

# What is Generative AI Worth?

Erik Brynjolfsson, Avinash Collis, Felix Eggers, Sophia Kazinnik, and David Nguyen\*

April 13, 2026

## Abstract

We estimate the consumer welfare gains from the rapid adoption of generative AI tools like ChatGPT, Gemini, Claude, or Copilot in the United States. Using online choice experiments, we elicit willingness to accept (WTA) compensation for giving up access to these AI chatbot tools for one month from representative samples of US adults, fielded in two waves, July 2025 and March 2026. We find that mean willingness to accept (WTA) increased from \$98 in 2025 to \$124.50 in 2026, a 27% increase, while the median value rose from \$3.4 to \$11.40. Combined with growth in the adult user base from 98 million to 115 million, these estimates imply that aggregate consumer surplus increased from \$116 billion to \$172 billion. This surplus substantially exceeds estimated revenues from generative AI in the United States, suggesting that consumers capture most of the welfare gains from these tools. We find substantial heterogeneity in valuations: usage frequency is the strongest predictor of WTA, followed by workplace use, and paid subscription status, with additional differences by gender, age, and ethnicity. Overall, the results suggest that generative AI is already generating substantial and rapidly growing welfare gains, even before its full effects on measured productivity and GDP are reflected in official statistics.

**Keywords:** Generative AI; consumer surplus; consumer welfare; willingness to accept; digital goods; online choice experiments.

**JEL Codes:** D12, D61, O33, C83, L86.

---

\*Brynjolfsson: Stanford University and NBER; Collis: Heinz College of Information Systems and Public Policy, Carnegie Mellon University; Eggers: Copenhagen Business School; Kazinnik: Stanford University; Nguyen: Stanford University and Economic Statistics Centre of Excellence. Contact: (davidngu@stanford.edu). We thank members of the Stanford Digital Economy Lab and Stanford Institute for Human-Centered AI (HAI) for valuable discussions. All remaining errors are our own. Authors listed alphabetically.

# 1 Introduction

Despite their rapid and widespread adoption, the effects of generative AI tools like ChatGPT, Gemini, Claude, or Copilot are not yet fully reflected in official economic statistics like GDP and productivity. This is a modern version of the Solow paradox: as with earlier waves of computing, the benefits of a transformative technology may take time to appear in aggregate statistics. Some of this delay is expected. The J-curve associated with general-purpose technologies implies that measured productivity gains often lag initial adoption (Brynjolfsson et al., 2021). But the gap also reflects a deeper measurement problem. Because most consumers pay little or nothing for generative AI services, market transactions capture only a small share of the value these tools create. The result is a widening gap between their contribution to GDP and the welfare they generate. Consumer surplus is a more direct way to measure those benefits and GDP-B provides a methodology for assessing its value (Brynjolfsson et al., 2025a).

To estimate this surplus, we fielded surveys in July 2025 ( $N = 1,491$ ) and March 2026 ( $N = 2,000$ ). Each survey included a single binary discrete choice experiment designed to measure respondents’ willingness to accept compensation for giving up access to generative AI tools for one month. Both samples are representative of the US adult population by age, gender, and ethnicity, with additional controls for education, income, and place of residence. We focus on willingness to accept (WTA) rather than willingness to pay (WTP) because respondents already possess access to these tools. For goods people already possess, WTA is the appropriate welfare measure, and the divergence between WTA and WTP is well established (Horowitz and McConnell, 2002).

Comparing the 2025 and 2026 surveys allows us to trace how consumer welfare evolved as adoption expanded and model quality improved. In early 2026, the average AI chatbot user would require about \$124 in compensation to give up access for one month, up from \$98 in 2025. Median WTA rose from \$3.40 to \$11.40 over the same period. In other words, in 2025, 50% of the user base would have been willing to give up access to AI chatbots for a compensation of \$2.27, whereas in 2026 the corresponding threshold had increased by roughly a factor of five. In parallel, Bick et al. (2026), found that the share of US adults using generative AI rose from 48% in early 2025 to 56% in late 2025 (see Figure 1).<sup>1</sup> While their dates don’t match up precisely with the dates of our surveys, combining these aggregate adoption estimates with our valuation estimates implies that total annual consumer surplus increased from roughly \$116 billion to roughly \$172 billion, or by about 50% in less than a year.

