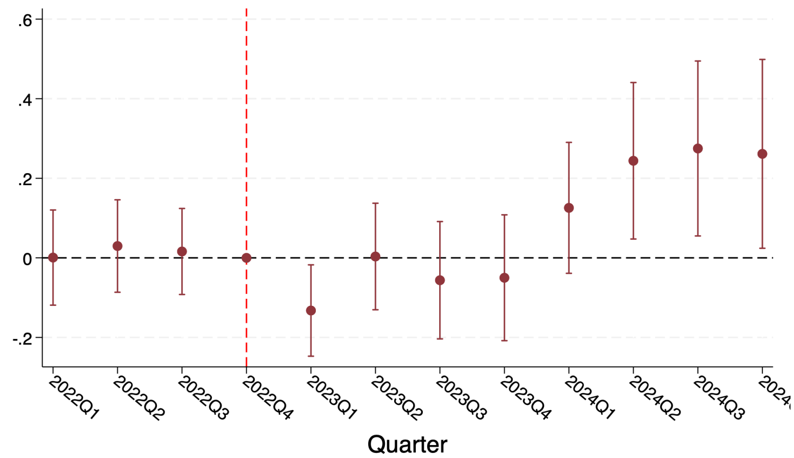

Panel B: Response of log browsing duration of productive websites

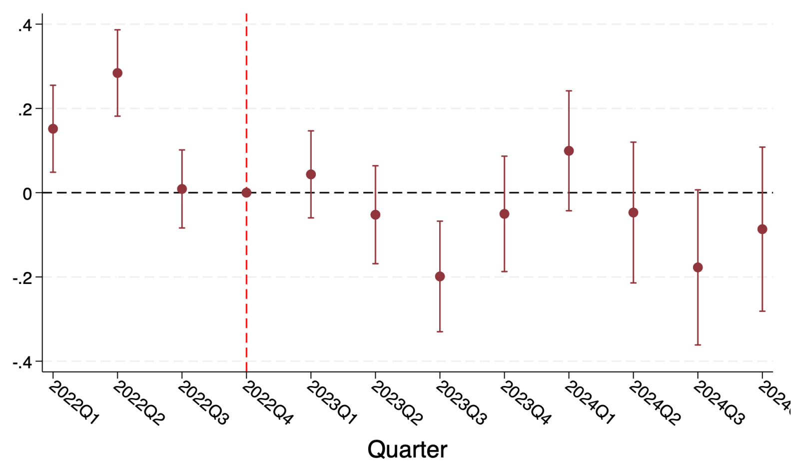

Figure 7: Household browsing activities during ChatGPT session

This figure plots the difference between households' browsing activities in leisure, productive, and other categories during the ChatGPT session relative to outside of the ChatGPT session. See more details in Section V.

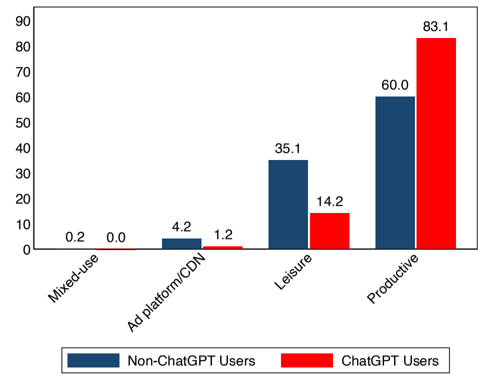

(A) GPT window: browsing shares

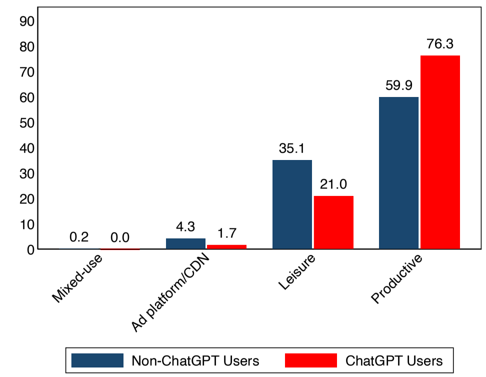

(B) GPT day, outside window: browsing shares

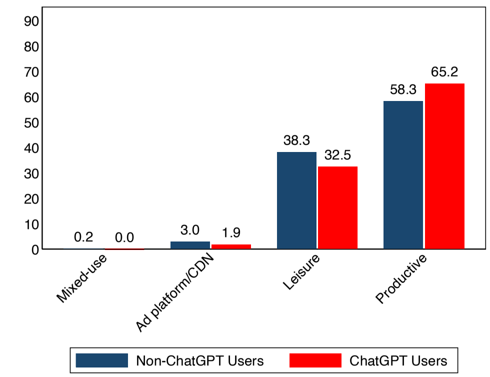

(C) Pre-GPT use: browsing shares

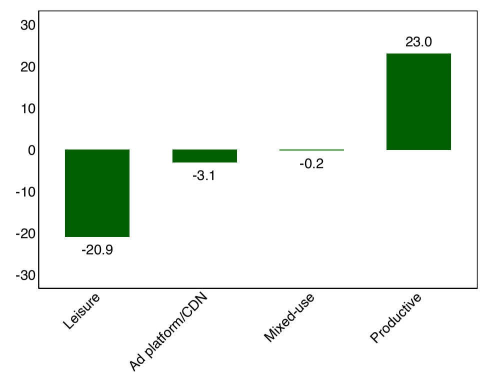

(D) GPT window: differences in shares

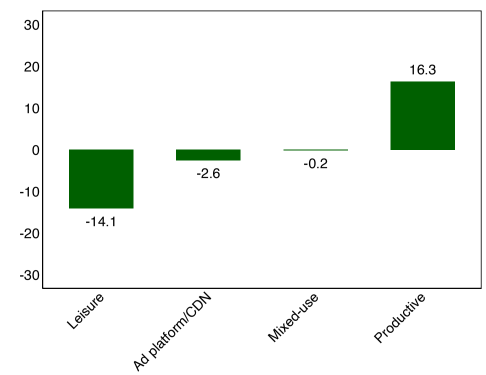

(E) GPT day, outside window: differences in shares

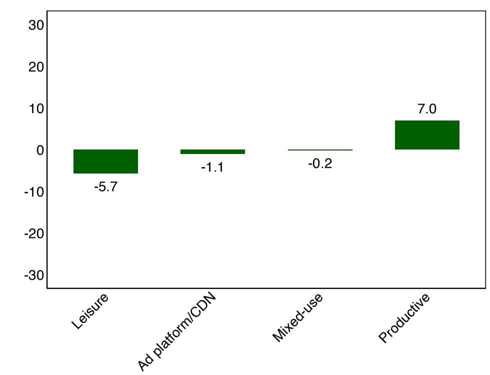

(F) Pre-GPT use: differences in shares

Figure 8:
Illustration: Generative AI and Household Production

This figure illustrates how a change in the productivity of productive household activities would be expected to change household time allocation: the solid budget line represents the household's time budget before generative AI. As generative AI makes it faster to complete productive tasks, this expands the household's effective time budget to the dotted line. The impact on relative time allocation depends on the elasticity of substitution between leisure time and the output of productive time spent browsing. In the illustrated case, the iso-utility curves are such that leisure behaves as a "normal" good with regard to the household's time budget, and leisure consumption increases in the new equilibrium with generative AI at point B, relative to the old equilibrium at point A.

