

OpenAI

ChatGPT and the price of work



March 2026

Contents

Executive summary	01
Introduction	02
Do people have accurate ideas about potential earnings?	04
Worker wage inquiries on ChatGPT	07
Introducing WorkerBench	16
Conclusion	19
References	20
Appendix	21

Executive summary

Wage information is hard to obtain, unevenly distributed, and critical for job search, negotiating, training, and career choice.

In January and February 2026, on average, nearly 3 million messages each day on consumer ChatGPT in the US involve wages and earnings. Role-specific and entrepreneurship questions account for the largest shares of these queries

Occupation-related wage search is more prevalent in higher-skill or less transparent fields, while entrepreneurship search centers on creative work and small service businesses.

Across occupation groups, wage search rises with both wage dispersion and wage level, consistent with stronger demand where returns to information and uncertainty are both higher.

We introduce a new effort – WorkerBench – to help understand how accurate ChatGPT is at providing national-level wage information in the U.S; GPT-5.4 can provide estimates that closely match national occupation benchmarks.

Introduction

Wages are the price of work. However, unlike many goods that are bought and sold, the price tag on labor – that of a worker’s time – is often unknown, ambiguous, or difficult to interpret for workers who aren’t already employed in a given role. Yet, obtaining accurate wage and salary information is critical for economic empowerment and career mobility. That’s apparent in ChatGPT aggregated usage data: on average, nearly 3 million messages each day on ChatGPT in the US ask about wages, compensation, or earnings. To preserve user privacy, we de-identify and aggregate a sample of user conversations. We then run classifiers across these conversations, so that OpenAI researchers can understand patterns of usage seeing the underlying conversation data. No member of the research team ever saw the content of user messages.

A large literature argues that labor markets have many information frictions. Search is costly; asking about pay can be socially awkward or risky; and informal networks distribute wage information unequally. Recent pay-transparency reforms reduce some of these frictions, but they remain partial and heterogeneous across occupations, firms, and places. The result is a setting in which even modest reductions in the cost of forming a credible wage belief may matter for worker welfare.

AI is a new type of labor-market intermediary. Rather than requiring a worker to search across multiple websites, interpret scattered salary pages, or ask a socially costly question, a model can synthesize wage information and return a benchmark in seconds. ChatGPT is already being used in this way. That revealed demand motivates the central question of this report: what kinds of information frictions are workers actually trying to solve with AI? Below, we describe why wage and earnings information matters for workers, what kinds of earnings information problems workers are trying to solve with ChatGPT, and provide early, suggestive evidence that ChatGPT is useful in obtaining better information for workers.

Millions of messages each day on ChatGPT in the US ask about wages, compensation, or earnings.

Why wage information matters

Misunderstanding potential earnings can keep workers stuck in lower-paying jobs, prevent students from investing in worthwhile education or training, and weaken workers' wage-negotiating power. Overall, difficult-to-obtain information can generate wage inequalities. The core intuition of this concept is simple: if learning about alternatives potential wages takes time and money, not everyone finds the best-paying option ([Stigler, 1961](#))¹.

Search costs are anything that makes it hard to learn what a job pays or what you could earn elsewhere: time, effort, lack of transparency, and even social risk (e.g., fear of retaliation or embarrassment for asking). Workers have fairly high search costs for obtaining accurate information about wages. For example, even with recent increases in information, a large share of job ads still lack pay information. The Indeed Hiring Lab reports that the share of US job postings with employer-provided salary information more than doubled from 18.4% in February 2020 to 57.8% in September 2024, but that still leaves many postings without pay info, with lots of differences by region and occupation ([Indeed Hiring Lab, 2024](#))².

Additionally, even when wage information is present for a given role, it's often unclear what the market wage may be and if that wage offer is negotiable. When the wage is "negotiable" but nobody says so, workers face an extra information problem: "is it okay to negotiate?" In general, wage negotiation is surprisingly rare: only one in three job-seekers negotiate ([Hall and Krueger 2012](#)). A large-scale field experiment showed that simply stating that pay is negotiable in the job ad substantially changes behavior: it reduces the gender gap in applications by about 45% and nearly triples negotiation initiation ([Leibbrandt and List, 2015](#)).

¹ Later work built on this insight with formal job-search models (e.g., [McCall 1970](#); [Mortensen 1986](#)). Empirical research using modern labor-market data—such as [Card, Heining, and Kline \(2013\)](#) and [Card, Cardoso, Heining, and Kline \(2018\)](#)—shows that wage differences across firms play a large role in overall inequality, consistent with the idea that search frictions and imperfect information affect where workers end up.

² Recent research on state pay-transparency laws, including [Arnold et al. \(2025\)](#) and [Cullen and Pakzad-Hurson \(2023\)](#), shows that mandatory salary ranges increase wage information in job ads and can influence application behavior, wage bargaining, and pay dispersion.

<p>Feb. 2020 Sep. 2024</p> <p>18.4% → 57.8%</p> <p>Share of US job postings with employer-provided salary information more than doubled</p> <p><u>Indeed Hiring Lab, 2024</u></p>	<p>Stating that pay is negotiable in a job ad:</p> <ul style="list-style-type: none">• 45% reduction in gender gap in applications• Near triple negotiation initiation <p><u>Leibbrandt and List, 2015</u></p>
---	--

When you need a job now, you can't spend weeks gathering wage intel or shopping offers. Accordingly, evidence suggests larger impacts from search costs for wage information when workers are more likely to hold inaccurate beliefs and when job acceptance decisions are made "under pressure" (Frimmel et al., 2024).

Wage information is especially helpful for vulnerable workers. One example: workers with fewer professional networks (new labor market entrants, recent migrants, career changers, people in small labor markets) could have fewer informal channels for wage information. That makes public or platform-based wage data more valuable (Harris 2018)³.

³ A related literature shows that referrals and social networks play a major role in transmitting wage information (e.g., Montgomery 1991; Dustmann, Glitz, Schönberg, and Brücker 2016), implying that workers without strong networks may face systematically higher information frictions.

Do people have accurate ideas about potential earnings?

Accurate potential wage information affects critical decisions: where to apply for a job, whether to accept an offer, whether to negotiate, and whether to stay or leave a career path.

Perceived future earnings matter even earlier in the job search process as students weigh the financial costs and time investment of different training and educational programs against what they can potentially earn in a given field. Many workers often have an inaccurate understanding of the pay of other potential jobs or opportunities, even within similar roles close to where they live. Researchers have also found that people tend to anchor on what they earn right now. Because of that, they find it difficult to guess how much higher their pay could be if they moved to a similar job at another company (Jäger et al., 2024). In their headline example, workers who would actually see about a 10% wage change from switching jobs only expect about a 1% change based on their current pay.

Even before people enter the labor market, beliefs about earnings matter. In one study, students' perceived returns to schooling were extremely low despite high measured returns to education, demonstrating how inaccurate beliefs about earnings can shape investment in skills (Jensen, 2010). The same logic shows up for unemployed workers deciding what jobs are "worth taking." Krueger and Mueller study the minimum wage unemployed workers would accept and find that these workers' reservation wages decline only 0.05–0.14 percent per week, implying remarkably slow downward adjustment in wage expectations over the unemployment spell. If a worker's internal compass is off for what kind of wage they should accept, workers may search for longer than they need to, or hold out for offers that are unlikely to arrive.

Do people have accurate ideas about potential earnings?

What can shape this?

Pay transparency reduces labor market friction by lowering the cost for workers to find wage data, empowering them to negotiate and filter jobs more effectively. AI models can extend this by providing near-instant, targeted salary estimates. By synthesizing vast data and generating tailored predictions, AI can democratize pay information and elevate workers' ability to navigate the job market.