Our estimates suggest that consumer welfare gains from AI chatbots exceeds current producer revenues by a wide margin. global revenues (including US) from generative AI tools like ChatGPT, Gemini, Claude, or Copilot are estimated at around \$14.2 billion, suggesting that much of the social value created by these tools accrues to consumers rather than firms. This is consistent with Nordhaus (2004), who finds that innovators captured only around 3% of the total social returns from major technological advances in the second half of the twentieth century. We also document

---

<sup>1</sup>Similarly, Hartley et al. (2026) report that “LLM adoption among U.S. workers has increased rapidly from 30.1 percent as of December 2024 to 38.3 percent as of December 2025.”




Figure 1: Share of U.S. adults that use Generative AI. Source: Bick et al. (2026)

significant heterogeneity in surplus across users: usage frequency is the strongest predictor of valuation, followed by work use, number of products used, and paid subscription status, with additional differences by gender, age, and ethnicity.

This paper relates to several strands of literature. Most directly, it builds on the GDP-B framework (Brynjolfsson et al., 2025a, 2019, 2023), which use massive online choice experiments to value free digital goods (Brynjolfsson et al., 2019). We contribute to the broader literature on measuring consumer surplus from the internet and digital services, including Goolsbee and Klenow (2006), Brynjolfsson and Oh (2012), and Nakamura and Soloveichik (2015), by providing the first such estimates for generative AI. Our discrete choice design also sits within the contingent valuation and stated preference literature (Hanemann, 1984; Carson, 2012).

Additionally, our paper complements the growing literature on the productivity effects of generative AI, including Noy and Zhang (2023), Dell’Acqua et al. (2023), and Brynjolfsson et al. (2025b). Those studies document gains in output quality and task completion speed in specific occupational settings. We provide the welfare counterpart to this literature by estimating consumer surplus directly. We show that the social value of generative AI is substantial and likely to grow as adoption expands and capabilities improve further. Finally, we relate to recent work documenting the patterns and evolution of generative AI use. Chatterji et al. (2025) document how consumers use ChatGPT across task categories. In addition, Bick et al. (2026) track adoption trends across the adult population in the U.S. while Hartley et al. (2026) track adoption in the labor force. We add the welfare dimension to this picture, and document how heterogeneity in usage patterns are related to heterogeneity in surplus.

The rest of the paper is organized as follows. Section 2 describes the data, survey design, and econometric framework. Section 3 presents the main results on aggregate consumer surplus and the drivers of heterogeneity. Section 4 discusses implications and Section 5 concludes.

## 2 Data

Our main surveys among users of generative AI are conducted in March 2026 (N=2,000; 1,908 after data cleaning) and July 2025 (N=1,500; 1,491 after data cleaning). Both surveys are representative of the U.S. Census in terms of age, gender, and ethnicity. We include additional controls for education, income, and place of residence.<sup>2</sup>

To estimate consumer surplus we ask the survey respondents the following: “*Would you give up access to all AI tools like ChatGPT, Gemini, Claude, or Copilot for one month starting tomorrow morning in exchange for [USD]?*” The monetary offer (*[USD]*) is drawn randomly from 7 price points between \$1 and \$500. The survey also collects detailed information on usage intensity, subscriber status, products used, use cases, and workplace adoption. Among respondents, 15.4% report using generative AI rarely, 40.7% occasionally, and 44% regularly (Table 1).

On average, respondents use 2.6 different products each (median = 2), while one in ten reports using five or more. ChatGPT is the most widely used product (81.4%), followed by Gemini (55%), Claude (18.7%), and Grok (16.4%). The most common use cases are information seeking (83%), practical guidance (55%), and writing (54%). The use cases in our survey broadly match Chatterji et al. (2025), who use confidential ChatGPT conversations and find that these categories account for more than three quarters of consumer messages on ChatGPT.

Almost half of all respondents (48.6%) report using generative AI for work, while 1 in 4 has access to a paid subscription. To capture intensity of work use we ask respondents to estimate the share of their work that is “supported” by these tools, defined as “*completing a task directly involves the use of Generative AI*”. On average respondents report around 20% of their work as supported by AI (median of 10%), while 12% estimate that it is more than half of their work. A Gallup survey finds that 12% of workers use Generative AI daily.<sup>3</sup> This aligns with Bick et al. (2026) who also find that between 1% to 7% of work hours are supported by Gen AI based on self-reported hours spent using Gen AI per day.