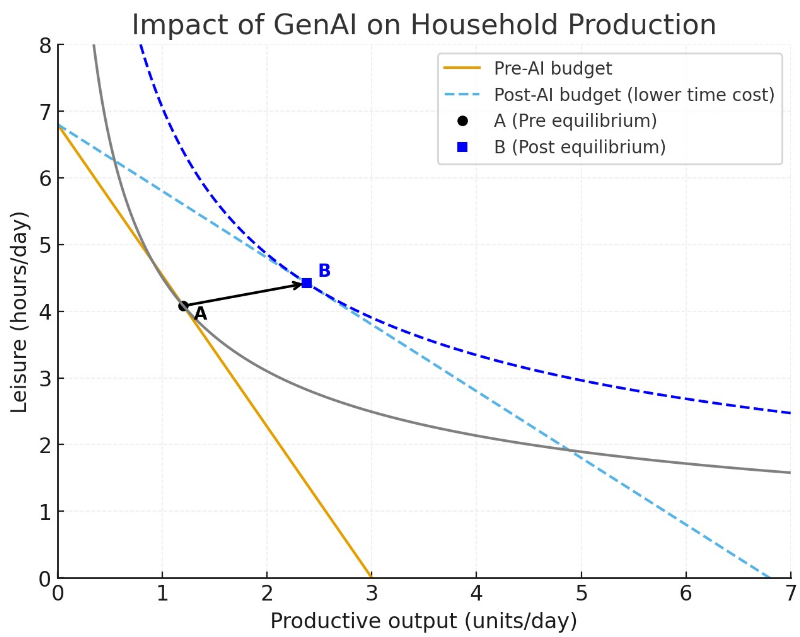

Table I:
Distribution by income and age, Comscore sample vs. ACS

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 adoption 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 IV.

		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 (YTD June 2022)	11.3	10.7	11.1	11.8	12.3	11.8	10.4
Panel B: % Tried ChatGPT							
Q4 2023	7.6	5.4	6.8	9.9	11.2	8.7	5.2
Q2 2024	8.9	6.1	8.0	11.6	12.8	10.5	6.0
Q4 2024	14.0	10.5	12.4	17.6	19.6	17.2	9.2
Panel C: % ChatGPT Regular User							
Q4 2023	3.0	2.1	2.8	3.8	4.6	3.4	1.9
Q2 2024	3.8	2.5	3.4	5.0	5.9	4.5	2.3
Q4 2024	6.3	4.6	5.5	8.0	9.5	7.8	3.8

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:
Household exposure effects on ChatGPT adoption

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

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

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 adoption: 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 adoption, 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.

	ChatGPT adoption measure			
Dependent variable:	Ever tried (%)	Regular user (%)	Share of Browsing Time (%)	Visits per Month
	(1)	(2)	(3)	(4)
Panel A: December 2023				
HH GenAI Exposure	6.56***	3.86***	0.18***	0.45***
	(0.72)	(0.49)	(0.05)	(0.08)
Observations	131,283	131,283	103,217	131,283
Dep. var. mean	8.114	3.153	0.157	0.271
Indep. var. SD	0.149	0.149	0.144	0.149
Panel B: June 2024				
HH GenAI Exposure	7.72***	4.23***	0.18***	0.44***
	(0.98)	(0.69)	(0.06)	(0.10)
Observations	77,554	77,554	60,998	77,554
Dep. var. mean	9.167	3.859	0.149	0.345
Indep. var. SD	0.151	0.151	0.148	0.151
Panel C: December 2024				
HH GenAI Exposure	9.62***	6.09***	0.21*	0.64
	(1.51)	(1.12)	(0.11)	(0.51)
Observations	47,383	47,383	47,383	47,383
Dep. var. mean	14.815	6.673	0.454	0.962
Indep. var. SD	0.147	0.147	0.147	0.147
Income x Age x HH Size x Children FEs	X	X	X	X

Table V:
Heterogeneity in household exposure effects on ChatGPT adoption: by age & income

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

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

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 adoption: 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	Q2 2024	
Dependent variable:	Ever tried (%)	Regular user (%)
	(1)	(2)
Panel A: By Income Group		
HH GenAI Exposure	6.59***	3.43***
	(1.04)	(0.67)
HH GenAI Exposure x Middle Income	1.24	1.68
	(2.00)	(1.42)
HH GenAI Exposure x High Income	1.84	0.53
	(2.16)	(1.49)
Observations	77,554	77,554
Panel B: By Age Group		
HH GenAI Exposure	11.28***	9.00***
	(2.99)	(2.25)
HH GenAI Exposure x Middle Age	-3.37	-5.87**
	(3.35)	(2.50)
HH GenAI Exposure x Old Age	-5.63*	-6.19***
	(3.20)	(2.37)
Observations	77,554	77,554
Panel C: By Tech Knowledge (proxied by Stack Overflow use)		
HH GenAI Exposure	5.77***	2.94***
	(0.94)	(0.68)
HH GenAI Exposure x Stack Overflow User	21.57**	15.37**
	(8.57)	(6.25)
Observations	77,554	77,554
Income x Age x HH Size x Children FEs	X	X

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

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

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

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															
Q4 2023	4.54	0.07	0.25	0.02	**7.65**	5.98	0.13	0.33	0.06	**8.75**	16.73	0.23	1.62	0.51	**14.13**
Q2 2024	5.46	0.08	0.20	0.02	**5.47**	6.79	0.15	0.59	0.11	**12.59**	18.33	0.28	2.44	0.78	**19.09**
Q4 2024	7.03	0.10	0.03	0.00	**1.92**	10.50	0.21	0.12	0.02	**3.35**	23.98	0.31	3.88	1.10	**22.07**
Panel B: % ChatGPT Regular user															
Q4 2023	1.76	0.04	0.05	0.01	**5.63**	2.69	0.08	0.37	0.07	**18.99**	10.13	0.14	0.93	0.29	**13.43**
Q2 2024	2.42	0.04	0.06	0.01	**4.36**	3.61	0.08	0.65	0.12	**23.74**	10.84	0.15	1.74	0.56	**22.58**
Q4 2024	3.41	0.06	0.09	0.01	**4.63**	5.75	0.13	0.62	0.13	**15.26**	15.14	0.19	1.83	0.52	**16.83**