In this report we:

01 Summarize why better pay information helps workers

02 Assess what role LLMs can play in providing this information

Wage information can shift intentions in a targeted way (Jäger et al., 2024). When workers get information about wages of others in similar roles, they improve beliefs about their outside options and change job search and wage negotiation intentions. That points to a practical principle: occupation-level wage information is helpful, but targeted information may be much more powerful.

This is where ChatGPT and AI models more broadly may help by providing timely and accurate wage, earnings, and salary estimates. Workers are already turning to ChatGPT for this information: nearly 3 million messages each day on average revolve around wages, earnings, and compensation.

ChatGPT's value for workers goes beyond providing public data; it eliminates the high personal cost of asking some wage questions. Research shows workers fear asking about pay due to social norms or fear of retaliation, which worsens information asymmetry, especially for vulnerable workers. By offering an anonymous, non-judgmental interface, AI democratizes inquiry, letting users ask difficult, sensitive compensation questions without social risk, empowering those who would otherwise remain uninformed.

Worker wage inquiries on ChatGPT

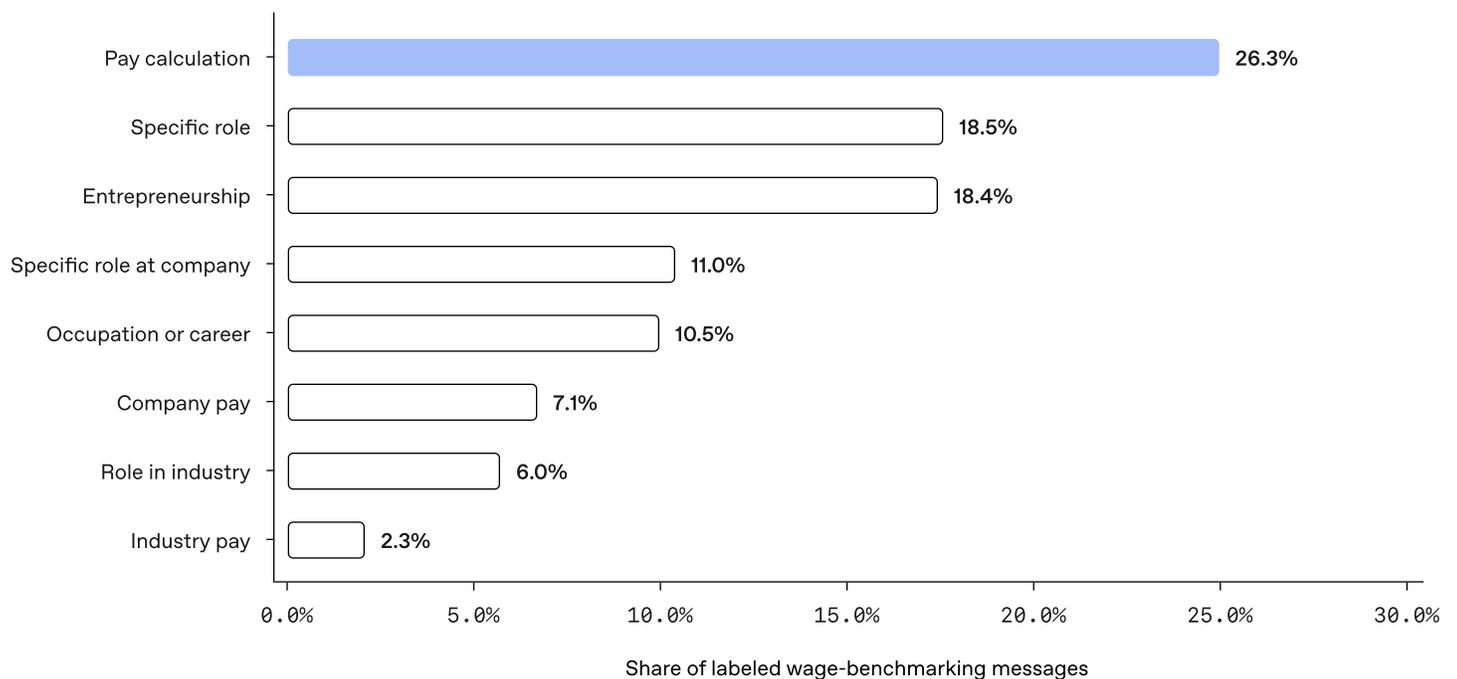
Users most often come to ChatGPT either to translate pay format calculations (e.g., convert hourly wages to a yearly salary) or to benchmark a concrete outside option. Figure 1 shows that pay calculation accounts for 26.3% of wage messages, followed by questions about a specific role (18.5%), entrepreneurship (18.4%), specific role at company (11.0%), and occupation or career (10.5%).

These messages capture two distinct worker problems. Pay-calculation messages reflect a basic but important budgeting need: workers often receive hourly or weekly pay information and need to translate it into an annualized number that can be compared across options. The role-specific and occupation-level categories, by contrast, reflect outside-option search: users are trying to understand what a concrete job, career path, or employer might plausibly pay before they apply, negotiate, or switch.

Fig. 1 Wage message breakdown

Pay calculation and role-specific questions dominate wage searches

Distribution of labeled wage-benchmarking messages by inquiry type



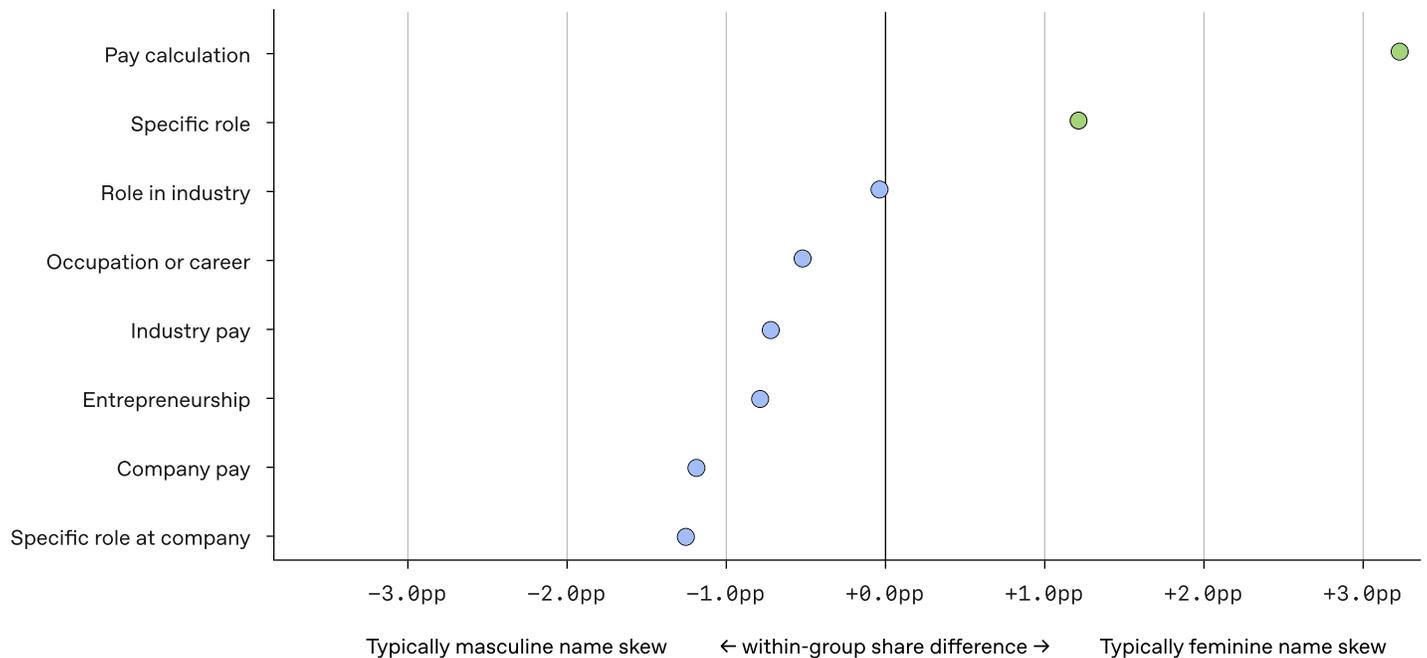
We investigated potential variation by gender by classifying a global random sample of ChatGPT users' first names using public, aggregated datasets of name-gender associations. Gender differences in the overall message mix are present but not large. Figure 2 shows that users with feminine names skew more toward pay calculation (+3.2 pp) and specific role (+1.2 pp), while users with masculine names skew more toward specific roles at a company (-1.2 pp) and company pay (-1.2 pp).