### 2.1 Econometric Framework

Each respondent is randomly assigned a single compensation amount  $p \in \{1, 10, 20, 50, 100, 200, 500\}$  and asked whether they would give up access to all generative AI tools for one month in exchange for \$ $p$ . Let  $Y_i \in \{0, 1\}$  denote whether respondent  $i$  accepts the offer. We model the acceptance

---

<sup>2</sup>Appendix A provides further information on survey design, sample construction, and robustness checks, while Appendix B reproduces the full questionnaire.

<sup>3</sup><https://www.gallup.com/workplace/701195/frequent-workplace-continued-rise.aspx>. Accessed March 13, 2026.

Table 1: Descriptive Statistics on GenAI Use

	N	%
<i>Usage frequency</i>		
Rarely	293	15.4
Occasionally	776	40.7
Frequently	839	44.0
<i>Use for work</i>		
Yes	927	48.6
No	981	51.4
<i>Estimated share of work supported by generative AI</i>		
0–25%	1,365	71.5
26–50%	312	16.4
51–75%	156	8.2
76–100%	75	3.9
<i>Paid subscriber</i>		
Yes	487	25.5
No	1,421	74.5
<i>Use cases (multiple responses allowed)</i>		
Writing	1,034	54.2
Practical guidance	1,042	54.6
Technical help	758	39.7
Multimedia	691	36.2
Seeking information	1,594	83.5
Self-expression	413	21.7
<i>Fear of job automation</i>		
Very likely	82	4.3
Somewhat likely	301	15.8
Not too likely	557	29.2
Not likely at all	668	35.0
Don't know	300	15.7

Notes: This table reports descriptive statistics for the March 2026 survey sample ( $N = 1,908$ ). Percentages may not sum to 100 because of rounding. For use cases, respondents could select more than one category.

decision using a logit specification:

$$\Pr(Y_i = 1 \mid p_i) = \Lambda(\alpha + \beta \ln p_i), \quad (1)$$

where  $\Lambda(\cdot)$  denotes the logistic CDF. The log transformation of the price reflects the well-documented concavity of valuation in money amounts (Hanemann, 1984). We report logit as our baseline specification; probit estimates yield nearly identical results.

The fitted model in Equation (1) traces out a demand curve: the share of users who would keep access at each compensation level. Median WTA is the price at which the predicted acceptance probability equals 0.5, recovered as  $\text{WTA}_{\text{med}} = \exp(-\hat{\alpha}/\hat{\beta})$ . Mean WTA corresponds to the area under the demand curve. We compute it as

$$\overline{\text{WTA}} = \int_0^{\bar{p}} [1 - \Lambda(\hat{\alpha} + \hat{\beta} \ln p)] dp, \quad (2)$$

where  $\bar{p} = 500$  is our highest offered amount. This truncation is conservative; we acknowledge that allowing valuations above \$500 would increase the estimate.

To compute aggregate consumer surplus, we multiply mean WTA by the number of US adult generative AI users and annualize:

$$\text{CS} = \overline{\text{WTA}} \times N_{\text{users}} \times 12, \quad (3)$$

where  $N_{\text{users}}$  is derived from Census population counts and the adoption rate reported in Bick et al. (2026).

To examine heterogeneity in valuations, we extend the logit model to include demographic and usage covariates:

$$\Pr(Y_i = 1 \mid p_i, \mathbf{X}_i) = \Lambda(\alpha + \beta \ln p_i + \mathbf{X}_i' \boldsymbol{\gamma}), \quad (4)$$

where  $\mathbf{X}_i$  includes age, gender, education, income, place of residence, race/ethnicity, usage frequency, number of products, work use, paid subscription status, and use case indicators. We report average marginal effects throughout, so that coefficients can be interpreted as changes in the probability of accepting the offer.

## 3 Results

### 3.1 How Large Is Consumer Surplus from Generative AI?

We estimate that the average generative AI user in March 2026 would require \$124.50 in compensation to give up access for one month. This represents a 27% increase relative to summer 2025, when the corresponding estimate was \$98 (see Table 2).

Table 2: Consumer Surplus from Generative AI

	2025	2026
Mean WTA (USD/month)	98.00	124.50
Median WTA (USD/month)	3.39	11.48
Adult population 18–64 (millions)	205.80	206.29
Usage rate	0.48	0.56
Users (millions)	98.78	115.33
Consumer surplus (USD billions/year)	116.2	172.3

Notes: Usage rates are from Bick et al. (2026). Consumer surplus is calculated by combining monthly WTA estimates with the number of users and annualizing.

