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

Michael Blank† Gregor Schubert‡ Miao Ben Zhang§

February 26, 2026

**Abstract**

This paper studies the impact of generative AI on U.S. households' task allocation at home, using detailed Internet browsing data from a large sample of home devices between 2021 and 2024. Leveraging pre-ChatGPT browsing patterns, we measure households' exposure to ChatGPT and use it as an instrument for ChatGPT adoption during the post-release period. Our IV estimates show that adopting generative AI substantially increases leisure browsing on home devices while leaving time spent on productive digital tasks unchanged. To examine mechanisms, we develop a novel method to infer the purpose of households' ChatGPT use and find that households primarily employ it for productive non-market tasks. Together, these results suggest that generative AI frees up leisure time by raising the efficiency of productive digital activities. A quantitative model implies economically sizable household productivity gains from generative AI.

**Keywords:** *Household Finance, Generative AI, Internet Browsing, Technology Adoption, Productivity, Leisure, Inequality, Digital Divide.*

*We thank participants at OpenAI economic seminar, UCLA finance brownbag, and conference participants at the 2026 AFA special session and the 2026 AEA meeting.

†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 (GenAI) diffuses rapidly into everyday life, it raises first-order economic questions: Which households adopt it, for which tasks, and with what consequences for household productivity, leisure, and inequality? While emerging evidence from surveys and laboratory experiments provides early insights into private GenAI use, it cannot link adoption to the broader set of activities GenAI substitutes for or complements, nor quantify resulting changes in non-market household productivity.

In this paper, we address this gap using detailed Internet browsing microdata from over 200,000 U.S. households' home devices between 2021 and 2024 to provide the first large-scale evidence on the adoption and impact of generative AI on households' non-market activity. Our focus on generative AI's impact "at home" complements existing empirical work on its impact "at work" (Brynjolfsson et al. (2023)) and highlights a central margin of economic adjustment, especially since Americans are more likely to use ChatGPT for personal purposes than for their jobs.[1]

We make five contributions. First, we document a widening "generative AI divide" across U.S. households by income and age, which motivates us to conduct our analysis within demographic bins. Second, we construct a household-level measure of ex-ante exposure to generative AI based on pre-ChatGPT Internet browsing patterns and show that it strongly predicts subsequent adoption of ChatGPT. Third, and most centrally, using the exposure as an instrument for household ChatGPT adoption in a long-difference empirical design, we show that ChatGPT adoption substantially increases leisure browsing on home devices while leaving time spent on productive digital tasks unchanged. Fourth, using the browsing context surrounding ChatGPT site visits, we infer the purpose of household generative AI use and find that private ChatGPT use is concentrated in productive online activities (e.g. education, job search, informational research) rather than leisure. Fifth, by embedding these findings in a quantitative model, we estimate sizable economic gains in non-market household productivity from generative AI adoption.

To study the impact of generative AI tools on private household behavior, we utilize a novel source of data for studying generative AI use by households: Comscore's web browsing panel that

[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.

allows us to comprehensively document the impact of generative AI on household online activities. We observe all browsing on household-owned computers. This data allows us to observe all browsing behavior of a large, diverse panel of U.S. households. A key advantage of this data over information on generative AI generated from surveys, lab settings, or chatbot conversation logs is that it allows us to track household-level changes in behavior over time—before and after the release of ChatGPT—and to observe changes in households' broader digital consumption patterns.

Our first key finding establishes that ChatGPT adoption among private households has been rapid — yet adoption is far from uniform. High-income and younger households adopt generative AI substantially faster than their low-income and older counterparts. Moreover, complementing the static adoption heterogeneity documented in prior survey evidence (Humlum and Vestergaard, 2025), we show that this gap is widening rather than converging over time, giving rise to a substantial and growing "generative AI divide" among U.S. households.

Second, we show that ChatGPT adoption can be in part predicted by households' pre-ChatGPT browsing patterns. To show this, we first develop a 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 January–December 2021, 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 expected benefit of using ChatGPT strongly predicts whether they actually adopt ChatGPT: a 1 SD higher exposure predicts a 2.5pp higher rate of having used ChatGPT by December 2024.

Our third and main finding shows 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" digital 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 tasks households spend more or less time

on after adopting ChatGPT. However, ChatGPT adoption may be endogenous to life events and household task demand. We address this concern by instrumenting adoption with a pre-ChatGPT exposure measure derived from 2021 browsing: households that previously spent more time on online activities that LLMs can substitute are more likely to adopt after November 2022. Conditional on detailed demographic-by-region fixed effects and controls for broad browsing composition, this exposure provides plausibly exogenous variation in adoption.

We find that ChatGPT adoption causes households to spend more time on digital leisure activities while leaving the total time spent on productive online activities (including any time spent using ChatGPT) unchanged. In long-difference IV estimates, ChatGPT adoption raises total leisure browsing time by roughly 150 log points and increases the leisure share of browsing duration by about 30 percentage points. This finding suggests that using ChatGPT for digital household production tasks may substitute for time spent on productive activities on other websites, while making households more efficient in completing those tasks. This leads to a reallocation of freed-up browsing time towards digital leisure activities, which suggests that ChatGPT adoption has positive welfare effects on households.

Fourth, to provide support for this mechanism, we analyze the high-frequency browsing context in which ChatGPT is used, based on the intuition that we can plausibly deduce what activities a person is doing inside ChatGPT by observing their online activities right before and after accessing the chatbot. Comparing the 30-minute browsing intervals around ChatGPT use to the browsing patterns of similar households that never use ChatGPT, we establish that households predominantly utilize ChatGPT in the context of online activities and websites that are related to productive tasks (rather than leisure-oriented). Moreover, we also document that the declines in browsing time as a result of ChatGPT adoption are concentrated in website categories, such as search or news, that are highly exposed to substitution by generative AI. This mechanism's evidence supports the idea that ChatGPT makes productive online tasks more efficient, which then allows households to reallocate freed-up browsing time to leisure.

Fifth, we provide a quantitative model, adapted from the approach in Aguiar et al. (2021), of time allocation by households across digital tasks, which maps our empirical estimates to changes in household productivity. This mapping requires Engel curve elasticities for leisure and productive

digital tasks as an additional parameter input, which we estimate using an IV approach based on exogenous shifts in total online browsing due to local changes in rainfall. We show that our empirical estimates of changes in time allocation in response to ChatGPT imply large increases in the efficiency of households' productive digital tasks as a result of the technology shock.

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 size of the implied productivity changes and the heterogeneity in adoption highlight the importance of better understanding the dynamic effects of using generative AI: if use of the technology outside of work complements workplace productivity or affects human capital development, inequality in initial private adoption can exacerbate disparities in who benefits from generative AI. Our findings highlight the importance of targeting technology literacy efforts, training, or subsidies at groups with low take-up rates (such as older households) to enable a more even distribution of both the market and non-market productivity benefits of the technology.

We note two limitations when inferring our findings. First, our data captures households' Internet activity on home devices but does not observe offline, non-Internet activities. Although Internet use likely accounts for a substantial share of at-home activities—particularly entertainment, information acquisition, and other digital tasks—one should take caution when extrapolating our findings on households' total activities at home. Second, our analysis focuses on generative AI use "at home" and does not connect to workplace exposure. While recent survey evidence suggests that generative AI use outside of work is at least as prevalent as use on the job, understanding how home and workplace adoption interact remains a highly promising avenue for future research.

**Related literature.** Analyzing the household setting complements the existing literature in several ways. First, it provides revealed-preference evidence on what households actually do, complementing early insights on reported generative AI use in surveys (e.g., Bick et al. (2024)) and on interactions with the technology in laboratory settings (e.g., Kosmyna et al. (2025)). Second, our focus on generative AI's impact "at home" complements a growing body of empirical work on its impact "at work" and within firms (e.g., Brynjolfsson et al. (2023); Schubert (2025)). While these studies document substantial disruption to jobs, including unequal productivity gains and job displacement,[2] households' responses to these shocks—such as rebuilding human capital or reallocating time toward leisure—likely occur to a large extent 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 content (Zhang et al., 2025), and household production tasks such as financial planning (Lo and Ross, 2024) and shopping (Lei and Liu, 2025).

As we document, ChatGPT affects household demand for digital services provided on other websites. Several recent studies examine changes in browsing behavior following the release of ChatGPT. Lyu et al. (2025) find that Wikipedia articles more similar to chatbot output experienced larger declines in views after ChatGPT's release. Padilla et al. (2025) show that informational Google searches decline after LLM adoption, while navigational searches do not; they also document declines in online ad exposure and visits to educational websites and Stack Overflow, with little change in visits to social media or Wikipedia. Declines in Stack Overflow engagement are similarly documented by del Rio-Chanona et al. (2024) and Burtch et al. (2024), who interpret these patterns as evidence that users resolve programming questions directly with ChatGPT. Our study complements this literature by examining households' overall browsing reallocation and documenting an increase in the leisure share of online activity.

Our findings on non-market household productivity complement evidence on generative AI's productivity effects in labor markets and firms. In an experiment granting workers access to ChatGPT for professional writing tasks, Noy and Zhang (2023) show that lower-performing workers

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

benefit disproportionately. Management consultants (Dell'Acqua et al., 2023), software developers (Peng et al., 2023), and call center employees (Brynjolfsson et al., 2023) all exhibit productivity gains, again with larger effects for initially less productive workers. Some studies document productivity effects outside traditional workplaces: Yu et al. (2024), for example, find improvements in academic writing proficiency among university students after ChatGPT's release. Jiang et al. (2025) show that generative AI exposure in *market* work may reduce leisure, as more-exposed workers work longer hours. In contrast, we study private household use and its implications for non-market productivity and leisure.

Concerns about unequal adoption also arise in earlier generations of AI. McElheran et al. (2024) document that firm-level AI adoption in 2018 was highly concentrated across cities and regions, describing an "AI divide" across geographies. Humlum and Vestergaard (2025) find in Denmark that younger and less-experienced workers are more likely to use ChatGPT, while women and lower-earning workers are less likely to adopt. We extend this literature by documenting a "generative AI divide" in private household adoption, identifying predictors of these gaps, and quantifying their implications for non-market productivity differences.

Our findings suggest that ChatGPT adoption improves household welfare by raising non-market productivity and enabling greater leisure consumption online. Our conceptual framework for time reallocation in response to task-level productivity shocks builds on Aguiar et al. (2021) and relates to empirical work by Brynjolfsson et al. (2025), who measure the consumer surplus generated by free digital goods.

The remainder of the paper proceeds as follows. Section I describes the Comscore browsing data and our LLM-based website classification and exposure measures. Section II documents adoption patterns, the predictive power of pre-ChatGPT exposure, and the generative AI divide. Section III estimates the causal impact of adoption on browsing behavior using an IV long-difference design, and Section IV provides mechanism evidence from high-frequency browsing context around ChatGPT use. Section V maps the reallocation into implied productivity gains in a simple time-allocation framework, and Section VI concludes.

# I. Data and Measurement

## A. *Comscore Internet browsing data*

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

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]

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 gener-

[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 its computer that tracks Internet usage. Comscore utilizes these data to produce estimates of online usage patterns as well as for other studies of online user behavior conducted internally or by their clients.

ative 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.

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, Comscore 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 January to December 2021 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.

[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 the 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 synthesized data does not.

[5]We define the benchmarking period as ending in December 2021 to ensure households' browsing behavior is not contaminated by any anticipation of the ChatGPT release. This also prevents short-run household shocks from affecting both the benchmarking period and the post-ChatGPT browsing activity.

**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 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.

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

This table shows the distribution with respect to household income and age in both our main Comscore sample (first column), where income and age are from the most recent month prior to November 2022 in which the panelist updated their demographic information, and in the U.S. Census American Community Survey (second column) among all Internet-using households in the 2022 survey.

| | **Comscore sample** | **ACS** |
|---|---|---|
| HH Income (% of Panelists) | | |
| <$25K | 21.25 | 13.94 |
| $25K - $40K | 21.38 | 10.45 |
| $40K - $60K | 13.99 | 14.33 |
| $60K - $75K | 4.67 | 9.89 |
| $75K - $100K | 7.34 | 13.38 |
| $100K - $150K | 9.96 | 17.50 |
| $150K - $200K | 2.49 | 9.19 |
| >$200K | 18.93 | 11.33 |
| Head of HH's Age (% of Panelists) | | |
| 18 - 24 | 7.19 | 4.09 |
| 25 - 34 | 9.78 | 16.14 |
| 35 - 44 | 14.06 | 18.50 |
| 45 - 54 | 38.63 | 17.69 |
| 55 - 64 | 16.68 | 18.64 |
| 65 and over | 13.53 | 24.93 |

## 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 a 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 focus on households' Internet browsing behavior at the monthly frequency. We focus on two measures of a household's browsing intensity on a website in a month. The first is the total number of seconds household $i$ 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, which capture the broad topic of the website's content (e.g., "social media").

**Categorizing productive vs. leisure websites.** A key focus of our study is on whether technology changes households' productivity in different *tasks* that can be accomplished through Internet browsing. In particular, in line with the common breakdown of time use into 'market work', leisure and 'non-market work' activities in time use studies (e.g., Aguiar et al. (2021)), we want to distinguish whether the home browsing activity that we observe is for 'productive' tasks (i.e., non-market work activity on the Internet) or 'leisure'. To the best of our knowledge, there is no existing approach for labeling websites based on the household tasks that they are associated with, so we develop a novel methodology for classifying domains using a large language model (LLM). We focus on the top 160K domains by visits between 2022 and 2024, which account for nearly all browsing activity during this period.

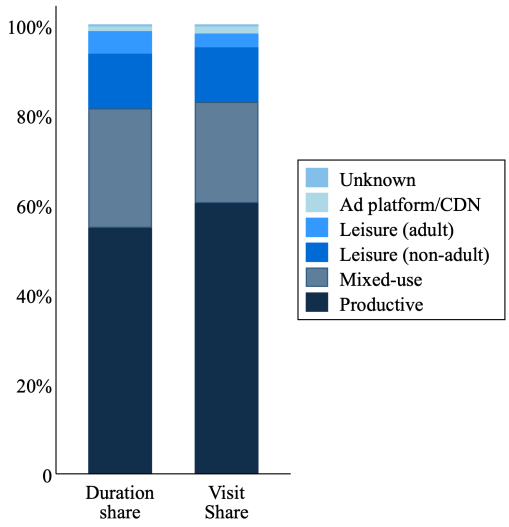
The methodology, as illustrated in Figure IA.1 and detailed in the Internet Appendix IA.1, involves the following three steps: First, for each domain, we obtain a description of each website by using a webscraper to try and access each domain address to retrieve the description and keywords

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 potentially 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.

**Figure 1: Website duration shares by assigned browsing purpose category**

This figure shows the share of browsing visits and duration in each assigned website purpose category in 2021 Comscore data.




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

[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.

ad platforms and content delivery networks as a separate category. These are not customer-facing sites in the sense that we are interested in, and they mechanically generate a large number of page views, so by putting them into a separate category, we ensure that they do not get mislabeled as productive or leisure uses of a household's time.

After scaling this process to the full set of 160K domains, we derived a final label for each domain as having a main use that is *Productive*, *Leisure*, *Mixed* or *Ad platforms/CDN*. Based on this categorization, we can 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. The distribution of total browsing activity over these purpose categories is shown in Figure 1: both by visits and by browsing duration, the majority of browsing activity is classified as 'productive', followed by 'mixed-use' activity. For illustration, we also split out adult vs. non-adult content leisure sites (but otherwise combine these in the rest of the paper). To illustrate what these labels mean, we break out the top 5 websites by duration for the 3 largest of these categories in Table II.

**Table II:**
**Top 5 domains by duration by assigned browsing purpose category in 2021**

This table shows the 5 biggest domains by duration in each of the three largest assigned browsing purpose categories in 2021 Comscore data.

| Domain | % of Category | % of All Sites | Automated description |
|---|---|---|---|
| **Productive Duration Top 5** | | | |
| google.com | 24.97 | 13.69 | widely used search engine |
| amazon.com | 3.74 | 2.05 | comprehensive online retail platform |
| bing.com | 3.50 | 1.92 | smart search engine |
| instructure.com | 1.95 | 1.07 | educational technology company |
| live.com | 1.93 | 1.06 | Microsoft-operated web portal |
| **Mixed-use Duration Top 5** | | | |
| youtube.com | 51.02 | 13.59 | video-sharing platform |
| msn.com | 33.78 | 9.00 | comprehensive web portal |
| yahoo.com | 14.24 | 3.79 | comprehensive web portal |
| playtika.com | 0.60 | 0.16 | gaming company |
| lycos.com | 0.15 | 0.04 | comprehensive web portal |
| **Leisure (non-adult) Duration Top 5** | | | |
| facebook.com | 15.48 | 1.89 | social networking platform |
| instagram.com | 4.09 | 0.50 | social media platform |
| twitch.tv | 1.91 | 0.23 | live streaming platform |
| espn.com | 1.87 | 0.23 | comprehensive sports media website |
| twitter.com | 1.48 | 0.18 | social media platform |

## C. *Measuring household exposure to GenAI*

**Website activity exposure to generative AI..** To evaluate the degree to which a household's Internet use patterns are 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]. The LLM assigns each activity a binary exposure label based on whether ChatGPT can perform or complement the website's tasks. The distribution of the types of activities that tend to be labeled as "exposed" by this method is shown in Figure 7.

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 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) are dominated by social media platforms like Facebook, Instagram, and Twitch; video streaming services like Netflix and YouTube, and other service sites like AOL, Craigslist, and Weather.com. The high end of the exposure score spectrum (4 or 5 out of 5) is dominated by education and knowledge sites, like Wikipedia, Quizlet, or Edgenuity. 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), or platforms for shopping and trading, such as Amazon, Robinhood, and eBay.

We show the distribution of these labels in Figure IA.3: 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

[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/

[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.

**Table III:**
**Top 5 websites by GenAI exposure**

This table shows the top 5 websites by browsing duration in each GenAI exposure category defined by how many of the top 5 website activities are exposed to generative AI.

| 0/5 Exposure | | 1/5 Exposure | | 2/5 Exposure | |
|---|---|---|---|---|---|
| Website | Share | Website | Share | Website | Share |
| facebook.com | 25.18 | youtube.com | 57.18 | msn.com | 33.84 |
| weather.com | 4.24 | aol.com | 4.34 | yahoo.com | 14.27 |
| twitch.tv | 3.11 | instagram.com | 2.10 | amazon.com | 7.71 |
| pluto.tv | 1.68 | craigslist.org | 1.14 | bing.com | 7.23 |
| netflix.com | 1.45 | veryfast.io | 0.89 | robinhood.com | 2.91 |

| 3/5 Exposure | | 4/5 Exposure | | 5/5 Exposure | |
|---|---|---|---|---|---|
| Website | Share | Website | Share | Website | Share |
| google.com | 57.03 | intuit.com | 6.83 | wikipedia.org | 7.93 |
| instructure.com | 4.46 | office.com | 2.93 | gingersoftware.com | 5.45 |
| live.com | 4.40 | edgenuity.com | 2.93 | wwnorton.com | 3.06 |
| ebay.com | 2.82 | indeed.com | 2.91 | quizlet.com | 2.81 |
| microsoft.com | 1.33 | reddit.com | 2.28 | mathxl.com | 2.34 |

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. Panel A of Appendix Figure IA.2 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.2 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.