The modest size of these skews is important. It suggests that demand for wage information is not limited to a narrow demographic of users; rather, the core need to benchmark pay appears widespread. The larger signal is common demand for wage information itself, with gender shaping the margin rather than the center of the distribution.

Fig. 2 Gender skew in overall wage message breakdown

Gender differences in overall wage inquiry mix are modest

Positive values indicate categories that skew toward typically feminine names; negative values indicate categories that skew toward typically masculine names



Occupation-related wage search

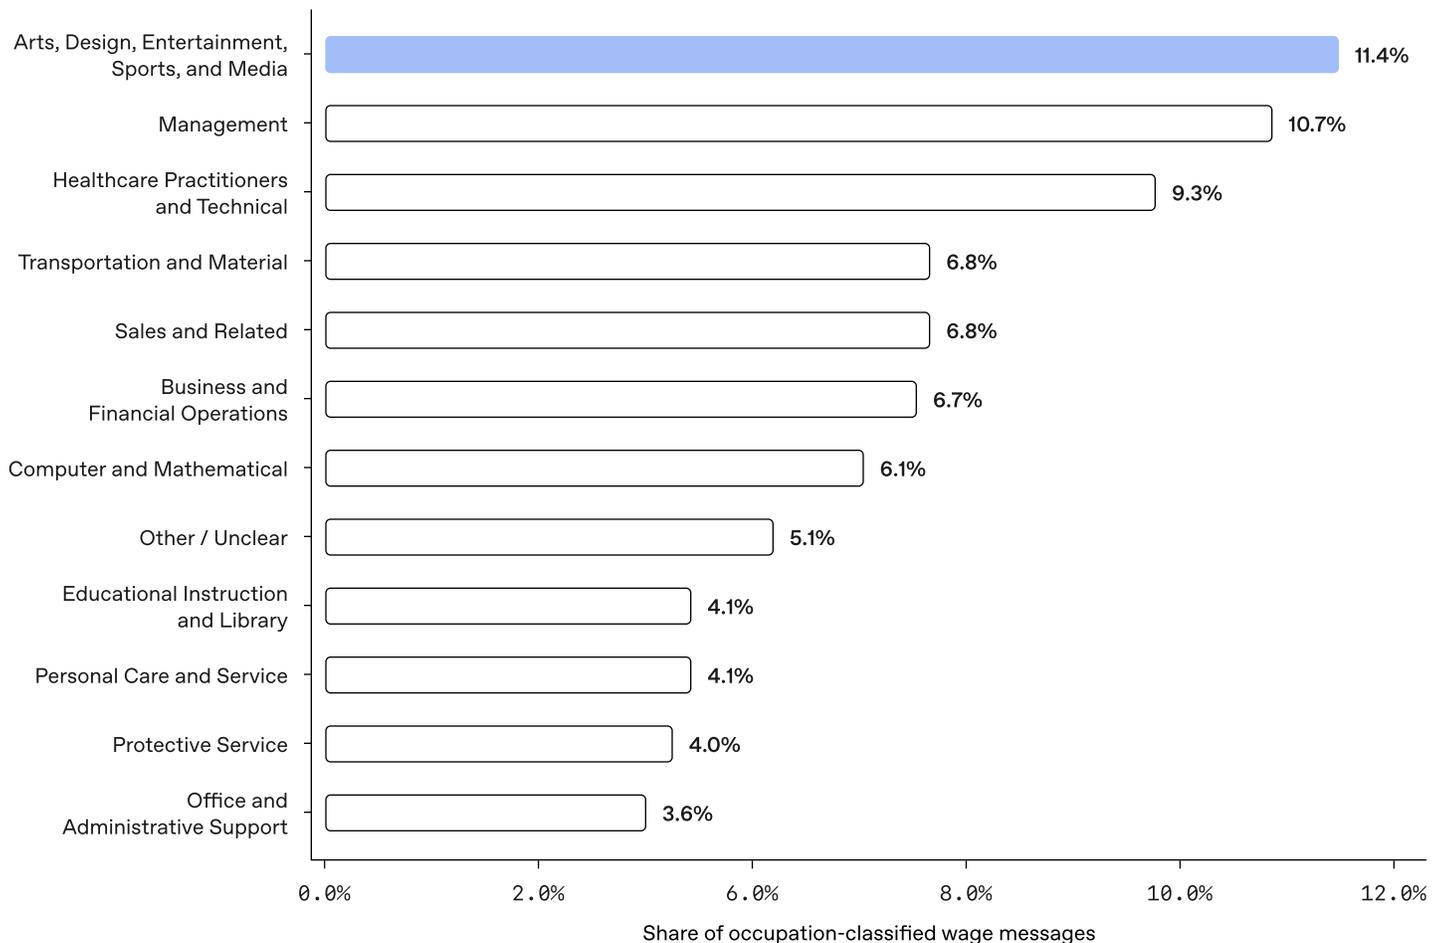
Occupation-related wage search is highly concentrated. To understand this, we classified a sample of messages that mention occupations or specific roles into SOC-2 occupation categories. Figure 3 displays the results of this exercise, showing the largest shares of messages in creative and athletic industries, management, healthcare, and transportation, with sales close behind. These are fields where pay is often less transparent, training investments are substantial, or wages vary sharply across employers and markets.

Healthcare, management, transportation, sales, and business-facing occupations are industries where wage beliefs have large consequences. In licensed healthcare and technical roles, workers make long training investments and need a clear sense of payoff. In management, sales, and other business roles, compensation often depends on location, firm, or incentive pay, making publicly posted benchmarks less informative.

Fig. 3 Distribution of occupation-related wage messages

Occupation-related wage searches concentrate in a few major groups

SOC-2 composition of occupation-related wage messages



Relative to employment, wage search over-indexes most in creative and athletic industries, management, and healthcare, and under-indexes most in office and administrative support, food service, and production (Figure 4). The composition of wage search is stronger in occupations where pay is more consequential or harder to benchmark.

This table is one of the clearest pieces of evidence that wage search follows information value, not headcount. Workers in large, standardized sectors such as office support or food service are numerous, but their wages are more likely to be posted, familiar, or have limited room for negotiation. By contrast, fields that over-index in wage search are the ones where compensation is more uncertain, more negotiable, or more consequential for mobility.