To translate these individual valuations into aggregate consumer surplus, we combine them with adoption rates from Bick et al. (2026) and population counts from the US Census. Conceptually, consumer surplus corresponds to the area under the demand curve. We recover this curve by plotting the share of users willing to forgo access to generative AI at each compensation amount (Figure 2).




Figure 2: WTA for giving up access to any Generative AI tool; March 2026

When thinking about the value of these tools to a typical user we can also look at the median, which is the point at which half of all users are indifferent between keeping access or not. The

median increased from \$2.27 to \$11.40 over the same period.<sup>4</sup> Looking at the median is also less susceptible to extreme valuations at the upper end of the distribution.

### 3.2 What Drives Consumer Surplus?

Our data include detailed demographic characteristics as well as self-reported measures of generative AI use along both the extensive and intensive margins. This allows us to examine which individual characteristics are most strongly associated with consumer surplus. Figure 3 reports median valuations that are statistically significant while controlling for all other factors included in the full logit model using average marginal effects (AMEs) calculated at observed covariate values (see Appendix C). As noted above, the median surplus for the average generative AI user is close to \$11 per month.

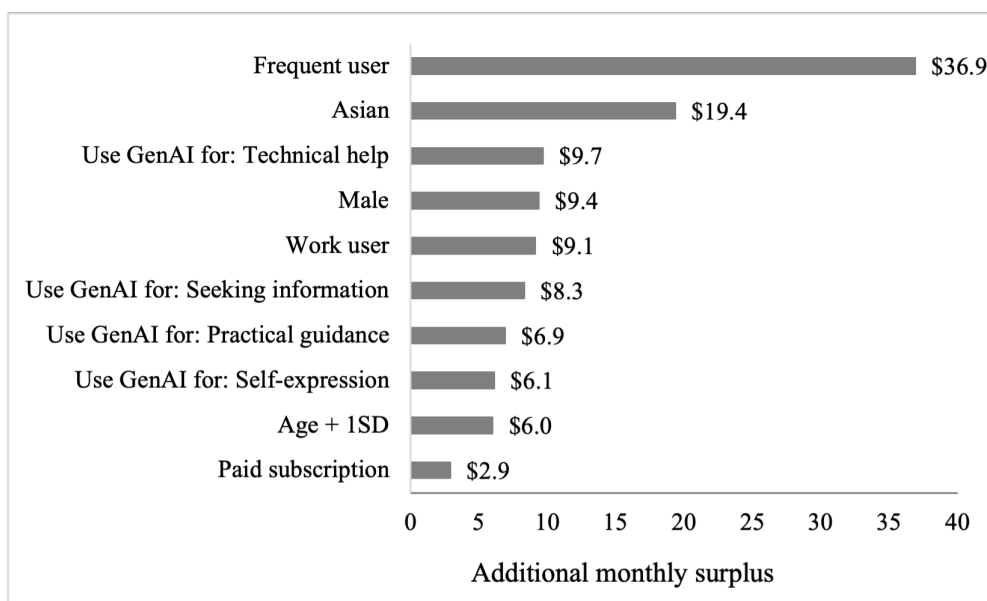


Figure 3: Additional surplus by individual characteristics

We find that usage frequency is the strongest predictor of consumer surplus as regular users have median monthly valuations that are \$36.9 higher than for infrequent users. Asian and male users also have higher valuations (+\$19.4 and +\$9.4 respectively). Work use is also a significant predictor of consumer surplus (+\$9.1) as are specific user cases such as using generative AI for technical help (e.g. coding, data formatting), seeking information, practical guidance, or self-expression. A one standard deviation increase in age is associated with \$6 of extra surplus per month while access to a paid subscription translates into \$2.9.

Other factors like having a college degree or living in a rural area are not statistically significant in our model and are not reported here. This also applies to reported use cases for writing or multimedia. Using multiple generative AI tools (“multi-homing”) is also insignificant, possibly

<sup>4</sup>We also collected data in December 2024 with a smaller sample ( $n = 502$ ). Median valuations were \$0.52, i.e., close to zero, suggesting that valuations in earlier periods were minimal.

hinting at the large degree of substitutability between tools for the average user. Similarly, we also do not find statistically significant differences in valuations for users of different tools, e.g. ChatGPT versus Gemini versus Claude.

## 4 Discussion

Our findings speak to a broader limitation of conventional welfare measurement. Consumer surplus captures the welfare gains from technological innovation more directly than GDP, which often understates the benefits of new goods, particularly in their early stages (Brynjolfsson et al., 2019). Since most people currently pay nothing or very little for generative AI products, there is a general wedge between what is captured in GDP and the value consumers derive from these tools (Brynjolfsson et al., 2025a).