**Household exposure to generative AI.** The exposure label for each household is then derived from the exposure of the domains they visit, weighted by the time spent on each site. 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 January to December 2021 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 of highly generative AI-exposed websites in its Internet browsing. That is, we compute

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

where $E_j$ is the count of activities (out of 5) on website $j$ that are generative AI exposed, and $\phi_{ij}$ 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 their 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*

To capture when a household first starts using generative AI technology, we focus on ChatGPT, the most popular LLM chatbot. ChatGPT use is a good 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 has *ever* used 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. We deliberately do not require a household to continuously use ChatGPT, as it is possible that households eventually either switch to other generative AI tools, or access ChatGPT through means other than browsing on a home computer (e.g., through a cellphone app)—but we would not want to label those households as not using generative AI.

Our second measure of interest is the intensity of the household's usage of ChatGPT, which is the number of seconds the household browsed the ChatGPT website within the time period.

# 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.

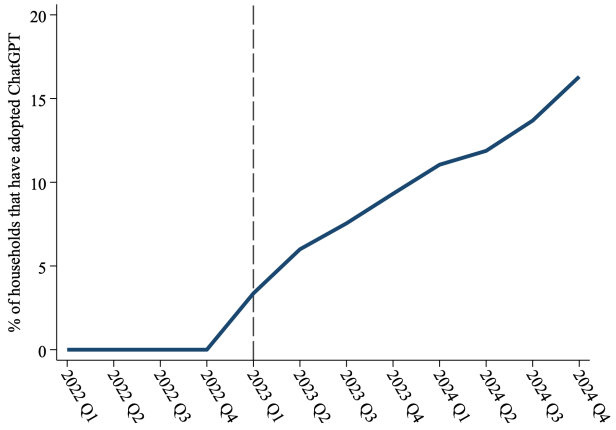
## A. *Household adoption of ChatGPT over time*

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.

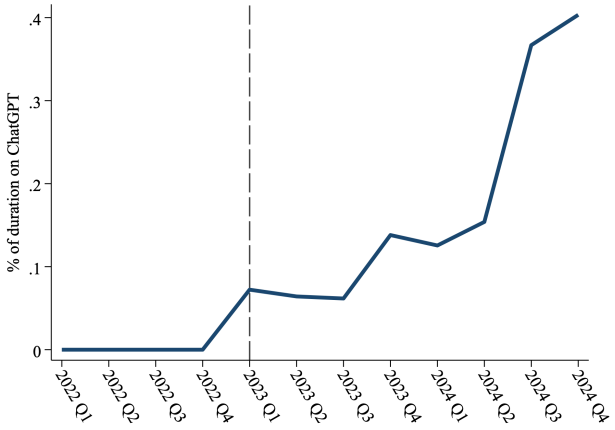
**Household ChatGPT adoption over time.** We plot the time series of ever using ChatGPT and the share of browsing time spent using the chatbot in Figure 2: the share of all households that have ever used ChatGPT rises from zero to 9.3 pp in Q4 2023 to 16.3 pp by Q4 2024, almost doubling within a year. The share of total browsing time spent on ChatGPT also captures the intensive use margin and shows an even steeper increase during 2024, rising from 0.14 pp in Q4 2023 to 0.40 pp by Q4 2024 (see Table IV for summary statistics).

**Figure 2: Household ChatGPT adoption over time**

This figure shows the share of households that have ever used ChatGPT (panel A) and the share of total browsing duration that is spent using ChatGPT (panel B) based on the Internet browsing activity in Comscore data.




**(A)** Share that has used ChatGPT




**(B)** Share of browsing duration on ChatGPT

**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 the measure of having used ChatGPT shows both similar rapid growth patterns and orders of magnitude of usage is reassuring.

**Table IV:**
**ChatGPT exposure and use over time in browsing activity: 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 Table VI.

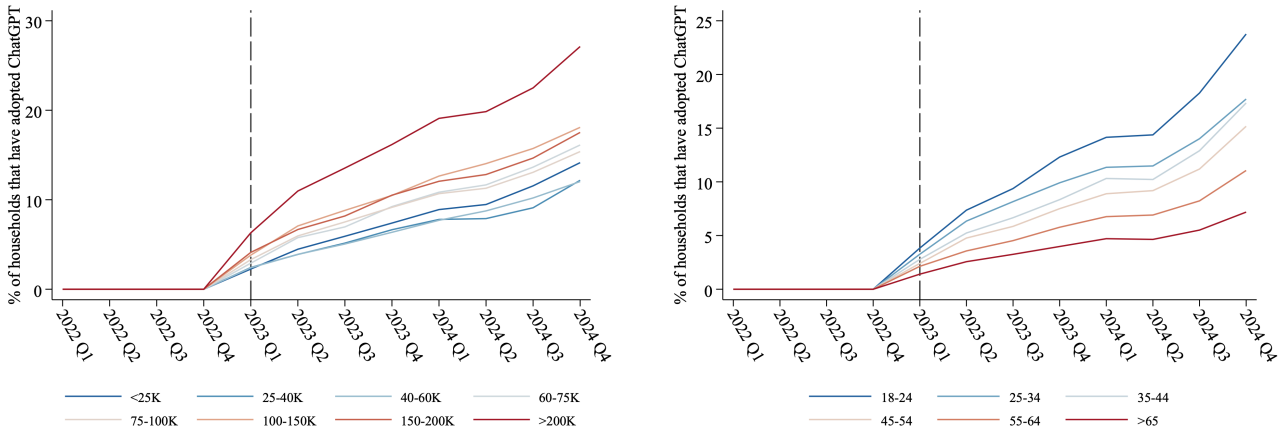
| | | By Income | | | By Age | | |
|---|---|---|---|---|---|---|---|
| | All | Low | Middle | High | Young | Middle | Old |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| HHGenAIExp % | 10.7 | 10.0 | 10.5 | 11.2 | 11.7 | 11.2 | 9.9 |
| *Used ChatGPT %* | | | | | | | |
| Q4 2023 | 9.3 | 7.1 | 8.1 | 12.0 | 14.1 | 11.3 | 6.3 |
| Q2 2024 | 11.9 | 8.8 | 10.4 | 15.3 | 18.0 | 14.5 | 7.9 |
| Q4 2024 | 16.3 | 13.3 | 14.3 | 20.3 | 24.1 | 21.0 | 10.6 |

**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). In panel A, we find that higher income households (> $200K) are substantially more likely to have used ChatGPT since its release. While the relationship is not precisely monotonic,

it is generally true that higher income categories tend to have higher rates of 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.

**Figure 3: ChatGPT adoption by household income and age**

This figure shows the share of households in each demographic group that have ever used ChatGPT.




**(A)** ChatGPT adoption by income bucket

**(B)** ChatGPT adoption by age bucket

**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 time series data is that a similar gap exists in the use of generative AI at the household level and low adoption is associated with other indicators of socioeconomic disadvantage, such as low income or old age. Large differences in generative AI adoption among households raise the specter of a persistent "Generative AI Divide". Moreover, the Comscore data allows us to track time trends in this adoption gap for a comparable population of households: the gap in adoption between high and low

income households, and between young and old households, is, if anything increasing over time. This increasing divide means that any productivity benefits of generative AI are also unequally distributed—which motivates our attempt to quantify the productivity effect in people's private lives.

## *B. Household exposure and ChatGPT adoption*

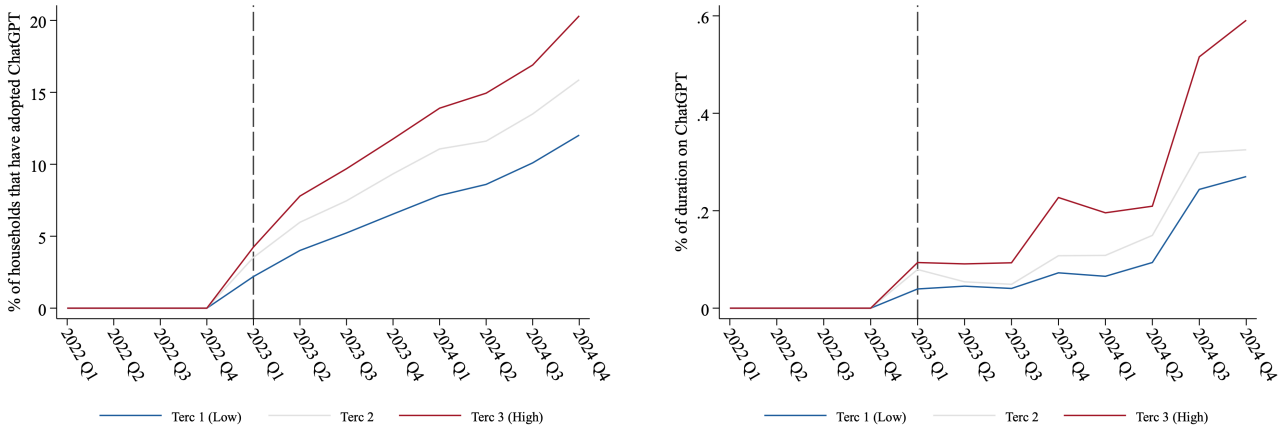
To understand whether households that are more likely to benefit from using ChatGPT are more likely to adopt the technology, we first consider whether greater exposure based on a household's pre-ChatGPT browsing behavior predicts greater adoption. This evidence motivates the empirical approach in the next section for estimating the causal effect of adoption on household behavior.

**Time patterns by exposure.** The relationship between predicted household benefits based on browsing behavior—"household GenAI exposure"—and adoption of ChatGPT can be seen in the raw time series: Figure 4 plots adoption rates by tercile 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 terciles aligning with the ordering of the terciles by exposure. The cumulative gap in having used ChatGPT is large: the likelihood of having used ChatGPT is 20pp by Q4 2024 in the highest exposure tercile households compared to 12pp in the lowest exposure tercile. Similarly, the share of browsing time spent using ChatGPT is 0.59pp for highly exposed households compared to 0.27pp among low exposure households in Q4 2024. That is, the adoption gap by exposure increases over time, with little sign of convergence: at least through 2024, it is true that households with greater exposure do not just adopt earlier than other households but also at a higher rate.

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

**Figure 4: ChatGPT adoption by household GenAI exposure**

This figure shows the share of households in each generative AI exposure tercile—with the 3rd tercile having the highest exposure—that have ever used ChatGPT (Panel A), and also the share of total browsing duration that is spent using ChatGPT in each tercile of exposure (Panel B).




**(A)** Share that has used ChatGPT

**(B)** Share of browsing duration on ChatGPT

AI among private households.

**Effect of household exposure on GenAI adoption.** To estimate the size of the effect of exposure on both the intensive and extensive margin of ChatGPT adoption, we estimate regressions of the form

$$\text{ChatGPT Use}_{j,t} = \theta \ln(\text{HHGenAIExp}_j) + X'_{j,t}\xi + \text{FE}_{j,t} + \varepsilon_{j,t} \tag{2}$$

where $\text{GenAIExp}_j$ is the household's pre-ChatGPT benchmark period (Jan.-Dec. 2021) exposure to generative AI, and $X_{j,t}$ is a vector of household characteristics. We include saturated interactions of income (8 categories), age (6 categories), and metro area as fixed effects in all cross-sectional household level regressions to capture the effect of different demographics and regional differences. We also control for the topic-based exposure to generative AI of the household - the component that is explained by the average exposure of websites in particular content categories, and the share of the household's total browsing activity in the benchmark period for which we can compute exposure scores.

We measure *extensive margin* adoption effects by using a dependent variable that reflects whether the household has ever used ChatGPT by 2024, such that the estimated coefficient on exposure indicates the degree to which exposure leads to adoption. To quantify the *intensive mar-*

*gin* of ChatGPT adoption, we estimate the effect on the average share of total browsing duration spent using ChatGPT during 2024.

The results are shown in Table V, where columns (1) and (2) show the extensive and intensive margin effects of household GenAI exposure on adoption. The coefficients in columns (1) and (2) show that a 1 SD (~1.3 unit) change in log household exposure is associated with about a 2.5pp increase in the likelihood of having used ChatGPT by 2024, and a 0.07pp increase in the share of total browsing time spent using ChatGPT. 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 adoption patterns across demographic groups that we previously documented.

The next section considers whether these changes in GenAI use as a result of a household's exposure lead to changes in a household's behavior that suggest productivity gains.

**Table V:**
**Effect of household GenAI exposure on ChatGPT adoption and browsing by purpose**

This table shows estimates of long difference reduced-form regressions comparing 2024 browsing characteristics to 2022 browsing characteristics, in an estimation of the form:

$$\Delta\text{BrowsingOutcome}_i = \gamma \ln(\text{HHGenAIExp}_i) + \lambda_{FEs} + X_i'\xi + \varepsilon_i.$$

The dependent variable in Column (1) is a dummy for whether household $j$ has ever used ChatGPT by the end of 2024. The independent variable is the log of the panelist household's predicted exposure to ChatGPT in 2021, based on the household's share of browsing on highly substitutable websites. The dependent variable in Column (2) is the share of the household's total browsing duration that is spent on ChatGPT; in Columns (3),(4), (5) and (6) it is the change in the log of browsing duration on all, leisure, productive, and mixed-use sites. All regressions include fixed effects for income bin X age bin X MSA. All regressions also include control variables for the generative AI exposure predicted by the households 2021 browsing composition by Comscore content category, and for the sum of the browsing share of websites with a GenAI exposure label in 2021. T-statistics based on heteroskedasticity-robust standard errors clustered at the income bin X age bin X MSA level are shown in parentheses: * p<0.10, ** p<0.05, *** p<0.01.

| *Specification:* | *Long differences: 2024 vs. 2022* | | | | | |
|---|---|---|---|---|---|---|
| | *GenAI adoption measures* | | *Browsing activity by purpose* | | | |
| *Dep. var.* | 𝟙[Used ChatGPT] | ChatGPT Duration Share % | Δ Log Duration All sites | Δ Log Duration Leisure | Δ Log Duration Productive | Δ Log Duration Mixed |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| ln(HHGenAIExp) | 0.019*** | 0.053*** | 0.024** | 0.029** | 0.000 | -0.006 |
| | (10.180) | (5.404) | (2.421) | (2.087) | (0.021) | (-0.452) |
| Observations | 42,886 | 42,886 | 42,886 | 42,886 | 42,886 | 42,886 |
| Estimation | OLS | OLS | OLS | OLS | OLS | OLS |
| Inc. X Age X MSA FEs | X | X | X | X | X | X |
| Bartik share control | X | X | X | X | X | X |
| Website category exposure | X | X | X | X | X | X |

# III. The Household Impact of Generative AI Adoption

The previous section established 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 do households spend more or less time on after they start using generative AI?

## *A. Empirical design*

We are 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. Conversely, 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.

**Empirical specification.** In general, we want to estimate empirical specifications similar to the following expression:

$$\Delta\text{BrowsingOutcome}_{i,t} = \gamma\, \text{ChatGPTUse}_{i,t} + \lambda_{FEs} + X'_{i,t}\xi + \varepsilon_i \tag{3}$$

Here, the dependent variable in each regression might, for example, be a measure of the household's browsing activity in different categories. We want to estimate the effect of a household $i$'s use of ChatGPT on the characteristics of the household's browsing activity.

**Instrumental variable identification.** We want to obtain an estimate of the effects of generative AI use in this empirical setting 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? The key identification concern is that adoption of ChatGPT may be endogenously driven by recent life events that may then also change online browsing behavior: for example, a recent job loss may lead an unemployed worker to use ChatGPT for drafting cover letters, while also devoting more time to online job search and to online education to brush up on in-demand skills, but the latter changes in browsing behavior should not be causally attributed to generative AI adoption.[9]

To arrive at a causal estimate of the effects of generative AI technology, we therefore need to identify variation in adoption that is plausibly exogenous with regard to other short-term changes in a household's life circumstances that might be driving browsing behavior.

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 exclusion restriction: Once we condition on observable sociodemographic characteristics and broad differences in pre-ChatGPT release browsing behavior, variation in GenAI exposure is the result of differences among otherwise similar households in the precise tasks done through online browsing, which are unlikely to drive short-run life events years later when the household starts using ChatGPT.

We therefore implement this IV approach by using the household-level GenAI exposure measure defined in equation (1) computed over the period from January 2021 to December 2021. This historical browsing behavior predicts which households utilize ChatGPT following its November 2022 release. Under the plausible assumption that the 2021 browsing behavior on which our measure of household GenAI exposure is based, is exogenous with regard to other post-2022 shocks to browsing behavior, this IV approach identifies the causal effects of generative AI on browsing

[9]Similarly, 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 the effects of ChatGPT use.

activity (conditional on household-level fixed effects and detailed demographic controls).

We estimate IV regressions corresponding to the two-stage system in which the second-stage equation is given by equation 3 and the first stage corresponds to equation 2 introduced in the previous section, which takes the form

$$\text{ChatGPT Use}_{j,t} = \theta \ln(\text{HHGenAIExp}_j) + X'_{j,t}\xi + \text{FE}_{j,t} + \varepsilon_{j,t}. \tag{4}$$

**Long difference sample.** In order to smooth out some of the noise from short-run fluctuations in browsing activity and account for the fact that Section II showed that ChatGPT use increased substantially only by 2024, we estimate long-difference regressions, comparing household browsing activity (i.e., machines) in 2024 to activity in 2022 before ChatGPT was released. This focuses our analysis on broad changes in behavior rather than short-run fluctuations and further reduces the possibility that the measurement period for our instrument (Jan.-Dec. 2021) is affected by household life event shocks that also drive digital activity in 2024.

**Control variables.** The long difference estimation includes age bucket×income bucket×region fixed effects. All regressions also include control variables for the generative AI exposure predicted by the households' 2021 browsing composition by Comscore content category, and for the sum of the browsing share of websites with a GenAI exposure label in 2021 - the 'Bartik share control'.

**Dynamic effect estimates.** To show how the effect of ChatGPT adoption changes over time and address the 'pre-trends' concerns often associated with this event study-style estimation, we also estimate dynamic quarterly panel IV regressions of the form

$$\text{Browsing outcome}_{j,t} = \Sigma_t \gamma_t^{IV} \ln\left(\text{HHGenAIExp}_{j,t}\right) \times \mathbb{1}[\text{Quarter} = t]_t + \lambda_j + \lambda_t + X'_{j,t}\xi + \varepsilon_{j,t},$$

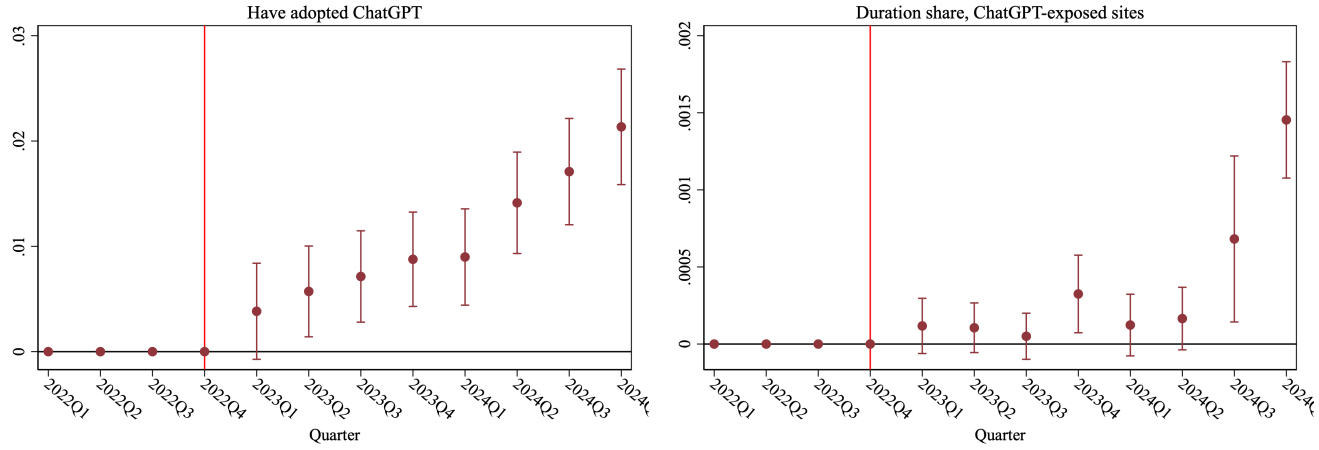
where the left-hand side is a time-varying browsing outcome of household $j$ in quarter $t$, $HHGenAIExp_{j,t}$ is the household's 2021 browsing exposure to GenAI. The dynamic effects estimation estimates a reduced-form effect because it would otherwise not be possible to estimate effects before ChatGPT is released (as the endogenous ChatGPT adoption variable is zero for everyone). We

**Figure 5: ChatGPT adoption effects of Household GenAI exposure over time**

This figure shows the results of estimating specifications of the form

$$\text{Browsing outcome}_{j,t} = \sum_{h=-4}^{8} \gamma_h^{RF} \ln\left(\text{HHGenAIExp}_j\right) \times \mathbb{1}\{t = 2022Q4 + h\} + \lambda_j + \lambda_t + X'_{j,t}\xi + \varepsilon_{j,t},$$

where the dependent variable in panel A is an indicator for whether the household has ever used ChatGPT by that time period, and in panel B it is the share of total browsing time spent using ChatGPT.




**(A)** Effect on share that has used ChatGPT

**(B)** Effect on duration share of ChatGPT use

use Q4 2022 as the pre-release reference period.[10] These panel regressions include household fixed effects, as well as quarter×age bucket×income bucket×region fixed effects.

## *B. Main findings: household browsing effects of generative AI adoption*

**Reduced-form dynamic effects.** To understand the time pattern of ChatGPT effects on households, we first estimate a version of the reduced-form dynamic effect estimation that corresponds to the 'first-stage' of the IV regression: Figure 5 shows the effects of household GenAI exposure on the use of ChatGPT over time. As the figure shows, households with greater exposure are increasingly more likely to have used ChatGPT, and also spend an increasing amount of time on the ChatGPT site. However, as the time patterns show—and in line with the raw adoption patterns in Section II—the effects only become more pronounced by the end of 2024. This is the reason why we will estimate the main behavioral effects parameters using the long difference design.

[10] 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.

How does household browsing activity change in response to household GenAI exposure? Figure 6 shows the effects of household GenAI exposure on the log of browsing time in different purpose categories, focusing on leisure vs. productive online activities. Panel (A) shows that there is an increase in total duration Panel (B) shows that households substantially increase their browsing time on leisure-oriented websites. This is especially the case starting in 2024, consistent with the pattern of when ChatGPT adoption in response to exposure surges as well. In comparison, Panels (C) and (D) show that exposure to generative AI does not lead to a similar increase in browsing on productivity-oriented websites and on mixed purpose sites. These dynamic effect estimates align with the reduced-form long-difference effects in columns (3)-(6) of Table V, which also show increases in total browsing and leisure browsing and no significant changes in time spent on productivity-related or mixed-purpose sites. To interpret the magnitudes, this suggests that a 1 SD change in log exposure increases total browsing by $0.024 * 1.3 \approx 3\%$ and leisure browsing by $0.029 * 1.3 \approx 4\%$.

**Pre-trends.** The dynamic effects are also reassuring with regard to the possibility that more GenAI-exposed households may have unobservable different trends in browsing behavior that precede the ChatGPT release: all 4 panels of Figure 6 show no significant differences in browsing behavior for exposed households in the year *before* ChatGPT is released. Moreover, the changes in behavior found after the ChatGPT release concentrate in the later quarters and do not appear in the initial quarters after the release when actual adoption rates were still low.

**Baseline effect estimates.** Table VI presents our long-difference IV regressions to estimate the effect of ChatGPT adoption on changes in browsing behavior between 2022 and 2024. In line with the dynamic reduced-form estimates in Figure 6, the amount of browsing activity devoted to leisure-oriented sites and to total browsing increases in an economically and statistically significant way in 2024. At face value, the estimates suggest that total leisure-related digital time increases by a factor of ~ 4.5 after households adopt ChatGPT.[11] We will discuss how to interpret this magnitude more in the next section. Here, the relevant takeaway is that ChatGPT adoption significantly

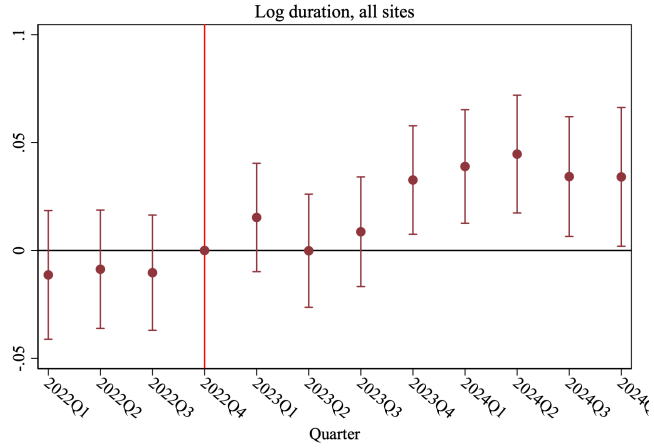
[11] Computed as $exp(1.512) \approx 4.5$.

## Figure 6: Browsing behavior effects of household GenAI exposure over time

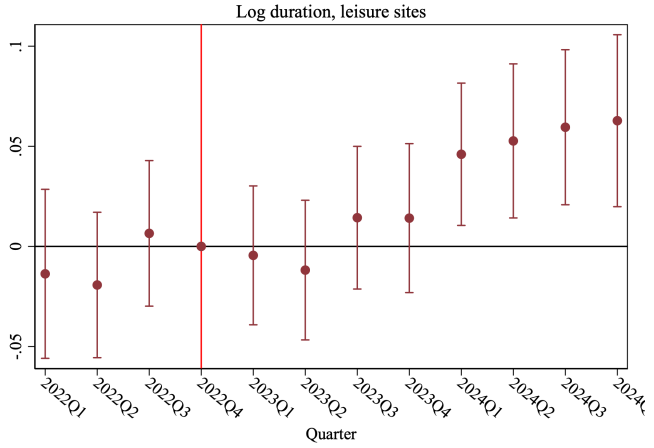
This figure shows the results of estimating specifications of the form

$$\text{Browsing outcome}_{j,t} = \sum_{h=-4}^{8} \gamma_h^{RF} \ln\left(\text{HHGenAIExp}_j\right) \times \mathbb{1}\{t = 2022Q4 + h\} + \lambda_j + \lambda_t + X'_{j,t}\xi + \varepsilon_{j,t},$$

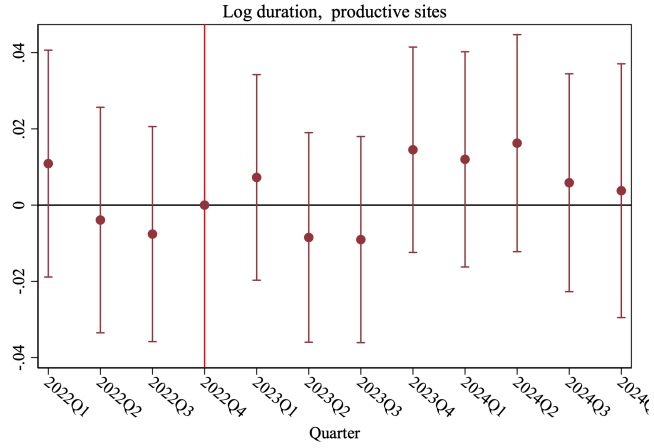
where the dependent variable in panel A is the log of the household's total browsing time, and in panels B, C and D, it is the log of total browsing time spent on leisure, productive, or mixed purpose sites. Confidence intervals shown correspond to 95% CIs based on heteroskedasticity-robust standard errors clustered at the age bucket×income bucket×region level. These panel regressions include household fixed effects, as well as quarter×age bucket×income bucket×region fixed effects.



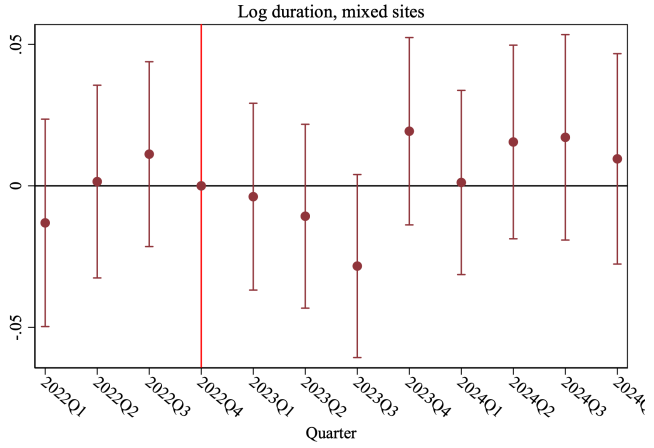

**(A)** Log browsing duration, all sites




**(B)** Log browsing duration, leisure sites




**(C)** Log browsing duration, productive sites




**(D)** Log browsing duration, mixed sites

increases leisure browsing and total browsing, but productive and mixed-purpose browsing levels (columns 3 and 4) do not change significantly. Columns (5)-(7) show that this implies a significant increase in the *share* of browsing time spent on leisure sites, a decrease in the share of time spent on productive sites, and a negative but non-significant decline in the share spent on mixed-purpose sites.

In the next section, we adapt a household time allocation model to show that these estimates can be used to infer the size of the productivity effect of ChatGPT on households' different digital activities.

**Table VI:**
**Effect of ChatGPT adoption on browsing activity by category**

This table shows estimates of long difference IV regressions comparing 2024 browsing characteristics to 2022 browsing characteristics, in an estimation of the form

$$\Delta \text{BrowsingOutcome}_i = \gamma \, \text{ChatGPTUse}_i + \lambda_{FEs} + X_i'\xi + \varepsilon_i.$$

The endogenous variable is a dummy for whether household $j$ has ever used ChatGPT by the end of 2024. The instrument is the panelist household's predicted exposure to ChatGPT in 2021, based on the household's share of browsing on highly substitutable sites. The dependent variable in each regression is a measure of the household's browsing activity in different categories. In Column (1) it is the quarterly log of browsing duration across all sites; in Columns (2),(3) and (4) it is the quarterly log of browsing duration on leisure, productive, and mixed sites; and in Columns (5), (6), and (7) the change in the share of the household's total browsing that is on leisure, productive, and mixed sites. All regressions include fixed effects for income bin X age bin X MSA. All regressions also include control variables for the generative AI exposure predicted by the households' 2021 browsing composition by Comscore content category, and for the sum of the browsing share of websites with a GenAI exposure label in 2021 - the 'Bartik share control'. T-statistics based on heteroskedasticity-robust standard errors clustered at the income bin X age bin X MSA level are in parentheses: * p<0.10, ** p<0.05, *** p<0.01.

| *Specification:* | *Long differences: 2024 vs. 2022* | | | | | | |
|---|---|---|---|---|---|---|---|
| *Dep. var.* | Δ Log Duration All | Δ Log Duration Leisure | Δ Log Duration Productive | Δ Log Duration Mixed | Δ Share Leisure | Δ Share Productive | Δ Share Mixed |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| ChatGPT Use | 1.243** | 1.512** | 0.011 | -0.285 | 0.307*** | -0.215*** | -0.092 |
| | (2.409) | (2.106) | (0.021) | (-0.449) | (4.103) | (-2.738) | (-1.609) |
| Observations | 42,886 | 42,886 | 42,886 | 42,886 | 42,886 | 42,886 | 42,886 |
| 1st-stage KP F-stat. | 104 | 104 | 104 | 104 | 104 | 104 | 104 |
| Estimation | IV | IV | IV | IV | IV | IV | IV |
| Inc. X Age X MSA FEs | X | X | X | X | X | X | X |
| Bartik share control | X | X | X | X | X | X | X |
| Website category exposure | X | X | X | X | X | X | X |

# IV. 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 while leaving the productive task share unchanged. We provide support for our interpretation of this empirical pattern as showing that ChatGPT primarily raises the efficiency of productive (non-market) digital tasks, freeing up time that is reallocated toward leisure digital tasks. Importantly, this interpretation assumes that the

leisure increase does not result from ChatGPT itself mainly being used to support leisure browsing.

We provide analyses to better understand this mechanism that answer the following questions: (1) Which types of websites are households using less or more after they adopt generative AI? (2) What types of websites and digital tasks (productive vs. leisure) are households engaging with in the context of generative AI use (right before or after using ChatGPT)?

## *A. Effects of ChatGPT adoption on browsing by website content category*

This section shows which categories of websites, based on their Comscore content category, experienced an increase or decrease in their share of browsing activity as a result of ChatGPT adoption. The results are shown in Table VII. As the table shows, households that adopt ChatGPT decrease their browsing time spent on search-related websites (column 1). They also decrease the time spent on news sites (column 11), while increasing the time spent on gaming sites (column 13). Note that these findings directly align with our proposed mechanism: categories of websites that are likely used for productive Internet activities, such as search or news consumption, see a decline in their relative share. Importantly, these are the types of websites that generative AI chatbots are most likely to substitute for. At the same time, the increase in website activity that we observe for gaming sites is unlikely to result from chatbots being directly useful for increasing the productivity of leisure activities related to gaming, as there is not an obvious overlap between basic ChatGPT capabilities and gaming tasks.

**GenAI-exposed activities.** For comparison, we also assign each activity on a website that is labeled as 'exposed' by our LLM classification pipeline into a broader category based on the type of activity that was labeled. The distribution of activities labeled as exposed over these categories is shown in Figure 7. As the graph shows, more than half of the activities that we determine ex ante to be most exposed to the substitution by generative AI fall into the category of 'Information & Research', which closely aligns with the search and news website categories for which we estimate that ChatGPT adoption actually causes a drop in browsing share.

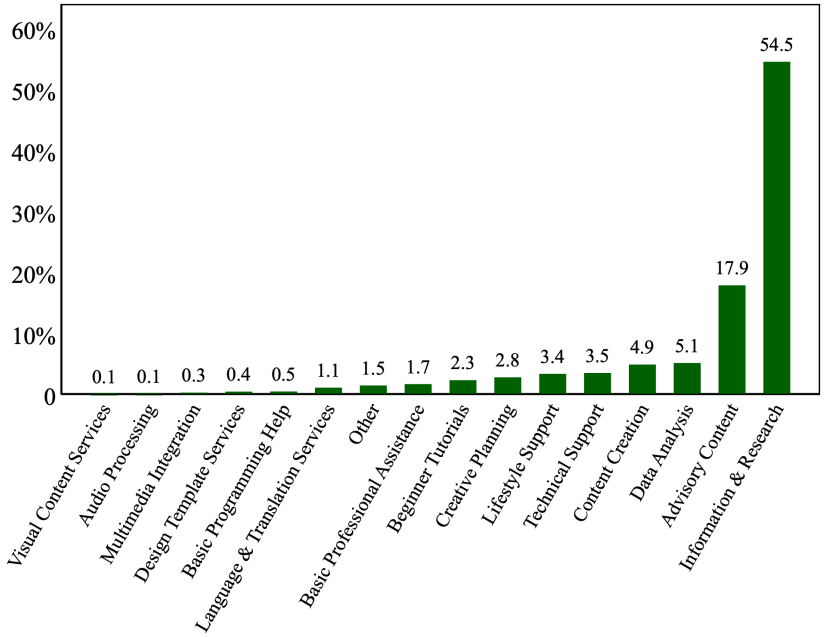
## Table VII:
## Effect of ChatGPT adoption on browsing activity by Comscore category

This table shows estimates of long difference IV regressions in the same specification as in Table VI, but with the change in the share of browsing in different Comscore content categories from 2022 to 2024 as the dependent variable. T-statistics based on heteroskedasticity-robust standard errors clustered at the income bin X age bin X MSA level in parentheses: * p<0.10, ** p<0.05, *** p<0.01.

| *Dep. var.:* | *Long difference change in browsing share: 2024 vs. 2022* | | | | | | |
|---|---|---|---|---|---|---|---|
| *Panel A: Productivity-related content categories* | | | | | | | |
| | (1) Search | (2) Educat. | (3) Shopping | (4) Technology | (5) Information | (6) Finances | (7) Job Search |
| ChatGPT Use | -0.177** | 0.053 | -0.024 | 0.029 | 0.035* | 0.052 | -0.022 |
| | (-2.200) | (1.330) | (-0.634) | (0.893) | (1.843) | (1.370) | (-1.250) |
| *Panel B: Leisure-related content categories* | | | | | | | |
| | (8) Communic. | (9) Social Media | (10) Promotional | (11) News | (12) Enter-tainment | (13) Games | (14) Other |
| ChatGPT Use | -0.031 | 0.014 | 0.242 | -0.039*** | 0.080 | 0.099*** | -0.041 |
| | (-1.280) | (0.397) | (1.451) | (-3.069) | (1.447) | (2.636) | (-0.587) |
| Observations | 42,886 | 42,886 | 42,886 | 42,886 | 42,886 | 42,886 | 42,886 |
| 1st-stage KP F-stat. | 104 | 104 | 104 | 104 | 104 | 104 | 104 |
| Estimation | IV | IV | IV | IV | IV | IV | IV |
| Inc. X Age X MSA FEs | X | X | X | X | X | X | X |
| Bartik share control | X | X | X | X | X | X | X |
| Website cat. exposure | X | X | X | X | X | X | X |

## Figure 7: Share of GenAI-exposed website activity by type

This figure shows the share by duration of different website activities among those that are classified as 'exposed' to generative AI by the methodology detailed in Appendix IA.1.



## B. *ChatGPT use context in 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 important empirical question in the context of our study is what *tasks* households actually choose to use ChatGPT for. While observing households' questions to ChatGPT would provide some information on this (see Chatterji et al. (2025) for evidence on this), we do not have access to chatbot interaction data. However, such interaction data could only provide a partial answer as it would lack the context of the broader task the household is performing on the Internet: any task that involves websites other than a chatbot would only be partially captured this way. For example, an inquiry like "*Can you give me a simple explanation of the Black-Scholes formula?*" inside ChatGPT may reflect a student preparing for an exam in finance, a financial analyst preparing for a client presentation, or a retail trader evaluating an investment. What is lacking for evaluating the task that is being done is the *context* of the chatbot interaction.

Conversely, in this section, we argue that one can use this intuition and infer households' purpose for using ChatGPT based on the browsing context of the chatbot interaction. We leverage our high-frequency browsing data to observe the websites that users visit in the minutes before or after they use ChatGPT (the "GPT-window"). This allows 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/CDN websites and a label of whether the household visited openai.com or chatgpt.com in the interval. On average, about 0.2% and 1.1% of the intervals are labeled as being a GPT-window in 2023 and 2024, respectively. For several of the analyses below, we want to compare browsing activity in the intervals around ChatGPT use to what is 'normal'. As a benchmark, we match ChatGPT use intervals to browsing activity by households that never use ChatGPT and are in the same age and income categories. Moreover, for general comparisons between users and non-users we also match browsing activity that occurs

on the same day of the week and at the same time of day, which avoids, for example, bias from comparing weekend browsing to work day evening browsing.

**Websites associated with ChatGPT use.** To provide some intuition for what kinds of Internet tasks are associated with ChatGPT use, we first consider which websites show the greatest difference in browsing shares in the intervals around ChatGPT use relative to matched browsing intervals by non-users. Table VIII shows the 30 sites with the largest gap. That is, these are the sites for which browsing time is unusually high around ChatGPT use. We find that search on google.com is unusually high. Moreover, unusually prevalent are also education sites such as Pearson.com or Instructure.com; the graphic design platform Canva; writing assistant platform Quillbot.com; and job search platforms/professional networks Linkedin and Indeed. This shows that ChatGPT use tends to be associated with sites commonly used for productive digital tasks. In Table IA.6 in the Appendix, we show that the sites that are unusually lacking from ChatGPT use intervals include general entertainment/streaming sites like YouTube, MSN.com, and Yahoo; the social network Facebook; and also some shopping sites like Amazon and eBay, as well as several adult content sites. Moreover, Wikipedia also has unusually low prevalence in ChatGPT use intervals, showing that some sites with high exposure are substituted by chatbot use. Overall, these simple descriptive patterns in our data align closely with our interpretation that ChatGPT use tends to occur in the context of websites used for productive digital tasks, and not in the context of leisure-oriented entertainment and social media sites. This supports our assumption that ChatGPT itself is more likely to be used for these productive tasks.

**Website content types associated with ChatGPT use.** We can compare ChatGPT use intervals and matched non-user browsing behavior more systematically by considering which types of Comscore website categories are over- or under-represented in browsing behavior around ChatGPT uses. Figure 8 shows that both ChatGPT use intervals and the matched browsing behavior by never-users show a large prevalence of browsing on multi-use sites, search engines, and on education sites (panel A). However, the ChatGPT use intervals show, in relative terms, a much larger prevalence of these three website categories, and a much lower prevalence of browsing on

**Table VIII:**
**Websites with highest relative prevalence around ChatGPT use.**

This table shows websites with the highest prevalence in 30 minute intervals with ChatGPT use compared to non-user browsing matched on time of day and age bin × income bin demographics.

| Rank | Domain | Δ Browsing % | Rank | Domain | Δ Browsing % |
|---|---|---|---|---|---|
| 1 | google.com | 16 | 16 | mheducation.com | .34 |
| 2 | instructure.com | 4.7 | 17 | grammarly.com | .29 |
| 3 | canva.com | 2.2 | 18 | udemy.com | .24 |
| 4 | quillbot.com | 1.3 | 19 | googleapis.com | .24 |
| 5 | discord.com | 1.3 | 20 | schoology.com | .24 |
| 6 | gist.build | .75 | 21 | edgenuity.com | .23 |
| 7 | quizlet.com | .48 | 22 | eesysoft.com | .22 |
| 8 | indeed.com | .45 | 23 | coinbase.com | .2 |
| 9 | pearson.com | .44 | 24 | edmentum.com | .2 |
| 10 | brightspace.com | .44 | 25 | github.com | .19 |
| 11 | blackboard.com | .42 | 26 | wwnorton.com | .17 |
| 12 | linkedin.com | .38 | 27 | chegg.com | .17 |
| 13 | cengage.com | .36 | 28 | character.ai | .14 |
| 14 | webassign.net | .35 | 29 | sharepoint.com | .13 |
| 15 | uworld.com | .34 | 30 | office365.us | .11 |

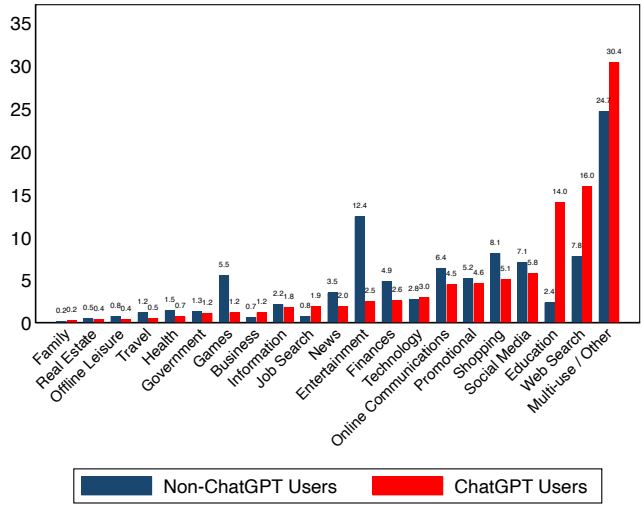
entertainment, gaming, and shopping sites (panel B).

**Task purposes associated with ChatGPT use.** Finally, we can map the behavior around ChatGPT uses to our distinction between sites associated with productive and leisure digital tasks (as well as mixed-use and ad/CDN sites for completeness). Figure 9 shows the distribution of the four types of website purposes during the GPT-window relative to time- and demographics-matched non-user browsing behavior. On average, during the 30-minute window when households use ChatGPT, the tasks they focus on are 80.1% productive, 9.5% mixed-use, 7.6% leisure, and 1.2% ad/CDN sites (panel A). Relative to never-user browsing (see panel D), GPT-window browsing focuses substantially more on productive tasks—a 25.2pp gap—and substantially less on leisure tasks, with a -13.7pp gap.

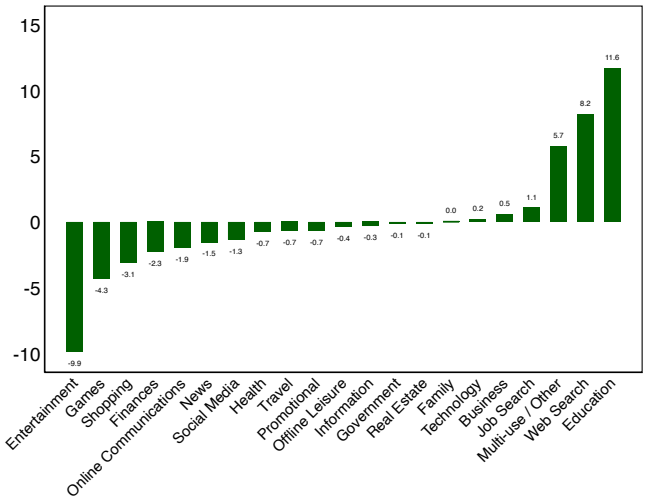
We explore further differences in the distribution of browsing task purposes between ChatGPT users and never-users in the other of Figure 9. Panels (B) and (E) show that ChatGPT users are generally more likely to do productive digital tasks than never-users on days when they use ChatGPT. However, this behavioral gap is not as large as the gap during the ChatGPT use interval itself. Panels (C) and (F) show that there is even a slight bias towards doing more productive digital tasks before ChatGPT is released among ChatGPT users—in fact, differences of this type underlie the

**Figure 8: Website content types during ChatGPT sessions**

This figure plots the level (panel A) and difference (panel B) of browsing duration shares by websites' Comscore content categories during a ChatGPT use session and matched never-user browsing sessions, matched on time and demographics.




**(A)** Browsing shares by website content category




**(B)** Difference between GPT windows and never-users

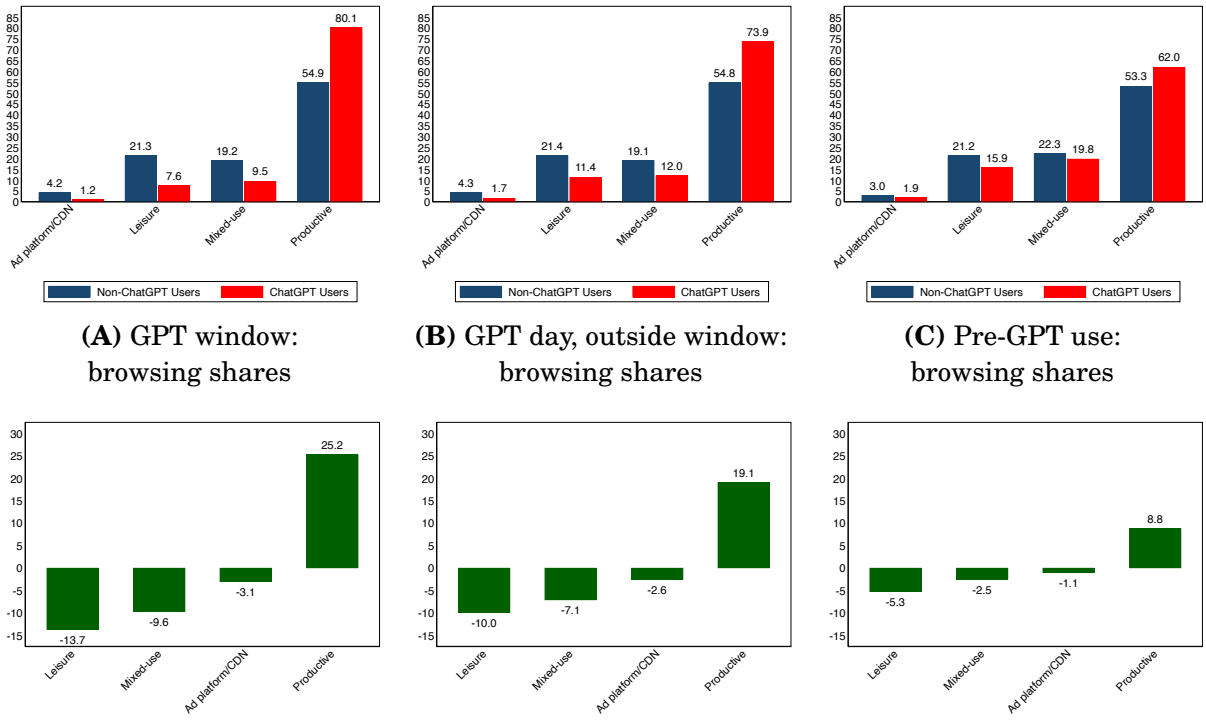
intuition why our pre-ChatGPT household exposure measure is able to predict ChatGPT adoption. However, this bias is much smaller than the latter gap in productive task browsing associated with ChatGPT use intervals. This evidence confirms that productive digital tasks are much more likely to be done in browsing sessions that involve ChatGPT use, and that this association cannot be explained by a selection effect with regard to on what days households use ChatGPT or which households choose to use ChatGPT at all.

Together, the mechanism evidence in this section provides an empirical justification for interpreting our baseline IV estimates as reflecting a productive digital task-biased productivity shock that frees up time which is re-allocated to leisure digital tasks.

**Methodological limitations.** We end this section by discussing the premises and limitations of this high-frequency browsing interval approach. Our approach is based on two premises: First, we view ChatGPT as a *tool* for fulfilling an ongoing 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

**Figure 9: 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 IV.




**(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

performs associated tasks 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 high-frequency data on browsing activity in short browsing intervals in which ChatGPT is used.

Our approach faces two potential sources of noise: First, households may resort to ChatGPT for tasks without visiting other websites. For example, as noted above, searching for factual information may no longer involve Wikipedia as ChatGPT becomes the first and only website to consult—in which case no trace of the completed task remains in the browsing context of ChatGPT. Second, households may be multi-tasking when visiting ChatGPT, making the surrounding websites less informative about the household's purpose of using ChatGPT. While this noise may lead to bias in the magnitudes found in our analysis, we would still expect our results to qualitatively reflect the average purpose of ChatGPT usage by households.

# V. Quantifying Household Productivity Gain

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 our empirical measures of changes in households' time spent on productive and leisure activities to the impact of generative AI as a productivity-enhancing tool. The conceptual framework shows that to translate these reallocations into an implied productivity improvement for digital tasks, we need to combine two ingredients:

1. Causal adoption effects on time use across categories. These are our empirical results that identify how adopting ChatGPT reallocates households' *digital time* across productive and leisure browsing (see the IV estimates in Table VI).

2. Time-demand ("Engel curve") elasticities that govern how category-level time responds to shocks to total digital time (estimated below in Table IX).

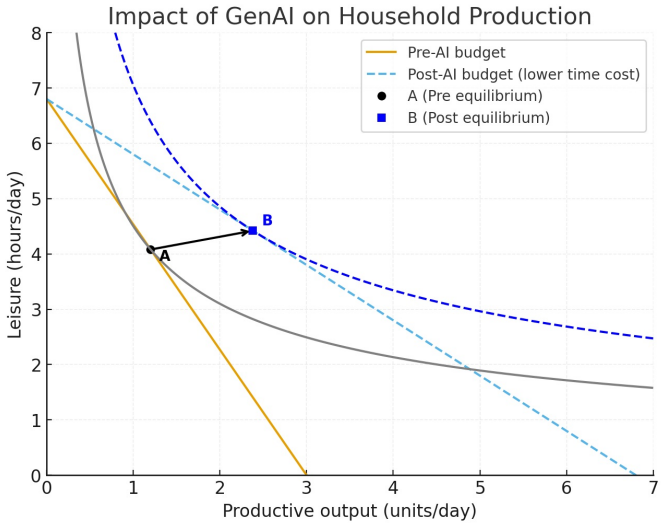
Below, we (1) derive expressions for changes in digital time allocation in response to ChatGPT adoption; and then (2) show how we can estimate "Engel curve" elasticities for digital time allocation from our Internet browsing data; (3) Finally, we combine these estimates into an implied magnitude of the household productivity change as a result of ChatGPT adoption.

**Model intuition.** Before presenting the full model below, the core intuition is illustrated in Figure 10: we consider a household that faces a budget constraint with regard to how to allocate effective browsing time between leisure purposes (effectiveness normalized to one and output shown on the vertical axis) and productive purposes (horizontal axis). The solid budget line represents the household's time budget before generative AI. The household chooses the time allocation where the effective time budget is tangent to the iso-utility curve. 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 productive time spent browsing. The illustrated case shows how the expansion of the effective time budget due to the effect of generative AI on productive tasks increases leisure browsing time. Note that the horizontal axis shows effective time units after accounting for productivity.

As less time is needed post-GenAI to produce a given effective productive digital task output, even the substantial increase in effective time output from productive digital tasks that is shown can be produced with a smaller 'clock time' input of browsing duration, so productive digital task time allocation falls. The full model below captures this overall mechanism in a more general setting.

**Figure 10: 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.




## A. *Model setup and household digital time allocation*

**Household utility and time allocation problem** We consider a household $i$ that, 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 $\ell_{it}$ 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

$$z_{it} + \ell_{it} \leq 1. \tag{5}$$

In the quantification below, we relax the normalization and allow total digital time to vary, so the model can map directly to observed variation in total browsing time.

The household optimally allocates time between these two digital activities to maximize flow

utility, given by the power utility function

$$v_{it} = \frac{(\theta^{\ell}_{it}\xi^{\ell}_{it}\ell_{it})^{1-(1/\eta^{\ell})}}{1-(1/\eta^{\ell})} + \frac{(\theta^{z}_{it}\xi^{z}_{it}z_{it})^{1-(1/\eta^{z})}}{1-(1/\eta^{z})}, \tag{6}$$

where $\theta^{\ell}_{it}$ and $\theta^{z}_{it}$ are taste shifters (preference weights) for leisure and productive digital activities, and $\xi^{\ell}_{it}$ and $\xi^{z}_{it}$ are activity-specific efficiency shifters which affect how much utility-relevant output per unit time the activity generates. The curvature parameters $\eta_{\ell}, \eta_{z} > 0$ govern diminishing marginal utility of time within each activity. When $\eta_a > 1$, activity $a$ is a *time luxury* (analogous to Aguiar et al. (2021) "leisure luxuries"): higher effective quality $\theta^a\xi^a$ increases the optimal time allocated to $a$. When $0 < \eta_a < 1$, activity $a$ is a *time necessity*: higher $\theta^a\xi^a$ reduces time allocated to $a$, because the household reaches diminishing returns quickly and reallocates time elsewhere. The knife-edge case $\eta_a = 1$ corresponds to log utility in activity $a$. The model generalizes to more than two digital activities (e.g., to incorporate the 'mixed' sites from our empirical design), but we simplify the exposition of the conceptual framework here by focusing on leisure and productive purposes. We consider the effect of adding mixed sites when calibrating the productivity shock estimates below.

**First order conditions** This optimization problem for each activity $a \in \{z, \ell\}$ yields a first-order condition which can be written as an activity time allocation equation (after taking logs of both sides) in the form

$$\ln(a_{it}) = (\eta^a - 1) \times \left(\ln(\theta^a_{it}) + \ln(\xi^a_{it})\right) - \eta^a \ln(\omega_{it}) \qquad a \in \{z, \ell\} \tag{7}$$

where $\omega_{it}$ is the shadow price of the time constraint. This equation shows 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—as long as the activity behaves like a "leisure luxury", i.e., $\eta > 1$. If $\eta < 1$, these effects reverse: households allocate less time to an activity when its taste/productivity parameter increases. If leisure and productive digital activities differ in their behavior—e.g., because $\eta^z < \eta^{\ell}$, as we show empirically below—then even an identical productivity change in both activities can lead to a relative change in the share of browsing time allocated to them. See the Internet Appendix

for details on the derivation.

## *B. The Release of ChatGPT and the choice to adopt*

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 $\xi^z$ while leaving $\xi^\ell$, $\theta^z$ and $\theta^\ell$ 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 $\xi^z_{it}$ of the household's digital home production. In particular, assume that this object's time $pre$ value equals $\xi^z_{i,pre} = \underline{\xi}$ while its time $post$ value is

$$\xi^z_{i,post} = \underline{\xi} \times \left(1 + \delta \cdot \text{Adopt}_{i,post}\right) \tag{8}$$

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. In what follows, it is also convenient to work with the log technology shock (conditional on adoption) $\Delta \ln \xi^z \equiv \ln(1+\delta)$, which is approximately $\delta$ for small productivity gains.

The household chooses to adopt ChatGPT if

$$v^{\text{Adopt}}_{i,post} - v^{\text{No adopt}}_{i,post} \geq c \tag{9}$$

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.

A first-order approximation assuming small productivity effects delivers a simple adoption rule: if adopting ChatGPT raises the efficiency of productive digital tasks by a factor $1+\delta$, and adoption entails a fixed time cost $c_{\text{time}}$, then adoption is optimal when the implied time savings relative to the counterfactual exceed the learning/setup time. Formally, adoption occurs when

$$\delta \cdot z_{i,\text{post}}^{\text{No adopt}} \geq c_{\text{time}},$$

with derivations in the Appendix. This maps directly to our empirical "exposure" measure: households with higher baseline intensity of GenAI-substitutable productive browsing have larger potential gains and should adopt at higher rates.

## C. *Browsing 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 (5) into the FOCs (7) and differencing 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. The relative impact on time allocated to productive tasks versus leisure tasks is

$$\frac{\beta_z^{GPT}}{\eta^z} - \frac{\beta_\ell^{GPT}}{\eta^\ell} = \frac{\eta^z - 1}{\eta^z} \times \ln(1+\delta_z) \tag{10}$$

while the impact on time allocated to productive tasks is

$$\beta_z^{GPT} = \frac{(\eta^z - 1)}{1 + \left(\eta^z/\eta^\ell\right) \cdot \left(z_{i,post}^{\text{No adopt}}/\ell_{i,post}^{\text{No adopt}}\right)} \times \ln(1+\delta_z) \tag{11}$$

Equation (11) 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, $\eta^z$, 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 $\eta^z$ is, due to the decreasing returns that the household gets from completing additional productive tasks.

If there is no direct impact on digital leisure efficiency as a result of ChatGPT use, the effect on the time allocated to leisure tasks takes a form that is similar but in which $\eta^z$ enters with the opposite sign:

$$\beta_\ell^{GPT} = \frac{(1-\eta^z)\cdot\frac{\eta^\ell}{\eta^z}}{1+\left(\eta^\ell/\eta^z\right)\cdot\left(\ell_{i,post}^{\text{No adopt}}/z_{i,post}^{\text{No adopt}}\right)} \times \ln(1+\delta_z) \tag{12}$$

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 $\eta^\ell$ itself goes down, the effect remains positive regardless of whether leisure is a luxury ($\eta^\ell > 1$) or an 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 a decrease in the share of time spent on productive browsing activities (see Table VI).

In the next section, we discuss how we obtain the "Engel curve" elasticities $\eta^a$ that we need

to translate the IV estimates of changes in browsing behavior into estimates of the size of the productivity shock.

## *D. Estimating Engel elasticities*

The previous sections showed how to derive an expression that relates the causal effects of ChatGPT adoption on digital task consumption to the size of the productivity shock and the parameters $\eta^a$. In this section, we show how the parameters $\eta^a$ can be estimated.

Let total digital time in period $t$ be $H_{it}$, allocated across productive, leisure, and mixed ("other") digital activity:

$$H_{it} = z_{it} + \ell_{it} + o_{it}. \tag{13}$$

Under the additively separable isoelastic aggregator in equation 6, optimal time allocations imply constant-elasticity "demands" for time in each category. A key implication is that each activity has an *Engel elasticity in time*,

$$\beta_a \equiv \frac{\partial \ln h^a_{it}}{\partial \ln H_{it}} = \frac{\eta^a}{\bar{\eta}}, \qquad a \in \{z, \ell, o\}, \tag{14}$$

that summarizes whether activity $a$ behaves as a *time luxury* ($\beta_a > 1$) or a *time necessity* ($\beta_a < 1$) as total digital time expands. The Engel elasticities $\{\beta^a\}$ identify $\{\eta^a\}$ only up to a common scale. To obtain levels of $\eta^a$, we need to impose (or calibrate) a value for the average curvature $\bar{\eta}$ to compute

$$\eta^a = \beta^a \bar{\eta}. \tag{15}$$

Conditional on a chosen $\bar{\eta}$, the Engel-curve estimates then deliver $\eta^\ell$, $\eta^z$, and $\eta^o$.

**Engel elasticity specification.** We estimate the elasticities $\{\beta^a\}$ using aggregated group panel regressions of the form

$$\ln h^a_{g,t} = \delta^a_t + \beta_a \ln H_{g,t} + u^a_{g,t}, \tag{16}$$

where $g$ indexes age–income–region cells and $\delta^a_t$ are quarter fixed effects. Here, $\ln H_{it}$ is the log of total browsing time, and $\ln h^a_{it}$ is the log of the browsing time of activity $a$ (e.g., leisure or productive

activity). All regressions include fixed effects that capture the interaction between income and age categories, and also fixed effects in the time dimension (at the year-quarter level).

**Instrument for total browsing time.** One concern in the Engel elasticity estimation is that total browsing time is likely endogenous with regard to other shocks to households' activity that may also change their need to complete productive vs. leisure tasks online. To address this endogeneity, we use an instrumental variable approach based on plausibly exogenous local precipitation shocks that discourage offline outdoor activity and increase total online browsing time.

**Engel elasticity estimates.** The Engel elasticity estimates are shown in Table IX. The estimates in panel B indicate that online leisure browsing is a time luxury ($\beta^{\ell} > 1$) while productive digital activity is a time necessity ($\beta^{z} < 1$), consistent with the reallocation patterns we observe after generative AI adoption. Intuitively, this suggests that households do not substantially expand the amount of productive online work they do in their private lives when they spend more total time online. Rather, they tend to allocate more total time spent online towards leisure activities such as social media or streaming. Moreover, to capture the full range of online activities we also estimate the response of other mixed-purpose browsing, which, based on our classification algorithm, combines leisure and productive tasks, and which, according to the estimates in panel C, has an Engel elasticity that lies about halfway between the estimates for leisure and productive browsing. The IV estimates result in somewhat more differentiated coefficients than the OLS estimates, but both tell a qualitatively similar story. Aguiar et al. (2021) conduct a similar estimation exercise using time shares as the dependent variables (rather than the log of activity time) to avoid dropping cells with zero activity, and then approximate the elasticity using the estimated semi-log coefficients. We show in IA.5 that a similar approach in our setting results in qualitatively and quantitatively comparable estimates as those shown in Table IX—likely because our estimation panel does not contain many cells with zero browsing activity.

## Table IX:
## Engel elasticity estimation results

This table shows estimates of age-income-region-level panel regressions of the form shown in equation 16. See Section V.D for details. T-statistics based on standard errors double clustered at the age-income-region level and at the year-quarter level in parentheses: * p<0.10, ** p<0.05, *** p<0.01.

| *Dep. var* | Log Activity Time | |
|---|---|---|
| Estimation | OLS | IV |
| | (1) | (2) |
| *Panel A: Productive digital activities* | | |
| Log(Total Browsing Time) | 0.991*** | 0.931*** |
| | (530.879) | (29.244) |
| Observations | 129,146 | 99,430 |
| 1st-stage KP F-stat. | | 54 |
| *Panel B: Leisure digital activities* | | |
| Log(Total Browsing Time) | 1.130*** | 1.374*** |
| | (241.949) | (15.823) |
| Observations | 120,785 | 94,691 |
| 1st-stage KP F-stat. | | 48 |
| *Panel C: Other digital activities* | | |
| Log(Total Browsing Time) | 1.026*** | 1.110*** |
| | (276.169) | (17.912) |
| Observations | 125,862 | 97,739 |
| 1st-stage KP F-stat. | | 51 |
| Age x Income Group FEs | X | X |
| Time FEs | X | X |

## E. *Implied generative AI efficiency gain*

With the estimates of $\beta^a$ in hand, we can use our estimates of the causal effect of ChatGPT use on browsing time allocation to infer the change in efficiency as a result of generative AI technology. Observed shifts in time allocation and the leisure Engel curves identify changes in technology up to the scaling parameter $\bar{\eta}$. Recall the assumption that generative AI adoption raises only the efficiency of productive digital tasks, and that we can therefore treat leisure browsing as our reference activity.

To estimate the implied scaled generative AI productivity change $(1-\eta_z)\delta$ for productive digital tasks,[12] We combine our estimates of Engel elasticities with our IV estimates of the causal effect of ChatGPT adoption on browsing activity. As before, let $\hat{\beta}_a^{GPT}$ denote our estimate of the causal effect of ChatGPT adoption on the log of activity time $a$:

$$\hat{\beta}_a^{GPT} \equiv \ln a_{\text{Adopt}} - \ln a_{\text{No adopt}}, \qquad a \in \{z, \ell, o\}, \tag{17}$$

which we obtain from our IV long-difference design (Table VI). Then, the difference in time allocation effects shown in equation 10 (using the exact version as $\delta$ may not be small) across the Adopt vs. No-adopt regimes eliminates the (unobserved) shadow value of digital time. Using the Engel curve estimates to substitute $\eta_z/\eta_\ell = \hat{\beta}_z/\hat{\beta}_\ell$, then yields:

$$(1-\eta_z)\ln(1+\delta_z) = (\beta_z/\beta_\ell)\hat{\beta}_\ell^{GPT} - \hat{\beta}_z^{GPT}, \tag{18}$$

if we are assuming that ChatGPT only impacts the efficiency of productive digital tasks and we are using leisure digital tasks as the unaffected reference category. Here, we have also multiplied both sides by $-1$, in order to ensure the technology shock on the left-hand side has a positive sign when $\eta_z < 1$, which we expect to be the case if productive browsing activity behaves as a "digital time necessity". Intuitively, this equation shows that the implied (scaled) efficiency shock due to ChatGPT use $(1-\eta_z)\ln(1+\delta_z)$ can be inferred from the *relative* change in browsing time. As

[12] This is analogous to the scaled productivity shock for computer time computed by Aguiar et al. (2021) under the assumption that leisure time spent on eating, sleeping, and personal care (ESP) can be treated as a reference category that does not experience a technology shock during the 2000s.

reported in column (2) of Table VI, leisure digital browsing time increased by 151 log points over the 2022-2024 period as a result of ChatGPT use. The estimates in Table IX give $\frac{\hat{\beta}^z}{\hat{\beta}^\ell} = 0.68$. This implies that, absent any technological change in productive digital tasks, the change in leisure browsing implies that productive browsing time should have increased by 102 log points as well. This is the term on the left-hand side of equation 18, which corresponds to the predicted movement along the Engel curve for productive digital tasks. However, productive browsing time actually rose by only 1.1 log points. The model interprets this gap relative to the expected productive digital time allocation as having to result from greater efficiency. We therefore estimate the change in $(1-\eta_z)\ln(1+\delta_z)$ to be a 101 log point, or 175%, increase in scaled efficiency.

**Allowing for leisure efficiency gains.** This calculation assumed that leisure digital tasks experience no change in efficiency due to ChatGPT use. This is likely too strong an assumption. While it seems plausible that leisure efficiency gains are weakly smaller than productive task efficiency gains—which is also consistent with ChatGPT being mainly used in the context of productive digital tasks—we can relax this assumption and allow leisure task efficiency gains to lie in a range between zero and being the same magnitude as productive task efficiency gains, i.e., we define a parameter $\psi$, the "efficiency gain ratio", defined by $\delta_\ell = \psi\,\delta_z$, where $\psi \in [0,1]$.

The relative impact on the observed changes in log browsing time for the two types of tasks is then given by the following modification of equation 10:

$$\frac{\beta_z^{GPT}}{\eta^z} - \frac{\beta_\ell^{GPT}}{\eta^\ell} = \frac{\eta^z - 1}{\eta^z}\ln(1+\delta_z) - \frac{\eta^\ell - 1}{\eta^\ell}\ln(1+\delta_\ell). \tag{19}$$

and equation 18 becomes

$$(1-\eta_z)\ln(1+\delta_z) - \frac{\beta_z}{\beta_\ell}(1-\eta_\ell)\ln(1+\psi\delta_z) = \frac{\beta_z}{\beta_\ell}\hat{\beta}_\ell^{GPT} - \hat{\beta}_z^{GPT}.$$

In the case where $\psi = 0$, this is identical to equation 18, but if $\psi > 0$, it shows that the size of the measured shock on the right-hand side can now be explained not just by a scaled efficiency shock to productive tasks (the first term on the LHS), but also potentially by a scaled efficiency shock to leisure tasks, as long as $\eta_\ell > 1$, which means that leisure tasks behave as a "time luxury", such that

the second term on the LHS is positive. Intuitively, as long as this will lead to smaller estimates of the scaled efficiency gain, because any gap between increases in leisure task time and increases in productive task time can no longer be explained only by a productivity shock lowering time spent on the (time necessity) productive tasks. Instead, part of the increase in leisure task time can now be due to its own efficiency gain and the fact that it behaves as a time luxury, such that the inferred efficiency gain for productive tasks is smaller.

As we do not know the *levels* of $\eta_z$ and $\eta_\ell$, we cannot generally solve this for the scaled efficiency gain. However, if we make assumptions about $\bar{\eta}$ in addition to assuming values of $\psi$, we can compute the scaled efficiency gain for productive tasks under these different scenarios. Aguiar et al. (2021) find $\bar{\eta} \approx 0.73$ and show robustness for a range $\bar{\eta} \in [0.73, 1.0, 1.19]$. In our setting, we estimated $\hat{\beta}_z = 0.931$ as the elasticity of productive browsing time with respect to total browsing time (see Table IX). Aguiar et al. (2021) show that $\beta_i = \eta_i/\bar{\eta}$, and we find in Table VI that, following adoption, total browsing rises, implying a proportional increase in productive browsing time, but our estimates in Table VI show no economically significant increase. This implies that $\eta_z < 1$, which is consistent with productive browsing behaving as a time necessity. As a result, we can infer that $\eta_z = \beta_z \bar{\eta} < 1$, which implies that $\bar{\eta} < 1/\beta_z \approx 1.07$. That is, $\bar{\eta}$ cannot be much above one if productive tasks are behaving as a time necessity. Coincidentally, the estimate of $\bar{\eta} \approx 0.73$ from Aguiar et al. (2021) also corresponds to the lowest $\bar{\eta}$ in our setting that is consistent with $\eta_\ell > 1$, which reflects the sensible assumption that leisure digital tasks behave as a time luxury good. We therefore consider values $\bar{\eta} \in [0.73, 0.9, 1.0, 1.07]$ as possible scenarios (where 1 is a salient benchmark and $\bar{\eta} = 0.9$ is the midpoint between the upper and lower bounds). Given the large observed increases in total browsing, without a change in productive browsing time, $\eta_z$ seems more likely to be well below 1, which implies that $\bar{\eta}$ is plausibly substantially smaller than the upper bound considered here.

**Calibrating the efficiency gain ratio.** We can determine plausible values for $\psi$ by assuming it ranges from leisure not being affected by ChatGPT ($\psi = 0$) to assuming that leisure and productive digital tasks see the same efficiency gains ($\psi = 1$), which provides plausible outer bounds for the relative effects. We can further narrow this range, by using our data to construct empirical proxies

for the share of activities on leisure sites that are exposed to ChatGPT relative to the share of productive site activities that are exposed. In Appendix Table IA.7, we show the duration-weighted average exposure during our benchmarking period in 2021 for the main usage categories. As the data show, leisure-oriented websites had an average activity exposure of 17.7%, while productive sites average 50.7% exposure, so the predicted efficiency gain ratio based only on this exposure proxy would be $\frac{0.117}{0.507} \approx 0.23$. We therefore consider efficiency gain ratios $\psi \in [0, 0.25, 0.5, 1]$, to include a value close to our empirical proxy and also consider what doubling this proxy or assuming the same shock for both task types would mean for our estimates.

**Table X:**
**Scaled productive efficiency gain calibration**

| Scaled productive efficiency gain (in % increase): $100\left(\exp\left((1-\eta_z)\ln(1+\delta_z)\right)-1\right)$ | | | | |
|---|---|---|---|---|
| $\bar{\eta} \setminus \psi$ | 0 | 0.25 | 0.5 | 1 |
| 0.73 | 175.52 | 174.46 | 174.12 | 173.76 |
| 0.90 | 175.52 | 85.16 | 75.59 | 66.45 |
| 1.00 | 175.52 | 33.60 | 28.80 | 24.22 |
| 1.07 | 175.52 | 1.73 | 1.47 | 1.21 |

*Note:* $(\bar{\eta}, \psi)$, $\delta_z$ solves $(1-\eta_z)\ln(1+\delta_z) - r(1-\eta_\ell)\ln(1+\psi\,\delta_z) = A_z$, with $\eta_z = 0.931\,\bar{\eta}$, $\eta_\ell = 1.374\,\bar{\eta}$, $\frac{\beta_z}{\beta_\ell} = 0.6776$, and $\frac{\beta_z}{\beta_\ell}\hat{\beta}_\ell^{GPT} - \hat{\beta}_z^{GPT} = \frac{0.931}{1.374} \cdot 1.512 - 0.011$.

**Discussion.** The resulting estimates of the scaled productive task efficiency shock in log points are shown in Table X. As the table shows, the implied scaled efficiency gains in percent for the productive digital task are large depending on the assumptions about the elasticity parameter and the efficiency gain ratio.

We consider the midpoint $\bar{\eta} = 0.9$ between the lower and upper bounds for $\bar{\eta}$ that would be consistent with productive time behaving as a necessity and leisure time as a luxury, as a pragmatic choice. We also assume that the actual efficiency ratio gain is no more than double our empirical proxy for the average ratio of exposure on leisure sites to that on productive sites, such that $\psi$ is plausibly between 0 and 0.5. Then, this table suggests a scaled efficiency gain range of 76%-176% for productive digital tasks from generative AI adoption.

Our estimated productivity gains of 76% to 175% from ChatGPT adoption on household productive Internet tasks are larger than, but within a similar order of magnitudes from, those found in

the recent literature on generative AI in workplace settings. For instance, Noy and Zhang (2023) found that ChatGPT reduced task completion time by 40% for professional writing tasks; Brynjolfsson et al. (2023) found that AI-based conversational assistance increased customer support agent productivity by 34% for novice workers; Dell'Acqua et al. (2023) found that BCG consultants using GPT-4 finished 25.1% faster on tasks within AI's capability frontier; Peng et al. (2023) found that GitHub Copilot users completed coding tasks 55.8% faster.

There can be several possible reasons for the effect to be larger for home setting. First, workplace experiments typically measure productivity on *structured professional tasks*—writing reports, resolving customer tickets, coding—where workers already possess specialized skills and established workflows. By contrast, household productive Internet tasks—such as researching product purchases, troubleshooting home appliances, looking up recipes, seeking tax or medical information, and helping children with homework—are often performed by individuals with little domain expertise, making them closer to the "novice" end of the skill distribution where productivity gains are largest (Brynjolfsson et al., 2023; Noy and Zhang, 2023). Second, our IV strategy identifies a local average treatment effect for compliers whose browsing patterns had the highest overlap with ChatGPT's capabilities, capturing the productivity gains among households most intensively exposed to the technology, which are known to amplify the average effect (Jiang (2017)). Third, and perhaps most importantly, many routine household Internet tasks involve extensive browsing—visiting multiple websites, comparing information across sources, reading forums, and watching tutorial videos—that ChatGPT can substitute almost entirely with a single conversational query.[13] Because our model aggregates across all productive browsing tasks—including those with more modest gains—the estimated range of 76% to 175% represents an average that is well within the plausible range implied by the mix of near-complete task substitution for some household activities and more moderate efficiency gains for others.

[13]For example, a household member researching how to fix a leaking faucet might previously spend 20–30 minutes browsing plumbing forums, watching YouTube tutorials, and comparing advice across websites, whereas ChatGPT can provide a step-by-step diagnosis and repair guide in under two minutes—a time savings exceeding 1,000%. Similarly, comparing features and reviews across consumer products, finding and adapting a recipe to dietary restrictions, understanding a health symptom before deciding whether to visit a doctor, or navigating government forms and tax questions are all tasks where a single ChatGPT prompt can replace what previously required 15–30 minutes of multi-site browsing, implying efficiency gains of 500% to well over 1,000% for these specific activities.

# 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.

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 to leisure browsing activities, consistent with a greater efficiency of online productive tasks that allows households to reallocate the resulting time savings to leisure browsing. 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 more 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. Moreover, accurately measuring the impact of generative AI on the economy and people's lives will require making further progress in measuring the effect on non-market activities—changes in private lives—that builds on our contribution.

# 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, 2017, Have Instrumental Variables Brought Us Closer to the Truth 6, 127–140.

Jiang, Wei, Junyoung Park, Rachel Jiqiu Xiao, and Shen Zhang, 2025, Ai and the extended workday: Pro-

ductivity, 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.

Schubert, Gregor, 2025, Organizational technology ladders: Remote work and generative ai adoption, *Available at SSRN* .

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.

# Internet Appendix for

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

Michael Blank Gregor Schubert Miao Ben Zhang

# IA.1. Methodology 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 160K 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 sites, o

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.

## *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 being do

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",
}}
```

## *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.

## *D. Rainfall instrument construction*

This section shows how we construct the measure of local precipitation used to instrument for total online browsing time in a quarter.

**Data Sources** We construct a metro region-level monthly weather panel from the grid-level daily weather dataset maintained by Wolfram Schlenker.[14] The raw data provide daily observations of precipitation (`prec`) for a fine grid of cells covering the contiguous United States, Alaska, and Hawaii. Each grid cell is pre-assigned to a county FIPS code via a grid-to-county crosswalk distributed with the data. Our sample covers the years 2019–2023. For each year, the data contain 3,105 county-level files, with each file containing daily observations for all grid cells assigned to that county. The typical county contains between 85 and 150 grid cells, depending on the geographic area.

**Geographic Mapping** Aggregation from grid cells to our target regions proceeds through a three-step geographic crosswalk:

1. **Grid cell → County.** This mapping is pre-assigned in the raw data. Each grid cell belongs to exactly one county, identified by its 5-digit FIPS code.
2. **County → MSA.** We use the Bureau of Labor Statistics Quarterly Census of Employment and Wages (QCEW) county-to-MSA/CBSA crosswalk (July 2023 vintage) to map county FIPS codes to Metropolitan and Micropolitan Statistical Area titles. This step successfully maps 1,843 counties to 924 unique CBSA titles. Counties not assigned to any MSA—primarily rural counties—are excluded from the regional aggregation.
3. **MSA → Region.** We map MSA titles to the 210 metro regions included in the Comscore data using a semi-manual approach. Because MSA titles (e.g., "Dallas-Fort Worth-Arlington, TX MSA") differ systematically from Comscore region names (e.g., "Dallas-Ft. Worth, TX"), simple fuzzy matching approaches alone are likely unreliable. Instead, we automatically match MSAs where (i) the MSA's primary state matches the Comscore region's primary state and (ii) the MSA's primary city name matches a city component in the Comscore region name. This automatic process matches 178 MSAs. We manually match an additional 95 of the

[14] Available at https://wolfram-schlenker.info/links.html.

Comscore regions to the QCEW MSA titles. In total, 273 of 924 MSA titles are mapped to one of 210 target Comscore regions. The remaining 651 MSAs correspond to areas outside our Comscore target regions. Three counties that lack QCEW MSA assignments but map to Comscore target regions are assigned directly (Aroostook County, ME → Presque Isle-Caribou, ME; Dawson County, MT → Glendive, MT).

The final crosswalk maps 1,013 unique counties to 210 Comscore regions.

**Aggregation Methodology.** Aggregation from the grid-level rainfall data then proceeds in two stages:

1. Grid cells → County-month: For each county and calendar day, we compute simple (unweighted) arithmetic means of precipitation across all grid cells assigned to that county. We then aggregate county-day observations to county-month observations by computing the mean daily precipitation rate in the month.

2. County-month → Region-month: For each region and month, we compute simple (unweighted) arithmetic means of weather variables across the counties assigned to that region. No population or area weighting is applied; each county within a region receives equal weight.

**Geographic coverage.** Of 210 target regions, 206 appear in the final panel with complete 60-month time series (12 months × 5 years). The four missing regions—Anchorage, AK; Fairbanks, AK; Juneau, AK; and Honolulu, HI—are absent because the weather data files for Alaska and Hawaii counties, while present in the raw data, do not match to MSA titles in the QCEW crosswalk for these regions.

# IA.2. 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:
## OLS estimates of ChatGPT adoption on browsing activity by category

This table shows estimates of long difference OLS regressions comparing 2024 browsing characteristics to 2022 browsing characteristics, in an estimatin of the form

$$\Delta \text{BrowsingOutcome}_i = \gamma \, \text{ChatGPTUse}_i + \lambda_{FEs} + \varepsilon_i .$$

The only RHS variable is a dummy for whether household $j$ has ever used ChatGPT by the end of 2024. The dependent variable in each regression is a measure of the household's browsing activity in different categories. In Column (1) it is the quarterly log of browsing duration across all sites; in Columns (2),(3) and (4) it is the quarterly log of browsing duration on leisure, productive, and mixed sites; and in Columns (5), (6), and (7) the change in the share of the household's total browsing that is on leisure, productive, and mixed sites. All regressions include fixed effects for income bin X age bin X MSA. T-statistics based on heteroskedasticity-robust standard errors clustered at the income bin X age bin X MSA level are in parentheses: * p<0.10, ** p<0.05, *** p<0.01.

| *Specification:* | *Long differences OLS: 2024 vs. 2022* | | | | | | |
|---|---|---|---|---|---|---|---|
| *Dep. var.* | Δ Log Duration All | Δ Log Duration Leisure | Δ Log Duration Productive | Δ Log Duration Mixed | Δ Share Leisure | Δ Share Productive | Δ Share Mixed |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| ChatGPT Use | 0.508*** | 0.392*** | 0.539*** | 0.407*** | -0.009** | 0.011** | -0.002 |
| | (14.566) | (8.480) | (14.836) | (9.443) | (-2.095) | (2.281) | (-0.608) |
| Observations | 43,874 | 43,874 | 43,874 | 43,874 | 43,874 | 43,874 | 43,874 |
| Estimation | OLS | OLS | OLS | OLS | OLS | OLS | OLS |
| Inc. X Age X MSA FEs | X | X | X | X | X | X | X |

## Table IA.3:
## Effect of ChatGPT adoption on browsing activity by Age

This table shows estimates of long difference IV regressions comparing 2024 browsing characteristics to 2022 browsing characteristics, in an estimatin of the form

$$\Delta \text{BrowsingOutcome}_i = \gamma \, \text{ChatGPTUse}_i + \lambda_{FEs} + X_i'\xi + \varepsilon_i.$$

The endogenous variable is a dummy for whether household $j$ has ever used ChatGPT by the end of 2024. The instrument is the panelist household's predicted exposure to ChatGPT in 2021, based on the household's share of browsing on highly substitutable sites. The dependent variable in each regression is a measure of the household's browsing activity in different categories. In Column (1) it is the quarterly log of browsing duration across all sites; in Columns (2),(3) and (4) it is the quarterly log of browsing duration on leisure, productive, and mixed sites; and in Columns (5), (6), and (7) the change in the share of the household's total browsing that is on leisure, productive, and mixed sites. All regressions include fixed effects for income bin X age bin X MSA. All regressions also include control variables for the generative AI exposure predicted by the households 2021 browsing composition by Comscore content category, and for the sum of the browsing share of websites with a GenAI exposure label in 2021 - the 'Bartik share control'. T-statistics based on heteroskedasticity-robust standard errors clustered at the income bin X age bin X MSA level are in parentheses: * p<0.10, ** p<0.05, *** p<0.01.

| *Specification:* | *Long differences: 2024 vs. 2022* | | | | | | |
|---|---|---|---|---|---|---|---|
| *Dep. var.* | Δ Log Duration All | Δ Log Duration Leisure | Δ Log Duration Productive | Δ Log Duration Mixed | Δ Share Leisure | Δ Share Productive | Δ Share Mixed |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| *Panel A: Young Households* | | | | | | | |
| ChatGPT Use | 0.314 | 2.375 | -1.590 | -1.023 | 0.366* | -0.514** | 0.147 |
| | (0.248) | (1.270) | (-1.130) | (-0.588) | (1.798) | (-2.183) | (0.967) |
| Observations | 6,924 | 6,924 | 6,924 | 6,924 | 6,924 | 6,924 | 6,924 |
| 1st-stage KP F-stat. | 11 | 11 | 11 | 11 | 11 | 11 | 11 |
| *Panel B: Middle Age Households* | | | | | | | |
| ChatGPT Use | 0.504 | 1.238 | -0.540 | 0.030 | 0.334*** | -0.305*** | -0.029 |
| | (0.873) | (1.416) | (-0.872) | (0.041) | (3.838) | (-3.124) | (-0.444) |
| Observations | 20,260 | 20,260 | 20,260 | 20,260 | 20,260 | 20,260 | 20,260 |
| 1st-stage KP F-stat. | 88 | 88 | 88 | 88 | 88 | 88 | 88 |
| *Panel C: Old Age Households* | | | | | | | |
| ChatGPT Use | 3.356*** | 1.351 | 2.302* | -0.277 | 0.206 | 0.190 | -0.395*** |
| | (2.692) | (0.922) | (1.830) | (-0.207) | (1.342) | (1.207) | (-2.972) |
| Observations | 15,702 | 15,702 | 15,702 | 15,702 | 15,702 | 15,702 | 15,702 |
| 1st-stage KP F-stat. | 34 | 34 | 34 | 34 | 34 | 34 | 34 |
| Estimation | IV | IV | IV | IV | IV | IV | IV |
| Inc. X Age X MSA FEs | X | X | X | X | X | X | X |
| Bartik share control | X | X | X | X | X | X | X |
| Website category exposure | X | X | X | X | X | X | X |

## Table IA.4:
## Effect of ChatGPT adoption on browsing activity by Income

This table shows estimates of long difference IV regressions comparing 2024 browsing characteristics to 2022 browsing characteristics, in an estimatin of the form

$$\Delta\text{BrowsingOutcome}_i = \gamma\,\text{ChatGPTUse}_i + \lambda_{FEs} + X_i'\xi + \varepsilon_i.$$

The endogenous variable is a dummy for whether household $j$ has ever used ChatGPT by the end of 2024. The instrument is the panelist household's predicted exposure to ChatGPT in 2021, based on the household's share of browsing on highly substitutable sites. The dependent variable in each regression is a measure of the household's browsing activity in different categories. In Column (1) it is the quarterly log of browsing duration across all sites; in Columns (2),(3) and (4) it is the quarterly log of browsing duration on leisure, productive, and mixed sites; and in Columns (5), (6), and (7) the change in the share of the household's total browsing that is on leisure, productive, and mixed sites. All regressions include fixed effects for income bin X age bin X MSA. All regressions also include control variables for the generative AI exposure predicted by the households 2021 browsing composition by Comscore content category, and for the sum of the browsing share of websites with a GenAI exposure label in 2021 - the 'Bartik share control'. T-statistics based on heteroskedasticity-robust standard errors clustered at the income bin X age bin X MSA level are in parentheses: * p<0.10, ** p<0.05, *** p<0.01.

| *Specification:* | *Long differences: 2024 vs. 2022* | | | | | | |
|---|---|---|---|---|---|---|---|
| *Dep. var.* | Δ Log Duration All | Δ Log Duration Leisure | Δ Log Duration Productive | Δ Log Duration Mixed | Δ Share Leisure | Δ Share Productive | Δ Share Mixed |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| *Panel A: Low-Income Households* | | | | | | | |
| ChatGPT Use | 3.132*** | 4.530*** | 0.887 | 1.624 | 0.407*** | -0.355** | -0.052 |
| | (2.996) | (3.459) | (0.853) | (1.310) | (2.873) | (-2.469) | (-0.471) |
| Observations | 16,489 | 16,489 | 16,489 | 16,489 | 16,489 | 16,489 | 16,489 |
| 1st-stage KP F-stat. | 55 | 55 | 55 | 55 | 55 | 55 | 55 |
| *Panel B: Middle-Income Households* | | | | | | | |
| ChatGPT Use | 0.992 | 0.716 | -0.190 | -0.123 | 0.257** | -0.227* | -0.030 |
| | (1.117) | (0.639) | (-0.195) | (-0.117) | (2.102) | (-1.685) | (-0.328) |
| Observations | 12,941 | 12,941 | 12,941 | 12,941 | 12,941 | 12,941 | 12,941 |
| 1st-stage KP F-stat. | 30 | 30 | 30 | 30 | 30 | 30 | 30 |
| *Panel C: High-Income Households* | | | | | | | |
| ChatGPT Use | 0.443 | 0.614 | -0.272 | -1.433 | 0.296** | -0.128 | -0.168* |
| | (0.549) | (0.488) | (-0.336) | (-1.335) | (2.407) | (-1.014) | (-1.752) |
| Observations | 13,456 | 13,456 | 13,456 | 13,456 | 13,456 | 13,456 | 13,456 |
| 1st-stage KP F-stat. | 39 | 39 | 39 | 39 | 39 | 39 | 39 |
| Estimation | IV | IV | IV | IV | IV | IV | IV |
| Inc. X Age X MSA FEs | X | X | X | X | X | X | X |
| Bartik share control | X | X | X | X | X | X | X |
| Website category exposure | X | X | X | X | X | X | X |

## Table IA.5:
## Full Engel Curve Estimation Results

This table shows estimates of age-income-region-level panel regressions of the form

$$s_{it}^{a} = \delta_{t}^{a} + \gamma^{a} \ln H_{it} + u_{it}^{a},$$

where $\ln H_{it}$ is the log of total browsing time, and $s_{it}^{a}$ is the browsing share of activity $a$ (e.g., leisure or productive activity). In columns 3 and 4, the dependent variable is instead the log of the browsing time for each activity category - the estimates are identical to those in Table IX. Implied Engel elasticities are computed in columns 1 and 2 as $\hat{\beta}^{a}_{\text{Engel}} = 1 + \frac{\hat{\gamma}^{a}}{\bar{s}^{a}}$. Columns (2) and (4) show IV estimates where the instrument is the log of total precipitation in the region in the year-quarter. All regressions include fixed effects that capture the interaction between income and age categories, and also fixed effects in the time dimension (year-quarter level). Heteroskedasticity-robust standard errors double clustered at the age-income-region level and at the year-quarter level in parentheses: * p<0.10, ** p<0.05, *** p<0.01.

| *Dep. var* | Browsing share | | Log Activity Time | |
|---|---|---|---|---|
| Estimation | OLS | IV | OLS | IV |
| | (1) | (2) | (3) | (4) |
| *Panel A: Productive digital activities* | | | | |
| Log(Total Browsing Time) | -0.015*** | -0.045*** | 0.991*** | 0.931*** |
| | (-20.871) | (-3.520) | (530.879) | (29.244) |
| Observations | 129,498 | 99,608 | 129,146 | 99,430 |
| 1st-stage KP F-stat. | | 55 | | 54 |
| Implied Engel Elasticity | 0.98 | 0.93 | 0.99 | 0.93 |
| *Panel B: Leisure digital activities* | | | | |
| Log(Total Browsing Time) | 0.017*** | 0.038*** | 1.130*** | 1.374*** |
| | (28.441) | (3.663) | (241.949) | (15.823) |
| Observations | 129,498 | 99,608 | 120,785 | 94,691 |
| 1st-stage KP F-stat. | | 55 | | 48 |
| Implied Engel Elasticity | 1.09 | 1.20 | 1.13 | 1.37 |
| *Panel C: Other digital activities* | | | | |
| Log(Total Browsing Time) | -0.002*** | 0.008 | 1.026*** | 1.110*** |
| | (-3.749) | (0.857) | (276.169) | (17.912) |
| Observations | 129,498 | 99,608 | 125,862 | 97,739 |
| 1st-stage KP F-stat. | | 55 | | 51 |
| Implied Engel Elasticity | 0.99 | 1.04 | 1.03 | 1.11 |
| Age x Income Group FEs | X | X | X | X |
| Time FEs | X | X | X | X |

## Table IA.6:
## Websites with lowest relative prevalence around ChatGPT use.

This table shows websites with the lowest prevalence in 30 minute intervals with ChatGPT use compared to non-user browsing matched on time of day and age bin × income bin demographics.

| Rank | Domain | Δ Browsing % | Rank | Domain | Δ Browsing % |
|---|---|---|---|---|---|
| 1 | youtube.com | -5 | 16 | fanfiction.net | -.46 |
| 2 | msn.com | -1.8 | 17 | xhamster.com | -.44 |
| 3 | yahoo.com | -1.7 | 18 | espn.com | -.38 |
| 4 | facebook.com | -1.7 | 19 | crunchyroll.com | -.33 |
| 5 | aol.com | -1.2 | 20 | bing.com | -.33 |
| 6 | amazon.com | -1.1 | 21 | centurygames.com | -.31 |
| 7 | capitaloneshopping.com | -1.1 | 22 | walmart.com | -.28 |
| 8 | pornhub.com | -1 | 23 | tubitv.com | -.27 |
| 9 | chaturbate.com | -.9 | 24 | intuit.com | -.27 |
| 10 | ebay.com | -.7 | 25 | xnxx.com | -.26 |
| 11 | live.com | -.65 | 26 | weather.com | -.25 |
| 12 | sprig.com | -.56 | 27 | searchmask.com | -.24 |
| 13 | xvideos.com | -.51 | 28 | instagram.com | -.23 |
| 14 | popin.cc | -.5 | 29 | wikipedia.org | -.23 |
| 15 | deviantart.com | -.48 | 30 | kick.com | -.23 |

## Table IA.7:
## Generative AI exposure by website category.

This table shows the average 2021 duration-weighted mean exposure score (based on website activities) for websites in each usage category.

| Usage type | Activity exposure score |
|---|---|
| Leisure | 0.177 |
| Mixed-use | 0.299 |
| Productive | 0.507 |

# IA.3. Additional Figures

**Figure IA.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.

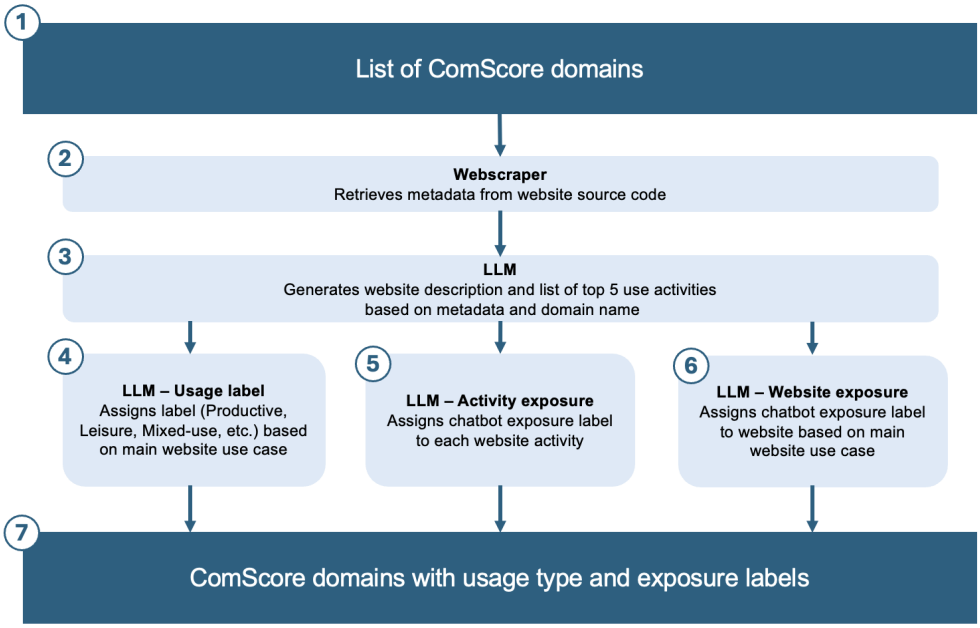
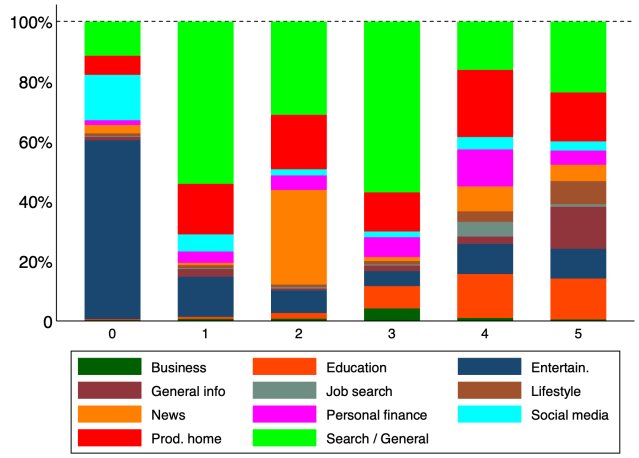
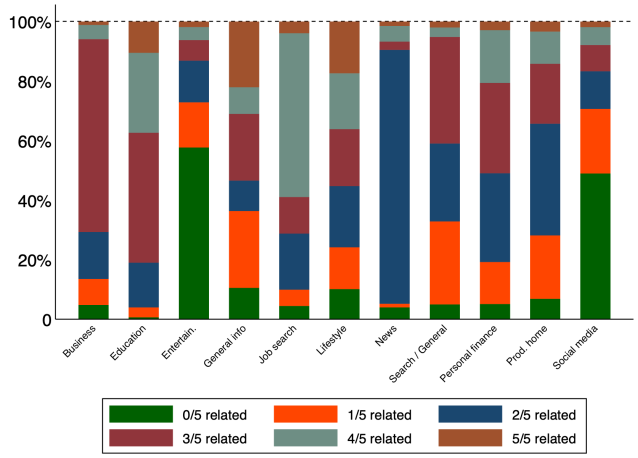



**Figure IA.2:** Duration-weighted double sorts of website overlap fractions and Comscore activity categories



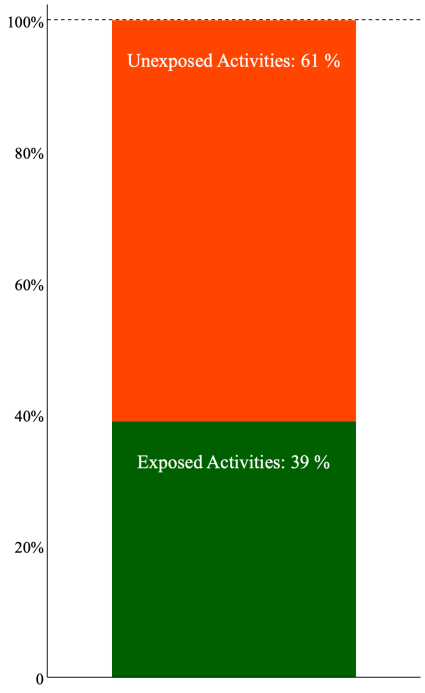

**(A)** Duration shares of overlap fractions



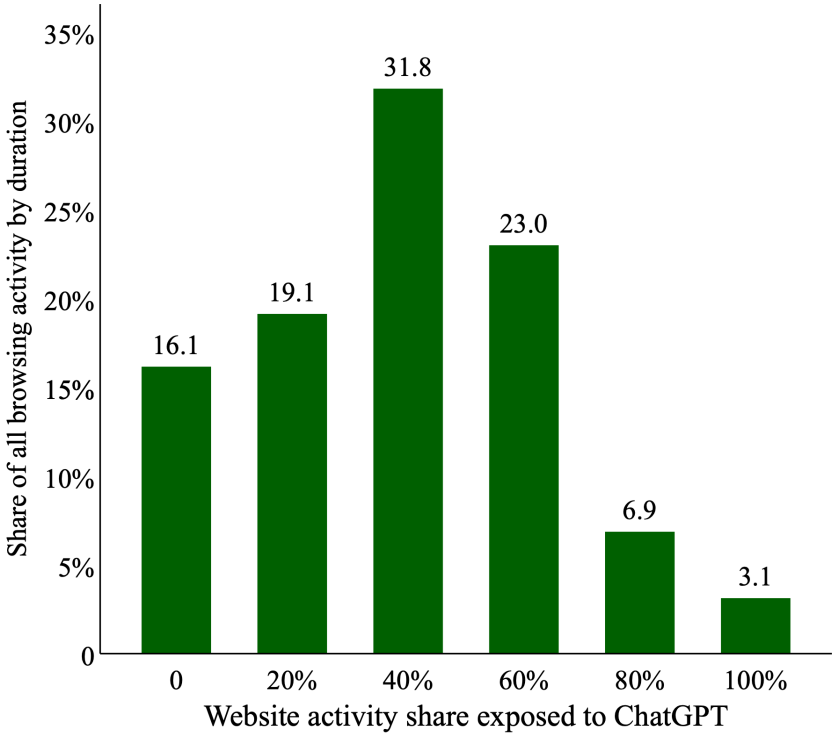

**(B)** Duration shares of site activity categories

**Figure IA.3:**
**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.4:** Pre-ChatGPT release duration shares for different site categories, by household income bucket

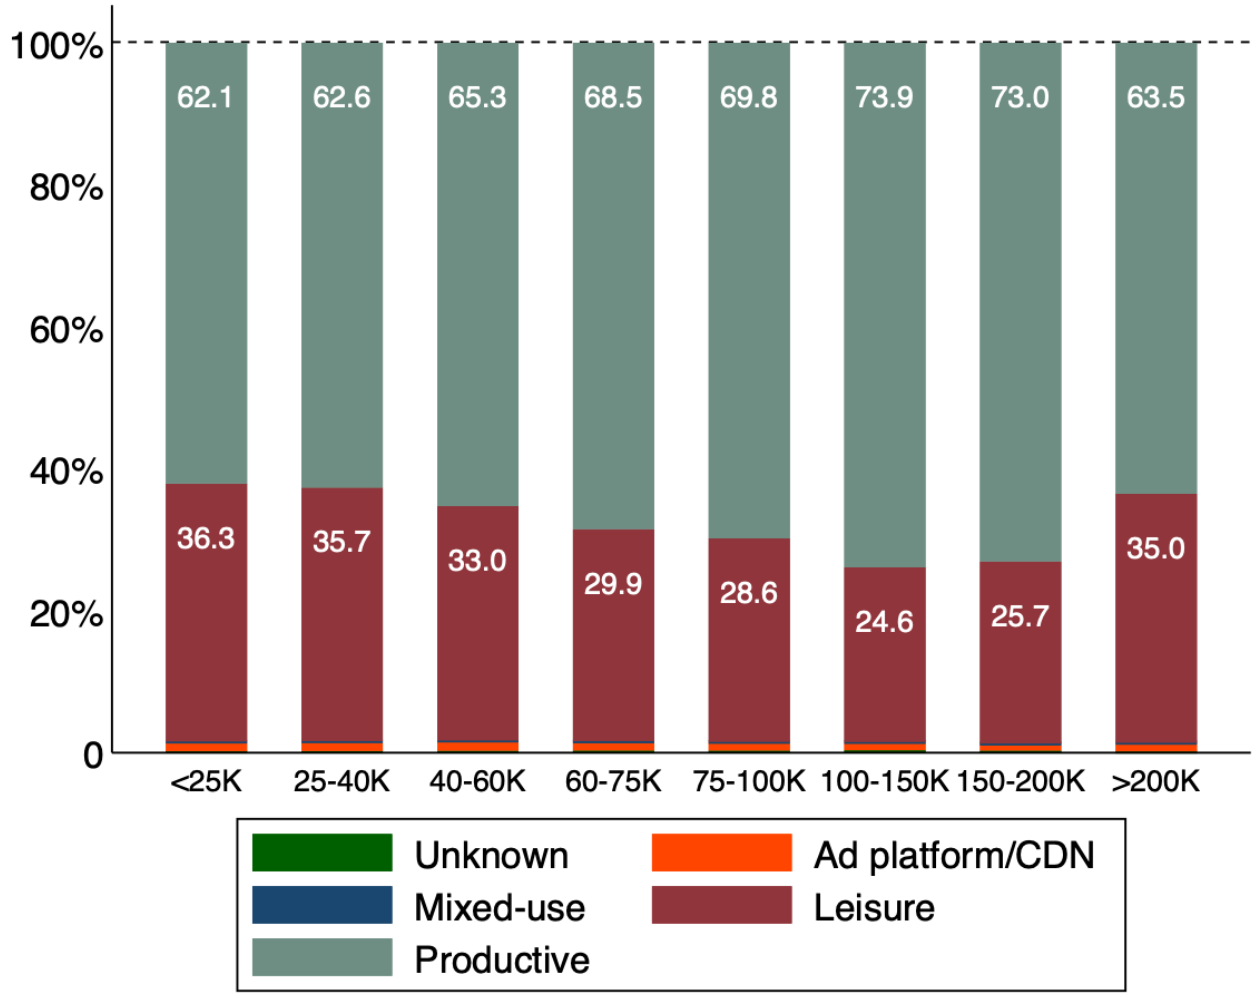
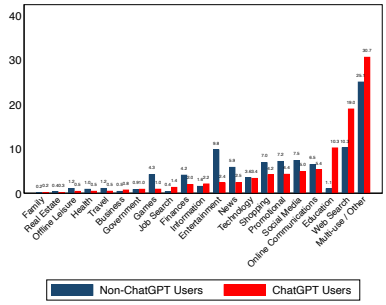
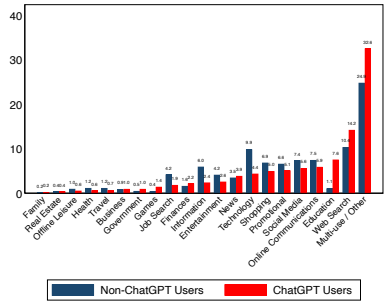


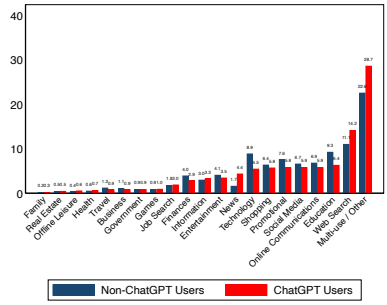
**Figure IA.5: Usage of sites associated with different activities in vs. outside ChatGPT adoption windows**



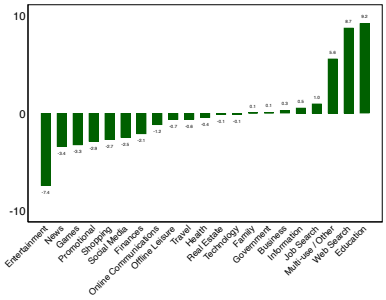

**(A)** GPT window: browsing shares



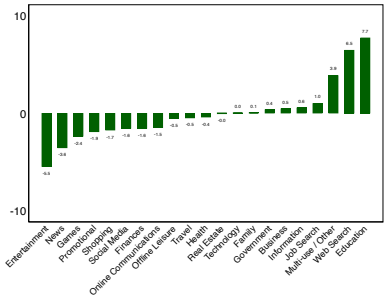

**(B)** GPT day, outside window: browsing shares



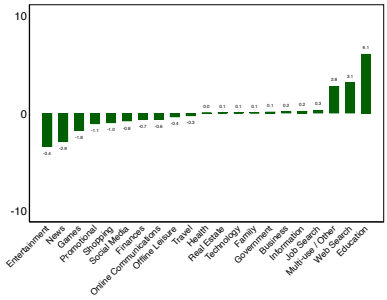

**(C)** Pre-GPT use: browsing shares



**(D)** GPT window: differences in shares

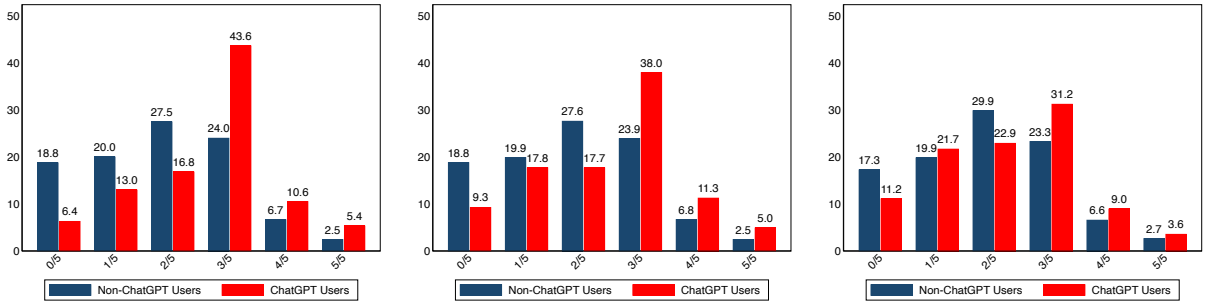


**(E)** GPT day, outside window: differences in shares



**(F)** Pre-GPT use: differences in shares

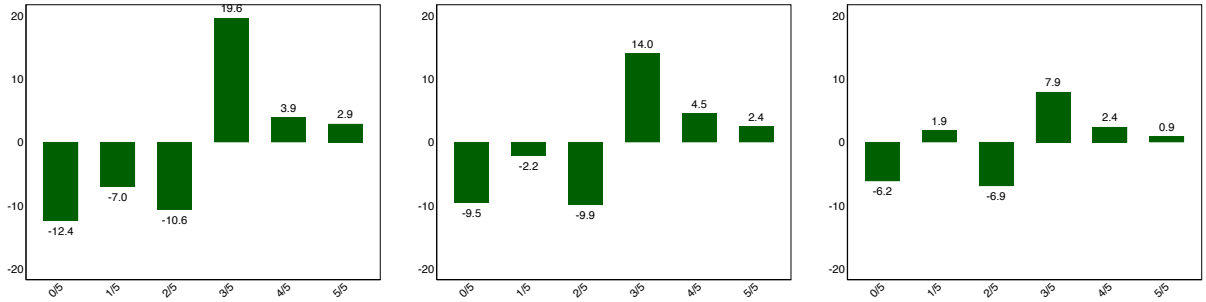
**Figure IA.6: Usage of sites with different GenAI overlap in vs. outside ChatGPT adoption windows**




**(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

# IA.4. Derivations

This appendix provides the step-by-step derivations underlying the conceptual framework in Section V.

## A. *Browsing activity demand system derivation*

**Connection to Aguiar et al. (2021) demand system** Aguiar et al. (2021) introduce an additively separable aggregator over time allocated to activities,

$$v(h;\theta,\xi) = \sum_{j=1}^{J} \frac{\left(\theta_j \xi_j h_j\right)^{1-\frac{1}{\eta_j}}}{1-\frac{1}{\eta_j}},$$

embedded in an outer utility $U(c,v)$. The key static subproblem is: allocate a fixed time endowment across activities to maximize $v$. Our framework is the special case with $J=2$ activities, $(\ell,z)$, and we work directly with $v$ (i.e., no outer $U(\cdot)$). This produces the same demand system mapping activity "quality" $(\theta\xi)$ and the shadow value of time into optimal time allocations.

**Household problem.** At time $t$, household $i$ chooses $(z_{it}, \ell_{it})$ to maximize

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

subject to the time budget

$$z_{it} + \ell_{it} \leq 1. \tag{2}$$

For $\eta^{\ell}, \eta^{z} > 0$, each term is increasing in its time argument, so the constraint binds at the optimum:

$$z_{it} + \ell_{it} = 1, \qquad \omega_{it} > 0.$$

**First-order conditions** Suppressing $(i,t)$ subscripts for readability, the Lagrangian is

$$\mathscr{L} = \frac{\left(\theta^{\ell}\xi^{\ell}\ell\right)^{1-\frac{1}{\eta^{\ell}}}}{1-\frac{1}{\eta^{\ell}}} + \frac{(\theta^{z}\xi^{z}z)^{1-\frac{1}{\eta^{z}}}}{1-\frac{1}{\eta^{z}}} + \omega(1-z-\ell),$$

where $\omega$ is the Lagrange multiplier (shadow value of one unit of digital time).

Define the exponents $\alpha_z \equiv 1-\frac{1}{\eta^{z}}$ and $\alpha_{\ell} \equiv 1-\frac{1}{\eta^{\ell}}$. Differentiating the $z$-term and using the chain rule:

$$\frac{d}{dz}\left[\frac{(\theta^{z}\xi^{z}z)^{\alpha_z}}{\alpha_z}\right] = (\theta^{z}\xi^{z}z)^{\alpha_z-1}(\theta^{z}\xi^{z}).$$

Since $\alpha_z - 1 = -\frac{1}{\eta^{z}}$, we have

$$(\theta^{z}\xi^{z}z)^{\alpha_z-1} = (\theta^{z}\xi^{z})^{-\frac{1}{\eta^{z}}} z^{-\frac{1}{\eta^{z}}},$$

and the first-order condition for $z$ is

$$(\theta^{z}\xi^{z})^{1-\frac{1}{\eta^{z}}}\, z^{-\frac{1}{\eta^{z}}} = \omega.$$

Taking logs and solving for $\ln z$:

$$\left(1-\frac{1}{\eta^{z}}\right)(\ln\theta^{z} + \ln\xi^{z}) - \frac{1}{\eta^{z}}\ln z = \ln\omega,$$

$$\ln z = (\eta^{z}-1)(\ln\theta^{z} + \ln\xi^{z}) - \eta^{z}\ln\omega.$$

Thus,

$$\ln(z_{it}) = (\eta^{z}-1)\left(\ln\theta^{z}_{it} + \ln\xi^{z}_{it}\right) - \eta^{z}\ln\omega_{it}. \tag{3}$$

Analogously, the FOC for $\ell$ yields

$$\ln(\ell_{it}) = (\eta^{\ell}-1)\left(\ln\theta^{\ell}_{it} + \ln\xi^{\ell}_{it}\right) - \eta^{\ell}\ln\omega_{it}. \tag{4}$$

**Activity gap condition.** Dividing (3) by $\eta^z$ and (4) by $\eta^\ell$, then subtracting the second from the first to eliminate $\ln \omega_{it}$:

$$\frac{\ln(z_{it})}{\eta^z} - \frac{\ln(\ell_{it})}{\eta^\ell} = \frac{\eta^z - 1}{\eta^z}\left(\ln \theta^z_{it} + \ln \xi^z_{it}\right) - \frac{\eta^\ell - 1}{\eta^\ell}\left(\ln \theta^\ell_{it} + \ln \xi^\ell_{it}\right). \tag{5}$$

**Interpretation:** Holding $\omega_{it}$ fixed, a rise in an activity's effective "quality" ($\theta\xi$) increases time in that activity if $\eta > 1$ (a "time luxury") and decreases time if $0 < \eta < 1$ (a "time necessity"). The knife-edge case $\eta = 1$ corresponds to log utility in that activity.

## *B. First-order approximation of the benefit of ChatGPT adoption*

**Technology shift.** Let $\xi^z_{i,pre} = \underline{\xi}$ and, in *post*, if the household adopts,

$$\xi^z_{i,post} = \underline{\xi}(1+\delta),$$

while if it does not adopt, $\xi^z_{i,post} = \underline{\xi}$. All other objects are held fixed. Let $(z^N, \ell^N)$ denote the post-period optimum under No-adopt, and $(z^A, \ell^A)$ under Adopt.

**Envelope approximation.** For small $\delta$, approximate the utility gain from adoption by evaluating the direct effect of changing $\xi^z$ at the no-adopt choice $(z^N, \ell^N)$. The $z$-component of utility is

$$u_z(\xi^z) = \frac{(\theta^z \xi^z z^N)^{\alpha_z}}{\alpha_z}, \qquad \alpha_z = 1 - \frac{1}{\eta^z}.$$

Under adoption, $\xi^z$ becomes $\xi^z(1+\delta)$, giving a utility gain

$$\Delta u_z = \frac{((1+\delta)^{\alpha_z} - 1)(\theta^z \xi^z z^N)^{\alpha_z}}{\alpha_z}.$$

Using the first-order expansion $(1+\delta)^{\alpha_z} \approx 1 + \alpha_z \delta$ for small $\delta$:

$$\Delta u_z \approx \delta(\theta^z \xi^z z^N)^{\alpha_z} = \delta \cdot \left(\theta^z \underline{\xi} z^N\right)^{\frac{\eta^z - 1}{\eta^z}}.$$

From the $z$ FOC, we have $\omega^N = (\theta^z \xi^z)^{1 - \frac{1}{\eta^z}} (z^N)^{-\frac{1}{\eta^z}}$. Multiplying both sides by $z^N$:

$$z^N \omega^N = (\theta^z \xi^z z^N)^{1 - \frac{1}{\eta^z}} = (\theta^z \xi^z z^N)^{\frac{\eta^z - 1}{\eta^z}}.$$

Therefore the adoption gain can be written as

$$v^A - v^N \approx \delta \cdot z^N \omega^N. \tag{6}$$

**Adoption condition and time-cost interpretation.** If the (utility) cost of adopting is $c$, the adoption rule is

$$v^A - v^N \geq c \quad \Rightarrow \quad \delta \cdot z^N \omega^N \geq c.$$

If instead the cost is a time cost $c^{time}$ (which has utility cost $c^{time}\omega^N$), the adoption condition becomes

$$\delta z^N \geq c^{time.}$$

Thus, to first order, adoption compares "time saved" $\delta z^N$ to time required to learn/use the technology.

## *C. Time-allocation treatment effects of ChatGPT adoption*

**Closed-form demand system** From the FOCs, each activity's time choice can be expressed as

$$z = (\theta^z \xi^z)^{\eta^z - 1} \omega^{-\eta^z}, \qquad \ell = (\theta^\ell \xi^\ell)^{\eta^\ell - 1} \omega^{-\eta^\ell}. \tag{7}$$

**Log-differentials.** Let No-adopt be $N$ and Adopt be $A$. Define the log treatment effects:

$$\beta_z^{GPT} \equiv \ln z^A - \ln z^N, \qquad \beta_\ell^{GPT} \equiv \ln \ell^A - \ln \ell^N.$$

Adoption changes $\xi^z$ by factor $(1+\delta)$, so $d\ln \xi^z = \ln(1+\delta) \approx \delta$, while $\xi^\ell$ is unchanged.

Log-differentiating (3) and (4) with respect to $\xi^z$:

$$d\ln z = (\eta^z - 1) d\ln \xi^z - \eta^z d\ln \omega,$$
$$d\ln \ell = -\eta^\ell d\ln \omega.$$

The binding time constraint $z + \ell = 1$ implies $dz + d\ell = 0$, or $z\, d\ln z + \ell\, d\ln \ell = 0$. Substituting the differentials above:

$$z\left[(\eta^z - 1) d\ln \xi^z - \eta^z d\ln \omega\right] + \ell\left[-\eta^\ell d\ln \omega\right] = 0,$$

which yields

$$d\ln \omega = \frac{z(\eta^z - 1)}{z\eta^z + \ell\eta^\ell} d\ln \xi^z.$$

**Treatment effects.** Substituting back into the differential for $z$ and simplifying:

$$\begin{aligned} d\ln z &= (\eta^z - 1) d\ln \xi^z \left[1 - \frac{\eta^z z}{z\eta^z + \ell\eta^\ell}\right] \\ &= (\eta^z - 1) d\ln \xi^z \left[\frac{\ell\eta^\ell}{z\eta^z + \ell\eta^\ell}\right] \\ &= (\eta^z - 1) \frac{1}{1 + \left(\eta^z/\eta^\ell\right)(z/\ell)} d\ln \xi^z. \end{aligned}$$

Evaluating at $(z^N, \ell^N)$ with $d\ln\xi^z \approx \delta$:

$$\beta_z^{GPT} \approx \frac{\eta^z - 1}{1 + \left(\eta^z/\eta^\ell\right)\left(z^N/\ell^N\right)} \cdot \delta. \tag{8}$$

Similarly, for leisure:

$$\begin{aligned} d\ln \ell &= -\eta^\ell \cdot \frac{z(\eta^z - 1)}{z\eta^z + \ell\eta^\ell} d\ln \xi^z \\ &= (1 - \eta^z) \cdot \frac{\eta^\ell/\eta^z}{1 + \left(\eta^\ell/\eta^z\right)(\ell/z)} d\ln \xi^z, \end{aligned}$$

yielding

$$\beta_\ell^{GPT} \approx \frac{(1 - \eta^z) \cdot \left(\eta^\ell/\eta^z\right)}{1 + \left(\eta^\ell/\eta^z\right)\left(\ell^N/z^N\right)} \cdot \delta. \tag{9}$$

**Relative impact.** Dividing the log-differentials by their respective elasticities and subtracting eliminates $d\ln\omega$:

$$\frac{\beta_z^{GPT}}{\eta^z} - \frac{\beta_\ell^{GPT}}{\eta^\ell} \approx \frac{\eta^z - 1}{\eta^z} \cdot \delta. \tag{10}$$

or, without the first-order approximation:

$$\frac{\beta_z^{GPT}}{\eta^z} - \frac{\beta_\ell^{GPT}}{\eta^\ell} \approx \frac{\eta^z - 1}{\eta^z} \cdot \ln(1 + \delta). \tag{11}$$

## *D. Digital-time Engel Curves and the Implied GenAI Productivity Shock*

This appendix summarizes how we use an Engel-curve ("time demand system") approach—following Aguiar et al. (2021)—to (i) estimate relative curvature parameters for different digital activities and (ii) translate observed adoption-induced reallocations of digital time into an implied productivity shock from generative AI.

## *E. A digital-time demand system*

Let the household allocate *total* digital time $H_{it}$ across three mutually exclusive categories: digital leisure $\ell_{it}$, digital productive home-production $z_{it}$, and other digital time $o_{it}$:

$$\ell_{it} + z_{it} + o_{it} = H_{it}. \tag{12}$$

(In the stylized model in the main text we normalize $H_{it} = 1$; here we keep $H_{it}$ explicit because it is useful for Engel-curve estimation and welfare calculations.)

Preferences over the composition of digital time are summarized by the additively separable "digital activity aggregator"

$$v(\ell_{it}, z_{it}, o_{it}; \theta_{it}, \xi_{it}) = \sum_{a \in \{\ell, z, o\}} \frac{\left(\theta^a_{it} \xi^a_{it} h^a_{it}\right)^{1-\frac{1}{\eta^a}}}{1-\frac{1}{\eta^a}} \qquad \text{where} \qquad (h^\ell_{it}, h^z_{it}, h^o_{it}) \equiv (\ell_{it}, z_{it}, o_{it}). \tag{13}$$

The parameter $\eta^a > 0$ governs diminishing returns in activity $a$: for $\eta^a > 1$, $a$ is relatively "luxurious" in that higher efficiency can raise time allocated to $a$; for $\eta^a < 1$, $a$ behaves like a "necessity" in that efficiency improvements tend to reduce time allocated to $a$ (holding total digital time fixed), pushing time into other activities.

Given total digital time $H_{it}$, the household solves the within-digital-time allocation subproblem

$$\max_{\{\ell_{it}, z_{it}, o_{it}\}} v(\ell_{it}, z_{it}, o_{it}; \theta_{it}, \xi_{it}) \quad \text{s.t.} \quad \ell_{it} + z_{it} + o_{it} = H_{it}. \tag{14}$$

Let $\hat{\omega}_{it}$ denote the *normalized* shadow price of digital time (the Lagrange multiplier on (12) divided by the marginal utility of the aggregator if $v$ is embedded in a higher-level utility function).[15] The first-order condition for activity $a$ implies a constant-elasticity "demand" for time in activity $a$:

$$h^a_{it} = \left(\theta^a_{it} \xi^a_{it}\right)^{\eta^a - 1} \hat{\omega}_{it}^{-\eta^a}. \tag{15}$$

Taking logs yields

$$\ln h^a_{it} = (\eta^a - 1)\left(\ln \theta^a_{it} + \ln \xi^a_{it}\right) - \eta^a \ln \hat{\omega}_{it}. \tag{16}$$

Equation (16) is the key object: it shows how (i) activity-specific shifters $(\theta^a, \xi^a)$ and (ii) the common shadow price of digital time $\hat{\omega}$ jointly determine the allocation of time across activities.

[15] This normalization is convenient because it drops out of relative demands across activities. See Aguiar et al. (2021) for details.

## F. *Engel curves in time*

Define the share of digital time devoted to activity $a$ as

$$s_{it}^{a} \equiv \frac{h_{it}^{a}}{H_{it}}, \qquad \sum_{a} s_{it}^{a} = 1. \tag{17}$$

Let $\bar{\eta}_{it}$ denote the share-weighted average curvature parameter:

$$\bar{\eta}_{it} \equiv \sum_{a \in \{\ell, z, o\}} s_{it}^{a} \eta^{a}. \tag{18}$$

The model implies an Engel elasticity ("leisure Engel curve" in Aguiar et al. (2021)) for each activity:

$$\beta^{a} \equiv \frac{\partial \ln h_{it}^{a}}{\partial \ln H_{it}} = \frac{\eta^{a}}{\bar{\eta}_{it}}. \tag{19}$$

Thus $\beta^{a} > 1$ means that activity $a$ increases *more than proportionally* with total digital time $H$ (a "time luxury"), while $\beta^{a} < 1$ means it increases *less than proportionally* (a "time necessity").

The empirical advantage of (19) is that it identifies *relative* curvature parameters. In particular,

$$\frac{\eta^{a}}{\eta^{b}} = \frac{\beta^{a}}{\beta^{b}} \qquad \text{for any activities } a, b. \tag{20}$$

This log-log relationship can be used directly to estimate the Engel curve elasticities from

$$\ln h_{it}^{a} = \delta_{t}^{a} + X_{it}' \alpha^{a} + \beta^{a} \ln H_{it} + u_{it}^{a},$$

where we use an IV approach based on rainfall shocks to identify exogenous changes in $\ln H_{it}$.

**Estimating Engel curves with an AIDS-style share equation.** Following Aguiar et al. (2021), a convenient empirical specification uses shares $s_{it}^{a}$ and regresses them on $\ln H_{it}$ (plus controls and fixed effects). A first-order approximation of the demand system implies an equation of the form

$$s_{it}^{a} = \delta_{t}^{a} + X_{it}' \alpha^{a} + \gamma^{a} \ln H_{it} + u_{it}^{a}, \tag{21}$$

where $\delta_{t}^{a}$ are time fixed effects and $X_{it}$ collects additional controls (demographics, region/industry cells, etc.). The coefficient $\gamma^{a}$ maps into the Engel elasticity $\beta^{a}$ via

$$\gamma^{a} = \bar{s}^{a} (\beta^{a} - 1) \qquad \Longrightarrow \qquad \widehat{\beta}^{a} = 1 + \frac{\widehat{\gamma}^{a}}{\bar{s}^{a}}, \tag{22}$$

where $\bar{s}^{a}$ is the sample mean share for activity $a$ in the estimation sample.

With three activities, the add-up restriction $\sum_{a} s_{it}^{a} = 1$ implies $\sum_{a} \gamma^{a} = 0$ (and likewise for the fixed effects). In practice, one can (i) estimate (21) for two activities and recover the third by add-

up, or (ii) estimate all three equations and check that the estimated coefficients approximately satisfy add-up.

**From Engel elasticities to the curvature parameters $\eta^a$** The Engel elasticities $\{\beta^a\}$ identify $\{\eta^a\}$ only up to a common scale, because $\beta^a = \eta^a/\bar{\eta}$. To obtain *levels* of $\eta^a$, we would have to impose (or calibrate) a value for the average curvature $\bar{\eta}$:

$$\eta^a = \beta^a \bar{\eta}. \tag{23}$$

Operationally, this means that once we choose $\bar{\eta}$ from the literature (or from an auxiliary calibration), the Engel-curve estimates deliver $\eta^\ell$, $\eta^z$, and $\eta^o$.

## G. *Backing out the GenAI productivity shock $\delta$ from adoption-induced time changes.*

Suppose adopting generative AI increases the efficiency of productive digital tasks only:

$$\xi_{it}^{z,\text{Adopt}} = (1+\delta)\,\xi_{it}^{z,\text{No}}, \qquad \xi_{it}^{\ell,\text{Adopt}} = \xi_{it}^{\ell,\text{No}}, \qquad \xi_{it}^{o,\text{Adopt}} = \xi_{it}^{o,\text{No}}, \tag{24}$$

and that $\theta^a$ does not change at adoption. Let

$$\beta_a^{GPT} \equiv \ln h_{it}^{a,\text{Adopt}} - \ln h_{it}^{a,\text{No}} \qquad (a \in \{\ell, z, o\}) \tag{25}$$

denote the causal (adoption-induced) change in log time allocated to activity $a$ in the post period.

Take the log first-order condition (16) for productive time $z$ and subtract the log FOC for a *reference* activity $r \in \{\ell, o\}$ whose technology is unchanged by adoption. Differencing across the Adopt and No-Adopt regimes eliminates the common shadow price $\hat{\omega}$ and isolates the technology shock, after using Engel-curve estimates to substitute $\eta^z/\eta^r = \beta^z/\beta^r$:

$$(\eta^z - 1)\ln(1+\delta) = \beta_z^{GPT} - \beta^z/\beta^r\, \beta_r^{GPT}. \tag{26}$$

Equation (26) has a direct interpretation: the adoption-induced increase in productive efficiency shows up as a residual increase in $z$ *relative* to the change in the reference activity $r$, after accounting for the fact that the shadow price of time affects activities with different elasticities $\eta^a$ differently.