Fig. 4 Wage search over- and under-representation relative to employment

Wage search over-indexes in higher-skill and less transparent occupations

Most over- and under-represented SOC-2 groups relative to employment share

Rank	More represented in wage search	Less represented in wage search
01	Art, design, entertainment, sports, and media	Office and administrative support
02	Management	Food preparation and serving related
03	Healthcare practitioners and technical	Production
04	Computer and mathematical	Healthcare support
05	Personal care and service	Transportation and material moving
06	Protective service	Sales and related
07	Legal	Building and grounds cleaning and maintenance
08	Life, physical, and social science	Installation, maintenance, and repair
09	Architecture and engineering	Educational instruction and library
10	Community and social service	Construction and extraction

Entrepreneurship-related earnings search

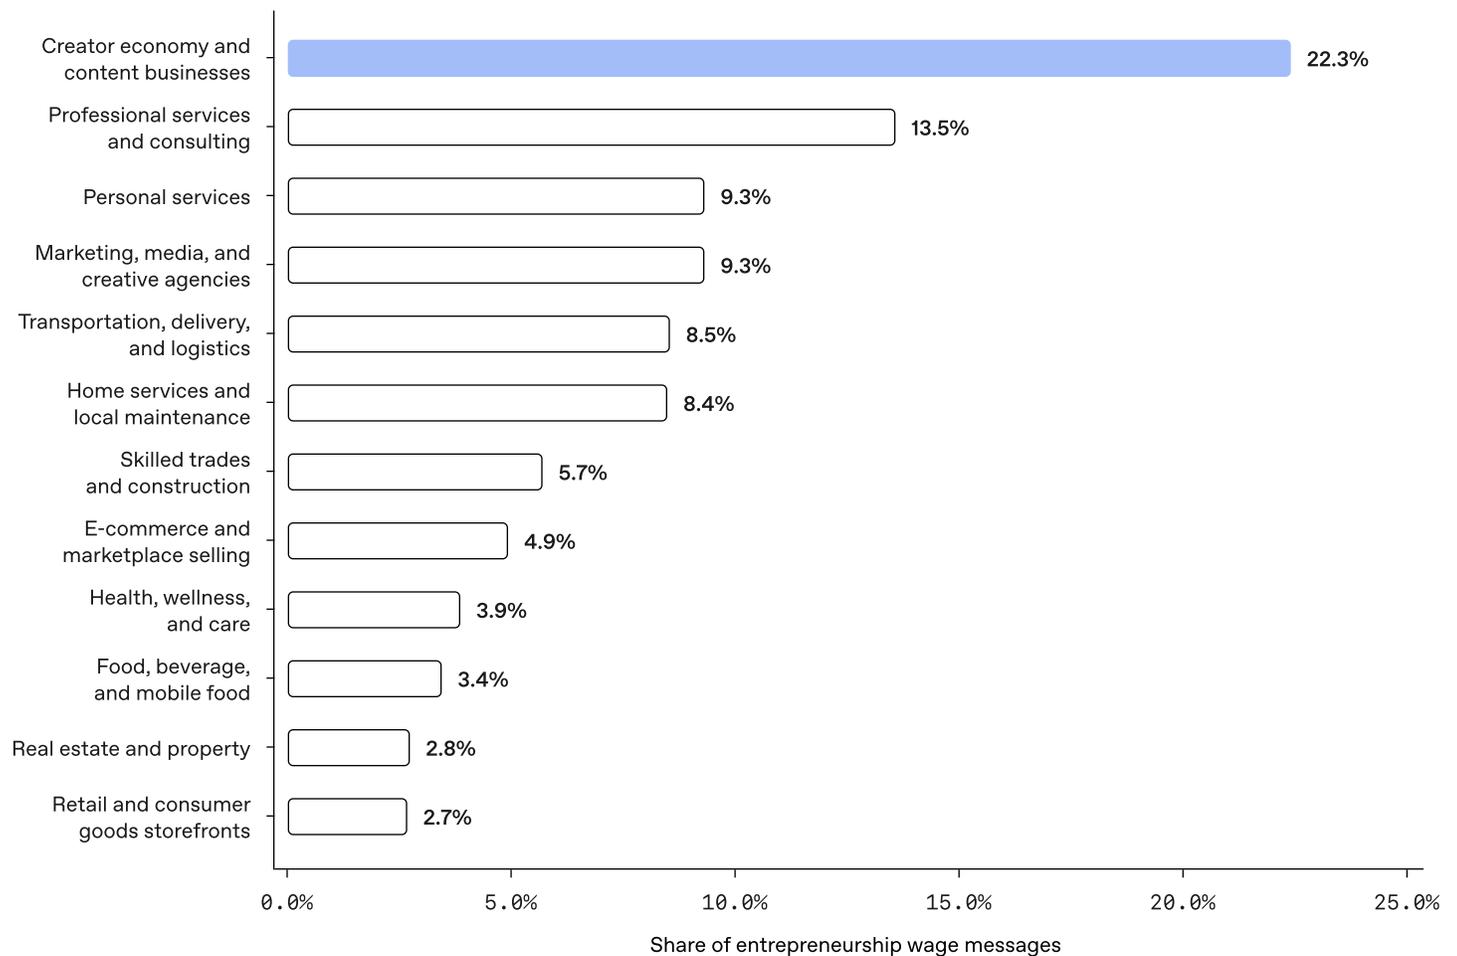
Entrepreneurship-related earnings questions are concentrated in creative work and small service businesses. Figure 5 shows the largest shares of questions in content and creator businesses, professional services and consulting, personal services, and marketing, media, and creative agencies.

That pattern is substantively important because self-employment rarely comes with a posted wage benchmark. For many small business ideas, the relevant question is not “what is the salary?” but “what could I realistically clear after costs and variable demand?” Creator businesses, consulting, local services, and logistics all fit that profile: entry looks feasible, but expected earnings are uncertain and depend heavily on pricing, utilization, and clients.