This gap between measured revenue and user value also helps place the distribution of returns in context. This is consistent with Nordhaus (2004), who finds this pattern for all major technological advances in the second half of the twentieth century. Global consumer revenue (not enterprise revenue) for OpenAI, Anthropic, Google, and Microsoft are estimated (not enterprises) are estimated to have reached \$14.2 billion in 2025.<sup>5</sup> While reliable estimates for the US alone are not available yet, we can use global revenues as an upper bound.

The distribution of these gains is uneven across users. More intensive use, measured by frequency, number of products used, and workplace adoption, is associated with higher surplus. Higher valuations are also associated with being male, older, and Asian or multiracial. By contrast, higher income is only weakly associated with valuation, and we find no systematic relationship between surplus and either college education or urban residence.

At the same time, these estimates should be interpreted with several limitations in mind. As with all stated preference studies, hypothetical bias remains a concern (Murphy et al., 2005; List and Gallet, 2001), and our results rely on self-reported behavior from an online sample. However, the anonymity of the survey environment may reduce social desirability bias, which has been documented in settings related to AI use (Ling et al., 2025). Any remaining bias is likely to be present across survey waves, making changes over time more informative.

Additionally, respondents who opt into online surveys may also be more digitally engaged than the broader population and therefore more likely to use AI, although this concern is mitigated by our focus on AI users and by our use of external adoption data to calibrate aggregate estimates.

A further limitation is the one-month time horizon: while natural for many consumption decisions, the compensation required to forgo generative AI for a year may exceed twelve times the monthly amount, implying that our estimates may understate the value of continued access.

More broadly, our study is not a full welfare analysis, as it does not account for producer surplus or for additional negative or positive externalities associated with generative AI, such as "digital

---

<sup>5</sup><https://www.gartner.com/en/newsroom/press-releases/2025-07-10-gartner-forecasts-worldwide-end-user-spending-on-generative-ai-models-to-total-us-dollars-14-billion-in-2025>. Accessed April 13, 2026

addiction”, loss or gain of expertise or changes in the rate of scientific discovery. Our calculation of total surplus requires capping compensation at a fixed amount. Because 12% of respondents reject even our highest offer (\$500), our aggregate surplus estimate may be an underestimate. We may also underestimate the value because generative AI is increasingly embedded in devices and services such as customer service chatbots and online search results, which respondents may not fully consider when reporting their valuations.

These limitations notwithstanding, our approach provides a starting point for measuring the consumer welfare gains from generative AI and points to several promising directions for future work, including longer time horizons, revealed preference validation, and broader welfare frameworks that incorporate spillovers and producer surplus.

## 5 Conclusion

We find that total consumer surplus from generative AI in the United States reached \$172.3 billion in early 2026. Compared to \$116.2 billion in the middle of 2025, this represents an increase of almost 50%. The increase is driven by wider adoption, as the total number of users grew from 98 million to 115 million (+21%), and by an increase in the average valuation from \$98 to \$124.50 (+27%). These trends are not surprising given the overall quality improvement of available models over this time period.

Our findings suggest that access to these tools appears to generate substantial welfare gains for a broad share of the population, far exceeding estimated consumer revenues by their vendors. In addition, the heterogeneity we document points to possible disparities in access, use, and realized benefits, suggesting that distributional questions are important for how economists and policymakers assess the broader impact of generative AI.

While the effects of generative AI are not fully reflected in traditional measures of productivity and GDP, they are clearly visible in consumer surplus using online choice experiments, and thus can be included in GDP-B. To the extent that not only generative AI tools, but also other digital goods and services are increasingly available at low or zero price, approaches like the one we use in this paper will be useful for understanding and assessing the level and distribution of their economic effects.