Table VII:
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

$$\text{exposed\_browsing}_{j,t} = \gamma^{RF}\ \text{HHGenAIExposure}_j \times \text{post\_gpt}_t + \lambda_j + \text{Char}_{j,t} + \varepsilon_{j,t},$$

and the first-stage equation is

$$\text{adopt}_{j,t} = \gamma^{FS}\ \text{HHGenAIExposure}_j \times \text{post\_gpt}_t + \lambda_j + \text{Char}_{j,t} + \eta_{j,t},$$

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 Try	-5.732***		-1.435***	
	(0.675)		(0.194)	
Post Short Run (2023) X Try		-6.400***		-1.578***
		(0.812)		(0.232)
Post Medium Run (2024) X Try		-4.742***		-1.220***
		(0.606)		(0.188)
Machine FEs	Yes	Yes	Yes	Yes
Quarter X Demographics X Region FEs	Yes	Yes	Yes	Yes
N	1,401,613	1,401,613	1,389,802	1,389,802
IV F-stat	73.608	32.839	73.493	33.056

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

$$\text{browsing\_outcome}_{j,t} = \gamma^{RF}\, \text{HHGenAIExposure}_j \times \text{post\_gpt}_t + \lambda_j + \text{Char}_{j,t} + \varepsilon_{j,t},$$

and the first-stage equation is

$$\text{adopt}_{j,t} = \gamma^{FS}\, \text{HHGenAIExposure}_j \times \text{post\_gpt}_t + \lambda_j + \text{Char}_{j,t} + \eta_{j,t},$$

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 Try	-2.661***		-0.091		0.266***	
	(0.790)		(0.595)		(0.087)	
Post Short Run (2023) X Try		-3.655***		-1.293*		0.244**
		(0.994)		(0.729)		(0.102)
Post Medium Run (2024) X Try		-1.093		2.090***		0.300***
		(0.768)		(0.749)		(0.095)
Machine FEs	Yes	Yes	Yes	Yes	Yes	Yes
Quarter X Demographics X Region FEs	Yes	Yes	Yes	Yes	Yes	Yes
N	1,401,794	1,401,794	1,202,423	1,202,423	1,401,794	1,401,794
IV F-stat	73.536	32.832	82.931	38.292	73.536	32.832

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

$$\text{browsing\_outcome}_{j,t} = \gamma^{RF}\,\text{HHGenAIExposure}_j \times \text{post\_gpt}_t + \lambda_j + \text{Char}_{j,t} + \varepsilon_{j,t},$$

and the first-stage equation is

$$\text{adopt}_{j,t} = \gamma^{FS}\,\text{HHGenAIExposure}_j \times \text{post\_gpt}_t + \lambda_j + \text{Char}_{j,t} + \eta_{j,t},$$

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 Try	0.424***	0.471***	-0.903***	0.733***	-0.074**	-0.141**	0.685***	-0.151**	0.029	0.069	-0.046	-0.248***	-0.194***	-0.204***
	(0.118)	(0.095)	(0.132)	(0.106)	(0.037)	(0.056)	(0.105)	(0.068)	(0.056)	(0.054)	(0.030)	(0.075)	(0.038)	(0.052)
Post Medium Run (2024) X Try	0.176*	0.343***	-0.591***	0.745***	0.030	-0.170***	0.513***	-0.122*	-0.044	0.130**	-0.034	-0.212***	-0.154***	-0.060
	(0.094)	(0.078)	(0.092)	(0.107)	(0.037)	(0.049)	(0.083)	(0.063)	(0.046)	(0.056)	(0.026)	(0.069)	(0.030)	(0.043)
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	1,401,794	1,401,794	1,401,794	1,401,794	1,401,794	1,401,794	1,401,794	1,401,794	1,401,794	1,401,794	1,401,794	1,401,794	1,401,794	1,401,794
IV F-stat	32.832	32.832	32.832	32.832	32.832	32.832	32.832	32.832	32.832	32.832	32.832	32.832	32.832	32.832

Internet Appendix for

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

Michael Blank Gregor Schubert Miao Ben Zhang

Appendix A. Methodology Appendix

Appendix A. Website classification

We sort websites in the public use version of Comscore 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

Here is the prompt used for labeling websites as exposed:

```
    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 adoption.

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",
}}
```

Appendix C. Website activity exposure classification

Here is the prompt used for labeling website activities as exposed:

```
I need to classify a specific activity on a website based on whether users might be able to substitute this activity with ChatGPT adoption.

Website Domain: {domain}
Website Description: {website_description}
Specific Website Activity: {activity}

Please classify this specific website activity based on whether it could be partially or completely substituted by using ChatGPT.

Here are the website capabilities that ChatGPT might replace:

### 1. Language & Translation Services
- **Translation Platforms:** Converting text between languages
- **Grammar & Writing Tools:** Spell checking, grammar correction, style suggestions
- **Language Learning Resources:** Language explanations, practice exercises
- **Multilingual Customer Service:** Bridging language gaps in customer communications

### 2. Content Creation
- **Creative Writing:** Story, poem, or script generation
- **Business Documents:** Professional document creation and templates
- **Marketing Copy:** Advertising text, slogans, product descriptions
- **Email Writing:** Professional communications drafting
- **Content Generation:** Basic informational articles

### 3. Information & Research
- **General Knowledge:** Facts, explanations, encyclopedia-like information
```

- **Simple Q&A:** Straightforward informational questions and answers
- **Educational Summaries:** Simplified explanations of concepts
- **Product Comparisons:** Basic review information
- **Academic Resources:** Standard subject explanations

4. Basic Programming Help
- **Code Snippets:** Common programming solutions
- **Programming Tutorials:** Fundamental coding concept explanations
- **Simple Debugging:** Basic code error identification
- **Standard Documentation:** Common technical explanations

5. Data Analysis
- **Basic Visualization:** Simple chart and graph generation
- **Data Interpretation:** Statistical information explanation
- **Simple Reporting:** Standard data report creation
- **Algorithm Explanation:** Common computational approach descriptions

6. Technical Support
- **Common Troubleshooting:** Solutions to frequent issues
- **Software Usage Guides:** Popular program usage explanations
- **IT Knowledge Base:** Standard technical solutions
- **Command References:** Command structures and usage information

7. Basic Professional Assistance
- **Resume Building:** Standard resume format generation
- **Cover Letters:** Professional letter framework creation
- **Business Plans:** Basic business documentation
- **Simple Legal Documents:** Standard legal form assistance

8. Advisory Content
- **General Financial Advice:** Basic investment information
- **Career Guidance:** Common professional development tips
- **Business Strategy:** Standard business insights
- **Health Information:** General wellness and nutrition advice

9. Lifestyle Support
- **Relationship Advice:** General interpersonal guidance
- **Trip Planning:** Standard travel itineraries
- **Fitness Programs:** Common exercise routines
- **Study Aids:** Standard learning assistance

10. Creative Planning
- **Event Planning:** Standard organizational frameworks
- **Project Management:** Basic structural approaches
- **Brainstorming:** Ideation and creative process facilitation
- **Organizational Systems:** Standard management frameworks

11. Visual Content Services
- **Image Description:** Caption or explanation of visual content
- **Basic Visual Analysis:** Element identification in images
- **Chart Interpretation:** Graph and visualization explanation
- **Document Reading:** Text extraction from images

12. Audio Processing
- **Transcription:** Speech-to-text conversion
- **Basic Audio Analysis:** Element identification in audio
- **Content Summarization:** Text summaries of audio content
- **Voice-to-Text:** Speech recognition functionality

13. Multimedia Integration
- **Basic Media Conversion:** Translation between content formats
- **Simple Virtual Assistance:** Basic multimodal interactions
- **Content Accessibility:** Making content available across formats
- **Media Description:** Text descriptions of visual/audio content

14. Beginner Tutorials
- **Introductory Learning:** Fundamental knowledge provision
- **How-To Guides:** Step-by-step instructions
- **Basic Skill Development:** Entry-level capability teaching
- **Simplified Explanations:** Complex topic reduction to basics

15. Design Template Services
- **Basic UI/UX Resources:** Standard design patterns
- **Template Libraries:** Pre-made design frameworks

```
- **Simple Logo Generation:** Basic brand element creation
- **Design Guidelines:** Fundamental aesthetic principle explanation

Determine if this specific activity is exposed to ChatGPT competition based on the above categories.

Return your response in this exact JSON format:
{{
    "activity_type": "Categorize this activity into one of the 15 categories above, or 'Other' if none apply",
    "exposure_reasoning": "Brief explanation of why this activity is or is not exposed to ChatGPT competition",
    "substitution_examples": "If exposed, provide 1-2 specific examples of how users might substitute this activity 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 IA.1: Sample Restriction Criteria

Exclusion Criteria	Number of Machines		Number of Machine - Quarters	
	Number	Percent	Number	Percent
None	303,244	100	2,439,710	100
- Machines for work	291,649	96	2,333,246	96
- Machines w/o demographic data	243,650	80	1,953,678	80
- Machines w/ < 1 hour of browsing from 7/2021 to 6/2022	199,783	66	1,741,563	71

Table IA.2:
Heterogeneity in household exposure effects on ChatGPT adoption: by age & income, Addl. time periods

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

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

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 adoption: 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	Q4 2023		Q4 2024	
Dependent variable:	Ever tried (%)	Regular user (%)	Ever tried (%)	Regular user (%)
	(1)	(2)	(3)	(4)
Panel A: By Income Group				
HH GenAI Exposure	5.27***	3.17***	10.34***	5.51***
	(0.74)	(0.48)	(1.79)	(1.18)
HH GenAI Exposure x Middle Income	1.13	1.37	-2.22	0.75
	(1.42)	(1.00)	(3.13)	(2.32)
HH GenAI Exposure x High Income	2.36	0.50	0.30	0.79
	(1.61)	(1.06)	(3.36)	(2.37)
Observations	131,283	131,283	47,383	47,383
Panel B: By Age Group				
HH GenAI Exposure	8.96***	6.69***	10.60**	10.63***
	(2.16)	(1.49)	(4.34)	(3.34)
HH GenAI Exposure x Middle Age	-2.91	-3.70**	-1.38	-5.63
	(2.42)	(1.68)	(4.94)	(3.79)
HH GenAI Exposure x Old Age	-3.21	-3.54**	-1.12	-5.92*
	(2.34)	(1.59)	(4.73)	(3.57)
Observations	131,283	131,283	47,383	47,383
Panel C: By Tech Knowledge (proxied by Stack Overflow use)				
HH GenAI Exposure	4.94***	2.88***	6.87***	4.58***
	(0.70)	(0.47)	(1.46)	(1.10)
HH GenAI Exposure x Stack Overflow User	14.58**	8.38*	34.39***	16.24*
	(6.09)	(4.41)	(10.28)	(8.70)
Observations	131,283	131,283	47,383	47,383
Income x Age x HH Size x Children FEs	X	X	X	X

Appendix C. Additional Figures

Figure IA.1: ChatGPT regular use 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.

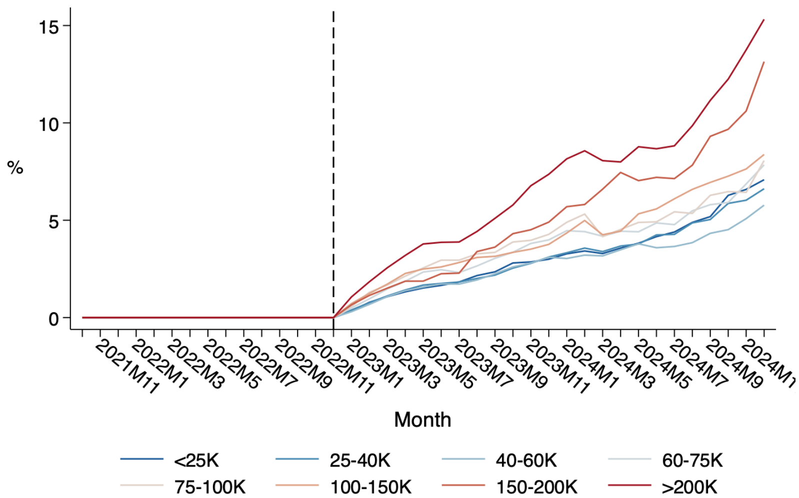

(A) ChatGPT regular use by income bucket

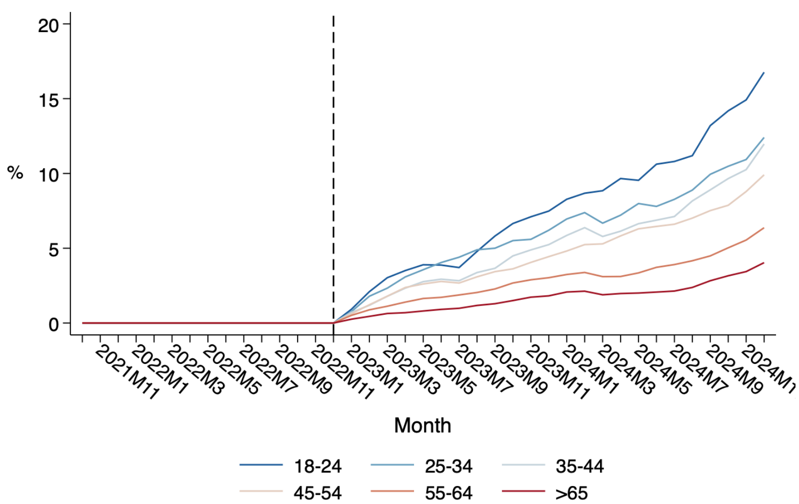

(B) ChatGPT regular use by age bucket

Figure IA.2: 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. The light-blue confidence intervals shown correspond to 2

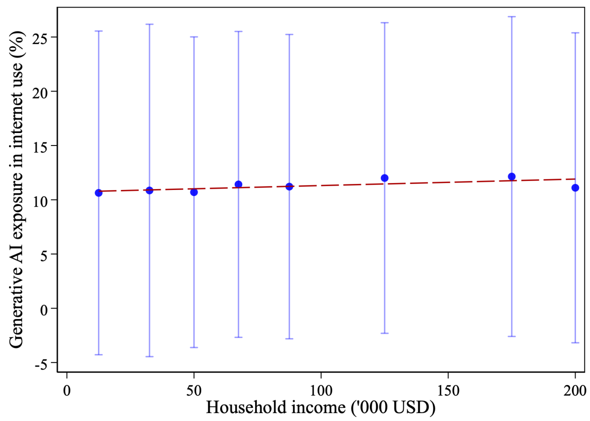

(A) Generative AI exposure by income bucket

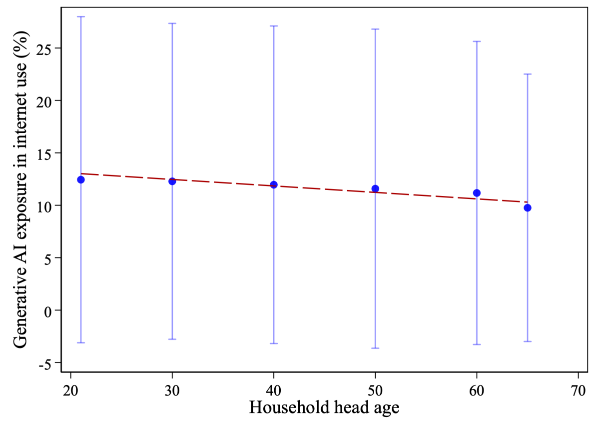

(B) Generative AI exposure by age bucket

Figure IA.3: Duration-weighted double sorts of website overlap fractions and Comscore activity categories

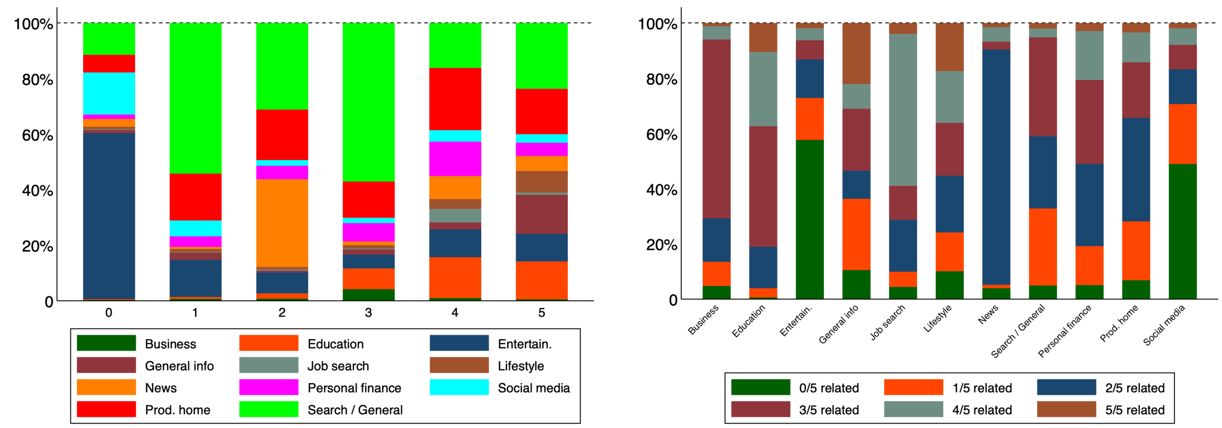

(A) Duration shares of overlap fractions

(B) Duration shares of site activity categories

Figure IA.4:
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.

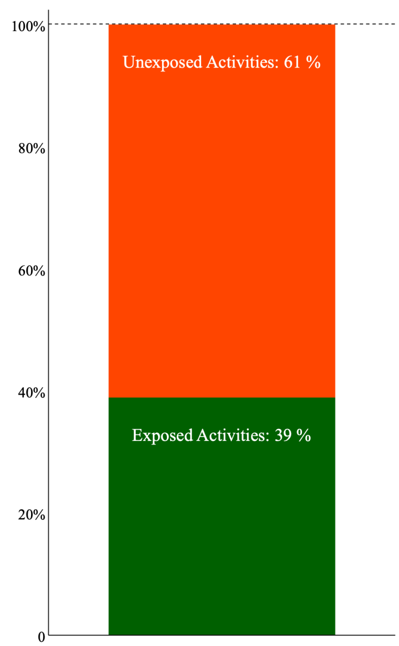

(A) Duration share of exposed activities

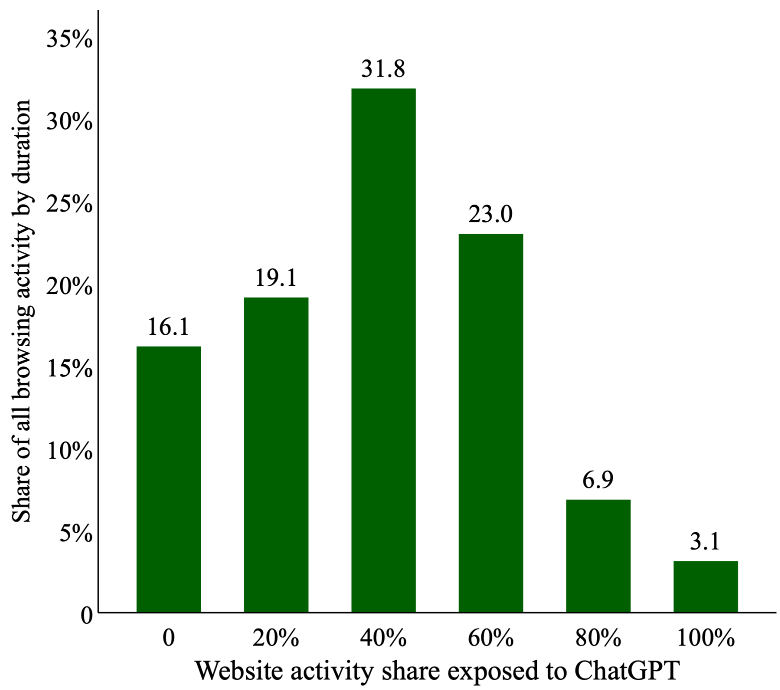

(B) Duration distribution over website exposure levels

Figure IA.5: Pre-ChatGPT release duration shares for different site categories, by household income bucket

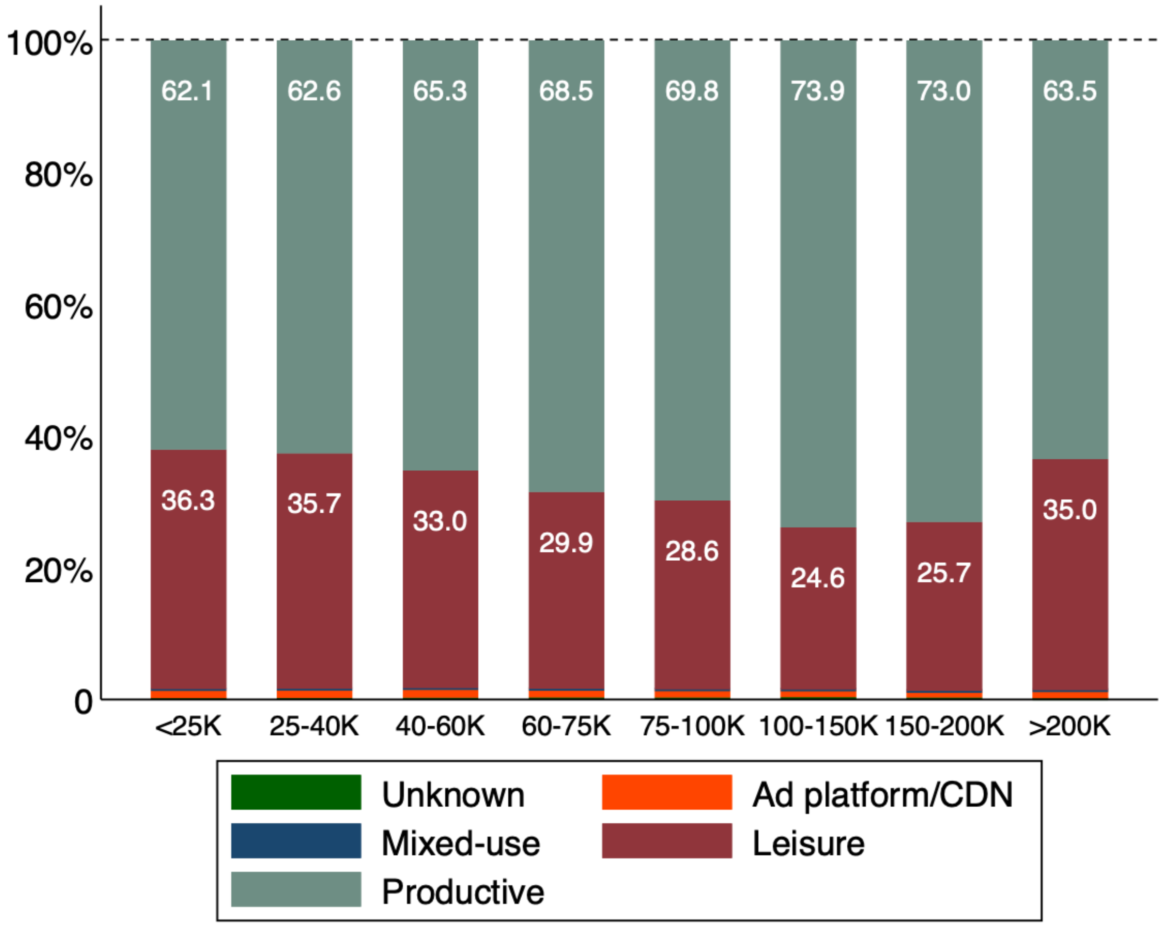

Figure IA.6: Estimated effect of placebo pre-ChatGPT adoption on GenAI-exposed browsing intensity

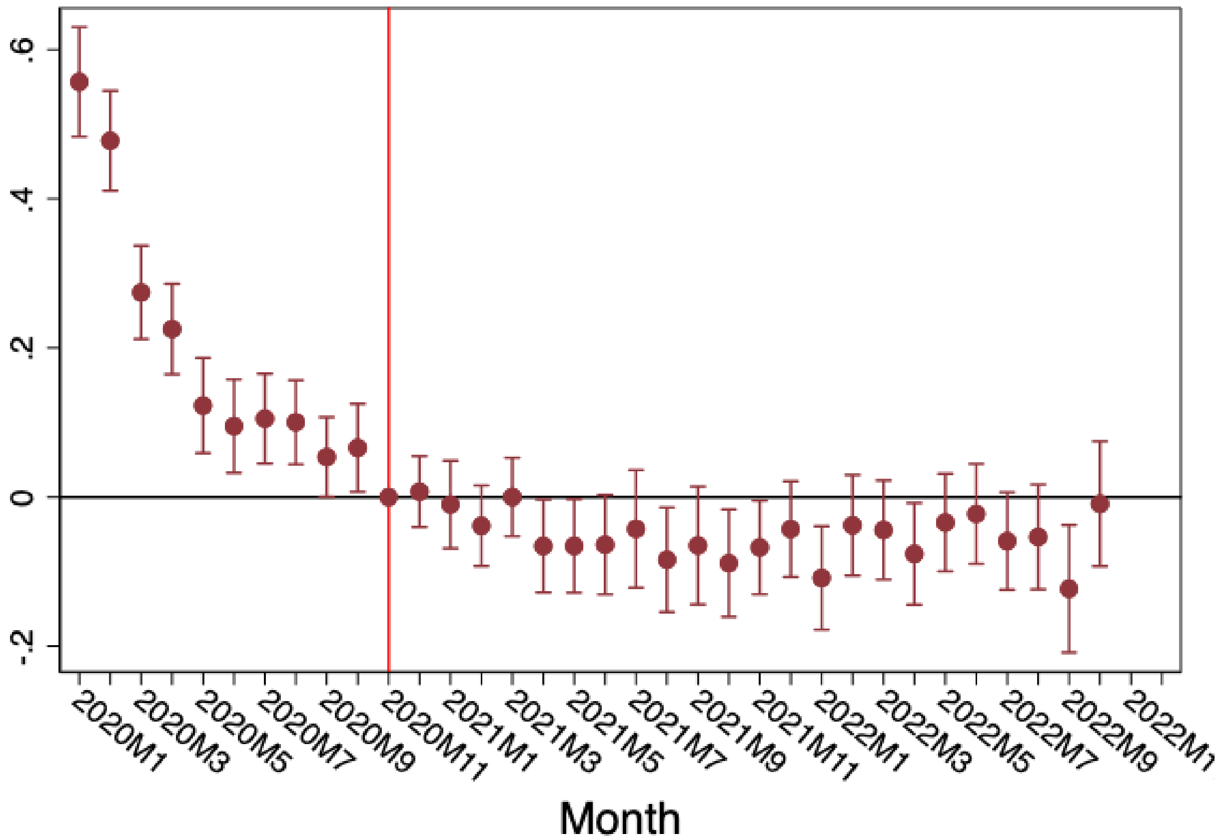

Figure IA.7: Share of households in the sample that have ever used ChatGPT, by whether they used Stack Overflow (or other coding-related subsites of Stack Exchange) or not.

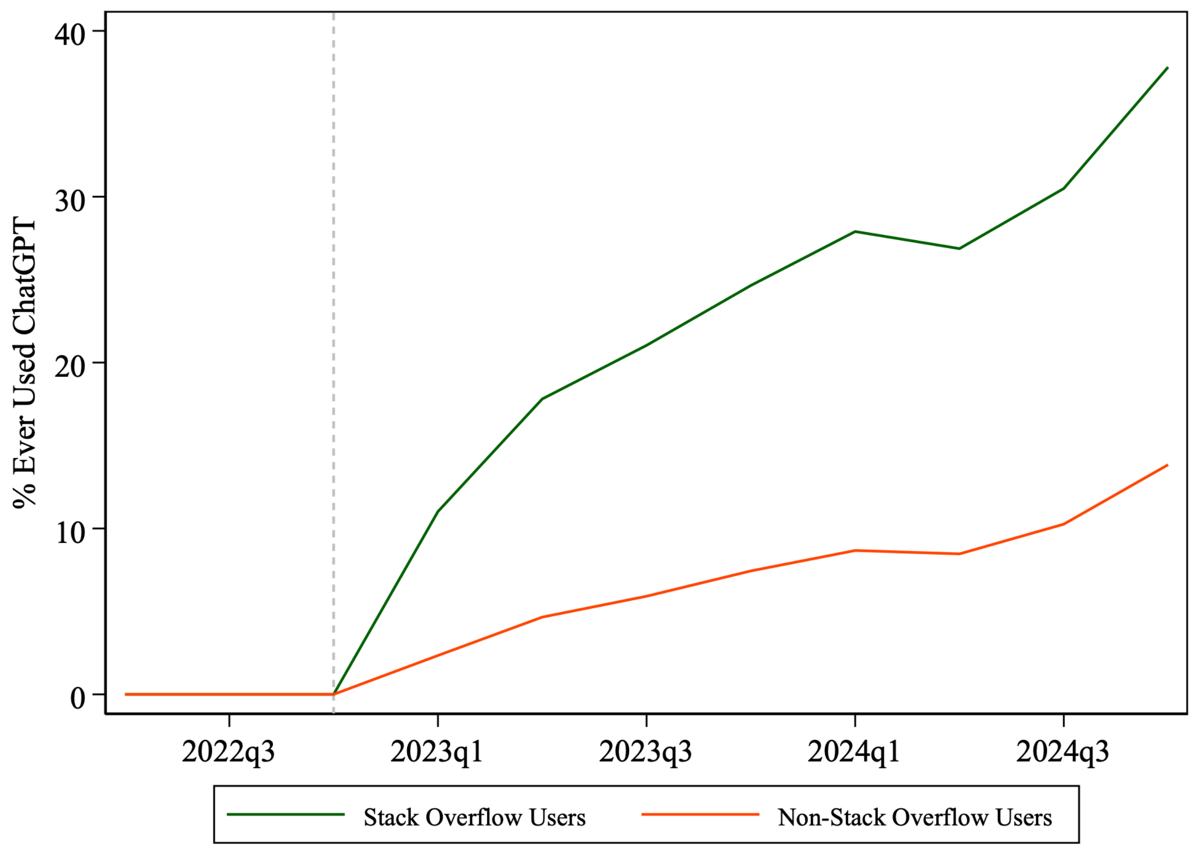

Figure IA.8: OLS estimates of effect of ChatGPT adoption on amount and composition of browsing activity.

.

Panel A: Response of log browsing duration of leisure websites

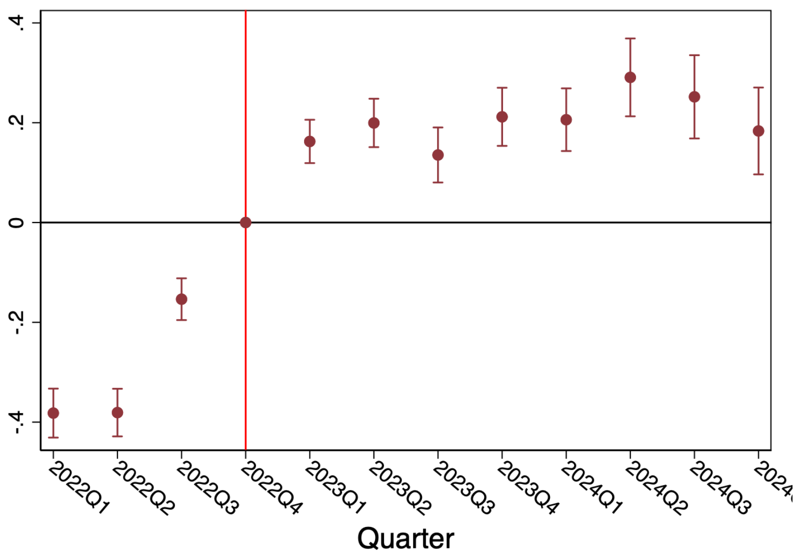

Panel B: Response of log browsing duration of productive websites

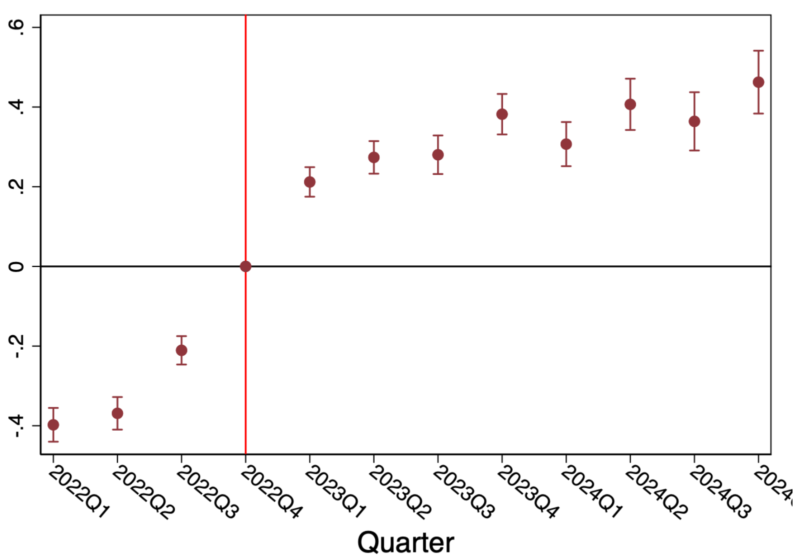

Figure IA.9: Usage of sites associated with different activities in vs. outside ChatGPT adoption windows

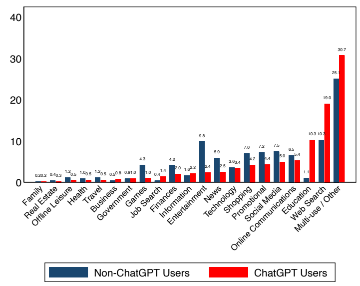

(A) GPT window: browsing shares

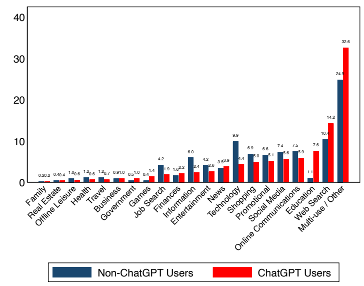

(B) GPT day, outside window: browsing shares

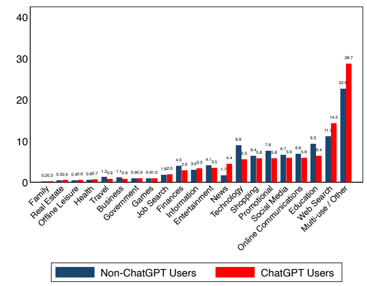

(C) Pre-GPT use: browsing shares

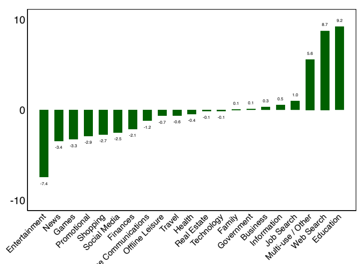

(D) GPT window: differences in shares

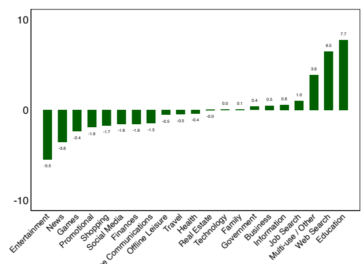

(E) GPT day, outside window: differences in shares

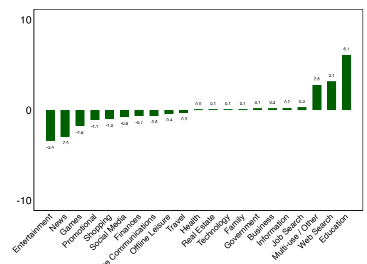

(F) Pre-GPT use: differences in shares

Figure IA.10: Usage of sites with different GenAI overlap in vs. outside ChatGPT adoption windows

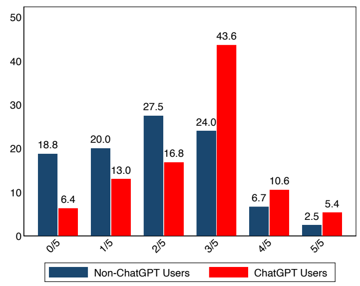

(A) GPT window: browsing shares

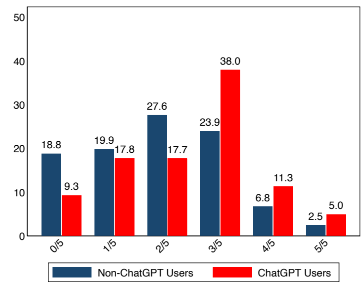

(B) GPT day, outside window: browsing shares

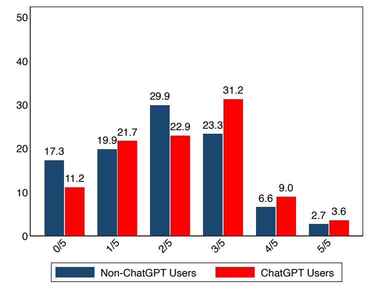

(C) Pre-GPT use: browsing shares

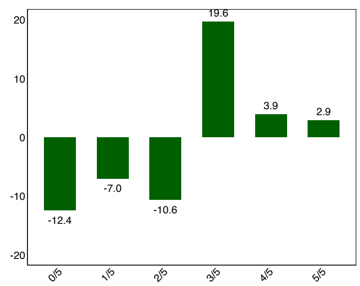

(D) GPT window: differences in shares

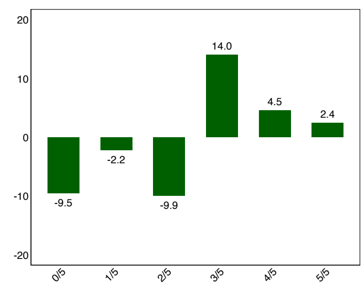

(E) GPT day, outside window: differences in shares

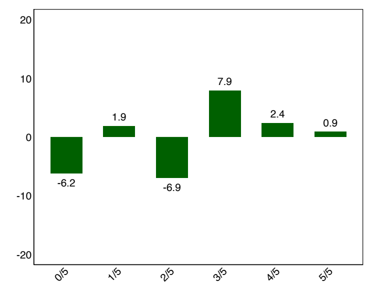

(F) Pre-GPT use: differences in shares