Fig. 5 Entrepreneurship wage search by business type

Entrepreneurship wage searches skew toward creator work and small services

Distribution of entrepreneurship-related wage messages by business type



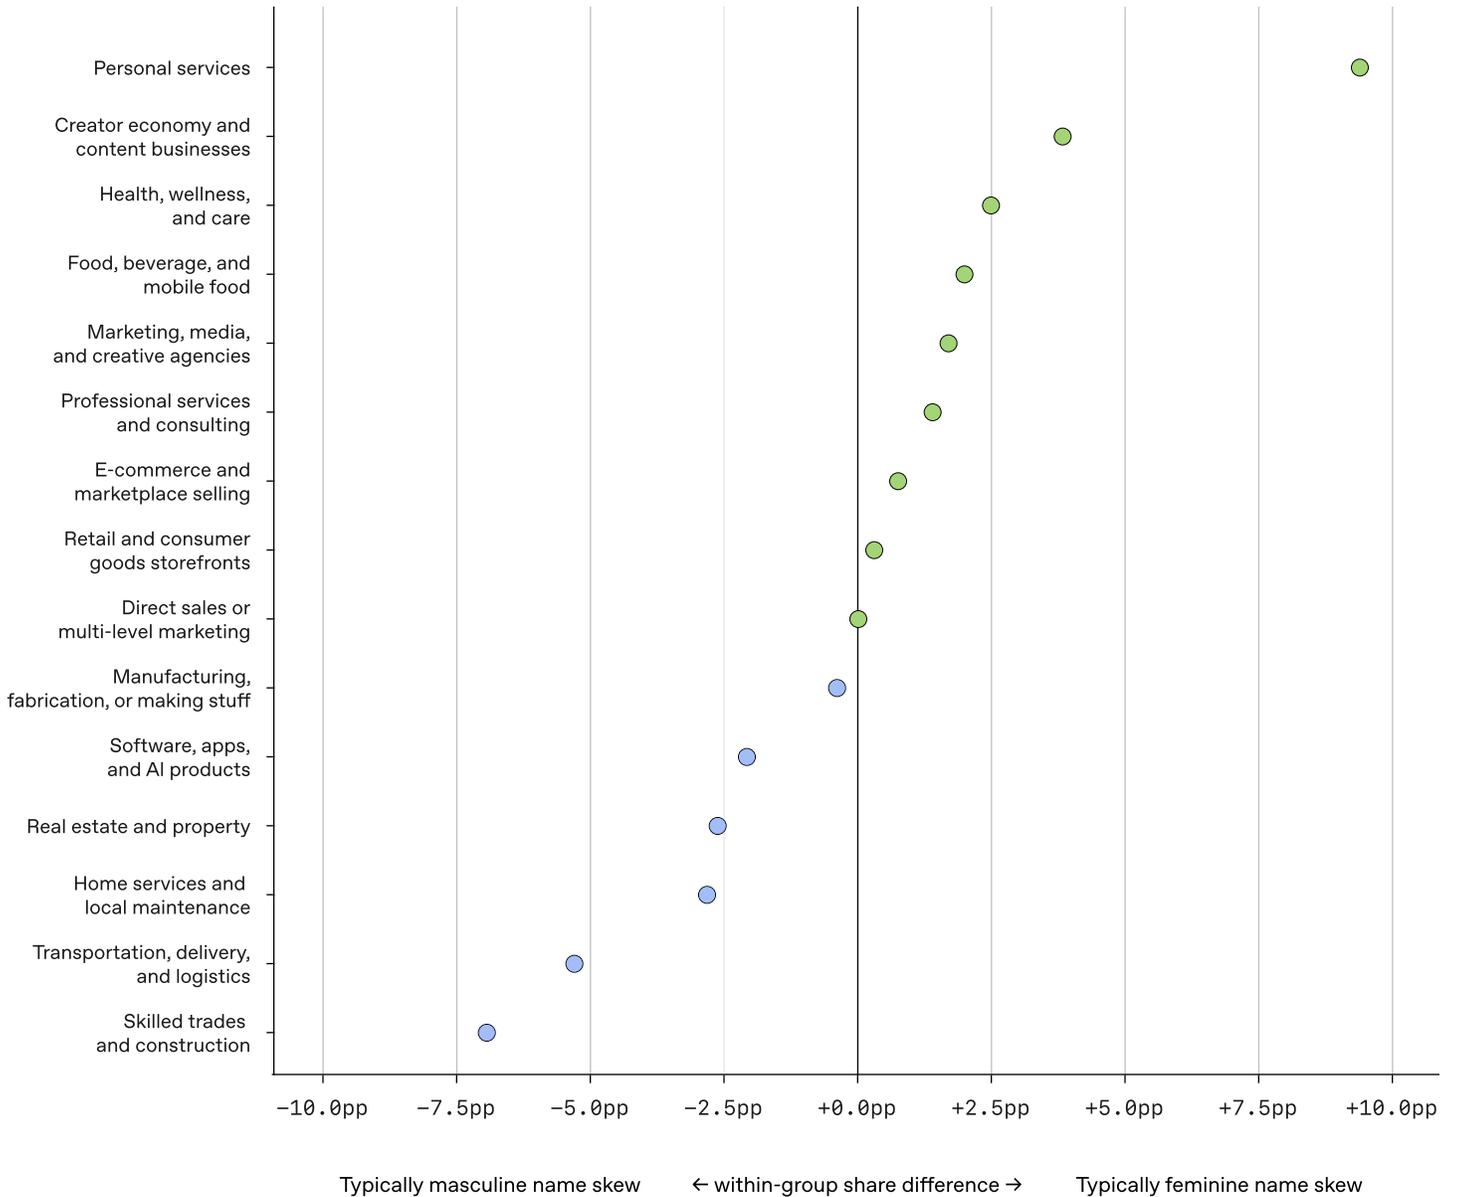
Gender differences are much larger in entrepreneurship search than in the overall wage mix. In terms of the difference between share of entrepreneurship messages, Figure 6 shows that users with feminine names skew toward personal services (+9.2 pp), content and creator businesses (+3.9 pp), and health, wellness, and care (+2.5 pp), while users with masculine names skew toward skilled trades and construction (-6.8 pp), transportation, delivery, and logistics (-5.4 pp), and home services and local maintenance (-3.2 pp).

Those differences imply that entrepreneurship search is more segmented than standard wage benchmarking. That matters for product design: a useful wage assistant should not treat entrepreneurship as a single class of question. The structure of uncertainty is different for creator businesses, home services, logistics, health and wellness, or personal services, and the relevant advice should adapt accordingly.

Fig. 6 Gender skew in entrepreneurship wage search

Entrepreneurship wage search shows sharper gender patterning

Positive values indicate categories that skew toward typically feminine names; negative values indicate categories that skew toward typically masculine names



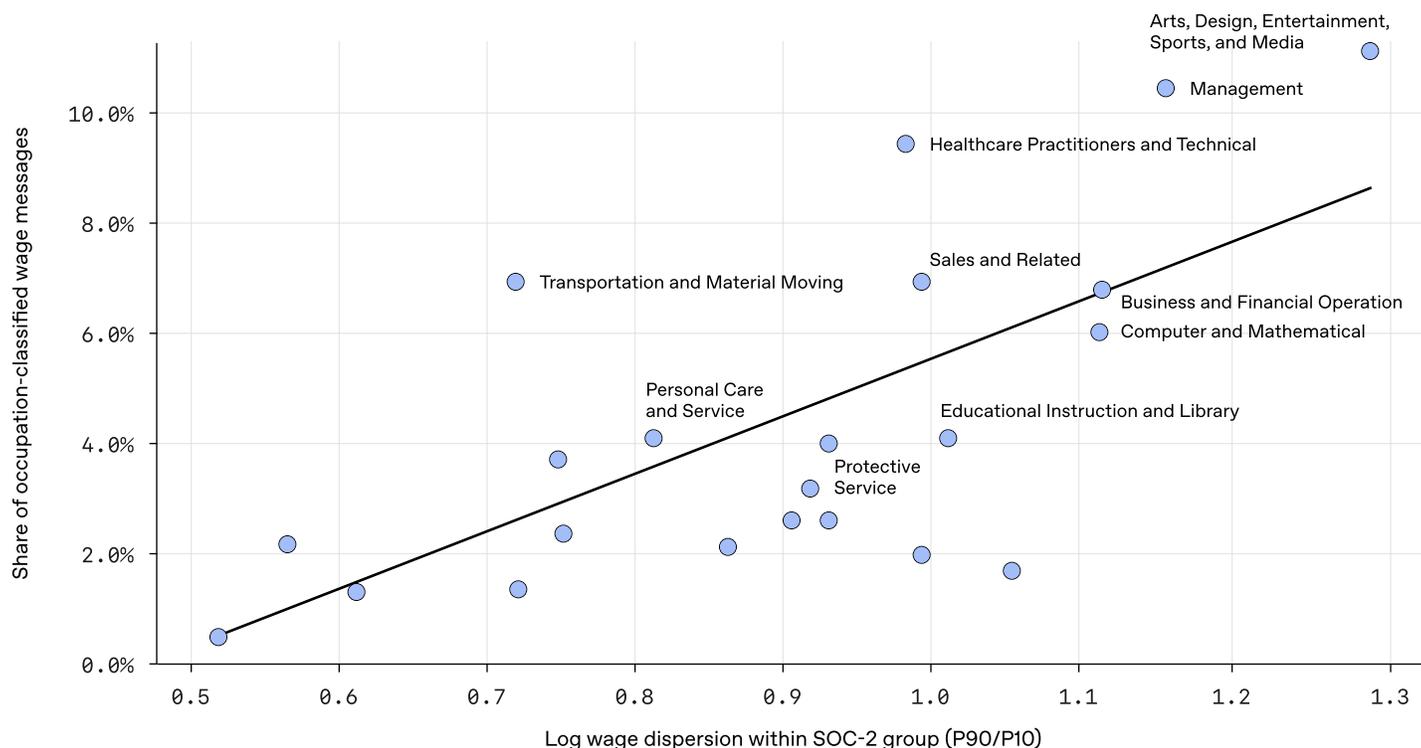
Where wage search concentrates

The distribution of wage search is consistent with the idea that workers search more where the stakes are higher and the underlying pay signal is noisier. Across industries, wage-query share is positively correlated with within-occupation wage differences ($r = 0.69$; Figure 7) and with median wages ($r = 0.44$; Figure 8). That pattern is directionally consistent with the core hypothesis for this paper: wage search is strongest where likely earnings are more uncertain and returns to better information are larger.