## References

- Bick, A., Blandin, A., and Deming, D. J. (2026). The Rapid Adoption of Generative AI. *Management Science*.
- Brynjolfsson, E., Collis, A., Diewert, W. E., Eggers, F., and Fox, K. J. (2025a). GDP-B: Accounting for the Value of New and Free Goods. *American Economic Journal: Macroeconomics*, 17(4):312–344.
- Brynjolfsson, E., Collis, A., and Eggers, F. (2019). Using massive online choice experiments to measure changes in well-being. *Proceedings of the National Academy of Sciences*, 116(15):7250–7255.
- Brynjolfsson, E., Collis, A., Liaqat, A., Kutzman, D., Garro, H., Deisenroth, D., and Wernerfelt, N. (2023). The digital welfare of nations: New measures of welfare gains and inequality. Working Paper 31670, National Bureau of Economic Research.
- Brynjolfsson, E., Li, D., and Raymond, L. (2025b). Generative AI at Work. *The Quarterly Journal of Economics*, 140(2):889–942.
- Brynjolfsson, E. and Oh, J. (2012). The Attention Economy: Measuring the Value of Free Digital Services on the Internet. In *Proceedings of the International Conference on Information Systems (ICIS)*.
- Brynjolfsson, E., Rock, D., and Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1):333–372.
- Carson, R. T. (2012). Contingent valuation: A practical alternative when prices aren’t available. *Journal of Economic Perspectives*, 26(4):27–42.
- Chatterji, A., Cunningham, T., Deming, D. J., Hitzig, Z., Ong, C., Shan, C. Y., and Wadman, K. (2025). How people use chatgpt. Working Paper 34255, National Bureau of Economic Research.
- Dell’Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraymer, L., Candelon, F., and Lakhani, K. R. (2023). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. Working Paper 24-013, Harvard Business School.
- Goolsbee, A. and Klenow, P. J. (2006). Valuing consumer products by the time spent using them: An application to the Internet. *American Economic Review*, 96(2):108–113.
- Hanemann, W. M. (1984). Welfare Evaluations in Contingent Valuation Experiments with Discrete Responses. *American Journal of Agricultural Economics*, 66(3):332–341.

- Hartley, J., Jolevski, F., Melo, V., and Moore, B. (2026). The labor market effects of generative artificial intelligence. Technical report, Social Science Research Network. Working paper, January 25, 2026. Previous version dated December 18, 2024.
- Horowitz, J. K. and McConnell, K. E. (2002). A Review of WTA/WTP Studies. *Journal of Environmental Economics and Management*, 44(3):426–447.
- Ling, Y., Kale, A., and Imas, A. (2025). Underreporting of AI use: The role of social desirability bias. Working paper, SSRN. Available at SSRN: <https://ssrn.com/abstract=5464215>.
- List, J. A. and Gallet, C. A. (2001). What Experimental Protocol Influence Disparities Between Actual and Hypothetical Stated Values? *Environmental and Resource Economics*, 20(3):241–254.
- Murphy, J. J., Stevens, T. H., and Weatherhead, D. (2005). Is Cheap Talk Effective at Eliminating Hypothetical Bias in a Provision Point Mechanism? *Environmental and Resource Economics*, 30(3):327–343.
- Nakamura, L. I. and Soloveichik, R. (2015). Capturing the Productivity Impact of the “Free” Apps and the Internet. Working Paper 15-25, Federal Reserve Bank of Philadelphia.
- Nordhaus, W. D. (2004). Schumpeterian profits in the american economy: Theory and measurement. Working Paper 10433, National Bureau of Economic Research.
- Noy, S. and Zhang, W. (2023). Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence. *Science*, 381(6654):187–192.

# A Survey Details

Our samples were recruited from Prolific with quotas representative to the US Census in terms of age, gender, and ethnicity. The median response time was 3 minutes 15 seconds. Data were collected in July 2025 (Wave 1, N=1,500), and March 2026 (Wave 2, N=2,000). The main analysis relies on wave 2 data, while we use wave 1 data to calculate the increase in consumer surplus over the 9-month period.

To improve data quality, we use the platforms advanced bot and LLM usage detection features and automatically screen out respondents that respond exceptionally fast. We also include implicit and explicit attention checks and randomize all response options wherever possible. We include a fake LLM as a first implicit attention check (none of respondents select it), while less than 0.3% of respondents fail our explicit attention check as shown in Figure 4. We removed them from the subsequent analysis.