Fig. 7 Wage dispersion and wage-query share

Wage search is higher where within-occupation pay is more dispersed

SOC-2 wage-query share vs log P90/P10 wage dispersion



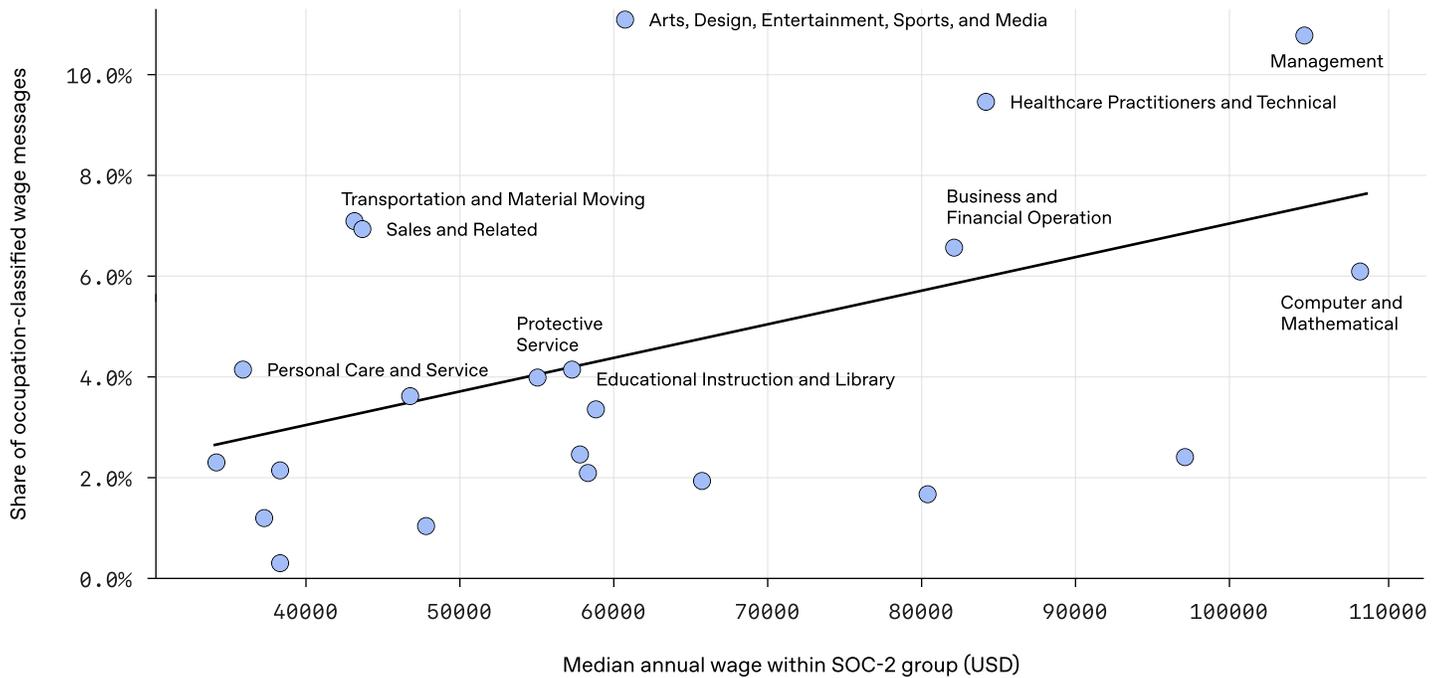
If users only searched the most common occupations, the distribution would mostly track employment. Instead, Figure 7 shows more search where within-occupation pay is spread out, which is precisely where a worker has more to gain from learning what a plausible outside option or local benchmark actually looks like.

Figure 8 adds a second channel: workers also search more in higher-paying occupations. Higher wage levels raise the payoff to getting the decision right, whether the decision is to invest in training, pursue a switch, negotiate harder, or stay put. Taken together, the two figures support the view that wage search is economically targeted rather than merely conversational.

Fig. 8 Median wage and wage-query share

Wage search is also higher in higher-paying occupation groups

SOC-2 wage-query share vs median annual wage



Introducing WorkerBench

Because of the importance of activities like wage benchmarking, we wanted to better understand how our models serve workers. WorkerBench is a project aiming to evaluate ChatGPT’s performance on tasks that are valuable for workers broadly, but that may not fit neatly into the context of an occupation . This is our first contribution to this research agenda, by attempting to broadly understand ChatGPT’s performance in providing accurate information. This is difficult for reasons likely related to the popularity of this topic: ground truth wage and earnings information is difficult to find. Here, we benchmark GPT-5.4 against 2024 OEWS mean annual wages at both the national occupation level and the metropolitan area level.

Results

National wage benchmark

The model does a solid job of accurately reporting occupation-specific official wages.⁴ When reporting national wage benchmarks, the model’s mean absolute error is \$101, mean absolute percent error is 0.1%, 99.8% of numeric estimates fall within 10% of the benchmark, and 99.8% fall within 20%. Mean bias is small at \$37. Figures 9 and 10 show the near one-for-one alignment with the benchmark and an error distribution concentrated at zero. This is the result of the following prompt: *“What was the mean wage in the US for [Occupation] in 2024? Use an official source.”*

The central empirical result is that the model is highly accurate in the observed sample: coverage is high, bias is small, and almost all numeric estimates fall very close to the benchmark. For users asking a national wage question, ChatGPT usually returns a benchmark-consistent annual wage rather than a loosely directional guess.

Substantively, these results imply that ChatGPT can serve as a reliable national wage look-up layer for workers asking about occupation-level pay. A user asking for a national occupation benchmark is therefore likely to receive a wage estimate that is very close to the accepted benchmark value in this observed sample.

⁴ Occupation is defined at the SOC-6 level.

Method

The evaluation pairs occupation titles with 2024 OEWS benchmark wages and asks the model for a mean annual wage estimate. The benchmark is limited to detailed 2024 national occupations with a directly available annual wage. The prompt gives the occupation title, specifies the United States as the geography, specifies the year 2024, and requires a JSON response with the annual wage, wage year, and source.

Fig. 9 National percent-error distribution

National wage estimate errors are tightly centered near zero

Distribution of percent error in annual wage estimates for observed detailed occupations

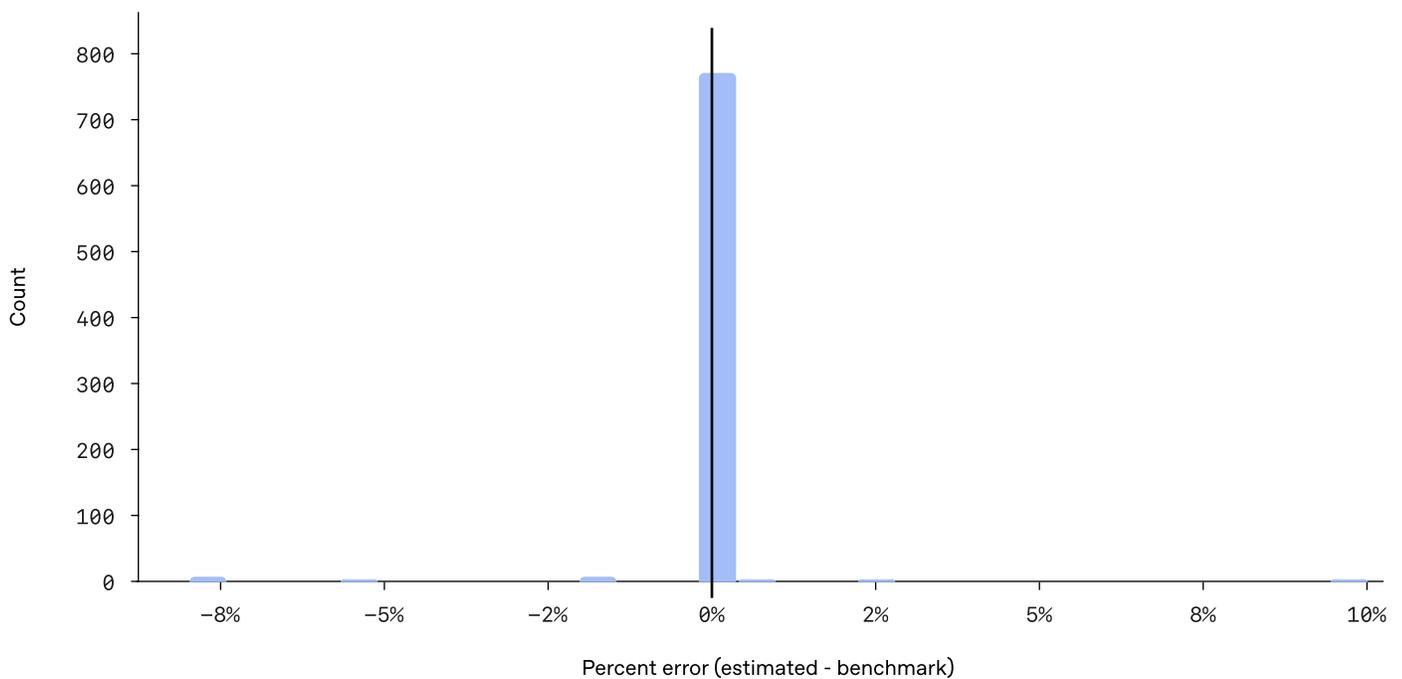
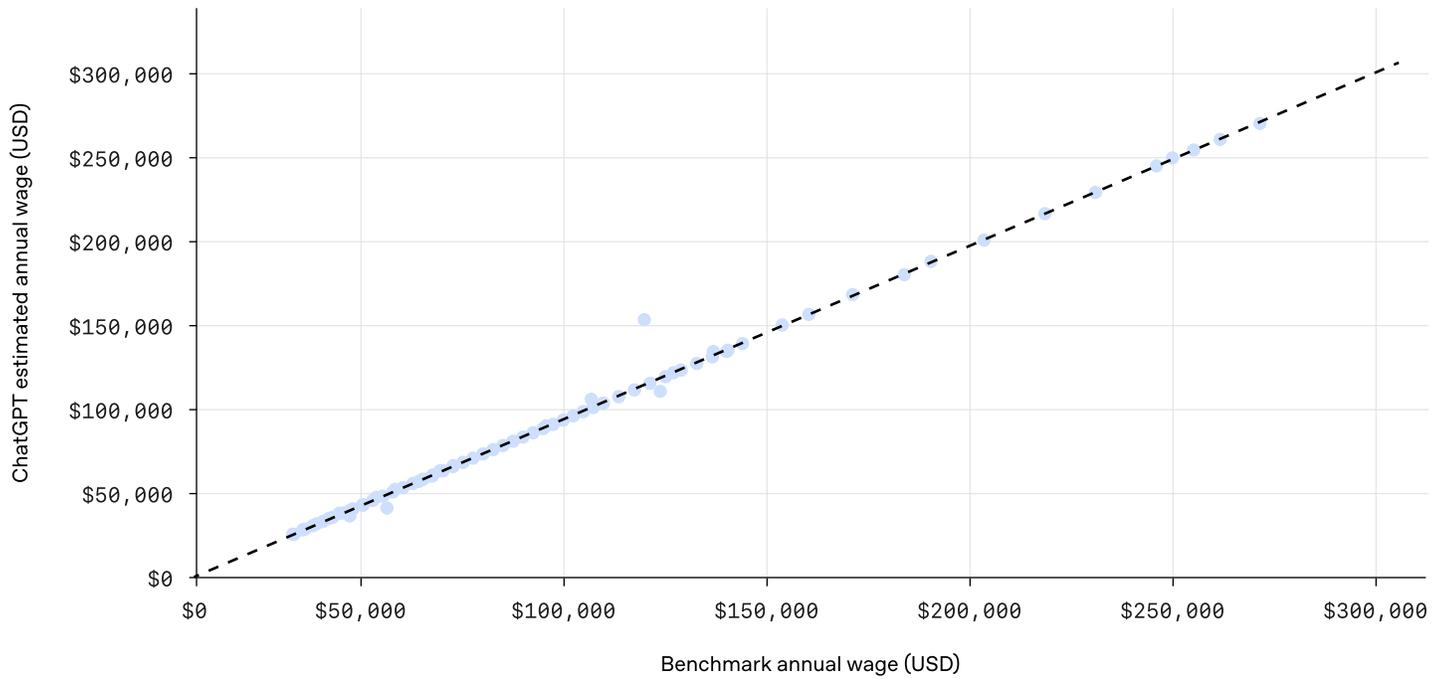


Fig. 10 National wage evaluation

ChatGPT's national wage estimates line up closely with benchmark wages

ChatGPT estimated annual wage vs 2024 benchmark wage for observed detailed occupations



Limitations

The benchmark accepts national BLS wage-release pages as ground truth, so the result should be described as agreement with the national BLS wage benchmarks, given the need to search for information for a given occupation on mean wages from an official source. To that end- this evaluation covers only benchmark annual mean wages. It does not cover local labor markets, firm-specific pay, bonuses, overtime, equity, benefits, or experience-specific wage ladders, which are often the quantities workers actually care about most.

Conclusion

Two facts stand out:

- 01 Users actively seek wage information on ChatGPT in categories where the financial stakes are large and the relevant market signal is uncertain

- 02 When asked for benchmark wages, the model is already highly accurate nationally

That combination matters because wage misperceptions affect job search, bargaining, occupational choice, and training investment. A tool that lowers the cost of forming a reasonable belief about pay can improve worker decisions even if it does not remove all uncertainty.

The next step for WorkerBench is to move beyond national occupation mean toward the settings workers actually ask about: geography, firm, level, and compensation structure. Those are also the settings where both user demand and model difficulty appear highest.

References

- Stigler, George J. (1961). "The Economics of Information." *Journal of Political Economy*. ([EconPapers](#))
- Dominitz, Jeff & Manski, Charles F. (1994/2006 posting). "Using Expectations Data to Study Subjective Income Expectations." NBER Working Paper w4937 ([SSRN page](#)). ([SSRN](#))
- Jäger, Simon; Roth, Christopher; Roussille, Nina; Schoefer, Benjamin (2024). "Worker Beliefs About Outside Options." *Quarterly Journal of Economics* (abstract page). ([OUP Academic](#))
- Caliendo, Marco; Mahlstedt, Robert; Schmeißer, Aiko; Wagner, Sophie (2024). "The Accuracy of Job Seekers' Wage Expectations." IZA Discussion Paper 17198 (abstract page). ([IDEAS/RePEc](#))
- Krueger, Alan B. & Mueller, Andreas I. (2016). "A Contribution to the Empirics of Reservation Wages." *AEJ: Economic Policy* (abstract page). ([American Economic Association](#))
- Arnold, David; Quach, Simon; Taska, Bledi (2025). "The Impact of Pay Transparency in Job Postings on the Labor Market." *SSRN abstract page*. ([SSRN](#))
- Indeed Hiring Lab (2024). "Pay Transparency in Job Postings Has More than Doubled Since 2020." (Indeed Hiring Lab)
- Belot, Michèle; Kircher, Philipp; Muller, Paul (2022). "How Wage Announcements Affect Job Search—A Field Experiment." *AEJ: Macroeconomics* (abstract page). ([American Economic Association](#))
- Belot, Michèle; Kircher, Philipp; Muller, Paul (2019). "Providing Advice to Jobseekers at Low Cost: An Experimental Study on Online Advice." *Review of Economic Studies* (abstract/results text). ([OUP Academic](#))
- Leibbrandt, Andreas & List, John A. (summary page). "Do Women Avoid Salary Negotiations? Evidence from a Large Scale Natural Field Experiment." Harvard Kennedy School Gender Action Portal. (gap.hks.harvard.edu)
- Jensen, Robert (2010). "(Perceived) Returns to Education and the Demand for Schooling." *Quarterly Journal of Economics* (abstract page). ([OUP Academic](#))
- Cullen, Zoe B. (2023). "Is Pay Transparency Good?" NBER Working Paper 31060 (RePEc abstract page). ([IDEAS/RePEc](#))
- Baker, Michael; Halberstam, Yosh; Kroft, Kory; Mas, Alexandre; Messacar, Derek (2019; revised 2021). "Pay Transparency and the Gender Gap." NBER Working Paper 25834. ([NBER](#))
- Gulyas, Andreas; Seitz, Sebastian; Sinha, Sourav (2023). "Does Pay Transparency Affect the Gender Wage Gap? Evidence from Austria." *AEJ: Economic Policy* (abstract page). ([American Economic Association](#))
- Frimmel, Wolfgang; Schmidpeter, Bernhard; Wiesinger, Rene; Winter-Ebmer, Rudolf (2024). "External Pay Transparency and the Gender Wage Gap." (RePEc abstract page). ([IDEAS/RePEc](#))
- Harris, Benjamin (2018). "Information is Power: Fostering Labor Market Competition through Transparent Wages" ([Brookings](#))
- Hall, Robert and Krueger, Alan (2012). "Evidence on the Incidence of Wage Posting, Wage Bargaining, and On-the-Job Search". *American Economic Journal: Macroeconomics* ([AEJ: Macroeconomics](#))

Appendix

Synthetic wage-benchmarking prompts

The examples below are synthetic sample exchanges written as if the prompt had been entered into ChatGPT. The prompts are invented, but where applicable the sample responses are anchored to the same benchmark data used in the evaluation and are written in a longer, more realistic style.

Example 1: National occupation benchmark

User prompt

I'm considering diagnostic medical sonography. What does a diagnostic medical sonographer typically make in the U.S., and is that enough to justify the training?

Sample response

A reasonable national benchmark is about \$89,340 per year for diagnostic medical sonographers. That is a useful starting point if you're asking whether the field can support a stable middle-income career, and the short answer is yes: for a two-year allied-health path, that is a meaningful earnings outcome.

The part I would be careful about is treating the national number as the number you personally will earn right away. Sonography pay moves around based on location, hospital versus outpatient setting, shift differentials, call requirements, and whether you later specialize. Someone working days in a lower-cost market may land below that benchmark at first, while a hospital-based role with evenings, weekends, or a tighter labor market can come in above it.

If you're deciding whether the training is worth it, I would frame it in payback terms: tuition, prerequisite time, certification costs, and the income you give up while training versus what you can plausibly earn once you finish. If you want, I can turn this into a simple break-even calculation using your program cost, how long the training takes, and the metro where you expect to work.

Example 2: Metro-specific occupation benchmark

User prompt

What does a registered nurse make around Chicago? I'm trying to decide if relocating there would be worth it.

Sample response

A reasonable benchmark for registered nurses in the Chicago-Naperville-Elgin area is about \$96,480 annually. That is the kind of number I would use as your anchor if you're comparing Chicago to another market, because it is much more relevant than the national RN median.

That said, relocation decisions should not treat the metro median as your offer. RN pay in Chicago can move meaningfully based on hospital system, unit, union coverage, shift mix, overtime availability, and years of experience. A med-surg staff role, an ICU role, a clinic job, and a travel-adjacent opening can all look very different even within the same metro.

If you're asking whether the move is worth it, I would compare three things side by side: expected base pay, expected schedule and differential pay, and your cost of living delta - especially housing, commuting, and licensing friction. If you want, I can help you compare Chicago against your current city using a rough after-housing budget instead of just headline salary.

Example 3: Role at a specific employer

User prompt

I have an interview for an assistant store manager role at a big-box retailer in Chicago. What salary range should I expect, and how should I think about negotiating it?

Sample response

A practical local anchor is the Chicago benchmark for first-line supervisors of retail sales workers, which is about \$48,530 a year. For an assistant store manager role at a big-box retailer, I would usually expect the base salary to land somewhere around that market level, with the exact number depending on store volume, weekend and closing expectations, prior people-management experience, and whether the role carries bonus eligibility.

If the company gives you a range, I would not focus only on the top-line base. Ask whether the role is exempt or non-exempt, whether bonuses are realistic or mostly theoretical, how staffing shortages affect hours, and whether you are effectively covering manager duties without manager pay. In retail, schedule burden and labor pressure can change the value of an offer a lot more than a few thousand dollars in base salary.

For negotiation, a good approach is: cite the local market benchmark, briefly tie your ask to your relevant experience, and then ask about the full package - bonus structure, PTO, promotion timing, and schedule predictability. If you want, I can help you script a salary ask based on the number they gave you.

Example 4: Local entry-level care job

User prompt

What does a CNA make in Houston? I only have about a year of experience, so I want a realistic benchmark, not the highest-end number.

Sample response

A fair Houston-area benchmark for nursing assistants is about \$36,940 per year. With about a year of experience, I would use that as a realistic market anchor rather than assuming you'll come in at the high end right away.

The most important caveat is setting. CNA pay can look noticeably different across hospitals, skilled nursing facilities, rehab centers, long-term care, and agency work. Some settings offer a lower base but steadier hours and better benefits; others may pay more because of nights, weekends, or heavier patient loads.

If you're comparing actual jobs, I would ask for the hourly rate, differential pay, guaranteed hours, and how often overtime is available or expected. For entry-level healthcare jobs, those details can change annual take-home more than the posted base rate. If you want, I can help convert Houston CNA offers into apples-to-apples annual numbers.

Example 5: Entrepreneurship income question

User prompt

Could a solo pressure-washing business in Atlanta realistically clear six figures, or is that internet hype?

Sample response

It is possible, but I would not treat six figures in owner take-home as the default case for a solo pressure-washing business. The biggest mistake in these conversations is confusing revenue with what the owner actually keeps after equipment, chemicals, fuel, insurance, marketing, taxes, and days when work falls through.

A more realistic way to think about it is by jobs, pricing, and utilization. If you're mostly doing residential work and you're still building a customer base, earnings can be decent but very inconsistent. Clearing six figures personally usually means some combination of strong repeat demand, good upselling, higher-ticket commercial accounts, or eventually adding labor so the business can do more work than one person alone can handle.

So I would call it achievable but not a safe base-case assumption. If you want, I can help you build a simple owner-earnings model for Atlanta using assumptions for jobs per week, average ticket size, seasonality, and overhead. That will give you a much better answer than generic internet revenue claims.

Example 6: Pay calculation and comparison

User prompt

If one job pays \$29 an hour and another offers \$61,000 salary, which one is better once I account for a normal full-time schedule?

Sample response

At a standard full-time schedule of 40 hours a week for 52 weeks, \$29 an hour works out to about \$60,320 per year. So on straight base pay, the hourly job and the \$61,000 salary are basically even.

The next question is what happens after that baseline calculation. If the hourly job offers overtime, shift differentials, or holiday pay, it could easily come out ahead in actual cash compensation. If the salaried job regularly expects 45 to 50 hour weeks, then the effective hourly rate may actually be lower than the hourly role even though the salary headline looks slightly better.

I would compare the two offers on four dimensions: total cash pay, expected hours, benefits, and schedule quality. If you want, send me the details on overtime eligibility, retirement match, health insurance cost, and expected weekly hours, and I can turn it into a side-by-side comparison.