Figure 4: Explicit attention check (pass rate: 99.7%)




Figure 5: WTA for giving up access to any Generative AI tool, Nov 2024

Table 3: Sample Demographics, March 2026, N=2,000

	N	%
All respondents	1,908	100.0
<i>Gender</i>		
Male	935	49.0
Female	973	51.0
<i>Age</i>		
18–24	225	11.8
25–34	342	17.9
35–44	341	17.9
45–54	303	15.9
55+	697	36.5
<i>Race and ethnicity</i>		
White	1,186	62.2
Black	256	13.4
American Indian or Alaska Native	37	1.9
Asian	153	8.0
Native Hawaiian or Other Pacific Islander	7	0.4
Hispanic, Latino, or Spanish Origin	247	12.9
Other	22	1.2
<i>Education</i>		
No high school	16	0.8
High school graduate	236	12.4
Some college, no degree	398	20.9
Associate’s degree	190	10.0
Bachelor’s degree	702	36.8
Graduate degree	366	19.2
<i>Marital status</i>		
Married, spouse in household	802	42.0
Married, spouse outside household	19	1.0
Widowed	58	3.0
Divorced	225	11.8
Separated	21	1.1
Never married	783	41.0
<i>Household income</i>		
\$0–\$25,000	173	9.1
\$25,000–\$50,000	377	19.8
\$50,000–\$75,000	376	19.7
\$75,000–\$100,000	352	18.4
\$100,000–\$125,000	212	11.1
\$125,000–\$150,000	163	8.5
\$150,000 or more	255	13.4
<i>Household size</i>		
1	405	21.2
2	625	32.8
3	389	20.4
4	293	15.4
5 or more	196	10.3
<i>Place of residence</i>		
Large city	407	21.3
Suburb near a large city	832	43.6
Small city or town	448	23.5
Rural area	219	11.5
Don’t know	2	0.1

Notes: This table reports demographic characteristics for the March 2026 survey sample. Percentages may not sum to 100 because of rounding.

## B Survey questionnaire

**Q1.** We will now ask about YOUR experiences with Generative AI. Generative AI is a type of artificial intelligence that creates text, images, audio, or video in response to prompts. Some examples of Generative AI include ChatGPT, Gemini, Claude, and Midjourney. Have you heard of Generative AI tools? *Note: This is not a check of whether you are using AI to fill in surveys. We want to better understand your experience with Generative AI.*

- a) No, I have never heard of them
- b) Yes, I have heard of them

**Q2.** Do you use Generative AI tools?

- a) I have never used them
- b) I have used them once or twice
- c) I use them occasionally
- d) I use them regularly

*Screen-out: Respondents who report that they have never heard of Generative AI (Q1 = "No") or have never used it (Q2 = "I have never used them") are screened out and do not continue the survey after this point. We collect 2,000 responses after screen-out.*

**Q3.** Which of the following Generative AI services have you used in the last month? **Check all that apply.** *[Answer options presented in random order.]*

- a) ChatGPT
- b) Claude
- c) Gemini
- d) GitHub Copilot
- e) Midjourney
- f) Scribe
- g) NeuroVista AI
- h) Generative AI tools embedded in existing software (for example, Microsoft Copilot)
- i) DALL-E
- j) Perplexity
- k) Grok
- l) Meta AI
- m) Canva
- n) DeepSeek
- o) Character.ai
- p) Other

**Q4.** What do you use Generative AI for? **Check all that apply.** *[Answer options presented in random order.]*

- a) Writing (editing, writing, translation, summary generation)
- b) Practical guidance (how-to, tutoring, creative ideation)
- c) Technical help (math, programming, data analysis)
- d) Multimedia (creating images, videos, or other media)
- e) Seeking information (specific information, products, recipes)
- f) Self-expression (chitchat, relationships, reflection, games, role play)
- g) Other

**Q5.** People use digital technologies such as smartphones, apps, or websites differently. To help us ensure that participants are reading carefully, please select “Never” for this question.

- a) Never
- b) Rarely
- c) Often
- d) Sometimes
- e) Daily

**Q6.** Do you use Generative AI tools for your job?

- a) No
- b) Yes

**Q7.** Think about how Generative AI tools supported your work during the PAST WEEK. Here, “supported” means that completing a task directly involves the use of Generative AI. Please provide your best estimate of the share of your work that is supported by using Generative AI.

- 0%–100% [slider]

**Q8.** Would you give up access to any AI tool like ChatGPT, Gemini, Claude, or Copilot for one month starting tomorrow morning in exchange for \$1/\$10/\$20/\$50/\$100/\$200/\$500?

- a) No
- b) Yes

**Q9.** Are you currently a paid subscriber to any Generative AI tool? Please also respond with “Yes” if your work pays for your subscription.

- a) No
- b) Yes

**Q10.** How likely is it that the job you have now will be eliminated within the next five years as a result of new technology, automation, robots, or artificial intelligence?

- a) Very likely
- b) Somewhat likely
- c) Not too likely
- d) Not at all likely
- e) Don't know / Doesn't apply

**Q11.** What is your sex?

- a) Male
- b) Female

**Q12.** How old are you in years?

- \_\_\_\_\_ [open box]

**Q13.** What is the highest level of school you have completed or the highest degree you have received?

- a) Did not complete 12th grade (no high school degree)
- b) High school graduate or the equivalent (for example, GED)
- c) Some college but no degree
- d) Associate's degree in college
- e) Bachelor's degree (for example, BA, AB, BS)
- f) Graduate degree (for example, Master's, professional, or doctorate degree)

**Q14.** Here is a list of categories of race and origin. Please select all that apply to you. You may choose more than one.

- a) White
- b) Black or African American
- c) American Indian or Alaska Native
- d) Asian
- e) Native Hawaiian or Other Pacific Islander
- f) Hispanic, Latino, or Spanish origin
- g) Other

**Q15.** What is your marital status?

- a) Married, spouse currently lives in this household
- b) Married, spouse currently lives outside this household
- c) Widowed
- d) Divorced
- e) Separated
- f) Never married

**Q16.** Which category represents the total combined income of all members of your family during the past 12 months? This includes money from jobs, net income from business, farm or rent, pensions, dividends, interest, Social Security payments, and any other money income received by members of your family who currently live with you and are 15 years of age or older.

- a) \$0–\$25,000
- b) \$25,000–\$50,000
- c) \$50,000–\$75,000
- d) \$75,000–\$100,000
- e) \$100,000–\$125,000
- f) \$125,000–\$150,000
- g) \$150,000 or more

**Q17.** Including yourself, how many people currently live in your household?

- \_\_\_\_\_ [open box]

**Q18.** Which of the following best describes the place where you currently live?

- a) Large city
- b) Suburb near a large city
- c) Small city or town
- d) Rural area
- e) Don't know

## C Regressions

Table 4 reports average marginal effects from the logit model in equation (4). Column (1) includes only demographic controls. Column (2) adds usage variables (frequency, number of products, work use, and paid subscription status). Column (3) further adds use-case indicators. Negative coefficients indicate a lower probability of rejecting the compensation offer, consistent with higher WTA.

Table 4: Logit Regression Results:  
Determinants of WTA (Average Marginal Effects)

	Dependent variable: Accept compensation offer		
	(1)	(2)	(3)
	Demographics	Full model	Full model + tasks
<i>Price</i>			
ln(price)	0.087*** (0.004)	0.086*** (0.004)	0.084*** (0.004)
<i>Demographics</i>			
Age	-0.002*** (0.001)	-0.002*** (0.001)	-0.002*** (0.001)
Female	0.110*** (0.020)	0.073*** (0.019)	0.065*** (0.019)
College	-0.023 (0.022)	0.023 (0.021)	0.018 (0.021)
Income	-0.035** (0.014)	-0.017 (0.013)	-0.016 (0.013)
Non-urban	0.063** (0.025)	0.027 (0.023)	0.028 (0.023)
<i>Race and ethnicity (omitted category: White)</i>			
Black	-0.079** (0.034)	-0.038 (0.030)	-0.039 (0.031)
Asian	-0.104** (0.045)	-0.106*** (0.041)	-0.101** (0.040)
Hispanic	-0.027 (0.041)	-0.031 (0.038)	-0.032 (0.038)
Other	-0.063 (0.085)	-0.124 (0.079)	-0.123 (0.077)
Multiracial	-0.066* (0.035)	-0.072** (0.032)	-0.079** (0.032)
<i>Usage</i>			
Usage frequency		-0.190*** (0.014)	-0.156*** (0.016)
Number of products		-0.011* (0.006)	0.000 (0.006)
Work use		-0.067*** (0.021)	-0.058*** (0.022)
Paid subscriber		-0.038*** (0.011)	-0.033*** (0.011)
<i>Use cases</i>			
Writing			-0.029 (0.021)
Practical guidance			-0.046** (0.020)
Technical help			-0.064*** (0.021)
Multimedia			-0.016 (0.020)
Seeking information			-0.074*** (0.028)
Self-expression			-0.043* (0.023)
Observations	1,908	1,908	1,908
Pseudo $R^2$	0.123	0.242	0.254
Log likelihood	-1120.71	-968.97	-952.73

Notes: Entries report average marginal effects from logit regressions. The dependent variable = 1 if respondent accepts compensation. Negative coefficients indicate a lower probability of accepting the offer, consistent with higher valuations. Standard errors are reported in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .