Empowering Inclusive Work*

Yanyou Chen
University of Toronto

Mitchell Hoffman
UC Santa Barbara and NBER

Huilan Xu
Zhejiang University

Zhe Yuan
Zhejiang University

August 2026

Abstract

How do workers with disabilities perform in the same job as non-disabled workers, and can AI-enabled tools narrow any gaps? We examine the relative performance of deaf or hard-of-hearing (DHH) workers on a major Chinese food-delivery platform, providing the first large-scale productivity comparison of disabled and non-disabled workers, and we evaluate the effects of an AI-based text-to-speech calling tool. Pre-AI, DHH workers are slower and receive worse customer ratings, but work more hours, quit less, and generate higher platform profits. After the tool's sudden introduction mid-sample, difference-in-differences estimates show faster deliveries; substantially fewer bad customer ratings; and higher labor supply, retention, and profits. The AI tool closes one-third of the disability hourly-pay gap.

Keywords: Disability, productivity, accommodation, personnel economics, pro-worker AI, gig work, deaf or hard of hearing

JEL Codes: M50, J14

*We thank Ach Adhvaryu, Amanda Agan, Naoki Aizawa, Heski Bar-Isaac, Gordon Dahl, Chiara Farronato, Kory Kroft, Peter Kuhn, Rafael Lalive, Imke Reimers, Kathryn Shaw, Joe Vecci, and especially David Autor for helpful comments. We also thank numerous seminar participants, especially those at NBER Labor Studies and NBER SI Personnel, for useful suggestions. We thank Xiaonan Zhu, Weiyue Zhang, Ruibo Sun, Yaqi Du, Yu Xu, Yibing Xia, and Sunyu Ye for excellent research assistance. We thank Schmidt Sciences and the National Natural Science Foundation of China [Grants 72141305, 72192803] for support. This project has been supported by the Gates Foundation (investment ID INV-044136), through a subgrant from IDinsight's Digital Economy Research Impact Initiative (DERII). We received ethics approval from the University of Toronto. One author has previously received compensation from the Platform for work unrelated to the research reported in this paper. The paper was reviewed to ensure that proprietary or confidential information is not revealed. We thank the Platform and its employees for their collaboration and for providing data. We thank delivery workers for participating in the qualitative data collection. Any opinions, findings, conclusions and recommendations expressed are those of the authors and do not necessarily represent the views of the Platform.

Around the world, workers with a disability face challenges in the labor market (Baldwin and Johnson, 2006; L'Horty et al., 2022). They have lower rates of employment and face substantial discrimination (Bellemare et al., 2023), and there are large fiscal costs from high use of disability insurance (Autor and Duggan, 2003; Milligan, 2012). Challenges faced by disabled workers[1] are especially acute in middle-income countries like China where government social insurance is often less developed (Banerjee et al., 2024).

Unfortunately, relatively little is known regarding "what works" in improving outcomes for disabled workers. A central policy to increase opportunity for disabled workers is civil rights laws, a key example being the Americans with Disabilities Act (ADA). The ADA prohibits discrimination against disabled workers and mandates firms to provide reasonable accommodations, and similar laws exist in many countries. However, research indicates that the ADA *decreased* the employment of disabled workers, potentially due to increased costs or legal risks for employers (Acemoglu and Angrist, 2001; DeLeire, 2000). Another common policy in many countries (e.g., Austria, Brazil, China) is to impose quotas, requiring larger firms to employ a minimum share of workers with disabilities. However, firms may comply with quotas by hiring disabled workers into lower-paid jobs (Szerman, 2022; Lalive et al., 2013). Rather than imposing compliance costs through civil-rights enforcement or quotas, an alternative is to rely on technologies like artificial intelligence (AI) that directly raise productivity—potentially increasing inclusion while aligning with firms' incentives.

Our study provides the first evidence from a large-scale setting on disabled workers' productivity and the effects of a targeted AI-enabled tool. We address two questions:

1. In the same job, how do disabled workers compare to non-disabled workers in terms of productivity, labor supply, output, wages, and profit for firms?

2. Can AI-enabled tools affect these outcomes for disabled workers, and how?

To address these questions, we leverage a unique collaboration with a major food delivery platform in China (hereafter, the "Platform"). The Platform offers an ideal setting for three reasons. First, the Platform operates on an immense scale: it is one of the two dominant food-delivery apps in China, directly engaging millions of active workers and completing hundreds of millions of deliveries per month. Second, the Platform follows an open hiring policy: all applicants who meet basic criteria are accepted, with no differential standards for disabled workers (Ameri et al., 2018). Thus, disabled workers are employed under the same conditions as their non-disabled counterparts, allowing for a cleaner comparison of on-the-

[1]While some use the *person-first* "worker with a disability," many disability scholars prefer the *identity-first* term "disabled worker" (Gernsbacher, 2017; Grech et al., 2024). We use both, but often the latter. About 1 in 6 people worldwide have a significant disability (World Health Organization, 2023).

job performance. Third, work on such platforms has become a key source of employment for disabled workers in China.

We focus on workers who are deaf or hard of hearing (DHH), who constitute the majority of disabled workers on the Platform, and compare their outcomes to those of non-disabled workers. People with different disabilities face different challenges and may be affected by AI-enabled tools in different ways. Focusing on a single, common disability allows us to more clearly examine the relevant challenges. Moreover, hearing impairment is a disability that gig platforms around the world increasingly seek to accommodate, yet important challenges remain—making our findings broadly relevant.[2]

Our analysis draws on a rich administrative dataset from the Platform (Section 1), covering all newly registered DHH workers and a random sample of newly registered non-disabled workers over a three-year period (2021–2023). This dataset allows us to directly compare workplace outcomes of DHH workers to non-disabled workers on many dimensions. The data include worker productivity variables capturing both the quantity of work, such as speed at different stages of order fulfillment, and work quality, as captured by customer ratings. The data also include labor supply variables on both the intensive and extensive margins, as well as information pertaining to worker compensation and Platform profits. In the middle of our observation period in mid 2022, the Platform introduces an AI-based "Intelligent Outbound Calling" feature designed to help DHH workers communicate with customers. We propose a simple conceptual framework about differences in workplace outcomes between disabled and non-disabled workers, and about how those differences are affected by AI tools.

Prior to the new AI tool, DHH workers are slightly slower than non-disabled workers (Section 2). This gap is especially pronounced in the final stage of order fulfillment—delivery (i.e., leaving the restaurant to customer drop-off), which involves interaction with customers—where DHH workers are 5% slower. DHH workers are substantially more likely to experience "left tail" outcomes, with 9% more late deliveries and 31% more bad customer ratings (though bad customer ratings are rare). Despite this, DHH workers have substantially higher labor supply: they work 5% more hours per week and are 41% less likely to quit. The higher labor supply of DHH workers more than offsets their small efficiency gap, with DHH workers completing more orders per week and generating higher weekly profits for the Platform than non-disabled workers. However, because DHH workers work slower, receive more bad ratings, and work more hours relative to income earned, they earn 10% less per

[2]About 8% of Chinese adults have moderate to severe hearing loss (Wang et al., 2025). China's rate is broadly in line with the World Health Organization (2025) global prevalence rate of 5% for disabling hearing loss (i.e., hearing loss of more than 35 decibels). Chen et al. (2024a) report that there are nearly 100,000 registered DHH couriers in China. Rates of US hearing impairment vary, from 3.7% of all Americans (Erickson et al., 2026) to 13% of US adults with some impairment (Madans et al., 2021).

hour than non-disabled workers, broadly similar to disability wage gaps in other settings.

As detailed in Section 3, the AI tool we study is a text-to-speech (TTS) phone calling system offering preset and customized audio prompts, and that uses AI to make calls sound natural. Pre-AI TTS systems sound robotic and are frequently confused for spam calls, so AI is critical for getting customers to listen to and act on the calls. Rolled out suddenly to all DHH workers—but to no other workers—we exploit this as a natural experiment to estimate, via difference-in-differences (DiD), how the AI tool affects workplace outcomes for DHH workers. The AI tool causes DHH performance to substantially improve. DHH work speed improves, driven by faster completion of the trip segment involving customer communication, where they lagged the most previously. They also have far fewer late deliveries and bad ratings, with the tool fully eliminating the pre-AI gap in late deliveries and cutting the gap in bad ratings by two-thirds. Orders increase, and so do hours and retention. These improvements translate into higher profitability of DHH workers, with Platform gains far larger than costs. The AI tool draws more DHH workers to the Platform, and it narrows the disabled worker pay gap from 10% to under 7%, a decrease of one-third. Sufficient-statistics welfare calculations imply that DHH workers gain the equivalent of roughly 4–12% of pre-AI earnings, so the tool's benefits are shared between the Platform and its workers. While DHH workers with non-profound (i.e., less severe) hearing loss also benefit, gains are largest among those with profound hearing loss. Overall, the results support that the key mechanism for effects is improved communication with customers, as opposed to alternatives like increasing customer sympathy for DHH workers. We also tend to find larger gains for DHH workers in stronger local labor markets, particularly on labor supply and total output.

This paper contributes to several literatures, the most important being the economics of disability. A large body of work studies the determinants and consequences of disability insurance and other social safety programs for disabled workers (Bound and Burkhauser, 1999; Deshpande, 2016; Autor et al., 2019; Deshpande and Lockwood, 2022). There is also work on disabled worker pay gaps (Forth and Jones, 2025), as well as estimating the impact of civil rights laws and quotas for disabled workers, but relatively little is known about what works in improving labor outcomes for disabled workers.[3] An important exception is Bloom et al. (2026), who find that the large rise in work from home has increased the employment of disabled workers. Accommodations are a natural candidate for improving disabled worker outcomes, but a central concern is their cost (Acemoglu and Angrist, 2001; Aizawa et al., 2025), and there is little research measuring the returns to particular accommodations (as

[3] Lise et al. (2023) find the ADA reduced employment for workers with physical limitations but increased it for those with non-work limitations like mental illness. Derbyshire et al. (2024) survey the impact of employer-focused interventions on disability employment.

we do). AI may alter the economics of accommodation, as well as which activities can be accommodated. Traditional accommodations often carry costs that scale with the number of workers or worksites (e.g., human assistance for communication challenges, ramps for mobility challenges). In contrast, the tool we study required only a modest one-time development cost: AI made the accommodation effective—producing speech natural enough for customers to engage with—while the near-zero replication cost of software let it reach every DHH worker on the Platform.[4] The tool paid for itself within days, making accommodation profitable rather than merely mandated. We are not aware of prior work in any academic field that compares disabled and non-disabled workers on a large scale in terms of individual productivity data.[5] The intervention we study, a simple AI-based communication tool, has the largest impact on closing the disability wage gap of any intervention we are aware of.

Second, it contributes to work on the economics of AI, especially the emerging discussion about pro-worker AI (Acemoglu et al., 2026). While AI use is accelerating, causal evidence on its impact for disadvantaged groups is scarce. As surveyed by Fruits and Stout (2026), work examines how AI affects outcomes for different types of workers, including accountants (Choi and Xie, 2025), translators (Merali, 2024), call center agents (Brynjolfsson et al., 2025), professionals in laboratory settings (Noy and Zhang, 2023), and software engineers (Cui et al., 2025; Hoffmann et al., 2024). In contrast, our study is the first to evaluate how a targeted AI tool affects workplace outcomes for disabled workers.[6] To our knowledge, our paper is also the first to estimate the effects of a specific "pro-worker AI" tool, defined by Acemoglu et al. (2026) as an AI tool that expands worker capabilities by increasing the value of worker skills and expertise (see also Autor 2024). The tool we study is a natural example of pro-worker AI because it enables DHH workers to communicate with customers—a simple and inexpert task for non-disabled workers—thus allowing them to potentially expand or improve their work.[7] Our heterogeneous treatment effects suggest that disability may critically differ from

[4]As we discuss further in the Conclusion, AI can provide numerous forms of potentially novel accommodations, e.g., serving as "eyes" for blind or low-vision people by describing their surroundings—awareness that fixed aids like braille markings cannot provide.

[5]This reflects data constraints rather than lack of broad scholarly interest in understanding disabled worker productivity and challenges. Prior evidence generally uses surveys and supervisor ratings, or uses small samples of administrative data. Appendix B reviews this literature and discusses why such comparisons have not previously been feasible. Human–computer interaction scholarship uses interviews to examine challenges faced by DHH workers on gig platforms, including Uber (Lee et al., 2019) and food delivery in China (Chen et al., 2024a). The Uber study is based on interviews with five drivers, supplemented by qualitative analysis of in-app feedback, while the China study relies on interviews with six couriers and four station managers. Neither analyzes large-scale administrative data nor examines the impact of AI tools.

[6]While generally not based on administrative data and usually lacking quasi-experimental variation in technology, research on personal computers suggests some role of technology in improving disabled worker outcomes (Kruse et al., 1996; Krueger and Kruse, 1995).

[7]On the basis of our paper, Acemoglu et al. (2026) describe the tool we study as an example of pro-worker AI. Their other examples are based on company accounts, agency statistics, or tools under development rather

other dimensions of workplace inequality like lower productivity or lower tenure, as we find heterogeneous effects by disability severity but not for those dimensions.

Third, our paper contributes to work in personnel and organizational economics (Oyer and Schaefer, 2011; Bloom and Van Reenen, 2011) on firm returns from disadvantaged workers, a topic that has received relatively scant attention (Hoffman and Stanton, 2025). One recent exception is Benson et al. (2024), who analyze productivity differences between white and minority workers in a large US retail firm. There is also work on firm returns from engaging workers with a criminal record (Minor et al., 2018). To our knowledge, ours is the first paper to examine the *profitability* of disabled workers for an organization or platform. Pre-AI, while DHH workers are less efficient than non-disabled workers, they are still more profitable overall. This highlights the importance of capturing a broad set of outcome variables in evaluating disabled worker profitability. An analyst who focused only on efficiency variables, and ignored the labor supply variables of hours and retention, would draw mistaken conclusions about the relative profitability of DHH workers, a point we return to in the Conclusion (Section 4).

1 Context, Conceptual Framework, and Data

The Platform is one of the two leading companies in China's massive online food-delivery market. It operates similarly to US apps like Uber Eats or DoorDash but on a much larger scale. Together, the two main Chinese platforms engage around 10 million delivery workers at a time (Lin, 2025). The overall Chinese gig economy engages around 200 million workers, larger than the entire US labor market (The Economist, 2025).

Delivery workers on the Platform use motorbikes or e-bikes to transport food, groceries, and other items from restaurants, convenience stores, and pharmacies to customers. The Platform is available nationwide, with a higher concentration of workers in urban and eastern regions of China where food delivery is most prevalent.

Our engagement with the Platform. Our engagement with the Platform on this project began in late 2023. One of the coauthors, who had received an AI call from a DHH worker, reached out to the Platform about the possibility of evaluating the impact of the AI tool. While the Platform frequently evaluates the impacts of changes to new customer features or other app improvements, it had never evaluated the impact of the AI tool for DHH workers, but seemed keen to better understand its impact.[8]

than academic studies.

[8]The Platform's driver operations team that we collaborate with does much less policy evaluation than other teams within the Platform.

"Hiring" policy. The Platform maintains an open registration policy (referred to hereafter as "hiring," though drivers are independent contractors). To begin working, an applicant must meet three requirements: be at least 18 years of age, possess valid identification, and have no disqualifying criminal record. These standards are uniform across applicants, with no separate criteria for workers with disabilities.

Disability status on the Platform. Across both major Chinese food delivery platforms, most disabled workers are DHH.[9] The Platform defines a worker as being disabled based on the driver uploading their *certificate of disability* during registration. The certificate is a formal, government-issued document requiring medical evaluation by doctors and is used to determine disability insurance benefits. Appendix G provides a screenshot of the disability question during the registration process.

During registration, both pre- and post-AI, drivers are asked (1) what type of disability, if any, they have, (2) the disability level, (3) whether they consent to the Platform sharing their disability with customers so that there is more understanding in unexpected situations. Disability level for DHH workers is coded based on the Chinese national four-grade system (GB/T 26341-2010), with further details on Chinese disability levels in Appendix F:

- Level 1, which is hearing loss of greater than 90 decibels, abbreviated as $>$ 90 dB HL. This is profound hearing loss, where someone is unable to hear even a shouted voice, and spoken speech is distorted (Olusanya et al., 2019). Extremely serious barriers occur regarding participation in social life (GB/T 26341-2010).
- Level 2 or 81-90 dB HL. This is severe hearing loss.
- Level 3 or 61-80 dB HL. This is moderately severe hearing loss.
- Level 4 or 41-60 dB HL. This is moderate hearing loss; individuals at this level can typically hear words only when spoken in a raised voice at close range, making unaided phone communication difficult (Olusanya et al., 2019).

Based on our qualitative work and understanding of the Platform, essentially all certified DHH workers disclose their disability to the Platform, a point we discuss further in Appendix G. Drivers know their level since it is listed on their certificate of disability.

[9]For one of the two major platforms, it was publicly reported that, as of end-2025, 80% of its roughly 15,000 active disabled workers were DHH (Minsheng Weekly, 2026). Two factors help explain why DHH workers predominate. First, since food delivery requires riding an e-bike through traffic, it may be hard for people with some disabilities (e.g., visual, mobility), whereas hearing loss leaves the core tasks intact. Second, conditions that account for a large share of US disability rolls, such as chronic pain or depression, may not qualify for a Chinese disability certificate.

Both before and after the introduction of the AI tool, customers are informed via the app that their driver is DHH, for those drivers who consented to this disclosure during registration. Though we do not have precise data on this, based on our conversations with the Platform, we believe consent rates are high. Despite most DHH workers electing to disclose their disability to customers, customers may not always attend to this information. There is a general expectation in China that drivers call customers upon arrival, and this can pose communication challenges for DHH workers.

Work models: team vs. gig. Workers on the Platform can choose between two work models: a *team* model, in which workers are assigned to stations and are paid under an accelerating piece-rate schedule, with unit prices that increase stepwise in the number of monthly completed orders, and a *gig* model, in which workers operate independently and are paid per delivery. Exact pay is determined algorithmically and varies across deliveries, with longer deliveries typically paying more. The team model is intended to be full-time while still allowing workers to choose their hours, whereas work in the gig model can be part time. Both gig and team workers are included in our data (about 2/3 are gig). We control for work model in our analyses, and present key results separately for gig and team workers.

Work model is also referred to as registration type. If a worker wants to switch work models, they need to re-register. Based on our conversations with the Platform, we believe that relatively few workers switched work models during the period of our data.

In both models, drivers tend to work for one delivery platform at a time (Lin, 2025).[10]

Matching of workers to deliveries. Gig and team workers also differ in order assignment. Gig workers primarily rely on independently selecting orders from the order pool.[11] Gig workers operate across the entire city, often covering multiple commercial districts. In contrast, team workers almost entirely depend on system-assigned orders. The typical assignment range for team workers is within 6km of their designated station.

The orders that appear for gig workers in the order pool, and the order that a team driver is assigned, depend on a driver's location, demand conditions, and the driver's contemporaneous *matching score*. The exact matching score is proprietary, and depends on a driver's (1) service score, (2) "Rider Level," and (3) order acceptance rate. The service score depends on customer reviews and late deliveries. The Rider Level is calculated based on tenure and the cumulative number of completed orders. Refusing system-assigned orders directly

[10]Team workers are prohibited from simultaneously working for other food-delivery platforms. Gig workers face no explicit prohibition, but working for another platform reduces their order history on this Platform, lowering their matching score.

[11]While gig workers do receive some system-assigned orders, these are preferentially assigned to team workers under the same conditions. On average, about 20% of gig workers' orders come from the system, while the remaining 80% are manually selected by the workers themselves.

reduces a worker's Acceptance Rate. Workers report that this factor has the greatest impact on assignment score, so they usually do not reject system-assigned tasks. The algorithm's formula does not include disability status as an explicit input.[12]

In sum, gig workers generally choose their orders and team workers are assigned orders. Both the order pool for gig and the assignments for team depend on workplace outcomes.

1.1 Conceptual Framework

We present a simple framework of labor supply on the Platform to ground our empirical analysis and to make predictions about differences between disabled and non-disabled workers, as well as how disabled worker productivity is affected by AI tools.

Workers have a productivity or efficiency level on the Platform, z, which may differ between disabled and non-disabled workers, and based on access to disability-focused AI tools. The share of total work time allocated to Platform work is denoted by h, while $1-h$ represents the share of time spent on other work pursuits. Output on the Platform is zh. Workers get paid a simple piece rate w, so their hourly wage is wz.

The return to work off the Platform is $r \cdot \frac{\beta}{\beta-1}(1-h)^{\frac{\beta-1}{\beta}}$, where r measures the value of one's outside option and $\beta > 1$. The parameter, β, captures how sharply the worker substitutes between work on and off the Platform. Worker utility is thus given by:

$$U = w \cdot zh + r \cdot \frac{\beta}{\beta-1}(1-h)^{\frac{\beta-1}{\beta}}$$

In our empirical analysis, we measure productivity broadly, incorporating speed at different points, customer reviews, and average total time per order. The firm makes no decisions and earns profits $A \cdot zh$, where A is a piece rate measuring the Platform's commission.

Workers make a decision about the share of their work time to spend on the Platform. Using the first order condition of $\frac{\partial U}{\partial h} = 0$, we get $h = 1 - r^{\beta}(wz)^{-\beta}$. Taking partial derivatives yields $\frac{\partial h}{\partial r} = -\beta r^{\beta-1}(wz)^{-\beta} < 0$ and $\frac{\partial h}{\partial z} = \beta r^{\beta} w (wz)^{-(1+\beta)} > 0$, i.e., people work more on the Platform when they have worse outside options and higher Platform efficiency.

Based on the literature and our qualitative data, we posit two main differences between disabled (denoted by D) and non-disabled workers (denoted by N):

Assumption 1. Disabled workers have worse outside options than non-disabled: $r_D < r_N$.

[12]More than assignments, the main thing affected by the matching score is the number of orders a driver is allowed to carry. The max number of concurrent orders starts at a baseline of 3, with an upper limit of 13 for gig workers and 20 for team workers. A higher assignment score results in a greater allowable maximum. Typically, experienced workers handle an average of six concurrent orders, while high-performing workers can manage even more, often reaching around ten concurrent orders.

This assumption (A1) is consistent with a large literature, cited above, on the challenges and discrimination faced by disabled workers.

Assumption 2. In the absence of AI tools, disabled workers have lower efficiency: $z_D < z_N$.

A2 is consistent with what we gleaned from conversations with DHH workers about work before the AI tool, namely, that there were often bottlenecks while making deliveries.[13]

Outcomes on the Platform can be grouped into those related to efficiency (z), labor supply (h), total output (zh), and income, wages, and profits. The base framework and the assumptions yield four implications (I1-I4) that we examine using the data.

Implication 1. Before the new AI tool, disabled workers have lower efficiency and wages than non-disabled workers.

Implication 2. Before the new AI tool, differences in labor supply between disabled and non-disabled workers are ambiguous. Disabled workers will work more on the Platform if the difference in outside options between disabled and non-disabled workers is large relative to differences in efficiency.

Implication 3. Before the new AI tool, average differences in total output, income, and profits between disabled and non-disabled workers are ambiguous.

Implication 4. Suppose an AI tool increases the efficiency of disabled workers. This will increase their labor supply, total output, wages, and profits on the Platform.

I1 follows directly from A2: lower efficiency implies lower wages since the hourly wage is wz. For I2, the difference in labor supply is ambiguous because A1 and A2 push in opposite directions: A1 makes disabled workers work more and A2 makes disabled workers work less. For I3, total output is zh, and since h is ambiguous, so is zh; income wzh is likewise ambiguous. For I4, by increasing z_D, an AI tool raises wages wz_D directly and also encourages disabled workers to increase labor supply, and thus two forces together work toward increasing total output and profits.

Framework extension: multiple severities of disability. While the framework has one severity of disability, it is easily extended to have multiple severities. Suppose, as in our setting, that there are workers who are profoundly disabled ("PD") and those who are non-profoundly disabled ("NPD"). Extending the logic of A1 and A2, it is natural to assume that $z_{PD} < z_{NPD}$ and $r_{PD} < r_{NPD}$. These assumptions yield pre-AI implications mirroring

[13] Disability scholarship (e.g., Krieger, 2010; Lane, 1992) emphasizes that differences between disabled and non-disabled people reflect both individual traits and environmental factors (e.g., discrimination, accommodations). Thus, A2 is a joint statement about capabilities and the environment, not inherent ability.

those between disabled and non-disabled. Turning to the impact of the AI tool, suppose communication situations vary in difficulty. Workers with non-profound hearing loss may be challenged primarily in difficult communication situations, whereas workers with profound hearing loss may face challenges in "simpler" interactions. If the AI tool helps resolve simpler communication situations but not the most difficult ones (Autor, 2015), then productivity gains from the AI tool will be larger for profoundly disabled workers.

Implication 5. Before the new AI tool, profoundly disabled workers have lower efficiency and wages than non-profoundly disabled workers, but differences in labor supply, output, income, and profits are ambiguous. The AI tool will have larger effects on profoundly disabled workers relative to non-profoundly disabled workers.

Framework extension: multiple local labor markets. Since the Platform operates across China, we observe AI impacts across a range of local labor markets, from the "rust belt" of Northeast China to the richer provinces in the Southeast. Disabled workers likely have stronger outside options when the local labor market is stronger, which can be proxied by a lower local unemployment rate, i.e., $r_D^S > r_D^W$, where r_D^S and r_D^W denote disabled workers' outside options in strong and weak local labor markets, respectively. Since $\frac{\partial^2 h}{\partial r \, \partial z} > 0$, an AI-induced increase in productivity leads to larger increases in hours for disabled workers in stronger local labor markets than for disabled workers in weaker local labor markets.[14]

Implication 6. The AI tool will have larger effects on hours, output, income, and profits for disabled workers in stronger local labor markets compared to disabled workers in weaker local labor markets.

Framework simplicity. This framework is deliberately simple and abstracts from several features of reality. First, true worker "productivity" is multidimensional—encompassing speed at different stages of order fulfillment, late deliveries, and bad ratings—whereas the model summarizes these dimensions with a single parameter z. Second, the model does not explicitly distinguish between hours worked and retention. In our setting, it is useful to think of "labor supply" broadly as including hours on the Platform, timing of work, and continued participation, e.g., workers with worse outside options may both work more hours and be less likely to quit the Platform. The empirical analysis examines these margins separately

[14]The positive cross-derivative follows from concavity of the outside option. Workers with stronger outside options spend more time outside the Platform, placing them where the marginal return to outside time is relatively insensitive to further increases in outside time. When AI raises platform productivity, these workers therefore find it relatively attractive to shift time back to the Platform, generating a larger labor supply response. This concavity is natural if outside opportunities exhibit decreasing returns: the most productive outside hours are used first, whereas jobs on the Platform get paid with a relatively constant piece rate.

and documents pre-AI differences between DHH and non-disabled workers, as well as how the AI tool affects those differences.

1.2 Data

We obtained a proprietary administrative panel from the Platform, organized at the worker-week level and containing detailed measures of worker activity and performance. The panel covers workers who register on the Platform in 2021-2023 and complete at least one week of work, and includes all 4,284 team or gig workers who report DHH as their only disability, as well as nearly 40,000 non-disabled workers randomly sampled from the same population.[15] Drivers with multiple disabilities were excluded from the panel shared with us.

The panel allows us to track each worker from the start of their tenure on the Platform and to analyze how their performance evolves over time. By focusing on newly registered workers, we compare DHH and non-disabled workers from their point of entry, avoiding stock-sampling concerns related to differential tenure.

The AI tool was rolled out suddenly at t_0 in mid 2022 to all DHH workers. To avoid identifying the Platform, we do not specify the exact month of t_0, but it is a date during April 2022-September 2022, thus giving over one year of pre-AI data to compare the performance of DHH and non-disabled workers, and we can track workers for more than a year to study the AI tool's impact. All variables are available for the full 2021-2023 period unless otherwise stated. We examine a wide range of worker outcomes, which we divide into four groups:

Efficiency: z in our framework. Efficiency or productivity refers to the speed (or quantity) and quality of work. Beyond overall speed, our data enable us to look at time spent during the three steps of the work process. At each step, the worker must perform an action on the app, which provides a clear time stamp for when each step starts and finishes.[16]

The three steps are as follows: First, once an order is initiated by the customer, it is assigned to a worker, via direct assignment by the Platform or via acceptance by the worker. The time for this first step is called the *order acceptance time*. Second, once the worker receives the order, they pick it up from the restaurant. The waiting time for food preparation is included in this step, which concludes when the worker actually receives the food from the restaurant. The time for this second step is called the *pickup time*. Lastly, the worker delivers the order from the store to the customer, the time for this being the *delivery*

[15] Beyond team and gig, a small share of workers on the Platform (we believe less than 5%) fall into a third work model called preferred gig. This group was excluded from the data shared with us, as essentially none of these workers are DHH and their work is different from team and regular gig.

[16] The relevant buttons can only be activated when GPS confirms the worker is sufficiently close to the relevant location, thus preventing manipulation. We do not have stage times by order, but rather as average time per order at the worker-week level.

time. Beyond overall speed, we also observe Platform-defined late delivery rate, which refers to deliveries that are late beyond the delivery window, indicating a significantly late order.

Quality is measured using customer ratings. Rather than use 1–5 stars, customers on our Platform click a button to indicate a Bad Rating, a Good Rating, or nothing (most common). Bad ratings are at the discretion of the customer and occur for different reasons, including late deliveries, not receiving the delivery, and food damage during delivery (e.g., spilled soup). Drivers incur a monetary penalty from the Platform for each bad rating received, giving bad ratings direct financial consequences.

Drivers can also receive Good Ratings, but we emphasize these much less, as they do not affect worker compensation or order assignment, and because the Platform emphasizes them much less. We therefore report results for Good Ratings only in the Appendix.

We cannot analyze accidents or injuries, as the Platform does not comprehensively record them in the administrative data we received.

Labor supply: h in our framework. Our two main ways of measuring labor supply are (1) a worker's weekly hours spent online (i.e., available for orders) and (2) whether the worker quits the Platform. Beyond total hours, we also analyze days worked per week (i.e., so we can see if hours effects are coming from more days or more hours per day).

On (1), hours data are measured based on the worker logging into the app or the app is working in the background. Hours data are available for most of the sample period, though missing in 28% of worker-weeks (the biggest determinant being that hours data are not available prior to $t_0 - 46$, i.e., not available prior to mid 2021); we discuss the missingness pattern further in Section 3.6.1, where we provide evidence that hours missingness does not drive our findings.

On (2), the data does not contain a formal flag for attrition. Instead, we proxy whether someone exits the Platform based on having 0 weekly orders for 8 weeks in a row.[17]

Total output: zh in our model. Measured by a worker's orders completed per week.

Income, wages, and profits. The administrative data we received does not contain fields for driver income, driver wages, and Platform profits. Instead, we use parameter information provided by the Platform, combined with our administrative data, to estimate these variables for each driver-week.

As on US delivery platforms, drivers are generally paid a piece rate per order. While the true piece rate varies by order, we estimate income using a single piece rate for gig workers and accounting for the accelerated piece rate profile of team workers. Hourly wages are

[17] Our conclusions are similar under other reasonable definitions, e.g., 0 orders for 7 or 9 weeks in a row.

defined as total income divided by hours online. To isolate the role of bad ratings in wage differences, we additionally compute wages excluding any pay losses due to bad ratings.

Tipping is effectively absent in Chinese food delivery, so tips are neither included in our administrative data nor used in calculating driver income and wages.

Platform profits from a driver in a week, π_{it}, equal per-order profit times the number of orders completed, minus the cost of any onboarding, minus the cost of any bad ratings:

$$\pi_{it} = \pi_0(zh_{it}) \cdot zh_{it} - c \cdot 1[\text{t is driver i's first week}] - k \cdot \text{Bad Ratings}_{it} \tag{1}$$

Profit per order, $\pi_0(zh_{it})$, is allowed to vary with the number of orders completed per week by the driver, reflecting the fact that driver compensation per order depends on weekly order volume. Due to confidentiality restrictions, we do not report the exact assumed per-order profit schedule. Second, the recruiting and onboarding cost of a new worker, c, is assumed based on our discussions with the Platform to be 400 CNY (Chinese yuan; 7 CNY $\approx$ 1 USD during our data period).[18] The presence of this term accounts for any value DHH workers provide by having higher retention. Third, k is the Platform's cost of a bad customer rating. Based on discussions with the Platform, we assign k to be 100 CNY, with k intending to capture the full costs of negative reviews, including effects on customer retention and reputational spillovers, net of the penalty paid by drivers to the Platform.

Control variables. The data also include rich control variables, detailed below in Section 2.1, about workers and the type of work performed.

Qualitative data. We conduct qualitative data collection to better understand DHH worker experiences, building on recent studies of workers in organizations (Blader et al., 2020). This includes: (1) interviews with 7 DHH workers about AI tool usage, outside job opportunities, daily routines, and challenges; and (2) on-the-ground observations in which one coauthor rode along with 3 DHH workers over several days and another delivered orders wearing earplugs to simulate the DHH experience. These data help us interpret quantitative patterns and identify barriers faced by DHH workers. Appendix H provides further detail; we reference qualitative findings when discussing quantitative results.

1.3 Sample and Summary Statistics

Sample restriction. Our analysis sample is based on excluding weeks where drivers have 0 orders, as these are weeks where drivers are not working.

[18] The biggest determinant of this is a bonus the Platform pays to workers after their first week.

Summary statistics. Table 1 presents summary statistics for our 43,605 food-delivery workers. Compared to non-disabled workers, DHH workers are slightly older (mean age 33 vs. 30 for non-disabled) and more likely to be male (93% vs. 91%), but are similar in work models: among both DHH and non-disabled workers, about 35% are team workers.

Panel C shows the distribution of disability severity among DHH workers, with 54% having profound hearing loss (Level 1 in the Chinese system).[19] Panel D of Table 1 compares DHH and non-disabled workers in terms of geography and distance variables. DHH workers are slightly more likely to be located in big cities than non-disabled workers.[20] DHH workers' orders have slightly higher travel distance than non-disabled workers.

Table 2 provides summary statistics on outcome variables. Workers do an average of 152 deliveries per week. The standard deviation is large at 142, indicating considerable variation in total deliveries. Peak-hour orders, which occur during high-demand times, account for 91 deliveries per worker-week, indicating that a substantial share of total orders happens during peak times. Team-based workers complete significantly more deliveries on average (212) compared to gig workers (73), reflecting the full-time nature of team-based work.

Average order duration in our data is about half an hour, with mean times of 3.4 minutes for order acceptance, 12 minutes for pickup, and 14 minutes for delivery. Per 10,000 orders, the late delivery rate is 416, meaning 4% of orders are substantially delayed.

Bad ratings, our measure of customer satisfaction, are much rarer. The bad rating rate per 10,000 orders is 9.6, meaning only about 1 in 1,000 orders yields a bad rating. Across driver-weeks, the median bad rating share is 0: most driver-weeks do not have bad ratings.[21]

Workers spend an average of 40 hours online per week. The high standard deviation of 29 hours reflects wide dispersion across drivers. The median number of workdays per week is 6, though the average is lower at 5.3, suggesting that while many workers work full weeks, others participate part-time or intermittently. Average hours are 58 for team workers and 23 for gig workers, as seen in Appendix Table A.3.[22]

The weekly attrition rate on the Platform is 8% for the new drivers we study, with a high share quitting in the first few months. While such an attrition rate is high, it is comparable to that in other low-skill jobs, e.g., the call-center workers studied in Burks et al. (2015). The

[19]The share of drivers with profound hearing loss is higher in our sample than the overall Chinese population. A natural explanation is lower barriers to entry in gig work (e.g., no interview requirement) for people who are profoundly DHH, rather than differential disclosure to the Platform. Appendix G discusses further.

[20]One explanation is greater opportunities for DHH workers in large cities. Another is the historical concentration of specialized deaf schools there. We do not observe schooling histories.

[21]The low incidence of bad ratings may reflect a cultural norm in China where customers generally do not give bad ratings unless something has gone wrong, rather than as an expression of mild dissatisfaction.

[22]While hours on the Platform are higher than for US Uber drivers (Cook et al., 2021), they are very typical for hours in middle-income and developing countries, where gig work often serves as a primary, full-time source of income (Chen et al., 2024b; Adhvaryu et al., 2026).

exit hazard decreases quickly, so there is a sizable share of workers who stay much longer.

Appendix Table A.3 presents summary statistics separately for gig and team workers. There is less variability in outcomes within each work model separately.

2 Pre-AI Differences by DHH Status

Our analysis proceeds in two stages. First, in Section 2, we measure differences in outcomes between DHH and non-disabled workers prior to the AI tool. Second, in Section 3, we evaluate the impact of the AI tool on outcomes of DHH workers (versus the comparison group of non-disabled workers) using a DiD approach.

2.1 Empirical Strategy

To compare outcomes between DHH and non-disabled workers in the pre-AI period, we regress each outcome Y_{it} on a DHH indicator and controls:

$$Y_{it} = \alpha + \beta \cdot DHH_i + \mathbf{X}'_{it}\gamma + \varepsilon_{it}, \tag{2}$$

where i is a worker and t is a week; DHH_i is a dummy for being DHH; $\mathbf{X}_{it}$ are controls; and ε_{it} is an error. With the exception of heterogeneity analysis of AI tool effects by regional unemployment (where we block-bootstrap cluster by region), throughout the paper, standard errors are clustered by worker, as DHH status varies by worker.

Our controls $\mathbf{X}_{it}$ include year-week fixed effects (capturing common shocks during the outcome period), city-of-work fixed effects (roughly 300 cities), age fixed effects, gender, team status (gig vs. team), and registration year-month fixed effects (capturing cohort differences at entry). In our main specifications, we do not control for a driver's average distance between restaurant and customer—a measure of task difficulty—or for worker tenure, as these variables represent mechanisms by which DHH and non-disabled workers may differ.[23] In the Appendix, we show that controlling for travel distance has little effect on results. However, not controlling for tenure is substantively important: as we discuss below, worker efficiency and output improve with tenure, and DHH workers are less likely to quit, so higher tenure is a way by which DHH workers benefit themselves and the Platform.

The purpose of equation (2) is to provide descriptive evidence on differences between DHH and non-disabled workers. It is not to estimate the "causal effect" of being DHH.

[23] These variables can be argued to be "bad controls" if they are mechanisms by which the workers differ (Angrist and Pischke, 2009). Still, it can be interesting to address the question of whether DHH and non-disabled workers differ in efficiency at a given level of tenure or for particular routes.

In addition to overall differences, we also present outcome-tenure curves separately for DHH and non-disabled workers. This allows us to analyze how outcomes vary with tenure, as well as to understand how DHH differences vary by tenure, e.g., whether differences are persistent or whether there is convergence. To do this, we regress outcomes in levels (non-transformed) on interactions between DHH status and worker tenure month indicators, controlling for the same controls listed in equation (2). The excluded group is non-disabled workers in month 1, and we plot the relevant regression coefficients added to the excluded group mean, giving estimated average outcomes for each group at each tenure month. Outcome-tenure curves are similar with no controls.

Functional form of dependent variable. For each variable besides late deliveries, bad ratings, and attrition, we use an inverse-hyperbolic sine transform on y, so that coefficients approximate percentage differences. Since we restrict our sample to weeks with positive orders, most of our outcome variables do not have zeros, and we avoid the "log of zero"-like problem with inverse hyperbolic sine (Chen and Roth, 2023). In line with Chen and Roth (2023), we estimate non-transformed OLS models for late deliveries, bad ratings, and attrition, and divide by a pre-AI DHH mean to calculate percentage differences. Late deliveries is analyzed per 10,000 orders, bad ratings as a weekly count, and attrition is binary (so the non-transformed OLS model is a linear probability model).[24]

2.2 Results

Figure 1 visually summarizes gaps between DHH and non-disabled workers across work outcomes prior to the introduction of AI tool. Each bar reports the percentage difference associated with the DHH indicator estimating equation (2) with full controls. Our findings are generally similar without control variables, as can be seen in Figure A.1, which plots raw means over time for several outcomes for DHH and non-disabled workers.

Productivity: speed variables. Figure 1 shows that DHH workers take longer per delivery than non-disabled workers, though differences vary considerably by stage of order fulfillment. For the initial stage of order acceptance (i.e., from order initiation by the consumer to order assignment or acceptance by the worker), DHH workers take 3.4% longer than non-disabled workers. For the second stage of pickup (i.e., from order acceptance to

[24]We analyze bad ratings as counts per week rather than per 10,000 orders: bad ratings are rare (most driver-weeks have zero), and the per-order rate can produce outliers when a worker completes few orders in a week. Figure 7 shows bad ratings per 10,000 orders evolving with tenure, before and after the AI tool. For late deliveries, which are more common, per-10,000-orders is the natural metric and a standard industry measure of delivery performance. In both cases results are highly robust: bad ratings results are similar using the per-10,000-orders rate (Table A.7), and late delivery results are similar using counts (Table A.8).

picking up the order at the restaurant), there is no difference between DHH and non-disabled workers. DHH workers are actually 0.6% faster, but it is highly statistically insignificant. The largest gap occurs in the final stage, from leaving the restaurant to delivering the order to the customer: DHH workers take 4.5% longer, which is about 40 seconds more. Overall, DHH workers do 1.9% fewer orders per service hour than non-disabled workers.[25]

Figure 2 presents the tenure-productivity relationship for each of the three order stages. Time spent in each stage declines with tenure for both groups, consistent with learning by doing (Shaw and Lazear, 2008).[26] For the first stage, order acceptance, differences are present in the first three months of tenure, but the two lines converge starting in the fourth month. In contrast, for the final stage, customer delivery, the gap is more persistent. That is, not only is the DHH difference largest for delivery time, but it also takes much longer to converge, and the average remains higher for DHH workers throughout.

The time numbers on each stage of the order fulfillment process are averages across worker-orders. The Platform may be more concerned with more extreme events, like whether an order is substantially late. Figure 1 shows that DHH workers have 9% more late deliveries, defined as deliveries beyond the expected window, than non-disabled workers. This works out to DHH workers having an extra 33 late deliveries per 10,000 orders.

A natural explanation for DHH workers being slower is communication challenges with customers. Other possibilities could include communication challenges with restaurants or differences in general driver behavior (e.g., DHH workers driving slowly or more cautiously). However, communication challenges with customers directly explain why DHH drivers are slowest on the final stage of the order.

This mechanism dovetails with what we observe in our qualitative fieldwork, especially riding along drivers. On most rides, DHH workers seem to go as fast as non-disabled workers. However, there are occasional situations that require extensive communication with customers, leading to substantial delays. In China, delivery workers are generally expected to call the customer upon arrival, e.g., to coordinate the handoff or confirm delivery if the food is left in a delivery cabinet or at the front desk. Prior to the AI tool studied, DHH workers who couldn't speak had to rely on in-app messaging, which customers sometimes

[25] Orders per service hour is calculated as $\frac{60}{t_1+t_2+t_3}$, where the denominator is the sum of average minutes spent in each stage of an order. This measure doesn't use our hours data, which capture time available for work rather than time completing orders. Because drivers can handle multiple orders concurrently, the sum of average stage times can exceed actual clock time spent delivering orders, so orders per service hour should be interpreted as a processing-rate measure rather than literal orders per hour doing deliveries.

[26] This pattern is also consistent with differential attrition, which we discuss in Section 3.6. However, the main pattern of order acceptance, pickup, and delivery time decreasing with tenure is also present when we analyze outcomes in the first 12 months of workers who stay at least 12 months. This supports that apparent first-year learning by doing is not driven by differential attrition.

don't see, contributing to occasional but more pronounced delivery delays.

Productivity: quality. Our measure of work quality is bad customer ratings. DHH workers receive more bad ratings than non-disabled workers: 0.15 per week versus 0.10, a gap of 47% relative to the non-disabled mean. Expressed as a fraction of the DHH pre-AI mean—consistent with how we scale our DiD estimates in Section 3—this corresponds to 31%, as shown in Figure 1. These differences are similarly strong for both gig and team workers separately (Figure A.3). Although bad ratings are rare in absolute terms, these are sizable percentage differences between DHH and non-disabled workers in the pre-AI period. They reinforce the above discussion on mechanisms pointing to challenges DHH drivers face in communication with customers. Table A.7 provides additional context and robustness checks on bad ratings results.

Qualitative work supports that communication challenges with customers are a natural explanation for why DHH workers get more bad ratings. Customers may get frustrated if communication is difficult, e.g. if the worker can't place a call to locate the customer, the customer may misinterpret the lack of spoken communication as rudeness. DHH workers recounted situations where customers left bad ratings citing "not calling" or "not speaking clearly." Overall, results on speed and bad ratings support A2 in Section 1.1 of $z_D < z_N$.

The presence of bad ratings is a concern for the Platform. If customers are dissatisfied, they may order less frequently, and the Platform could get a bad reputation. The Platform itself noted this concern, which was one motivation for investing in calling technology.

Labor supply. Figure 1 shows that DHH workers work 4.9% more hours per week than non-disabled workers, conditional on controls, equivalent to an extra two hours per week. Additional hours could come from DHH workers working more days per week or more hours per day, and our data indicate both. Specifically, DHH workers work an extra 4.3% days per week than non-disabled workers, equivalent to an extra 0.2 days per week.

Beyond working more hours and days, Figure 1 shows that DHH workers are 41% less likely to quit the Platform than non-disabled workers, conditional on observables, in a linear probability model. To dig further into this overall difference, panel (a) of Figure 8 presents Kaplan-Meier survival curves for disabled and non-disabled workers separately. The raw gaps in survival emerge starkly within the first month and remain even after a year.

These are substantial differences in hours and driver retention. To benchmark magnitudes for retention, research across numerous firms shows that workers hired via employee referral are less likely to quit than workers hired through other means (Hoffman and Stanton, 2025), with differences around 10-15% for low-skill workers (Burks et al., 2015). Thus, the retention difference by DHH status is roughly 3-4 times that associated with being referred. Appendix Figure A.3 shows that the tendencies of DHH workers to work more hours, work more days,

and quit less are all present when one looks separately at gig and team workers.

Our conceptual framework implies that DHH workers should work more if differences in efficiency are small relative to differences in outside options. We believe this is exactly what is occurring in our setting, where Figure 1 indicates that differences in efficiency are small, at least for efficiency as defined in terms of overall speed. In our qualitative work, it was not uncommon for DHH workers to work 6-7 days a week on the Platform, and these workers also indicated having limited outside job opportunities. Our findings of DHH workers quitting less broadly aligns with Minor et al. (2018), who find that workers with criminal records are less likely to quit, also seemingly because of worse outside options.

Total output. Combining the two offsetting forces—DHH workers being slightly slower while working more hours—the net effect is that DHH workers deliver more orders per week than non-disabled workers. DHH workers do 5.3% more orders than non-disabled workers, equivalent to an extra 8 orders per week.

Panel (a) of Figure 6 shows orders as a function of worker tenure in the pre-AI period. The tenure-productivity curves in Figure 6 are obtained from the same specification used to construct Figure 2. As with the speed variables, performance increases with tenure. However, for total output, DHH and non-disabled workers have very similar numbers of orders at a given tenure level. The advantage of DHH workers therefore arises not from higher output conditional on tenure, but from their having longer tenure and lower quit rates. In other words, DHH workers tend to operate at higher tenure levels of the learning curve than non-disabled workers.

The output gap is largest in the team model. Among team workers, DHH workers complete 8.7% more orders per week than non-disabled workers, about 18 additional orders. Among gig drivers, the difference is 0.7% and statistically insignificant. A likely explanation is greater support in the team model, where DHH workers can seek help from station managers and teammates.

Income, wages, and profits. Figure 1 shows that, in line with doing more orders per week, DHH workers earn 5.4% more in weekly gross income relative to non-disabled workers.[27] This increase is driven entirely by greater labor supply rather than higher pay per unit of time: hourly wage rates for DHH workers are lower than those of non-disabled workers, as also seen in Figure 1.

Not accounting for bad rating penalties, disabled workers earn 9.4% lower hourly wages. When penalties from bad ratings are incorporated, the hourly wage gap widens slightly to 10.2%. These estimates are broadly in line with those in Forth and Jones (2025), who es-

[27] Weekly gross income is based on earnings received, and does not include penalties for bad ratings or costs from e-bike rental, maintenance, or charging. We account for bad rating penalties in the hourly wage.

timate a disability hourly pay gap of 6.8% within firm for a representative sample of firms in England and Wales. Overall, that DHH workers have more bad ratings only accounts for about one-tenth of the 10% wage gap. Most of the pay gap reflects DHH workers having similar income to non-disabled workers (in the sample with observed hours used to compute wages) despite working more hours. To further benchmark these patterns, Appendix Table A.1 reports DHH employment and wage gaps in the U.S. general labor market using CPS ASEC microdata from 2008–2023. The specification most analogous to our within-platform comparisons controls for occupation and industry fixed effects in addition to demographics, year, and state (column (4)), since our Platform estimates compare workers in the same job. This yields an hourly wage gap of about 4%, somewhat smaller than, but broadly consistent with, the wage gap we estimate on the Platform. DHH workers also face an employment gap of 25 percentage points relative to non-disabled workers (column (3) of Panel A).

Turning last to Platform profits, DHH workers generate 12% more profit per week than non-disabled workers. While DHH workers are slower and earn lower hourly wages, Platform profits depend on total orders rather than speed per se, as shown in equation (1). Thus, slower work reduces DHH workers' own hourly wages rather than Platform profits. The Platform also benefits from DHH workers' longer tenure: a smaller share of their worker-weeks occur in the first week, when onboarding costs are incurred.

Sources of profit differences. Panel A of Table 3 decomposes average pre-AI profit differences between DHH and non-disabled workers into its three components: orders, bad ratings, and onboarding costs. For each component, we multiply its estimated percentage difference between DHH and non-disabled workers by the component's average value in CNY, and report its share of the total CNY difference. The main source of DHH profitability is they do more orders per week, accounting for 85% of the profit gap. Getting more bad ratings works against DHH profitability, but bad ratings are rare. Higher DHH retention means the Platform incurs onboarding costs in a smaller share of weeks than for non-disabled workers, but this contribution to pre-AI profits is limited, as onboarding costs are modest.

Heterogeneity in pre-AI differences by DHH severity. Figure A.4 presents heterogeneity in pre-AI outcome differences according to disability severity. We split by DHH workers with profound HL (Level 1 in Chinese system; 54% of the sample) and DHH workers with non-profound HL (Levels 2-4). Profound HL workers tend to obtain worse outcomes than non-profound HL workers, though not all differences are statistically significant.

Starting with efficiency, profound HL workers get 48% more bad ratings than non-disabled workers, whereas non-profound HL DHH workers get 14% more bad ratings than non-disabled workers. There is no difference between non-profound HL workers and non-disabled workers in late deliveries, whereas profound HL workers have 16% more late deliveries. Turn-

ing to labor supply, the overall DHH hours advantage is driven by non-profound HL workers, who work 17% more hours per week than non-disabled workers. Profound HL workers work 12% fewer hours per week than non-disabled workers. Differences in total output, income, and Platform profits are likewise driven by non-profound HL workers.

Both groups earn lower hourly wages than non-disabled workers, but, interestingly, the wage gap is larger for non-profound HL workers (a 13% gap vs. a 7% gap). This reflects their substantially higher hours. Profound HL workers may concentrate work in more lucrative days or times, consistent with greater communication challenges.

3 Evaluating the AI-Enabled Communication Tool

3.1 Origins of the AI tool

The "Intelligent Outbound Call" (IOC) is a TTS calling feature on the app to help DHH workers communicate better with customers. Many DHH workers face challenges in speaking and being understood. Before the IOC feature was introduced, such DHH workers communicated with customers via text-based messages within the Platform's app. This posed challenges, as typed messages such as "Can you open the door for me?" or "I cannot find the place, please allow me an extra five minutes" often went unnoticed or unanswered by customers, resulting in delays, misunderstandings, and added stress for DHH workers.

Development of the tool. The tool was developed in-house at the Platform by around 3 software engineers during off time over several months. Both of the two major Chinese food delivery platforms emphasize the use of technology for good as part of corporate social responsibility, and engineers at the Platform are encouraged to spend time helping others. In developing the tool, the Platform emphasized helping DHH workers as part of its broader corporate social responsibility efforts. At the same time, the tool had the potential to increase profits (e.g., by increasing productivity, reducing costly delivery errors).

How the tool works. Under the IOC feature, DHH workers make natural-sounding, automated calls to customers. The IOC provides workers with a set of default communication prompts, as well as the option to create and customize prompts tailored to specific situations. Appendix Figure C.1 illustrates both the default prompts offered by the Platform and examples of customized prompts used by a DHH worker.[28] After the DHH worker

[28] From our observations, prompts often fall into two main categories: (1) difficulties locating the customer, such as being denied entry by security or requiring a buzzer code; and (2) notifications that the order has been placed at a specific location, like a food cabinet. One worker we worked with also made polite requests to avoid bad ratings.

selects a prompt, the app initiates an automated voice call to the customer. Each call begins with an introductory message stating that the caller is DHH followed by a specific message conveying the worker's issue or request, e.g., "Your food will be placed in a locker". The tool is strictly text-to-speech instead of speech-to-text, i.e., customers can respond in the app instead of on the call. The feature is available only to DHH workers; neither non-disabled workers nor workers with other disabilities can use it.

How the tool uses AI, and why AI matters. The IOC is built on AI-based TTS. The tool's effectiveness rests on a behavioral requirement: customers must answer an automated call, listen to the message, and act on it (e.g., open the door, or note where the food was left). What AI changes is whether customers are willing to listen and to act. Pre-AI TTS systems—which generated speech from hand-coded rules or by splicing together fragments of recorded speech—were intelligible but unmistakably robotic, lacking natural intonation, rhythm, and pauses. Robotic calls are widely associated with telemarketing and fraud, a salient concern in China, and customers routinely ignore them or hang up. Modern TTS instead uses deep learning: neural networks trained on large corpora of human speech learn to predict how natural speech should sound (Agrawal et al., 2018), producing voices nearly indistinguishable from human recordings in blind listening tests (Shen et al., 2018; Tan et al., 2021). Research suggests that AI voices are more pleasing to customers than traditional robocalls (Shah et al., 2023; Arcot, 2025). A TTS accommodation for DHH drivers was thus technically conceivable before AI, but became viable—natural enough for customers to listen to and engage with—only once AI-based synthesis became widely available. Appendix D provides further background on the evolution of TTS technology and the role of AI.

The advantages of the IOC tool include: first, enhanced customer responsiveness, as customers are more responsive to phone calls compared to app notifications, and voice communication is perceived by the Platform as "warmer" than text, ensuring prompt and higher-quality interactions; second, improved efficiency, since DHH workers can quickly address common issues through pre-set prompts without repeatedly typing messages, saving valuable time especially during peak delivery hours or adverse weather; and third, personalization and flexibility, as workers can customize messages to communicate unique situations.

AI implementation and usage. The Platform announced the AI-based tool to DHH drivers through multiple in-app notifications, making awareness of the tool widespread. Our administrative data do not record individual-level usage. However, the Platform informs us that, to their knowledge, all DHH workers use the AI tool and that the tool is popular. This is validated by our qualitative evidence: all 9 DHH workers whom we interviewed or worked alongside reported using the AI tool (details in Appendix H).

3.2 Overall Effects of the AI Tool

Estimation strategy and graphical presentation of results. To assess the impact of the AI tool, we estimate a DiD model that compares outcomes for DHH workers before and after the introduction of the AI tool, relative to changes over the same period for non-disabled workers, who do not receive access to the tool:

$$Y_{it} = \alpha + \gamma \cdot DHH_i \times Post_t + \lambda_i + \delta_t + \varepsilon_{it}. \tag{3}$$

where $Post_t$ is a dummy equal to 1 if the current week is on or after the introduction of the AI tool at time t_0 in mid 2022. γ captures the differential change in outcomes for DHH workers after the introduction of the AI tool relative to non-disabled workers. Under the parallel trends assumption, γ identifies the average treatment effect on DHH workers. Worker effects λ_i capture fixed differences across workers while time effects δ_t capture common shocks. We present our baseline DiD results without time-varying controls, and show robustness to time-varying controls in Section 3.6. Beyond overall DiD results, we (1) show event study results analyzing how outcomes change month-by-month before and after the AI tool and (2) compare outcome-tenure curves before and after the AI tool.

For analysis of employee attrition, we do not include worker fixed effects—as workers can only quit once—and instead use the same controls as for equation (2).

The identifying assumption underlying equation (3) is that, absent the introduction of the AI tool, outcomes for DHH and non-disabled workers would have evolved similarly over time, conditional on controls. This could be violated if there are other policies or procedures introduced for DHH workers around the event time window (for the event study analysis) or in the sample period in general (for overall effects). We examine concerns about differential trends using event studies.

Figure 3 summarizes DiD estimates of the AI tool's impact across all outcomes. Each bar shows an estimate of γ from equation (3) and represents the percentage change in the outcome for DHH workers. We analyze the same outcome as in Figure 1.

Event study setup. Our event study uses a window of 6 months before and 10 months after the introduction of the AI tool. We use all available data and bin observations outside this window into the endpoints—i.e., all months more than 6 months before or more than 10 months after t_0 are pooled into boundary coefficients (Schmidheiny and Siegloch, 2023). Because our analysis is based on new entrants to the Platform, the sample available at each pre-period month shrinks as we go further from t_0, mechanically widening confidence intervals in the early pre-period. We thus focus on this window to ensure a reasonably high degree of statistical power for all event-time months shown. Section 3.6 discusses longer-run

effects beyond the 10-month post-period and shows that results are highly robust.

Productivity variables: speed. Figure 3 shows that the AI tool led to an improvement in DHH worker speed, with orders per service hour increasing by 1.6%. Improvements are concentrated in the final stage of order fulfillment, delivery from restaurant to customer, where there was the largest initial disadvantage for DHH workers. Delivery time decreases by 3.5% for DHH workers, closing roughly 3/4 of the pre-AI gap of 4.5%. There is also a small 0.8% reduction in order acceptance time and no meaningful effect on pickup time.[29]

In terms of magnitudes, the reduction in delivery time of 3.5% is equivalent to moving about 6 months down the DHH delivery time tenure curve (panel (c) of Figure 2).

Panel (b) of Figure 4 shows event study estimates of the impact of the AI tool on delivery time (i.e., the final stage). There are not systematic pre-trends, consistent with the identifying assumption in DiD. In the first couple months of the AI tool, delivery time may actually increase (though not significantly so), consistent with drivers taking a little time to use the AI tool effectively (Brynjolfsson et al., 2021). Effects on delivery time gradually turn negative, with effects seeming to gradually increase in magnitude.

The impact of the AI tool on substantially late deliveries is even larger, with a 11% reduction for DHH workers. The IOC entirely closes the late delivery gap between DHH and non-disabled workers. This is seen in raw means in panels (e) and (f) of Figure A.1. As whether an order is classified as substantially late can hinge on a minute or two, modest improvements in delivery speed—and reductions in communication-related stoppages—can generate larger effects on late delivery rates.

As seen in Appendix Figure A.5, reductions in delivery time are pronounced and economically similar for both gig workers and team workers. Reductions in late deliveries are pronounced for team workers but not statistically significant for gig workers, possibly because gig workers take many more orders after the IOC and are therefore more likely to experience late deliveries, for example when carrying more orders simultaneously during peak hours, which increases the likelihood of delivering late.

Productivity variables: quality. Figure 3 shows the AI tool has a large effect on bad ratings, reducing them by 22%. This closes roughly 2/3 of the pre-AI gap of 31%.[30]

As seen in Figure 7—which plots bad ratings over tenure—prior to the AI tool, there are pronounced gaps between DHH and non-disabled workers that persist for most of a worker's

[29]The pre-AI DHH order acceptance time gap is 3.4%, so the AI tool closes about 20–25% of it, driven by gig workers. Since gig workers (unlike team workers) have discretion over which orders to accept, the AI tool may make them more willing to accept offers confidently.

[30]The sizable negative impact of the AI tool on bad ratings is robust across many methods, which is important given that bad ratings are rare. For example, the AI tool reduces the chance of receiving any bad ratings in a week by around 20%.

first year of tenure on the Platform. After the AI tool, there is still an initial gap in the first months of tenure, but the gap is much smaller, and the gap dissipates by around the 4th month of tenure. In terms of magnitudes, the impact of the AI tool is equivalent to moving about 3 months down the tenure curve in bad ratings.

Panel (c) of Figure 4 shows effects on bad ratings over time in an event study. As for delivery time (in panel (b)), there is no pre-trend. Effects appear to grow somewhat over time, potentially reflecting adjustment costs in learning about the AI tool, though confidence intervals are larger.

Overall, the productivity results suggest that the AI tool alleviates frictions at the delivery stage: pickup time is unchanged, while delivery is faster and bad ratings decrease. Qualitative interviews (Appendix H.1) support this interpretation, with DHH workers reporting fewer misunderstandings at delivery and greater comfort in interacting with customers.

Labor supply. Pre-AI, DHH workers have greater hours and retention than non-disabled workers, and Figure 3 shows that the AI tool further accentuates these differences. The AI tool increases DHH weekly hours online by 10.6%, roughly tripling the pre-AI hours gap of 4.9%. The AI tool increases work days per week by 3.5%, which is significant, but most of the increase in hours comes from an increase in hours per day. The AI tool further reduces DHH attrition from the Platform by 21%, further expanding the pre-AI gap of 41% by half. As seen in the comparison of worker survival curves in Figure 8, the AI tool noticeably increases the difference between the DHH and non-disabled survival curves. Effects on labor supply and retention are present in both gig and team workers, as seen in Figure A.5. Overall survival is somewhat lower for both groups in the post-period, possibly reflecting continued growth in China's delivery workforce, but the decline is smaller for DHH workers.

These effects are in line with the conceptual framework, which predicts that drivers increase labor supply in response to an improvement in efficiency. While our intentionally simple framework does not account for non-pecuniary work amenities, it is also possible that the AI tool increases labor supply "directly" by making Platform work more pleasant instead of simply by helping DHH workers make more money. Consistent with this, DHH workers we interviewed reported that the AI tool reduced misunderstandings with customers and lowered stress and anxiety about bad ratings and late deliveries (Appendix H.1). While bad ratings and late deliveries carry monetary consequences, which our wage and income measures capture, reduced worry about these outcomes is a distinct non-pecuniary benefit. We take up this possibility in Section 3.4 when analyzing the worker welfare impacts of the AI tool.

DHH worker entry. In addition to the hours and retention of workers on the Platform, we also examine the impact of the AI tool on entry of DHH workers onto the Platform, using

a group-by-week design (as entry cannot be analyzed in our worker-level panel). Appendix E.2 shows that the AI tool increases entry of DHH workers by 27.5%. Relative to the tool's effect on weekly income, this implies an entry elasticity of roughly 2.5, broadly in line with quasi-experimental estimates of firm-level labor supply and recruitment elasticities (see the Appendix for benchmarking details and caveats).

Total output. Figure 3 shows that the AI tool increases DHH weekly orders by 11%. Pre-AI, DHH workers do 5% more orders than non-disabled workers, so the AI tool extends this initial advantage. Effects are present for team workers (+8%), but are especially large for gig workers (+15%), though the difference by work model is not significant ($p = 0.29$). These results suggest that the AI tool may be especially helpful in the more independent gig mode. One possibility is that team DHH workers already had some support or workarounds in place in the pre-AI period, whereas gig DHH workers, with less support, benefit substantially.

As seen in panel (a) of Figure 4, effects are present directly after the AI tool is introduced, and seem to increase over time, consistent with effects on delivery speed and hours building over time, with panel (a) of Figure 5 showing the event study for hours.[31]

Income, wages, and profits. Following the introduction of IOC, weekly gross income earned by DHH workers increases by 11% relative to non-disabled workers. Hourly wages also rise for DHH workers, though by a more modest magnitude of 3.6%. These estimates imply that roughly one-third of the increase in weekly income is driven by higher hourly pay, with the remainder attributable to increased hours worked. Given a pre-AI disability pay gap of 10.2%, these gains imply that the AI tool closes around one-third of the wage gap between DHH and non-disabled workers. Wage gains tend to be concentrated among team workers; gig workers see smaller hourly wage gains as their large hours increases offset their income gains (Figure A.6). We are not aware of other workplace interventions or technologies that generate effects of comparable magnitude on disability wage differentials.[32]

From the Platform's perspective, the profitability of DHH workers rises substantially following the introduction of IOC. Net weekly profit generated per DHH worker increases by 17%, a sizable and statistically significant effect.

Sources of profit gains. Panel B of Table 3 applies the same approach—multiplying

[31] Unlike for hours, effects on total output are not statistically significant within a couple months, but take about 6 months to become statistically different from 0. This reflects, in part, that effects on delivery speed are also a little slower to manifest, but we believe mainly, that there is more noise in orders completed relative to hours, which are a pure choice variable. The event study coefficients on orders completed are estimated to be positive relatively quickly, even if they are not statistically significant.

[32] Evaluating the expansion of work-from-home arrangements during the COVID-19 pandemic, Bloom et al. (2026) show that disabled workers experienced larger wage declines in occupations with greater work-from-home exposure, consistent with increased labor supply exerting downward pressure on wages. Gentry (2025) finds that gaining access to Uber closes less than 2% of the disability pay gap.

each component's estimated AI-driven percentage change by its average CNY value—to decompose the source of the gains in profitability for DHH workers. The main source of the profit gains is making DHH workers do more orders as opposed to avoiding onboarding costs or getting fewer bad ratings. Order volume alone accounts for 93% of the profit impact, with reduced attrition and fewer bad ratings contributing only 5% and 2%, respectively.[33]

3.3 Benefits and Costs to the Platform

While Figure 3 shows that the profitability of DHH workers increases, this does not account for the costs of the AI tool itself. The Platform's average profit per order is confidential and varies with work type and order volume, which we incorporate directly in our worker-week profit calculations based on equation (1). For the purpose of translating our estimates into aggregate dollar magnitudes, we use a benchmark value of 7 CNY (≈ $1 USD) per order, based on conversations with the Platform. Under this assumption, the 17% improvement in DHH worker profitability implies roughly $163k USD per month in additional platform profit.[34] The Platform's cost of the tool consists entirely of engineer time during development: there are no per-use licensing or subscription costs, and the tool was built by three engineers working off-hours over several months. We conservatively assume that each engineer spent 10 hours per week for three months, implying a total time investment of 390 hours—less than 0.2 annual FTE—and thus no more than about $20,000 USD in labor cost. Under a simple comparison of benefits and costs, the AI tool pays for itself in less than a month: specifically, in 20/163 months, or about 4 days.

In summing up the Platform benefits as above, we implicitly assume that additional orders completed by DHH workers do not displace orders from non-DHH workers, corresponding to an infinitely elastic labor demand curve at the Platform level. To relax this assumption, suppose that a fraction δ of additional orders completed by DHH workers are displaced from non-DHH workers. In this case, the net increase in platform profits is scaled by $(1-\delta)$. Even under substantial displacement—e.g., $\delta = 0.5$, implying that half of additional DHH orders come at the expense of non-DHH workers—the tool would still generate

[33]The onboarding component (5%) captures only the direct savings from paying the onboarding bonus less often. It does not capture the contribution of retention to order volume: the AI-induced reduction in attrition also moves workers to higher-tenure points on the learning curve where they complete more orders per week. Table 3 thus understates the full profit contribution of retention. To analyze how much of the order effect is due to improved retention, one strategy is to compare the order effect in Figure 3 vs. the effect in a model with inverse probability weighting (IPW). Using IPW shrinks the effect on orders from 10.5% to 9.7% (Figure A.15), suggesting that increased retention drives about 1/10 of the effect on orders.

[34]There are on average about 1,600 DHH workers on the Platform per week, including both new and experienced workers. Assuming 600 orders per month, we have 1,600*600*0.17*$1 ≈ $163k per month.

roughly $82k USD per month in additional profits and pay for itself in about 1 week. Even with $\delta = 0.75$, the tool would still pay for itself in about 2 weeks.

3.4 Worker Welfare from the AI Tool

We use a sufficient statistics approach to estimate by how much the AI tool affects DHH welfare. Starting by ignoring amenities, the increase in DHH workers' hourly wage both (i) raises the value of their work on the Platform and (ii) allows substitution away from their outside option. As detailed in Appendix E.4, by the Envelope Theorem, the change in DHH welfare as a fraction of pre-AI earnings E_0 is:

$$\frac{\Delta V}{E_0} = \underbrace{\left(e^{\hat{\gamma}_w} - 1\right)}_{\text{inframarginal wage gain}} + \underbrace{\tfrac{1}{2}\,\hat{\gamma}_h\,\hat{\gamma}_w}_{\text{re-optimization gain}} . \tag{4}$$

where $\hat{\gamma}_w$ and $\hat{\gamma}_h$ denote the DiD effects on hourly wages and on weekly hours, respectively. The first term captures that hours the worker was already supplying now earn a higher wage. The second term is a Harberger-style triangle: at the pre-AI optimum, the marginal off-Platform hour was worth exactly the pre-AI Platform wage, so hours newly reallocated to the Platform earn only the wedge between the new wage and their (rising) opportunity cost. Under this calculation, we find that DHH welfare increases by 3.9% relative to pre-AI Platform earnings, primarily driven by (i), as shown in Panel A of Table 4. We can also additionally account for potential changes in the amenity value of working on the Platform. The main idea is to compare the implied labor supply elasticity to the Platform against an assumed "true elasticity," with the differential assumed to reflect changes in amenities. Doing so, we estimate that DHH welfare increases by 5.7% to 11.7% of pre-AI Platform earnings (Panel B). These gains strike us as substantial.

3.5 Heterogeneous Treatment Effects

A central question regarding the deployment of AI in the workplace is how treatment effects vary across different workforce segments. For example, some evidence indicates that generative AI may have larger effects on novice or lower-skilled workers (Brynjolfsson et al., 2025; Cui et al., 2025; Noy and Zhang, 2023). We focus our analysis on four key dimensions: pre-AI tenure, pre-AI productivity, disability severity, and local labor market conditions. For each dimension of heterogeneity, we present estimates separately for each half of the sample split at the median, and test for differences using interaction term models similar to those in equation (3). For late deliveries, bad ratings, and attrition, we divide by a pre-AI DHH

mean within each subgroup to calculate percentage differences.

Pre-AI tenure and productivity. We find no consistent evidence of heterogeneous treatment effects based on pre-AI tenure (Figure A.8) or pre-AI productivity (Figure A.9), with pre-AI productivity measured in orders per week. Excluding pickup time, the outcome where there was no effect of the AI tool, there is significant heterogeneity by pre-AI tenure or pre-AI productivity for only 3 of 26 outcomes. That is, the AI tool's effectiveness does not systematically vary with a driver's tenure or average productivity.

Disability severity. We observe substantial heterogeneity in effects based on severity of disability. Figure 9 plots the coefficients for workers with profound hearing loss (meaning Level 1, hearing loss of more than 90 decibels) versus those with non-profound hearing loss (Levels 2–4), whom we refer to as profound HL and non-profound HL workers.[35] Profound HL workers, who are 54% of DHH workers in our sample, cannot hear even shouted speech.

We observe substantial heterogeneity in effects based on severity of disability, usually in the direction that profound HL workers benefit more. Orders per week increase by 6% for non-profound HL workers, but by 16% for profoundly impaired, with a p-value of 0.08 on the difference. Bad ratings decrease by 7% for non-profound HL workers, but decrease by 40% for profound HL workers ($p = 0.02$). As seen in Figure A.4, pre-AI, the bad rating gap is largest for profound HL workers (and profound HL workers do 8% fewer orders than non-profound HL workers), so the AI tool also helps close the pre-AI gap by disability severity within DHH workers. This suggests that the AI tool is most effective at reducing communication frictions for those who face the steepest barriers.[36]

Relative to non-profound HL workers, profound HL workers also see substantially greater gains in income ($p = 0.08$) and wages ($p = 0.03$). While statistical significance is sometimes borderline on differences, the differences are economically quite sizable, e.g., wages go up by around 1.3% for non-profound but by 7% for profound HL workers.

For attrition, profound and non-profound HL drivers show similar effects. While Implication 5 predicts profound HL workers to benefit more, one explanation for similar effects is that profound HL workers have fewer outside work options regardless of the tool, whereas the AI tool effectively tips the scales for non-profound HL who might otherwise have exited.

We split by profound and non-profound HL both because this is roughly a 50/50 split

[35] Unlike our other heterogeneity dimensions, severity is defined only within DHH workers, so an indicator Profound_i for profound HL is absorbed by worker fixed effects λ_i and the standard four-term interaction specification reduces to two terms: $DHH_i \times Post_t \times \text{Profound}_i$ and $DHH_i \times Post_t$. The worker fixed effects also address the standard concern that Profound_i may be correlated with unobservables.

[36] Interestingly, non-profound HL workers exhibit larger improvements in delivery time and orders per service hour. Profound HL workers' gains instead appear in total output, as they complete more orders per week despite similar increases in hours online.

(and good for statistical power) and because profound HL can be qualitatively different from non-profound HL. However, results are broadly similar when we estimate dosage-response heterogeneous treatment effects (Figure A.10). Heterogeneity by severity is also robust to controlling for interactions of $DHH_i \times Post_t$ with age and gender, thus indicating that the heterogeneity is not driven by demographics correlated with severity (Figure A.11).

Local labor market conditions. Figure A.12 examines heterogeneity in AI effects by local labor market conditions. Exploiting that our Platform operates nationwide, we split China's 31 provincial-level administrative regions into high and low unemployment regions, where the 2021 average unemployment rates were 3.6% and 2.5% respectively.

On the 13 variables where we detect significant positive AI effects (i.e., excluding pickup time), in 7 of 13 variables, we observe a significantly more beneficial effect in low unemployment rate regions. The AI tool has no effect on late deliveries in high unemployment regions, while it reduces DHH late deliveries by 14% in low unemployment regions. Likewise, there is no effect of the AI tool on DHH hours online, work days, or orders in high unemployment regions, but sizable improvements in all these dimensions in low unemployment regions. DHH driver income increases by 17% in low unemployment regions, but is not significantly affected in high unemployment regions. Heterogeneity by regional unemployment is robust to controlling for interactions of $DHH_i \times Post_t$ with age and gender (Figure A.13).

3.6 Threats to Identification and Robustness Checks

3.6.1 Threats to Identification

Existence of concurrent policy changes. Our event study graphs show no clear pre-trends prior to the AI tool's introduction. A threat to identification would be if the Platform made other policy changes around the same time as the AI tool that differentially affected DHH workers. According to the Platform, there were no changes made close to the introduction of the AI tool, and in fact, no meaningful other changes for DHH workers over the sample period of 2021-2023.

Several of our quantitative results are strongly consistent with this. Suppose there was another policy launched by the Platform (or another entity like the Chinese government) to improve the productivity of DHH workers. Such a policy would likely show effects across all stages of order fulfillment: order acceptance, pickup, and delivery. In contrast, the effects of the AI tool are concentrated in the delivery stage—the only stage for which the tool is designed—where communication with customers is required and where we observe the largest improvements in speed. We do observe a smaller improvement in order acceptance, which reduces the pre-AI gap by roughly 1/4; this is consistent with DHH workers being more

willing or confident to accept orders when customer communication is easier. There is no effect on pickup, consistent with the tool not being used for communication with restaurants.

In addition, the observed heterogeneity by disability severity speaks to the plausibility of the identification assumption. If there was an alternative Platform intervention that benefited DHH workers, it would not necessarily bestow greater benefit on profoundly impaired DHH drivers relative to non-profoundly impaired DHH drivers.

Improvements to DHH workers or declines for non-disabled workers? Could our DiD estimates be driven by non-disabled outcomes getting worse instead of DHH outcomes getting better? Several pieces of evidence speak against this. First, non-disabled workers getting worse would not generate the observed heterogeneity by disability severity, with larger AI benefits on more severely disabled workers. Second, raw means of various outcomes (Figure A.1) suggest substantial improvements for DHH workers in levels.

One particular form of non-disabled workers getting worse would be if DHH workers take delivery opportunities away from non-disabled workers. This is unlikely to be quantitatively important for non-disabled worker outcomes because the share of DHH workers on the Platform is small (i.e., less than 1% of workers on the Platform).

Missing data on hours. As noted above, we have significant missing data on hours. Figure A.2 plots the missing rate by calendar time. Hours is missing for all workers before week t_0-46 where t_0 is the week of AI implementation. In t_0-46, hours become available for gig workers, and then become available for team workers in week t_0-4 of 2022. Missingness rates are similar between DHH and non-disabled workers and appear to reflect Platform data-keeping practices rather than endogenous selection.[37]

Because of the missing hours data, our estimates regarding hours-based outcomes, namely, hours online and wages, rely on a different sample than other outcomes.[38] To ensure that our constellation of results is not being driven by different samples, we redo our main results from Figure 3 while restricting the sample to observations where hours are non-missing. As seen in Figure A.14, our main conclusions are unchanged in this sample.[39] In the restricted sample, the AI tool drives significant improvements: orders per week increase by 13%, weekly income rises by 13%, and delivery time decreases by 2.6%.

Accounting for regional trends. DHH workers are more likely to live in Eastern

[37] Among team workers, hours missingness is extremely similar for DHH and non-disabled workers, whereas among gig workers, it is slightly higher for DHH workers than non-disabled workers (Figure A.2, panel (b)). However, as seen in Figure A.5, our main DiD results are robust when restricted to team workers, and we cannot reject that AI effects on hours and wages are the same for gig and team workers.

[38] Hours per service hour relies on the data on time spent on different parts of the production process.

[39] In the restricted sample, the AI tool has a negative effect on late deliveries and bad ratings, but the effect is no longer statistically significant, reflecting the smaller sample with hours data.

China and big cities than non-disabled workers. If regions in Eastern China experience different shocks than other areas (e.g., different lockdown policies), this could potentially confound our analyses. To address this, we re-do our analyses while adding region-month fixed effects (where we have 31 regions and 3 years of data). Our results are highly robust to this, indicating that such trends do not drive our results (Appendix Figure A.14).

Matching DiD. Beyond regions, DHH and non-disabled workers vary slightly in terms of observable characteristics. While our event studies strongly support the parallel trends assumption, for robustness (exploiting our large sample of non-disabled workers), we perform matching where we match each DHH worker to a single non-disabled worker, matching exactly on entry year, work model (gig or team), city, 5-year age bin, and gender. Using the resulting sample of 2,354 DHH and 2,354 non-disabled workers, we re-do our DiD in the matched sample. Our results are also highly robust to this (Appendix Figure A.14).

3.6.2 Alternative Mechanisms and Interpretations

Differential attrition. Since the AI tool increases DHH worker retention, a concern is that this could "bias" our estimates of the effect of the AI tool on DHH non-attrition outcomes like productivity variables, hours, and wages. Given that productivity increases with tenure (see Figures 2, 6, and 7), if the AI tool boosts DHH worker retention, then DHH workers will be more likely to be in higher-productivity parts of the learning-by-doing curve.

If we take the perspective of the Platform, our baseline DiD estimates should not be "corrected" for possible differential attrition, as it is an important *mechanism* by which the AI tool affects outcomes. The Platform cares about the overall effect of the AI tool on outcomes, not the impact of the AI tool on particular workers.

At the same time, we are still interested in how the AI tool affects outcomes separate from affecting retention, and such estimates might be more relevant for DHH workers themselves (i.e., how much will the AI tool affect my productivity). To shed light on this, and "shut down" differential attrition, we pursue two approaches. First, we perform our DiD analysis restricting the sample to a panel of "stayer" workers, who are observed working at least one week both before and after the introduction of the AI tool.[40] As shown in Figure A.15, the results from this panel are quantitatively and qualitatively consistent with our baseline findings. We continue to observe significant improvements across key metrics: total weekly orders rise by 10%, weekly income increases by 10%, and bad ratings decrease by 20%. These

[40]We cannot pursue a fully balanced panel approach over 2021-2023, as our sample consists exclusively of workers joining the Platform. A fully balanced panel would require analyzing only the DHH and non-disabled workers who join the Platform in January 2021 and then stay throughout the whole data period, and such a sample is too small for reasonable analysis.

strong effects are similar to our baseline estimates, indicating that differential attrition is not a central driver of the DiD estimates. Rather, the AI tool increases worker productivity directly in addition to increasing retention.

Second, we use inverse probability weighting (Wooldridge, 2002), which proceeds in two stages. In the logit first stage, we predict a worker's probability of remaining on the Platform each week as a function of time-invariant worker characteristics. Second, we re-do our DiD estimates weighting by the inverse of the predicted probability of remaining each week. Our non-attrition variables all show sizable effects from the AI tool under inverse probability weighting (Figure A.15, bottom panel). This supports that the AI tool increases productivity directly.

The role of late deliveries and bad ratings in load assignment. As noted in Section 1, late deliveries and bad ratings can affect the set of orders workers are assigned or able to select. Beyond direct productivity improvements, an alternative mechanism is that the Platform algorithm assigns treated workers easier or higher-quality orders as their performance metrics improve. In that case, observed gains in delivery speed, orders, hours, or income would reflect task reallocation rather than improved customer communication. This is a question about mechanisms rather than identification.

We address this concern in two ways. First, we re-estimate our main DiD specification while controlling for task difficulty, proxied by the distance between the restaurant and customer. As shown in Figure A.14, the estimated effects are very similar to the baseline results, indicating that changes in trip distance do not explain the observed DiD gains. Second, we conduct a mediation analysis controlling for late deliveries and bad ratings (Appendix E.3). The estimated effects of the AI tool on orders are largely unchanged, suggesting that reductions in late deliveries and bad ratings account for little of the overall increase in orders.

Overall, the results suggest that reductions in late deliveries or bad ratings, or the resulting reallocation of tasks, are unlikely to drive our other key DiD findings.

Salience or sympathy. Could it be that AI effects reflect salience or sympathy of hearing impairment as opposed to improved communication? It is highly unlikely that salience and sympathy could explain the full constellation of results. Pre-AI, customers were sent a notification that their driver was DHH, though more customers would be aware of this after a call. While sympathy could make customers less likely to give bad ratings, it is less clear how this reduces delivery times. Moreover, there is no reason why salience or sympathy would generate larger effects for DHH workers with profound HL relative to non-profound HL, as customers were not notified of drivers' level of hearing loss. This interpretation is also consistent with our qualitative observations: during a coauthor's three days riding alongside DHH workers and over a hundred orders, no customer appeared to

make any accommodation for the driver being DHH — for example, no one opened a door preemptively or waited downstairs after receiving an AI call.

Reference dependence / income targeting. Could the pre-AI difference simply reflect that DHH workers target income, doing more hours to make up for lower efficiency? This is inconsistent with our results by DHH severity, where profound HL workers have both lower efficiency and fewer hours than non-profound HL workers.

3.6.3 Long-run Effects and External Validity

Longer-run persistence of effects. The other major food delivery platform in China introduced a similar AI tool for DHH workers later in the same year, after our Platform. Because the market is largely dominated by these two platforms, the period in which only our Platform offers the tool raises the possibility of cross-platform spillovers (e.g., worker or customer reallocation). To address this, we re-estimate our DiD results excluding the period in which the AI tool is available on our Platform but not on the other platform. The estimates are very similar (Figure A.16). This suggests that our results are unlikely to be driven by the temporary presence of the tool on only one platform.[41] The subsequent adoption of a similar tool by the other platform is also consistent with the tool being broadly valuable for DHH workers.

Beyond considering adoption of the other Platform, we also assess whether effects persist further into the future. While our main event study graphs cover 10 months after the AI tool, we extend them to 16 months, pushing toward the edge of our data. As seen in Figure A.7, effects continue to grow and remain significant and robust.

Accounting for the transition period. Like many new technologies, usage of the AI tool was not instantaneous at 100% usage. Rather, according to the Platform, there was a transition period of a few months where DHH drivers learned about the new AI tool and how to best use it. This is supported by the event study graphs, which often show impacts that are more muted in the first few months, but that increase over time. Since we are primarily interested in the longer-run effects of the AI tool, we also re-do our DiD analyses but excluding the first 4 months immediately following the policy change. The results, presented in Figure A.16, confirm that our main findings are robust to this exclusion and often become larger. The reduction in bad ratings intensifies from 22% to 25%, the increase in output per hour grows from 1.6% to 2.1%, the increase in weekly orders grows from 11% to 13%, and the increase in income grows from 11% to 13%. Thus, our baseline

[41] Our DiD results are also qualitatively robust restricting to the months where only our Platform uses an AI tool for DHH workers.

DiD estimates, which include the first several months of the AI tool being available, may represent a more conservative lower bound of the tool's longer-run effects.

Could COVID outbreaks or lockdown policies limit external validity or threaten identification? Chinese cities instituted lockdowns at different times at various periods in 2021 and 2022. As shown in Figure A.16, our DiD results are robust to excluding worker-week observations from regions-weeks affected by COVID outbreaks. Results are also similar if we control for whether a COVID outbreak is occurring instead of excluding those periods.

Would the AI tool have equally benefited non-disabled workers? The AI tool is only introduced to DHH workers, raising the question of whether non-disabled workers would benefit as much. The heterogeneity results by disability severity suggest this is unlikely. DHH workers with non-profound HL benefit substantially less than those with profound HL, suggesting that non-disabled workers might also benefit less, or possibly not at all. The Platform's decision not to roll out the AI tool to non-disabled workers is also consistent with the view that benefits for non-disabled workers would be smaller or negligible.

3.7 What the AI Tool Does Not Address

Although the AI tool substantially improves communication between DHH workers and customers, some workplace outcome gaps remain. Qualitative evidence helps clarify the types of situations that allow such gaps to persist. Our qualitative data were collected in summer 2024, around two years after the AI tool was introduced. Consistent with what the Platform told us, in the workers we interviewed or worked alongside (7 interviewed and 3 accompanied during deliveries, with one worker appearing in both), all 9 of them used the AI tool. Thus, the remaining gaps do not reflect limited adoption of the tool.

Instead, the evidence points to communication frictions that the AI tool does not fully resolve. These include difficulty locating customers' exact locations, locations with complex layouts (e.g., buildings with multiple entrances), and customer instructions that remain ambiguous without verbal clarification. A coauthor working alongside a DHH worker experienced the following, as recorded in his field notes, where "I" denotes the coauthor:

> The [DHH] worker followed the customer's instructions to contact the apartment complex management via the intercom. After the system connected, no one spoke on the other side. Although I heard the door unlocking sound, the door did not automatically pop open, and there was no visible indication that it had unlocked. The worker continued waiting by the intercom until I eventually walked over and pulled the door open.

A second situation arose while the coauthor was accompanying the same DHH worker, illustrating how written instructions may remain unclear in the absence of verbal communication:

> One order was delivered to an urban village [chengzhongcun in Mandarin], but the address was inaccurate. Upon arriving, the worker used the intelligent outbound call to contact the customer, who messaged back, "I'm upstairs." However, it was unclear what the customer meant by "upstairs." I then heard the customer shouting from inside a factory behind us, but by that time, the worker had already started riding in the opposite direction. I stopped the worker, informing him that the customer was inside the factory, and we managed to find the customer. In such situations, a worker with a hearing disability could waste significant time locating the exact address. Later that afternoon, we encountered a similar issue with an order delivered to a construction site.

The AI tool helps initiate contact when customers might otherwise not notice or respond to messages. However, ambiguity in customers' directions and physical delivery environments can persist when clarification would normally rely on real-time verbal communication.

4 Conclusion

By 2050, roughly 10% of the world's population will have disabling hearing loss (World Health Organization, 2025). While DHH workers face intense labor market challenges including lower earnings and employment, research has often struggled to identify policies or tools that can benefit DHH workers or disabled workers in general (Derbyshire et al., 2024).

We comprehensively examine the relative workplace outcomes of DHH food-delivery workers in China, and analyze how DHH workers are affected by a disability-focused AI tool. Before the tool is introduced, DHH workers are more likely to be late and get bad customer ratings, though they also work more hours, deliver more orders, and are substantially less likely to quit, patterns consistent with our conceptual framework where disabled workers have worse outside options. While DHH workers earn greater profit for the Platform on a per-worker basis, they earn about 10% less per hour.

The introduction of a simple AI-based communication tool substantially enhances DHH worker performance. Work speed improves, driven by faster completion of the final stage of order fulfillment—delivery to customers—where the AI tool is used. Orders rise and customer complaints fall. Retention, already higher for DHH workers, increases further. The AI tool decreases the disability wage gap by one-third and increases DHH worker profitability.

These large benefits occur even though the new tool was relatively inexpensive and straightforward. Far from being a limitation, the simplicity of the tool illustrates that even relatively modest AI-enabled interventions can meaningfully reduce barriers faced by disabled workers, without requiring large language model or autonomous conversational capabilities. Further work can examine the impact of more advanced AI tools.

Relevance for other disabilities. While our results pertain specifically to DHH workers, AI tools may also benefit people with other disabilities. Many human faculties involve prediction, the core function of modern AI (Agrawal et al., 2018). In the context of deafness, AI predicts how "natural" speech should sound. For visual impairment, AI tools such as Seeing AI and OrCam MyEye use images and video to predict what is nearby and narrate the environment (Granquist et al., 2021). AI may also enable smart wheelchairs for people with mobility impairments (Giansanti and Pirrera, 2025), smart glasses to help people with autism spectrum disorder interpret emotional cues (Iannone and Giansanti, 2024), and brain–computer interfaces (e.g., Neuralink) that restore greater control to people with paralysis. Future work could examine the impact of such technologies on workplace outcomes.

We speculate that one reason the AI tool was especially beneficial in our setting is that hearing loss creates a relatively localized bottleneck. Delivery work in China does not require a lot of speaking, but speaking can be especially useful in one part, i.e., for communicating with customers. By contrast, in some jobs, disabilities may affect many complementary aspects of work simultaneously (e.g., DHH physicians), such that an AI tool targeting one task might generate smaller gains. Still, we believe there are many other jobs where a person's disability creates a localized bottleneck where AI might be useful.[42]

Relevance for other jobs. Our results are also specific to food delivery work, a job with short, frequent interactions with customers, decentralized work, limited educational requirements, and heavy use of smartphones. Other jobs sharing these features include ride-hailing workers (e.g., Uber, Lyft, Didi); package delivery and last mile logistics (Amazon Flex, DoorDash Drive, UPS seasonal contractors); home services and on-demand repairs (TaskRabbit, Thumbtack, Handy); field sales or merchandising work (gig workers stocking shelves for Instacart, Shipt, or retail merchandising platforms); and customer-facing retail or hospitality roles where interactions are brief and standardized.

Public policy implications. An important policy question is whether governments should subsidize AI tools that help disadvantaged workers. In our particular setting, we estimate that it was *privately beneficial* for the Platform to introduce the AI tool, with benefits much larger than costs. On the Platform, the sizable number of DHH workers increases the return to implementing the AI tool. More broadly, firms may have incentives to implement similar technologies, without subsidies, when the disability is common and per-worker gains are high. However, for rarer disabilities, or where productivity improvements

[42]For example, a blind or low-vision person might be fine with most aspects of knowledge work, but face a bottleneck in accessing visual information: she or he could benefit from an AI tool like Seeing AI. Likewise, someone with dyslexia may perform well in most aspects of a marketing job, but face a localized bottleneck related to spelling or drafting written text, and AI writing tools could be useful for this.

are smaller, public support may be warranted given the social welfare benefits of boosting disabled worker employment. Separately, while news stories often emphasize adverse aspects of AI like job loss, AI tools like ours may help reduce inequality (Acemoglu et al., 2026).

The value of data. Separate from AI, our paper illustrates the value of granular administrative data in understanding the contributions of disabled workers. In analyzing the performance of DHH workers, one might naturally focus on their speed of work. While it is true that DHH workers are slower, they also work more hours and quit much less, consistent with having a worse outside option. We speculate that some organizations might "under-notice" this "loyalty benefit" of disabled workers. Big data can illuminate the ways in which disabled workers do or do not benefit organizations.

References

Acemoglu, Daron and Joshua D Angrist, "Consequences of Employment Protection? The Case of the Americans with Disabilities Act," *JPE*, 2001, *109* (5), 915–957.

_ **, David Autor, and Simon Johnson**, "Building Pro-Worker Artificial Intelligence," Working Paper 34854, National Bureau of Economic Research February 2026.

Adhvaryu, Achyuta, Martin Atela, Valentina Brailovskaya, Priyanka Dua, Jenny Susan John, Pratibha Joshi, and Rivandra Royono, "Platform Work: Evidence from Drivers in India, Indonesia, and Kenya," 2026. NBER WP 34680.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb, *Prediction Machines: The Simple Economics of Artificial Intelligence*, Boston, MA: Harvard Business Review Press, 2018.

Aizawa, Naoki, Corina Mommaerts, and Stephanie L. Rennane, "Firm Accommodation After Workplace Disability: Labor Market Impacts and Implications for Subsidy Design," *Econometrica*, 2025, *94* (2), 341–374.

Ameri, Mason, Lisa Schur, Meera Adya, F. Scott Bentley, Patrick McKay, and Douglas Kruse, "The Disability Employment Puzzle: A Field Experiment on Employer Hiring Behavior," *ILR Review*, 2018, *71* (2), 329–364.

Angrist, Joshua D. and Jörn-Steffen Pischke, *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton University Press, 2009.

Arcot, Siva Venkatesh, "Comparative Analysis of AI-powered Outbound Dialer Campaigns vs. Legacy Outbound Dialer Campaigns (Contact Center)," *Asian Journal of Research in Computer Science*, Feb. 2025, *18* (3), 73–84.

Autor, David, "Polanyi's Paradox and the Shape of Employment Growth," in Lisa M. Lynch, ed., *Re-Evaluating Labor Market Dynamics*, Kansas City Fed, 2015, pp. 129–177.

_ , "Applying AI to Rebuild Middle Class Jobs," Technical Report, NBER WP 32140 2024.

_ **and Mark G. Duggan**, "The Rise in the Disability Rolls and the Decline in Unemployment," *Quarterly Journal of Economics*, 02 2003, *118* (1), 157–206.

_ **, Andreas Kostøl, Magne Mogstad, and Bradley Setzler**, "Disability Benefits, Consumption Insurance, and Household Labor Supply," *AER*, 2019, *109* (7), 2613–2654.

Baldwin, Marjorie L. and William G. Johnson, "A Critical Review of Studies of Discrimination Against Workers with Disabilities," *Handbook on Econ. of Discrim.*, 2006.

Banerjee, Abhijit, Rema Hanna, Benjamin A. Olken, and Diana Sverdlin Lisker, "Social Protection in the Developing World," NBER Working Paper 32382 2024.

Bellemare, Charles, Marion Goussé, Guy Lacroix, and Steeve Marchand, "Physical Disability and Labor Market Discrimination: Evidence from a Video Résumé Field Experiment," *American Economic Journal: Applied Economics*, 2023, *15* (4), 452–476.

Benson, Alan, Simon Board, and Moritz Meyer-ter-Vehn, "Discrimination in Hiring: Evidence from Retail Sales," *Review of Economic Studies*, 2024, *91* (4), 1956–1987.

Blader, Steven, Claudine Gartenberg, and Andrea Prat, "The Contingent Effect of Management Practices," *Review of Economic Studies*, 2020, *87* (2), 721–749.

Bloom, Nicholas and John Van Reenen, "Human Resource Management and Productivity," *Handbook of Labor Economics*, 2011, *1*, 1697–1767.

_ , Gordon B. Dahl, and Dan-Olof Rooth, "Work from Home and Disability Employment," *American Economic Review: Insights*, 2026, *8* (2), 179–195.

Bound, John and Richard V. Burkhauser, "Economic Analysis of Transfer Programs Targeted on People with Disabilities," *Handbook of Labor Economics*, 1999, *3*, 3417–3528.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, "The Productivity J-curve: How Intangibles Complement General Purpose Technologies," *American Economic Journal: Macroeconomics*, 2021, *13* (1), 333–372.

_ , Danielle Li, and Lindsey Raymond, "Generative AI at Work," *Quarterly Journal of Economics*, 2025, *140* (2), 889–942.

Burks, Stephen V., Bo Cowgill, Mitchell Hoffman, and Michael Housman, "The Value of Hiring through Employee Referrals," *QJE*, 2015, *130* (2), 805–839.

Chen, Jiafeng and Jonathan Roth, "Logs with Zeros? Some Problems and Solutions," *Quarterly Journal of Economics*, 2023, *139* (2), 891–936.

Chen, Shi et al., "Silent Delivery: Practices and Challenges of Delivering Among Deaf or Hard of Hearing Couriers," in "Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems" CHI '24 Association for Computing Machinery 2024.

Chen, Yanyou, Yao Luo, and Zhe Yuan, "Driving the Drivers: Algorithmic Assignment in Ride-Hailing," *Available at SSRN 4299499*, 2024.

Choi, Jung Ho and Chloe Xie, "Human+ AI in accounting: Early evidence from the field," 2025.

Cook, Cody, Rebecca Diamond, Jonathan V. Hall, John A. List, and Paul Oyer, "The Gender Earnings Gap in the Gig Economy: Evidence from over a Million Rideshare Drivers," *Review of Economic Studies*, 2021, *88* (5), 2210–2238.

Cui, Zheyuan, Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz, "The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers," aug 2025. SSRN working paper.

DeLeire, Thomas, "The Wage and Employment Effects of the Americans with Disabilities Act," *Journal of Human Resources*, 2000, *35* (4), 693–715.

Derbyshire, Daniel W., Emma Jeanes, Esmaeil Khedmati Morasae, Susan Reh, and Morwenna Rogers, "Employer-focused interventions targeting disability employment: A systematic review," *Social Science & Medicine*, 2024, *347*, 116742.

Deshpande, Manasi, "Does Welfare Inhibit Success? The Long-term Effects of Removing Low-income Youth from the Disability Rolls," *AER*, 2016, *106* (11), 3300–3330.

_ **and Lee M Lockwood**, "Beyond Health: Nonhealth Risk and the Value of Disability Insurance," *Econometrica*, 2022, *90* (4), 1781–1810.

Erickson, William, Camille Lee, and Sarah von Schrader, "Disability Statistics from the American Community Survey (ACS)," Cornell University Yang-Tan Institute (YTI), Ithaca, NY 2026. https://www.disabilitystatistics.org.

Forth, John and Melanie K. Jones, "The Disability Pay Gap Within and Across Firms," IZA Discussion Paper 17679, Bonn, Germany February 2025.

Fruits, Eric and Kristian Stout, "AI, Productivity, and Labor Markets: A Review of the Empirical Evidence," Technical Report ICLE Issue Brief February 5 2026.

Gentry, Melissa, "Driving Inclusion: The Effect of Improved Transportation for People with Disabilities," *Unpublished Manuscript, Texas A&M*, 2025.

Gernsbacher, Morton Ann, "Editorial Perspective: The Use of Person-first Language in Scholarly Writing May Accentuate Stigma," *J. Child Psychology & Psychiatry*, 2017, *58.*

Giansanti, Daniele and Antonia Pirrera, "Integrating AI and Assistive Technologies in Healthcare: Insights from a Narrative Review of Reviews," *Healthcare*, 2025, *13* (5), 556.

Granquist, Christina et al., "Evaluation and Comparison of Artificial Intelligence Vision Aids: Orcam MyEye 1 and Seeing AI," *Journal of Visual Impairment & Blindness*, 2021, *115* (4), 277–285.

Grech, Lisa B., Donna Koller, and Amanda Olley, "Person-first and Identity-first Disability Language: Informing Client Centred Care," *Social Science & Medicine*, 2024, *362*, 117444.

Hoffman, Mitchell and Christopher Stanton, "People, Practices, and Productivity: A Review of New Advances in Personnel Economics," *Handbook of Labor Economics*, 2025.

Hoffmann, Manuel, Sam Boysel, Frank Nagle, Sida Peng, and Kevin Xu, "Generative AI and the Nature of Work," Technical Report CESifo WP No. 11479 April 2024.

Iannone, Andrea and Daniele Giansanti, "Breaking Barriers—The Intersection of AI and Assistive Technology in Autism Care: A Narrative Review," *Healthcare*, 2024, *12* (2).

Krieger, Linda Hamilton, *Backlash Against the ADA: Reinterpreting Disability Rights*, University of Michigan Press, 2010.

Krueger, Alan B. and Douglas Kruse, "Labor Market Effects of Spinal Cord Injuries in the Dawn of the Computer Age," NBER Working Paper 5302 1995.

Kruse, Douglas, Alan Krueger, and Lisa Drastal, "Computer Use, Training, and Employment of People with Spinal Cord Injury," *Spine*, 1996, *21* (7), 891–896.

Lalive, Rafael, Jean-Philippe Wuellrich, and Josef Zweimüller, "Do Financial Incentives Affect Firms' Demand for Disabled Workers?," *JEEA*, 2013, *11* (1), 25–58.

Lane, Harlan, *The Mask of Benevolence: Disabling the Deaf Community*, AO Knopf, 1992.

Lee, Sooyeon, Bjorn Hubert-Wallander, Molly Stevens, and John M. Carroll, "Understanding and Designing for Deaf or Hard of Hearing Drivers on Uber," in "Proceedings of the 2019 CHI Conf. on Human Factors in Computing Systems" 2019, pp. 1–12.

L'Horty, Yannick, Naomie Mahmoudi, Pascale Petit, and François-Charles Wolff, "Is Disability More Discriminatory in Hiring than Ethnicity, Address or Gender? Evidence from a Multi-criteria Correspondence Experiment," *Social Science & Medicine*, 2022, *303.*

Lin, Ou, "What Do Platform Workers Think About the Law? The Ambiguous Legal Status and Legal Consciousness of On-Demand Food Delivery Riders in China," *Industrial Law Journal*, 2025.

Lise, Jeremy, Elena Pastorino, and Luigi Pistaferri, "Revisiting the Employment Effects of the Americans with Disabilities Act," 2023.

Madans, Jennifer H., Julie D. Weeks, and Nazik Elgaddal, "Hearing Difficulties Among Adults: United States, 2019," NCHS Data Brief 414, Centers for Disease Control and Prevention, National Center for Health Statistics July 2021.

Merali, Ali, "Scaling Laws for Economic Productivity: Experimental Evidence in LLM-Assisted Translation," *arXiv preprint*, December 2024. Last revised Dec 7 2024.

Milligan, Kevin, "The Long-Run Growth of Disability Insurance in the United States," in "Social Security Programs and Retirement around the World," UChicago Press, 09 2012.

Minor, Dylan, Nicola Persico, and Deborah M. Weiss, "Criminal Background and Job Performance," *IZA Journal of Labor Policy*, 2018, *7*, 1–49.

Minsheng Weekly, "Report on Disabled Delivery Riders," *People's Daily*, https://paper.people.com.cn/mszk/pc/content/202602/09/content_30142581.html 2026.

Noy, Shakked and Whitney Zhang, "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence," *Science*, 2023, *381* (6654), 187–192.

Olusanya, Bolajoko O., Adrian C. Davis, and Howard J. Hoffman, "Hearing Loss Grades and the International Classification of Functioning, Disability and Health," *Bulletin of the World Health Organization*, 2019, *97* (10), 725–728.

Oyer, Paul and Scott Schaefer, "Personnel Economics: Hiring and Incentives," *Handbook of Labor Economics*, 2011.

Schmidheiny, Kurt and Sebastian Siegloch, "On Event Studies and Distributed-Lags in Two-Way Fixed Effects Models: Identification, Equivalence, and Generalization," *Journal of Applied Econometrics*, 2023, *38* (5), 695–713.

Shah, Shariq, Hossein Ghomeshi, Edlira Vakaj, Emmett Cooper, and Shereen Fouad, "A Review of Natural Language Processing in Contact Centre Automation," *Pattern Analysis and Applications*, 2023, *26* (3), 823–846.

Shaw, Kathryn and Edward P. Lazear, "Tenure and Output," *Labour Economics*, 2008, *15* (4), 704–723.

Shen, Jonathan et al., "Natural TTS Synthesis by Conditioning WaveNet on Mel Spectrogram Predictions," in "Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)" 2018, pp. 4779–4783.

Szerman, Christiane, "The Labor Market Effects of Disability Hiring Quotas," *SSRN*, 2022.

Tan, Xu, Tao Qin, Frank Soong, and Tie-Yan Liu, "A Survey on Neural Speech Synthesis," *arXiv preprint arXiv:2106.15561*, 2021.

The Economist, "China's 200m Gig Workers Are a Warning for the World," 2025.

Wang, Yu et al., "Hearing Loss Prevalence and Burden of Disease in China: Findings from Provincial-level Analysis," *Chinese Medical Journal*, 2025, *138* (1), 41–48.

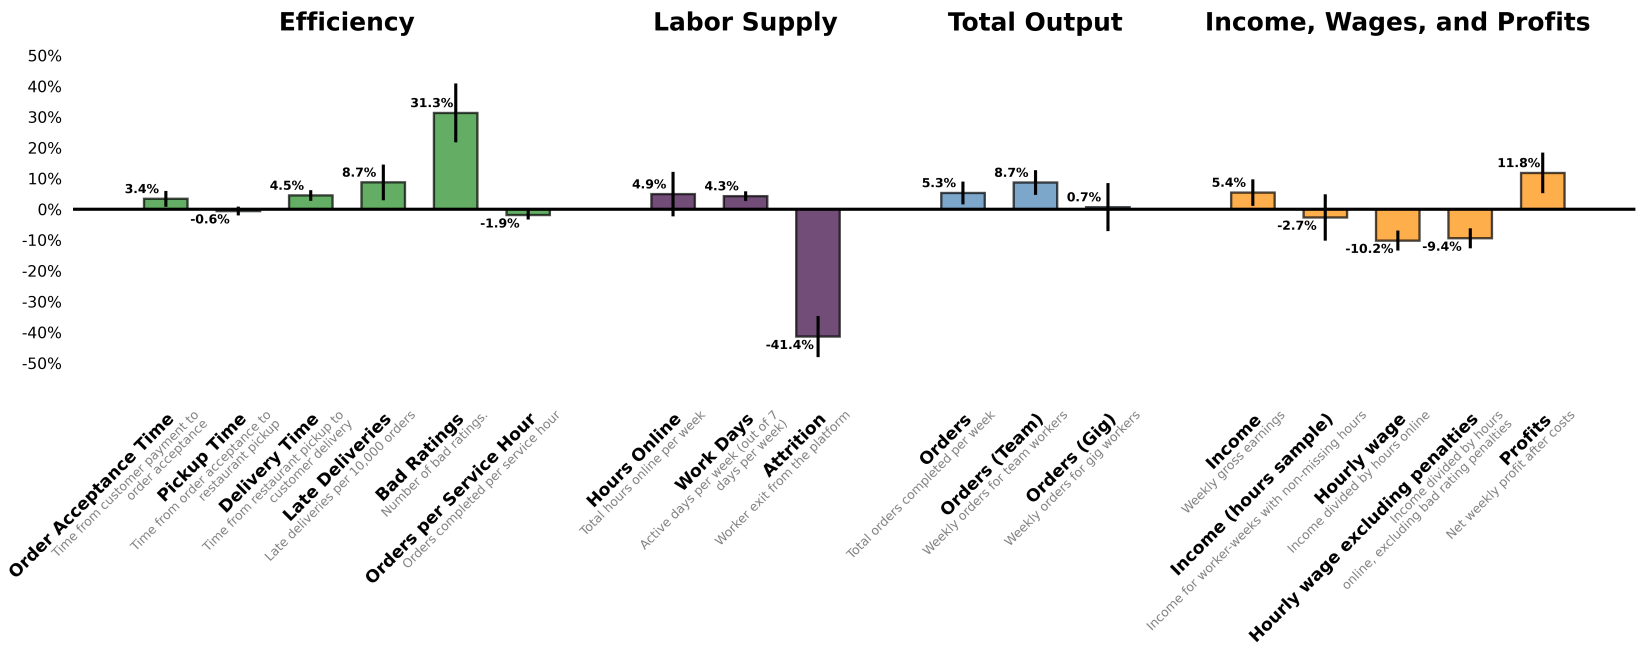
Wooldridge, Jeffrey, *Econometric Analysis of Cross Section and Panel Data* 2002.

World Health Organization, "Disability and Health," Fact Sheet 2023.

_ , "Deafness and hearing loss," https://www.who.int/en/news-room/fact-sheets/

detail/deafness-and-hearing-loss 2025. Accessed 02/26.

Figure 1: Pre-AI Differences in Workplace Outcomes Between DHH and Non-Disabled Workers



Main notes: This figure visually summarizes gaps between DHH and non-disabled workers during the pre-AI period. For each outcome, each bar represents a percentage difference associated with being a DHH worker. Each bar graph comes from estimating equation (2) with a different outcome, with the coefficient shown in the bar and the 95% confidence interval accounting for clustering by driver shown in the line. For all outcomes except late deliveries, bad ratings, and attrition, percentages are derived from an inverse hyperbolic sine transformation of the dependent variable; these coefficients approximate log differences and are interpretable as percent changes. For late deliveries, bad ratings, and attrition, the plotted value is a non-transformed OLS coefficient divided by the pre-AI DHH mean, giving the percent difference relative to the pre-AI DHH baseline. See column 1 of Table A.5 for the underlying regressions plotted here. The data are from 16,912 workers and have 180,375 worker-weeks.

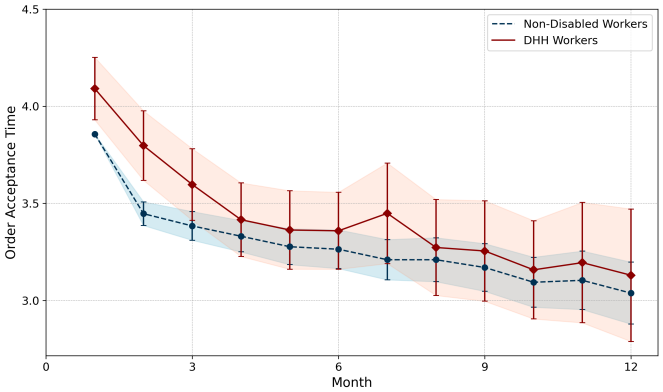
Efficiency notes: In the green bars, *Order Acceptance Time* is time from order placement to a worker accepting it, *Pickup Time* is time from accepting the order to receiving the food from the restaurant, and *Delivery Time* is time from pickup to customer delivery. These times are measured in minutes per order averaged over the week for each worker. *Late Deliveries* is the number of late deliveries per 10,000 orders. *Bad Ratings* is the number of bad ratings in a week. *Orders per Service Hour* measures orders completed per service hour, and is calculated as $60/(t_1 + t_2 + t_3)$, where the denominator is the sum of average minutes spent in each stage of an order.

Labor supply notes: In the purple bars, *Hours Online* is total hours the worker was online and available for orders, *Work Days* is number of days the worker was active in that week, and *Attrition* is a dummy for exiting the Platform.

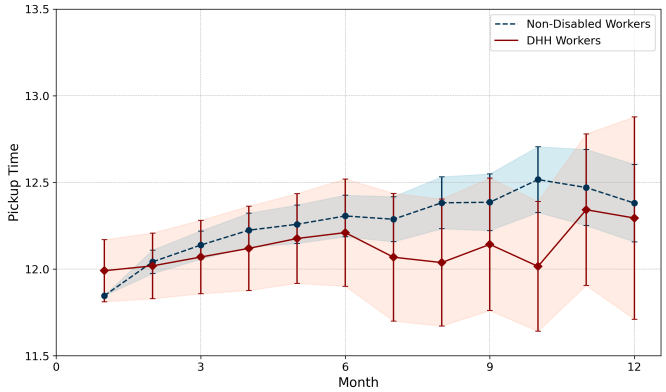
Total output notes: In the blue bars, *Orders* is the total number of delivery orders completed by a worker in a given week. We also show Orders results separately for team and gig workers.

Income, wages, and profits notes: In the orange bars, *Income* is a driver's weekly gross earnings. Since hours data are missing for some drivers, we also show results on income restricted to drivers with non-missing hours in *Income (hours sample)*. *Hourly wage* is equal to income divided by hours. *Hourly wage excluding penalties* is similar to *Hourly wage*, but excludes bad ratings penalties from income. *Profits* is profit for the Platform.

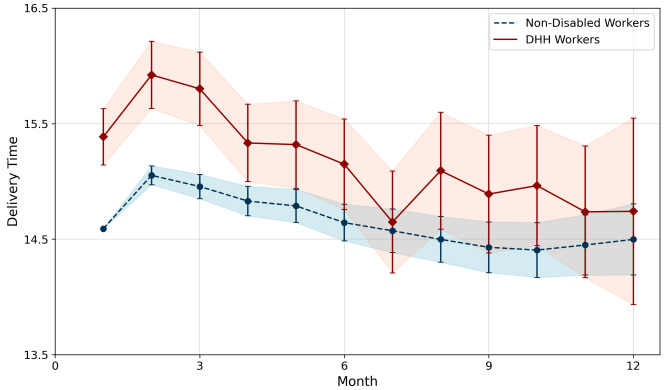
Figure 2: Pre-AI Tenure-Productivity Curves: Time Spent on Different Parts of the Order on Months of Tenure



(a) Order Acceptance Time



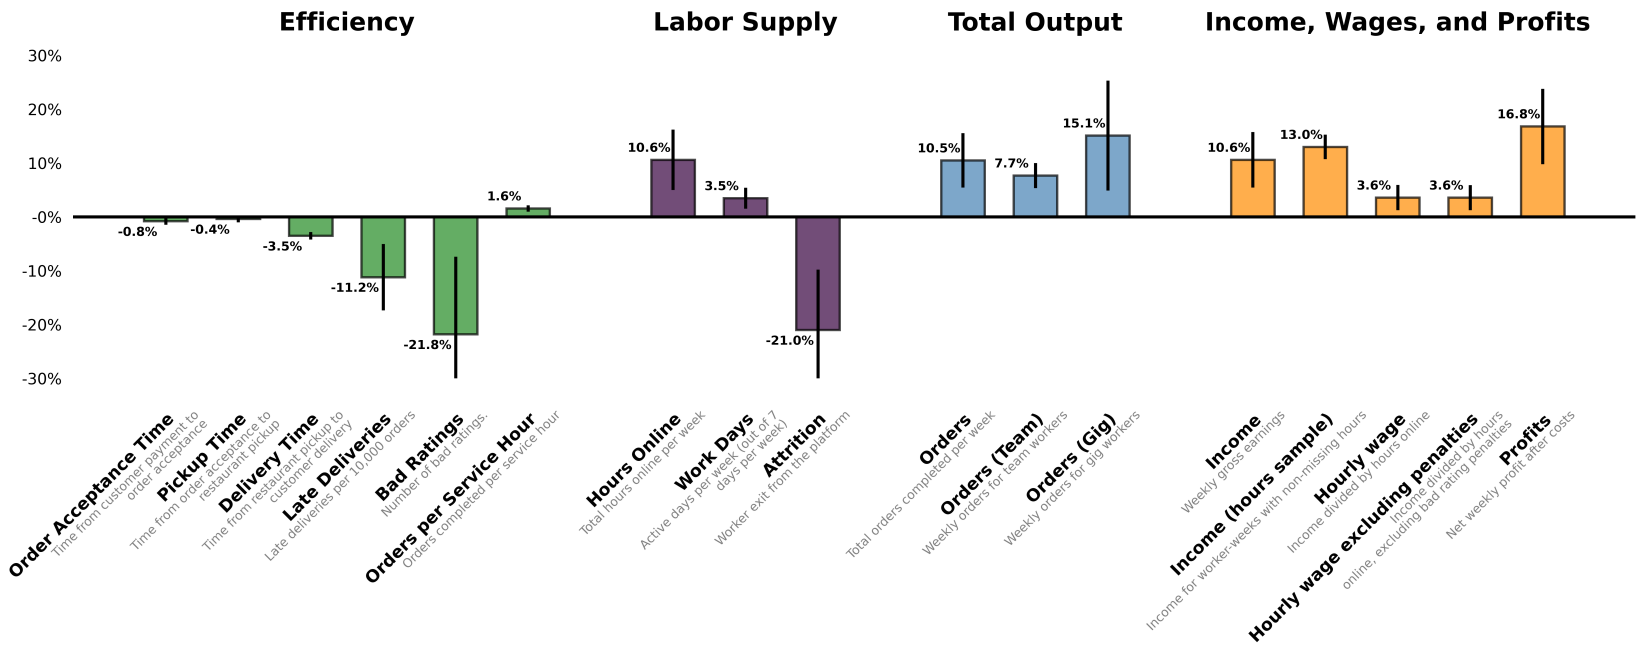
(b) Pickup Time



(c) Delivery Time

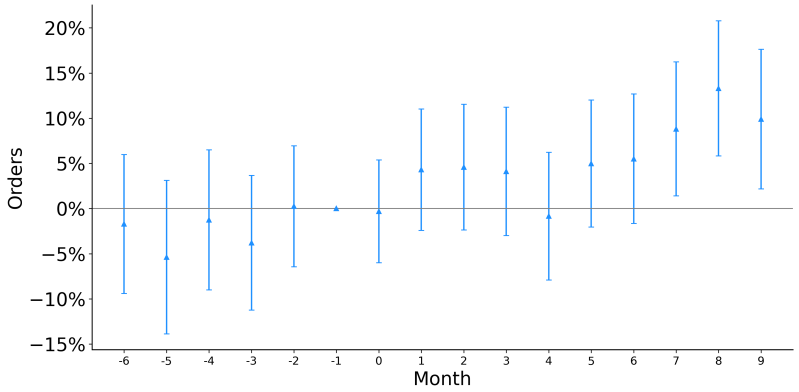
Notes: This figure examines how mean worker speed at different stages of the delivery process varies with worker tenure on the Platform. It plots time spent per order at different stages across each month of a worker's tenure, separately for DHH and non-disabled workers. The three time measures (order acceptance time, pickup time, and delivery time) are defined as in Figure 1. Each point on a curve represents the estimated average time at a given tenure month, obtained from the specification described in the main text. The red solid line corresponds to DHH workers, and the blue dashed line corresponds to non-disabled workers. The first month for non-disabled workers is normalized to zero; we then rescale all coefficients by adding back the average time for non disabled workers in their first month. Thus, at "Month 1," the figure shows the average per order time for workers in their first month of tenure. Error bars indicate 95% confidence intervals, with standard errors clustered at the worker level. The sample is workers in the pre-AI period, restricting to worker-weeks during the first 12 months of worker tenure.

Figure 3: The Impact of the AI Tool on the Outcomes of DHH Workers

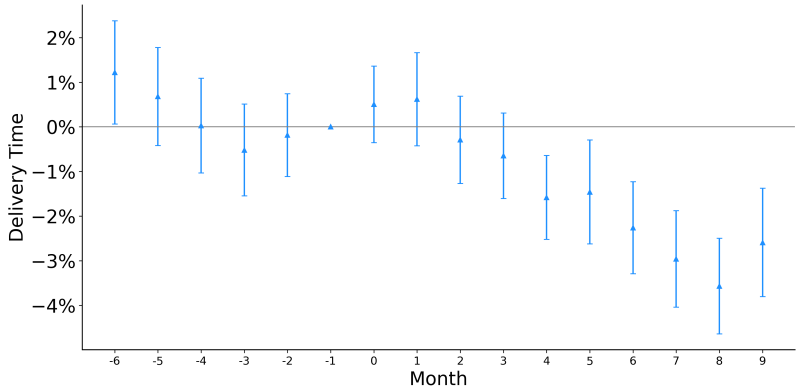


Notes: This figure shows DiD estimates of the impact of the AI-based Intelligent Outbound Call tool on DHH workers. Each bar reports the percentage change in outcomes for DHH workers relative to non-disabled workers following the introduction of the AI tool. Positive values indicate improvements for DHH workers relative to non-disabled workers (e.g., higher output or earnings). For duration-based outcomes and negative performance indicators (e.g., delivery times, bad ratings), negative values represent improvements. Error bars denote 95% confidence intervals, with standard errors clustered by workers. Outcome definitions and outcome categories follow Figure 1. Percentages are calculated as in Figure 1: inverse hyperbolic sine regressions (with coefficients shown) for most outcomes, and non-transformed OLS coefficients divided by the pre-AI DHH mean for late deliveries, bad ratings, and attrition. See column 1 of Table A.6 for the underlying regressions plotted here. The sample is the full sample of drivers summarized in Table 1, so the maximum sample size is 480,070 worker-weeks from 43,605 workers. Of these workers, 11,724 are singletons (i.e., appear in the data for only one week). For outcomes other than attrition, which are estimated with worker fixed effects using `reghdfe` in Stata 17, singleton observations are perfectly absorbed by the worker fixed effects and do not contribute to estimation, and `reghdfe` drops them. Attrition models do not include worker fixed effects, so singletons contribute there. Throughout the paper, we report sample sizes prior to any exclusions from `reghdfe`. As noted in Figure A15, our results are robust to restricting to drivers working at least one week pre- and post-AI.

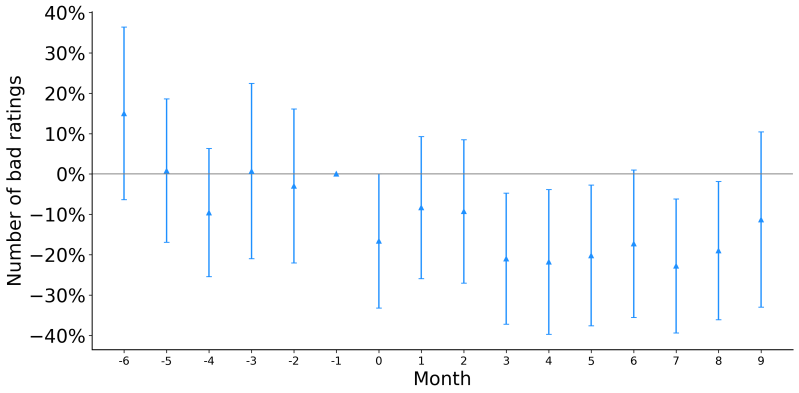
Figure 4: Event Study Analysis: AI Tool Effects on Productivity



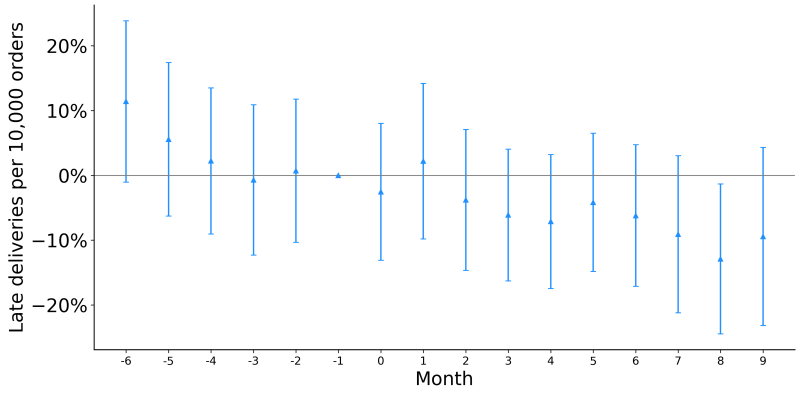
(a) Number of Orders Completed



(b) Delivery Time



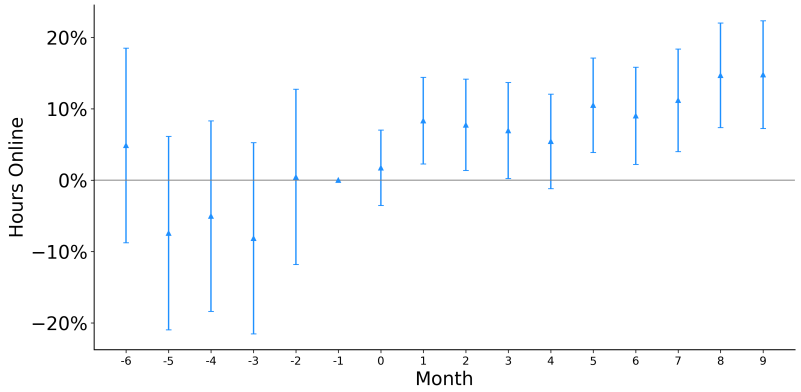
(c) Bad Ratings



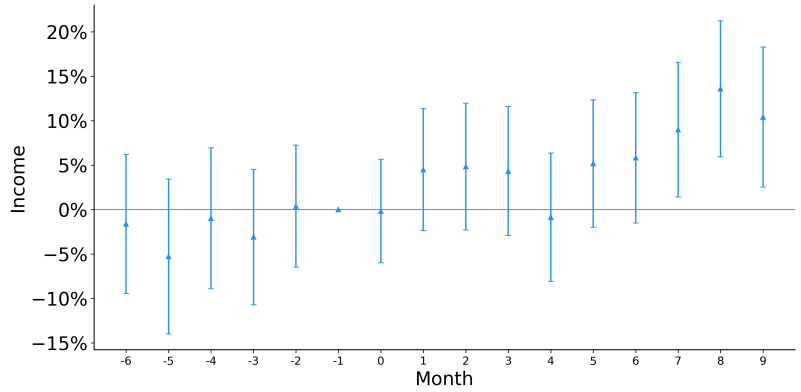
(d) Late Deliveries

Notes: This figure presents an event study analysis of the introduction of the AI-based Intelligent Outbound Call (IOC) system. The x-axis shows event time in months relative to the intervention, with 0 denoting the first month in which the tool is in effect. The underlying data are weekly; months are defined as 4-week periods, anchored on the exact week of implementation. The event study window covers 6 months before and 10 months after the intervention; observations outside this window are pooled into the boundary coefficients at −7 and +10. Dots show the event study coefficients, and whiskers show the 95% confidence intervals accounting for clustering by driver. The y-axis shows the estimated difference between DHH and non-disabled workers in percentage terms; percentages are calculated as in Figure 3. The sample used here is the full sample of drivers summarized in Table 1.

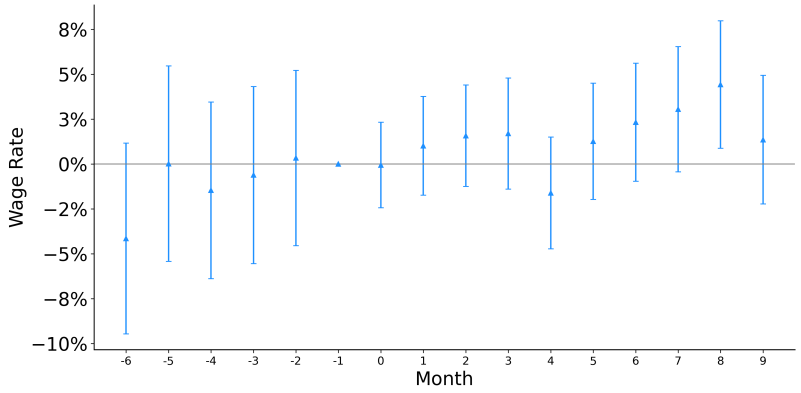
Figure 5: Event Study Analysis: AI Tool Effects on Hours, Income, Wages, and Profits



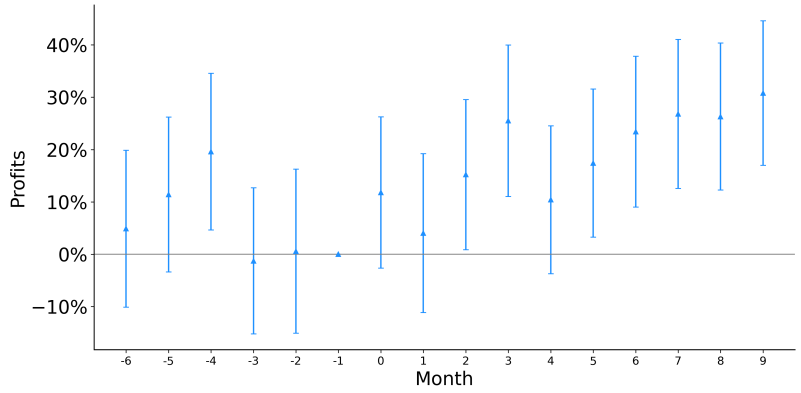
(a) Hours Online



(b) Income



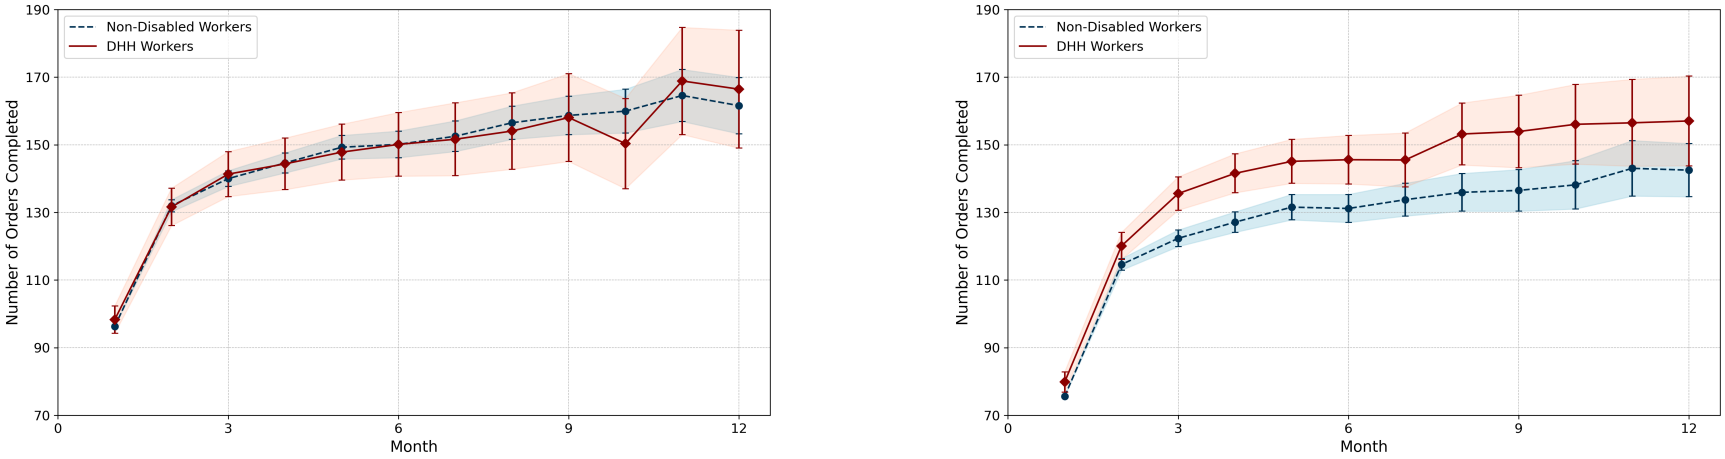
(c) Hourly Wage



(d) Profits

Notes: This figure is similar to Figure 4, but examines the effects of the AI tool on DHH workers' hours online, income, wages, and on platform profits from DHH workers. All these variables are defined in the main text and also described in Figure 1. Dots show the event-study coefficients, and whiskers show the 95% confidence intervals, accounting for clustering by worker. The y-axis shows the estimated difference between DHH and non-disabled workers in percentage terms, as in Figure 4. The sample used here is the full sample of drivers summarized in Table 1.

Figure 6: Tenure-Productivity Curves on Number of Orders Completed, Before and After the AI Tool

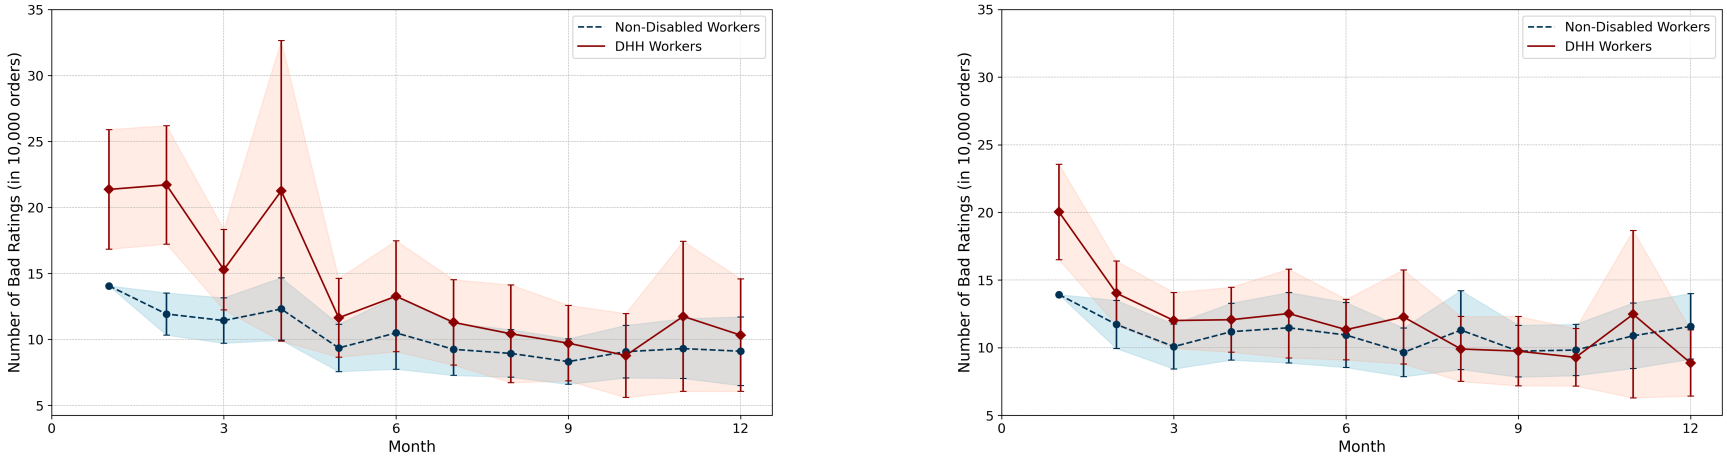


(a) Pre-AI tool

(b) Post-AI tool

Notes: This figure examines how the productivity trajectory of new workers over their first year differs before and after the introduction of the AI tool. The figure plots the average number of orders completed as a function of months of tenure. Error bars show 95% confidence intervals accounting for clustering by worker. The red solid line corresponds to DHH workers, and the blue dashed line corresponds to non-disabled workers. Tenure profiles are constructed following the same cohort-based approach as Figure 2, but are split by whether workers had access to the AI tool from the start of their tenure. The AI tool was implemented in mid 2022 at time t_0. Panel (a), *Pre-AI tool*, displays tenure profiles for workers who entered the Platform prior to the AI tool, and we include their observations until the time of AI implementation. For example, if someone joins the Platform at time $t_0 - 2$, we include two months of tenure for this worker, corresponding to months $t_0 - 2$ and $t_0 - 1$. Panel (b) consists of workers who enter the Platform after the AI tool is introduced. In Panel (a), the sample includes 176,221 worker-week observations from 16,912 workers, of whom 1,515 are DHH and 15,397 are non-disabled. In Panel (b), the sample includes 183,193 worker-week observations from 26,693 workers, of whom 2,769 are DHH and 23,924 are non-disabled. We restrict to worker-weeks during the first 12 months of worker tenure.

Figure 7: Tenure-Productivity Curves in Bad Ratings, Before and After the AI Tool

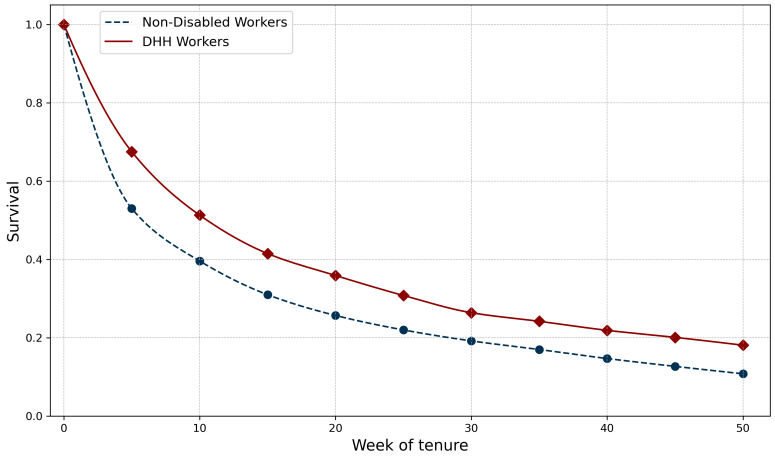


(a) Pre-AI tool

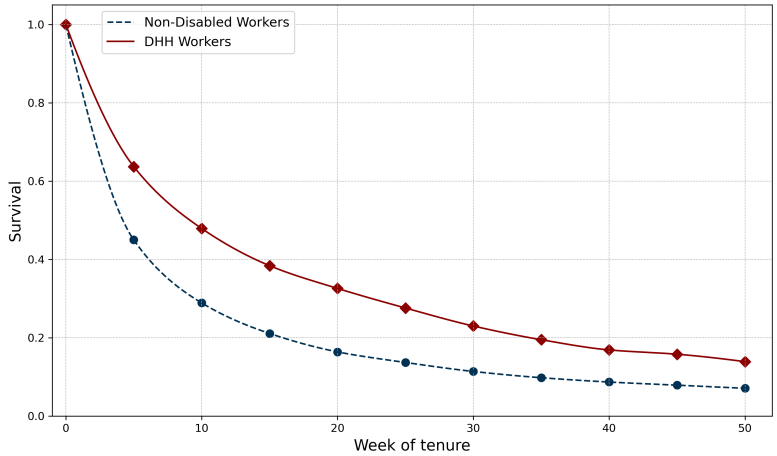
(b) Post-AI tool

Notes: This figure is similar to Figure 6, but the outcome is bad customer ratings, measured in terms of bad ratings per 10,000 orders. Error bars show 95 percent confidence intervals accounting for clustering by worker. The samples in each panel are the same as in Figure 6. Appendix Figure A.17 shows that results are similar when the dependent variable is a worker's weekly count of bad ratings instead of ratings per 10,000 Orders.

Figure 8: Survival Curves: Before and After the AI Tool



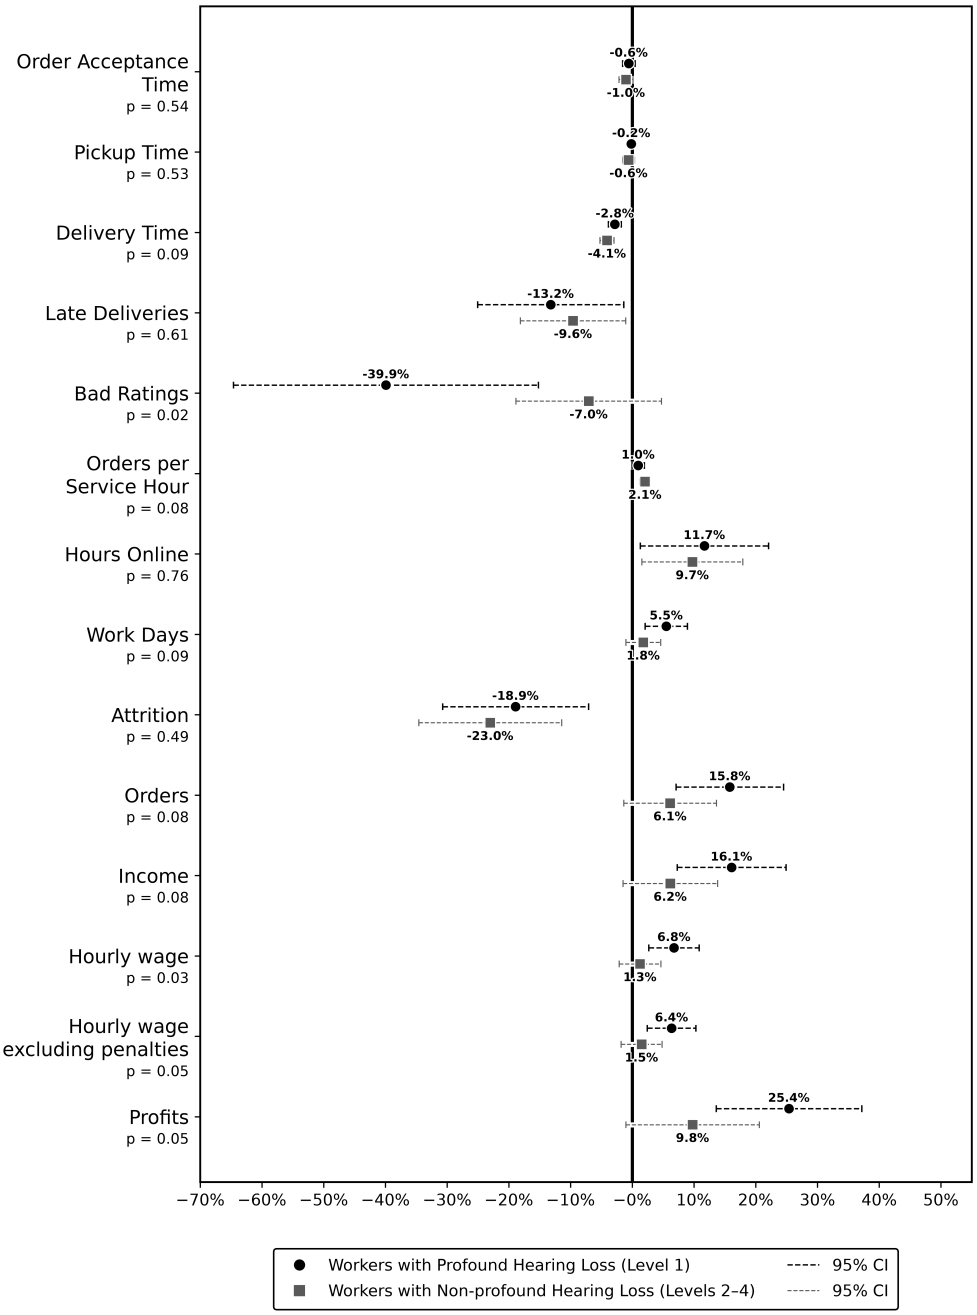
(a) Pre-AI tool



(b) Post-AI tool

Notes: This figure is similar to Figure 6, but shows survival curves as the outcome. The y-axis reports the fraction of workers remaining active on the Platform (i.e., not exited) and the x-axis reports the number of weeks since initial registration. The samples in each panel are the same as in Figure 6.

Figure 9: Heterogeneous Treatment Effects of the AI Tool by Disability Severity



Notes: This figure examines how impacts of the AI tool vary according to disability severity. Each dot-and-whisker plot is calculated using a regression where DHH workers of a given severity are compared to non-disabled workers. Each dot or square represents the treatment effect of the AI tool on DHH workers of a given severity. The whiskers show 95% confidence intervals accounting for clustering by driver. The p-values shown are from tests of whether there is a differential effect by disability severity. Each p-value is calculated from one regression with two key interaction terms, Profound DHH$_i \times Post_t$ and Non-profound DHH$_i \times Post_t$, testing whether the two interactions are different, where Profound DHH$_i$ and Non-profound DHH$_i$ are dummies for having profound hearing loss (HL) and non-profound HL, respectively. As described in Section 1, severity follows the official Chinese four-level classification (GB/T 26341-2010), where Level 1 is profound HL, corresponding to more than 90 decibels of HL, and Level 4 is moderate (41-60 decibels of HL). Severity level is based on a worker's government certificate of disability. The sample here is the full sample of drivers summarized in Table 1.

Table 1: Mean Characteristics of DHH and Non-disabled Workers

	DHH Workers (1)	Non-Disabled Workers (2)	Full Sample (3)	P-Value of Difference (4)
Panel A: Worker Demographics				
Age	32.8	30.3	30.5	0.00
Male	92.6%	90.8%	91.0%	0.00
Panel B: Work Model				
Team worker	35.4%	36.4%	36.3%	0.18
Gig worker	64.6%	63.6%	63.7%	0.18
Panel C: Disability Severity for DHH Workers				
Level 1: Profound hearing loss	53.6%			
Level 2: Severe hearing loss	16.5%			
Level 3: Moderate-Severe hearing loss	14.2%			
Level 4: Moderate hearing loss	15.7%			
Panel D: Distance & Geography Variables				
Driver is in a big city	36.6%	32.3%	32.7%	0.00
Driver is in Shanghai	7.27%	5.58%	5.75%	0.00
Driver is in Beijing	3.60%	3.36%	3.38%	0.40
Driver is in Eastern China	64.2%	58.2%	58.8%	0.00
Avg. distance bwn. restaurant and customer, in km	2.24	2.16	2.17	0.00
Unique restaurants served, per week	73	65	66	0.00
Number of workers	4,284	39,321	43,605	
Number of worker-weeks	70,301	409,769	480,070	

Notes: This table provides summary statistics on food-delivery workers on the Platform, both deaf or hard of hearing (DHH) and non-disabled. We use data from the whole sample period of 2021-2023. Column (1) shows DHH worker averages; (2) shows non-disabled worker averages; (3) is the full sample average. Column (4) reports the p-value for the difference in means between DHH and non-disabled workers. The unit of observation is a worker for demographics. In Panel A, age is age upon joining the Platform. In Panel B, Team worker and Gig worker are dummies for a driver's work model, also called their registration type. Panel C reports the distribution of disability severity among the DHH workers in column 1. Severity follows the official four-level classification (GB/T 26341-2010). Level 1 is profound hearing loss, meaning loss of greater than 90 decibels; Level 2 is severe hearing loss of 81-90 decibels; Level 3 is moderately severe hearing loss of 61-80 decibels; and Level 4 is moderate hearing loss of 41-60 decibels. Level 4 is the lowest level of severity to have a certificate of disability in China. Severity is based on government medical certificates of disability, which are provided by drivers to the Platform. Big cities cover the top 10 cities in China by population: Chongqing, Shanghai, Beijing, Chengdu, Guangzhou, Shenzhen, Wuhan, Tianjin, Xi'an, and Suzhou. Eastern China provinces include: Beijing, Tianjin, Hebei, Shanghai, Jiangsu, Zhejiang, Fujian, Shandong, Guangdong, and Hainan. Location information is missing for 2,904 workers in the full sample (about 7%).

Table 2: Descriptive Statistics for Key Outcome Variables, all at the Worker-Week Level

	Mean (1)	Median (2)	Std. Dev. (3)	Obs. (4)	Workers (5)
Panel A: Efficiency/productivity: Quantity and Quality					
Order Acceptance Time: minutes from customer payment to order acceptance	3.37	2.48	9.20	464,717	42,281
Pickup Time: minutes from order acceptance to restaurant pickup	11.97	11.47	6.86	464,850	42,293
Delivery Time: minutes from restaurant pickup to customer delivery	14.18	13.44	5.10	465,546	42,367
Late Deliveries per 10k orders	416.44	134.23	816.42	480,070	43,605
Bad Ratings per Week	0.093	0.00	0.365	480,070	43,605
Week where worker gets at least one Bad Rating, 0 or 1	0.076	0.00	0.265	480,070	43,605
Good Ratings per Week	9.32	6.00	12.30	480,070	43,605
Week where worker gets at least one Good Rating, 0 or 1	0.816	1.00	0.387	480,070	43,605
Panel B: Labor Supply on the Platform					
Hours Online, in hours per week	40.42	37.85	29.19	343,733	34,965
Work Days per week	5.31	6.00	2.13	480,070	43,605
Driver attrites from the Platform, 0 or 1	0.077	0.00	0.267	456,192	40,567
Panel C: Total Output					
Total Orders	152.39	127.00	141.52	480,070	43,605
Peak Hour Orders	90.92	76.00	88.42	480,070	43,605
Total Orders, restrict to Team drivers only	212.35	205.00	146.18	272,986	15,828
Total Orders, restrict to Gig drivers only	73.35	42.00	85.28	207,084	27,777
Panel D: Income, Wages, and Profits					
Driver Income, in CNY per week	Confid.	Confid.	Confid.	480,070	43,605
Driver Income, restrict to weeks with hours data, in CNY per week	Confid.	Confid.	Confid.	343,733	34,965
Hourly Wage, in CNY per hour	Confid.	Confid.	Confid.	342,061	34,909
Hourly Wage, excluding any bad rating penalties from income, in CNY per hour	Confid.	Confid.	Confid.	342,061	34,909
Platform Profits from a driver, in CNY per week	Confid.	Confid.	Confid.	480,070	43,605

Notes: This table summarizes weekly work outcomes for workers on the Platform. We use data from the whole sample period of 2021-2023. The unit of observation is a worker-week. Outcome definitions and grouping into the four categories are the same as in Figure 1. The "Confid." marking indicates that certain summary statistics cannot be shared due to confidentiality requirements. An observation in column 4 is a worker-week. The attrition outcome has fewer observations than other outcomes because it excludes the last 7 weeks of 2023: since attrition is defined as 8 consecutive weeks of no activity, the last week in which an attrition decision can be observed is the 8th-to-last week of the sample. Appendix Table A.2 presents summary statistics separately for DHH and non-disabled workers. Appendix Table A.3 presents summary statistics separately for gig and team workers.

Table 3: Sources of Pre-AI Profit Differences and of Profit Impact of AI Tool

Component	Share
Panel A: Sources of Pre-AI Profit Differences	
Orders	85%
Bad Ratings	-4%
Onboarding	19%
Panel B: Sources of Profit Impact	
Orders	93%
Bad Ratings	2%
Onboarding	5%

Notes: Profit has three components as in equation (1): profit from orders (order revenue minus wage payments), penalty from bad ratings, and onboarding cost. To assess the relative importance of each component, we convert all three to a common CNY scale. For each component j, we multiply the estimated percentage difference $\hat{\beta}_j$—from Figure 1 for Panel A and from Figure 3 for Panel B—by the component's signed average profit contribution $\bar{y}_j$ (negative for bad-rating penalties and onboarding costs) to obtain an implied CNY impact $\Delta y_j = \hat{\beta}_j \cdot \bar{y}_j$. Each component's share is its implied CNY impact divided by the sum across all three components.

Table 4: Worker Surplus from the AI Tool: Calculations Based on the Framework

	Worker surplus, % of pre-AI earnings
Panel A: Excluding amenity changes	
Baseline	3.9 (1.5)
Workers present pre- and post-AI	3.6 (1.6)
Panel B: Including amenity changes	
$\eta = 1$	11.7 (4.1)
$\eta = 1.5$	7.7 (2.7)
$\eta = 2$	5.7 (2.0)
Panel C: By disability severity, excluding amenity changes	
Profound HL	7.4 (2.4)
Non-profound HL	1.3 (1.8)

Notes: The table reports the welfare gain to DHH workers from the AI tool, expressed as a percent of their pre-AI weekly earnings, with delta-method standard errors in parentheses. Each entry applies equation (4) to the effects of the AI tool on hourly wages and on weekly hours online, estimated separately for each row; the overall effects are reported in Figure 3 and those by disability severity in Figure 9. Panel A assumes the tool benefits workers only through the higher hourly wage—i.e., that it does not change the amenity value of Platform work—so that the entire increase in hours is movement along a fixed labor supply curve to the Platform; this makes it a lower bound on the welfare gain. Panel B allows the tool to also raise the non-pecuniary value of Platform work, such as by reducing stress and difficult customer interactions. Here η denotes the elasticity of weekly hours on the Platform with respect to the hourly return to Platform work: our wage and hours estimates imply $\eta \approx 2.9$, above existing estimates for ride-hail drivers and other workers with flexible hours, which mostly lie between about one and two (see Appendix E.4), and Panel B attributes the excess increase in hours to improved amenities.

Web Appendix for "Empowering Inclusive Work"

by Yanyou Chen, Mitchell Hoffman, Huilan Xu, and Zhe Yuan

The Web Appendix is organized as follows. Appendix A provides additional figures and tables (with a guide at the start). Appendix B reviews prior work, largely from outside economics, on the productivity of disabled workers. Appendix C gives details on the AI tool. Appendix D provides background on the evolution of text-to-speech (TTS) technology, and describes the critical role of AI in TTS. Appendix E presents additional analysis: effects on good ratings (Section E.1), DHH worker entry (Section E.2), a mediation analysis (Section E.3), and welfare calculations quantifying how much DHH workers gain from the AI tool (Section E.4). Appendix F describes hearing loss levels in China. Appendix G discusses potential nondisclosure of DHH status. Appendix H presents qualitative data on DHH workers.

A Appendix Figures and Tables

Organization. Appendix A figures and tables are grouped into six themes. The themes roughly follow the order of the main text.

Gaps between DHH and non-disabled workers in the US	
Table A.1	DHH employment and wage gaps, U.S. labor market
Summary statistics	
Tables A.2–A.4	By DHH status; by work model (gig vs. team); by DHH severity
Figure A.1	Raw means over time, DHH vs. non-disabled workers
Figure A.2	Missingness rates for hours data over time
Pre-AI differences in outcomes (Section 2)	
Figure A.3	Pre-AI differences by work model
Figure A.4	Pre-AI differences by severity of hearing loss
Table A.5	Pre-AI differences in tabular form
AI treatment effects and event studies (Section 3.2)	
Figure A.5	Main DiD estimates by work model
Table A.6	Extended DiD results: subsamples and specifications
Figure A.6	Event studies by work model
Figure A.7	Event study extended through month 15 post-adoption (months 0–15)
Heterogeneous treatment effects of the AI tool (Section 3.4)	
Figure A.8	By pre-AI tenure
Figure A.9	By pre-AI productivity
Figure A.10	Robustness on HTE by disability severity: dosage-response model
Figure A.12	By regional unemployment
Figures A.11, A.13	Robustness on HTE by disability severity and regional unemployment: controlling for interactions with age and gender
Robustness checks on impact of AI tool (Section 3.5)	
Figure A.14	Hours sample, matched sample, and additional control variables
Figure A.15	Differential attrition
Figure A.16	Longer-run effects and external validity
Tables A.7, A.8, A.9	Additional results and robustness: bad ratings; late deliveries; attrition
Figure A.17	Tenure-productivity curves for bad ratings, counts instead of ratings per 10k orders

Table A.1: DHH Employment and Wage Gaps in the U.S. General Labor Market
(CPS ASEC 2008–2023)

	(1) No controls	(2) Demog. + Year FE	(3) Demog. + Year + State FE	(4) Demog. + Year + State + Occ. + Ind. FE
Panel A: Employment (Linear Probability Model)				
DHH	-0.2665^{***}	-0.2445^{***}	-0.2454^{***}	—
	(0.0041)	(0.0039)	(0.0039)	
Observations			1,549,931	
Panel B: Log Hourly Wage (OLS, conditional on employment)				
DHH	0.0000	-0.0619^{***}	-0.0567^{***}	-0.0369^{***}
	(0.0086)	(0.0077)	(0.0077)	(0.0069)
Observations			985,211	

Notes: CPS ASEC microdata from 2008–2023, accessed via IPUMS (Flood et al., 2024); sample restricted to working-age adults aged 18–64. Each observation is a worker (one per person per survey year). DHH equals 1 if the respondent reports hearing difficulty (`DIFFHEAR`), and 0 for the non-disabled comparison group (those reporting no difficulty on any of the six CPS disability questions, `DIFFANY` = No). Panel A estimates a linear probability model for employment. In Panel B, hourly wages are constructed as annual wage and salary income divided by weeks worked times usual weekly hours; the top and bottom 1% of the weighted hourly wage distribution are trimmed before taking logs. Column (1) has no controls; column (2) adds age, age squared, gender, education categories, race/ethnicity indicators, and year fixed effects; column (3) further adds state fixed effects; column (4) further adds occupation and industry fixed effects. Column (4) is omitted for Panel A because industry is only observed for employed workers, making industry fixed effects inapplicable in an employment regression. All regressions are weighted by CPS ASEC person weights (`ASECWT`) to produce population-representative estimates. The near-zero unadjusted wage gap in Panel B, Column (1) reflects DHH workers being substantially older on average (mean age 49.7 vs. 40.2 for non-disabled workers); once age is controlled for, a sizable gap emerges. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.2: Summary Statistics for Key Outcomes at Worker-week Level: DHH vs. Non-disabled Workers

	DHH Workers				Non-disabled Workers			
	Mean (1)	Median (2)	SD (3)	Obs. (4)	Mean (5)	Median (6)	SD (7)	Obs. (8)
Panel A: Efficiency/productivity: Quantity and Quality								
Order Acceptance Time, in minutes	3.96	2.98	10.34	69,444	3.27	2.38	8.97	395,273
Pickup Time, in minutes	11.47	10.94	4.13	69,462	12.05	11.55	7.23	395,388
Delivery Time, in minutes	14.42	13.85	4.93	69,498	14.14	13.36	5.13	396,048
Late Deliveries per 10k orders	426.81	147.06	811.01	70,301	414.67	131.93	817.34	409,769
Bad Ratings per 10k orders	11.53	0.00	83.23	70,301	9.30	0.00	95.88	409,769
Week with a Bad Rating, 0 or 1	0.09	0.00	0.29	70,301	0.07	0.00	0.26	409,769
Panel B: Labor Supply on the Platform								
Hours Online, in hours per week	45.02	46.07	29.05	54,984	39.54	36.16	29.13	288,749
Work Days per week	5.57	7.00	2.18	70,301	5.26	6.00	2.12	409,769
Driver attrites from Platform, 0 or 1	0.05	0.00	0.22	66,622	0.08	0.00	0.27	389,570
Panel C: Total Output								
Total Orders	161.10	141.00	126.66	70,301	150.90	124.00	143.86	409,769
Peak Hour Orders	96.19	85.00	77.27	70,301	90.02	74.00	90.16	409,769
Panel D: Income, Wages, and Profits								
Driver Income, in CNY per week	Confid.	Confid.	Confid.	70,301	Confid.	Confid.	Confid.	409,769
Hourly Wages	Confid.	Confid.	Confid.	54,675	Confid.	Confid.	Confid.	287,386
Platform Profits from a driver	Confid.	Confid.	Confid.	70,301	Confid.	Confid.	Confid.	409,769
Workers	4,284				39,321			

Notes: This table is similar to Table 2 in the main text, but provides summary statistics separately for DHH workers and for non-disabled workers. Summary statistics use the full sample period, including both pre- and post-AI observations.

Table A.3: Summary Statistics Comparing Gig vs. Team Workers

	Gig Workers				Team Workers			
	Mean (1)	Median (2)	SD (3)	Obs. (4)	Mean (5)	Median (6)	SD (7)	Obs. (8)
Panel A: Efficiency/productivity: Quantity and Quality								
Order Acceptance Time, in minutes	4.17	3.29	3.65	206,131	2.74	2.01	11.85	258,586
Pickup Time, in minutes	11.06	10.65	3.55	206,201	12.69	12.07	8.57	258,649
Delivery Time, in minutes	16.69	16.24	5.45	206,237	12.19	11.84	3.76	259,309
Late Deliveries per 10k orders	517.33	140.85	981.70	207,084	339.91	132.16	653.86	272,986
Bad Ratings per 10k orders	15.60	0.00	135.09	207,084	5.09	0.00	41.15	272,986
Week with a Bad Rating, 0 or 1	0.07	0.00	0.26	207,084	0.07	0.00	0.26	272,986
Panel B: Labor Supply on the Platform								
Hours Online, in hours per week	22.83	15.55	21.54	172,695	58.17	61.83	24.83	171,038
Work Days per week	4.25	4.00	2.23	207,084	6.11	7.00	1.65	272,986
Driver attrites from Platform, 0 or 1	0.12	0.00	0.32	195,834	0.05	0.00	0.21	260,358
Panel C: Total Output								
Total Orders	73.35	42.00	85.28	207,084	212.35	205.00	146.18	272,986
Peak Hour Orders	42.87	22.00	54.47	207,084	127.38	125.00	91.75	272,986
Panel D: Income, Wages, and Profits								
Driver Income, in CNY per week	Confid.	Confid.	Confid.	207,084	Confid.	Confid.	Confid.	272,986
Hourly Wages	Confid.	Confid.	Confid.	171,184	Confid.	Confid.	Confid.	170,877
Platform Profits from a driver	Confid.	Confid.	Confid.	207,084	Confid.	Confid.	Confid.	272,986
Panel E: Worker Characteristics								
Age		30.8				30.1		
Male		90.2%				92.3%		
Driver is in a big city		34.8%				26.4%		
Workers		27,777				15,828		
DHH Workers		2,769				1,515		
Non-disabled Workers		25,008				14,313		

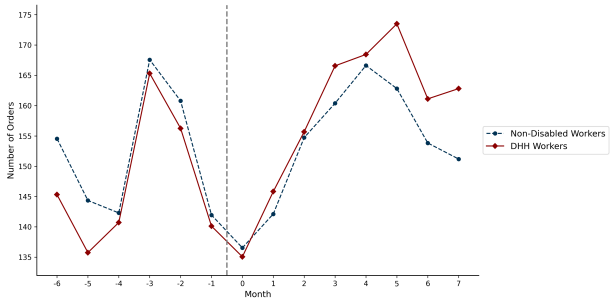
Notes: Panels A–D are similar to Table 2 in the main text, reporting worker-week-level summary statistics separately for gig and team workers. The "Confid." marking indicates that certain statistics cannot be shared due to confidentiality requirements. Panel E reports worker-level characteristics, similar to Table 1 in the main text, averaged across workers within each group. Summary statistics use the full sample period, including both pre- and post-AI observations.

Table A.4: Summary Statistics Comparing Profound DHH Workers vs. Non-Profound DHH Workers

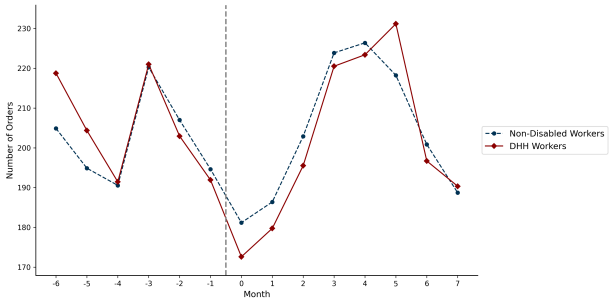
	Profound DHH Workers				Non-profound DHH Workers			
	Mean (1)	Median (2)	SD (3)	Obs. (4)	Mean (5)	Median (6)	SD (7)	Obs. (8)
Panel A: Efficiency/productivity: Quantity and Quality								
Order Acceptance Time, in minutes	3.72	3.11	2.97	34,218	4.19	2.85	14.22	35,226
Pickup Time, in minutes	11.69	11.15	3.74	34,224	11.25	10.75	4.46	35,238
Delivery Time, in minutes	14.98	14.30	4.83	34,244	13.87	13.33	4.97	35,254
Late Deliveries per 10k orders	448.67	161.29	827.42	34,444	405.80	132.45	794.38	35,857
Bad Ratings per 10k orders	12.29	0.00	91.38	34,444	10.80	0.00	74.55	35,857
Week with a Bad Rating, 0 or 1	0.10	0.00	0.30	34,444	0.08	0.00	0.28	35,857
Panel B: Labor Supply on the Platform								
Hours Online, in hours per week	46.47	50.20	28.88	26,430	43.68	41.98	29.15	28,554
Work Days per week	5.52	7.00	2.08	34,444	5.63	7.00	2.26	35,857
Driver attrites from Platform, 0 or 1	0.06	0.00	0.23	32,759	0.04	0.00	0.21	33,863
Panel C: Total Output								
Total Orders	170.65	160.00	126.29	34,444	151.92	122.00	126.35	35,857
Peak Hour Orders	104.88	100.00	78.88	34,444	87.85	71.00	74.75	35,857
Panel D: Income, Wages, and Profits								
Driver Income, in CNY per week	Confid.	Confid.	Confid.	34,444	Confid.	Confid.	Confid.	35,857
Hourly Wages	Confid.	Confid.	Confid.	25,068	Confid.	Confid.	Confid.	27,207
Platform Profits from a driver	Confid.	Confid.	Confid.	34,444	Confid.	Confid.	Confid.	35,857
Panel E: Worker Characteristics								
Age		31.1				34.9		
Male		90.4%				95.0%		
Driver is in a big city		42.2%				31.4%		
Workers		2,295				1,989		

Notes: Panels A–D are similar to Table 2 in the main text, reporting worker-week-level summary statistics separately for DHH workers with profound hearing loss (Level 1; >90 dB) and non-profound hearing loss (Levels 2–4), who together constitute the full DHH sample. The "Confid." marking indicates that certain statistics cannot be shared due to confidentiality requirements. Panel E reports worker-level characteristics, similar to Table 1 in the main text, averaged across workers within each group. Summary statistics use the full sample period, including both pre- and post-AI observations.

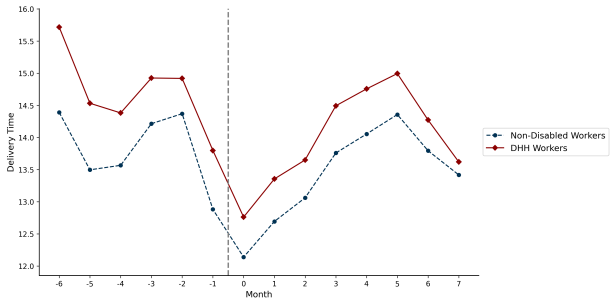
Figure A.1: Comparison of Raw Means Between DHH and Non-disabled Workers



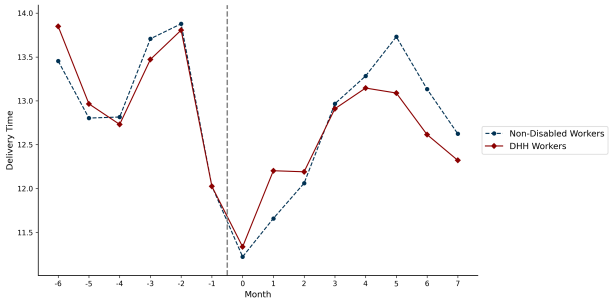
(a) Orders Completed (All Workers)



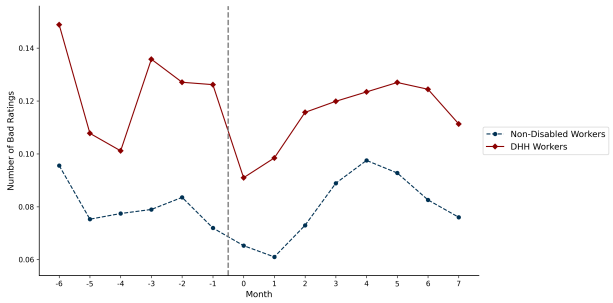
(b) Orders Completed (Balanced Panel)



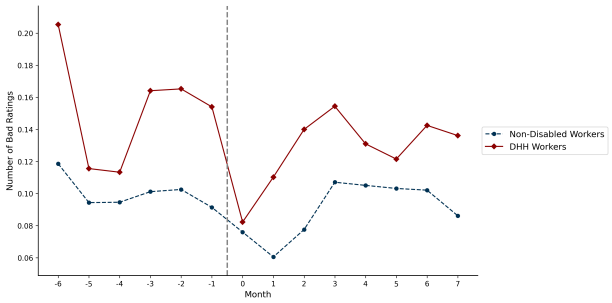
(c) Delivery Time (All Workers)



(d) Delivery Time (Balanced Panel)



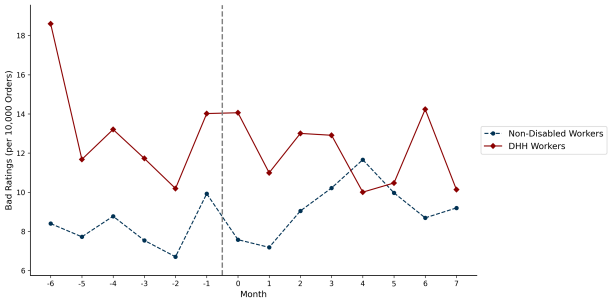
(e) Bad Ratings (Count) (All Workers)



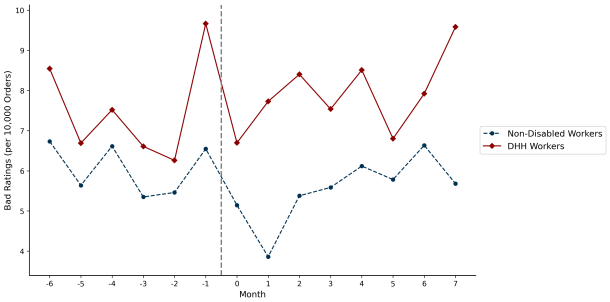
(f) Bad Ratings (Count) (Balanced Panel)

Figure continues on the next page.

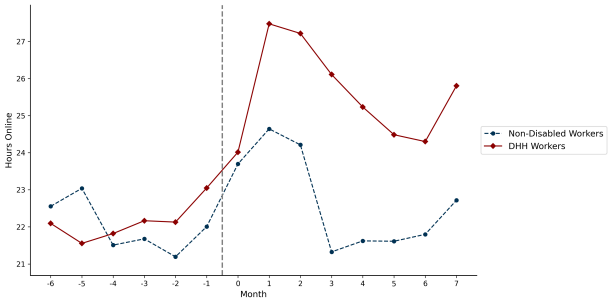
Figure A.1: Comparison of Raw Means Between DHH and Non-disabled Workers (*continued*)



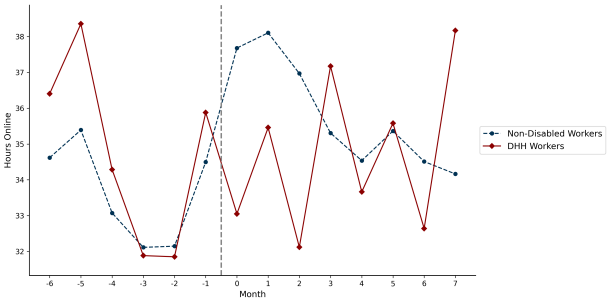
(g) Bad Ratings (per 10,000 Orders) (All Workers)



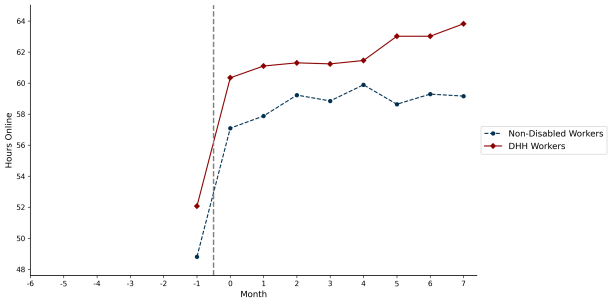
(h) Bad Ratings (per 10,000 Orders) (Balanced Panel)



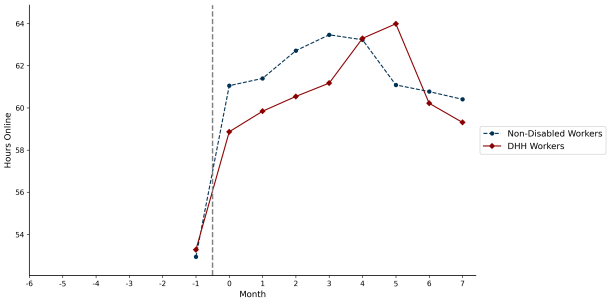
(i) Hours Online – Gig Workers (All Workers)



(j) Hours Online – Gig Workers (Balanced Panel)



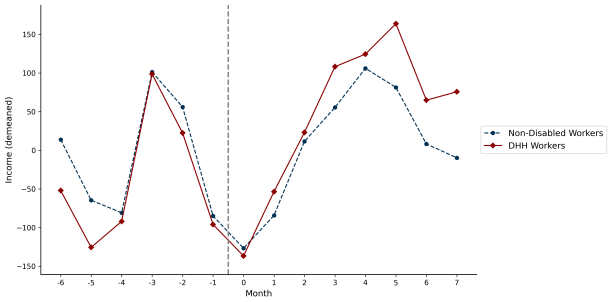
(k) Hours Online – Team Workers (All Workers)



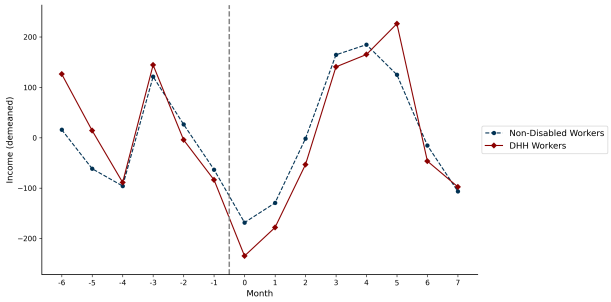
(l) Hours Online – Team Workers (Balanced Panel)

Figure continues on the next page.

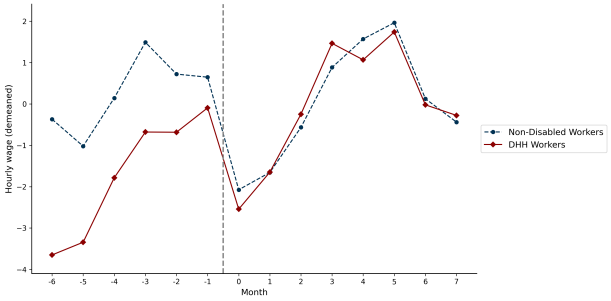
Figure A.1: Comparison of Raw Means Between DHH and Non-disabled Workers (*continued*)



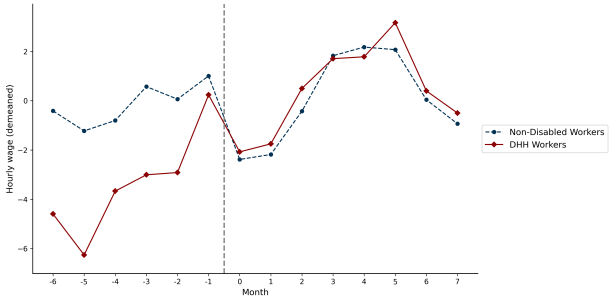
(m) Income
(All Workers)



(n) Income
(Balanced Panel)



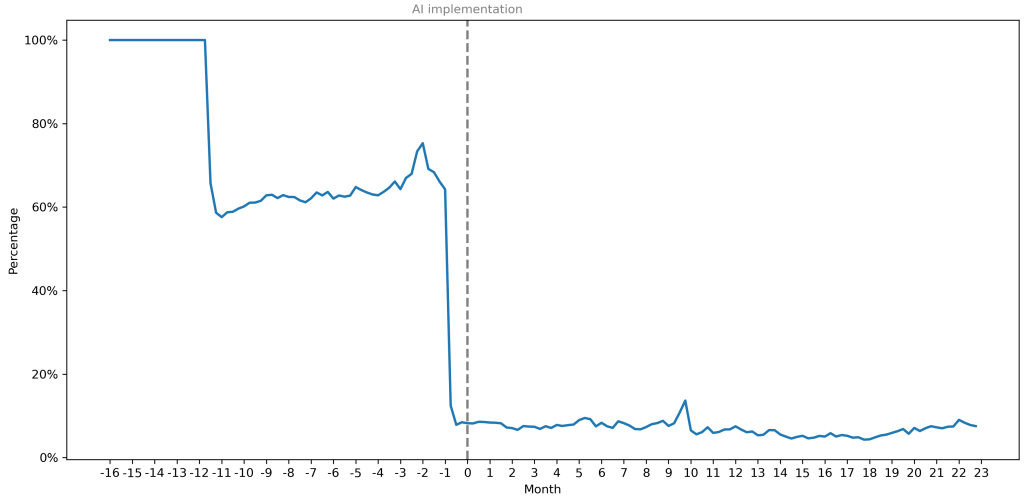
(o) Hourly Wage
(All Workers)



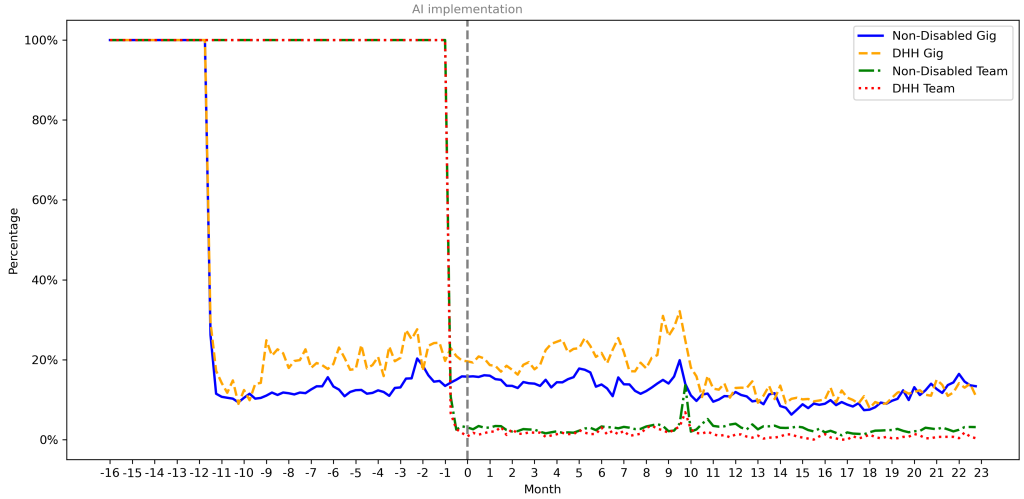
(p) Hourly Wage
(Balanced Panel)

Notes: This figure compares raw means of outcomes for DHH workers (red) and non-disabled workers (dark blue, dashed). The left panels show results for all workers. There are 43,605 workers in total, including 4,284 DHH workers and 39,321 non-disabled workers. The right panels show results for workers who remained continuously active on the platform (i.e., did not exit) from month -6 before the AI implementation to month 7 after the AI implementation. This sample includes 779 workers in total, with 107 DHH workers and 672 non-disabled workers. Bad ratings are shown in two ways (following footnote 20 in the main text) as a weekly count (panels (e) and (f)), consistent with the main analysis, and as bad ratings per 10,000 orders (panels (g) and (h)), which normalizes by order volume. Since these are raw means, we show Hours Online separately for gig and team workers (panels (i)–(l)). Hours data become available for gig workers at week $t_0 - 46$ and for team workers only at week $t_0 - 4$, where t_0 is the week of AI implementation (see Figure A.2). Showing the two groups in a single pooled series would therefore conflate differences in outcome levels with differences in coverage periods. This is not a concern for our DiD analyses, where work model fixed effects ensure the comparison is always made within work model. The apparent increase in hours around t_0 for team workers (panels (k) and (l)) reflects that hours data are only available for the final few weeks of month -1, so the pre-AI data point is an average over an incomplete month rather than a true behavioral change. For confidentiality, the income variable (panels (m) and (n)) and the hourly wage variable (panels (o) and (p)) are demeaned using the sample average, with the same value subtracted for both DHH and non-disabled workers. Income is in CNY per week and hourly wage is in CNY per hour.

Figure A.2: The Share of Worker-Weeks with Missing Hours Data Over Time



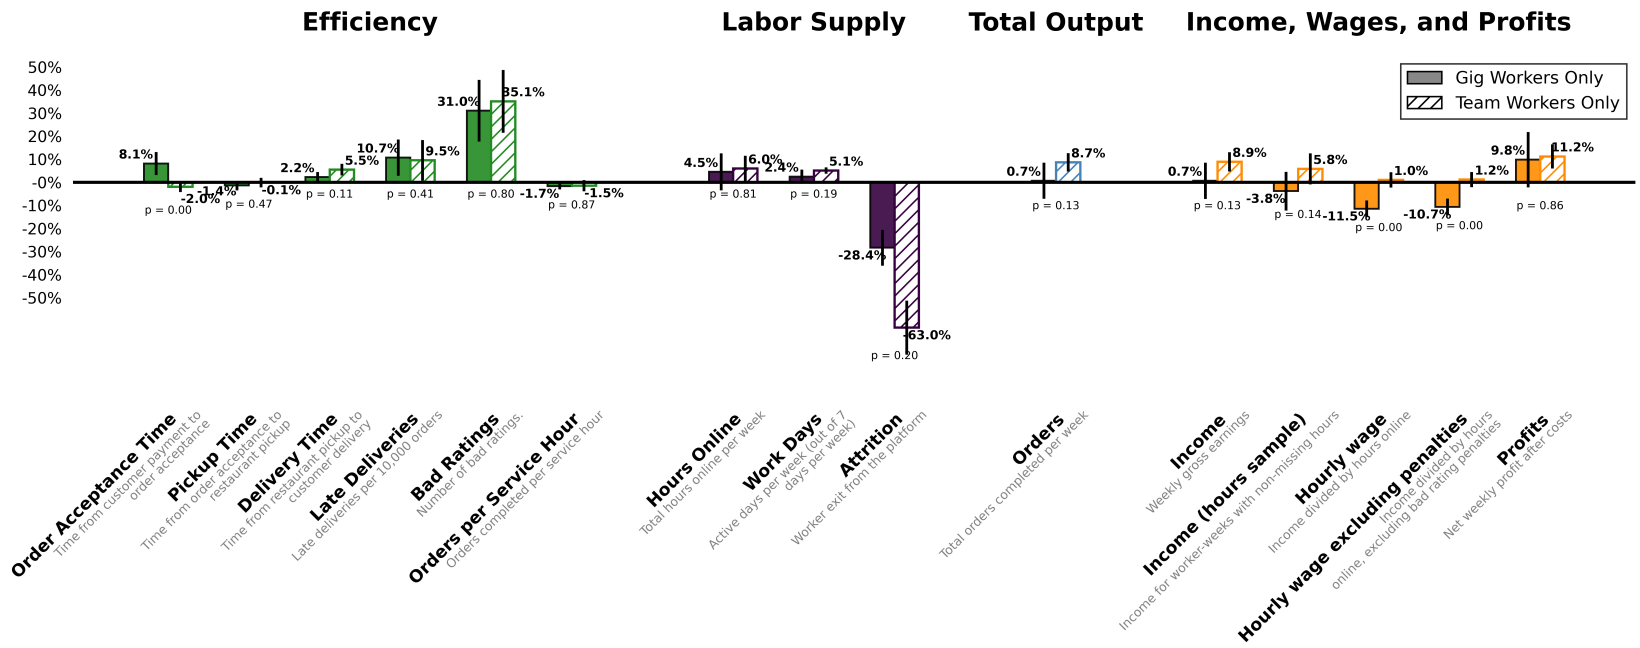
(a) Share of worker-weeks with missing hours over time



(b) Share of worker-weeks with missing hours, split by work model and DHH status

Notes: This figure shows the share of worker-weeks with missing hours data over time. The dotted vertical line shows when the AI tool is implemented. Panel (a) shows the overall share of missing hours data over time. As seen in panel (b), which shows missing rate broken down by work model (gig or team) and DHH status, hours data become available for gig workers at week $t_0 - 46$, and then for team workers at week $t_0 - 4$, where t_0 is the week of AI implementation.

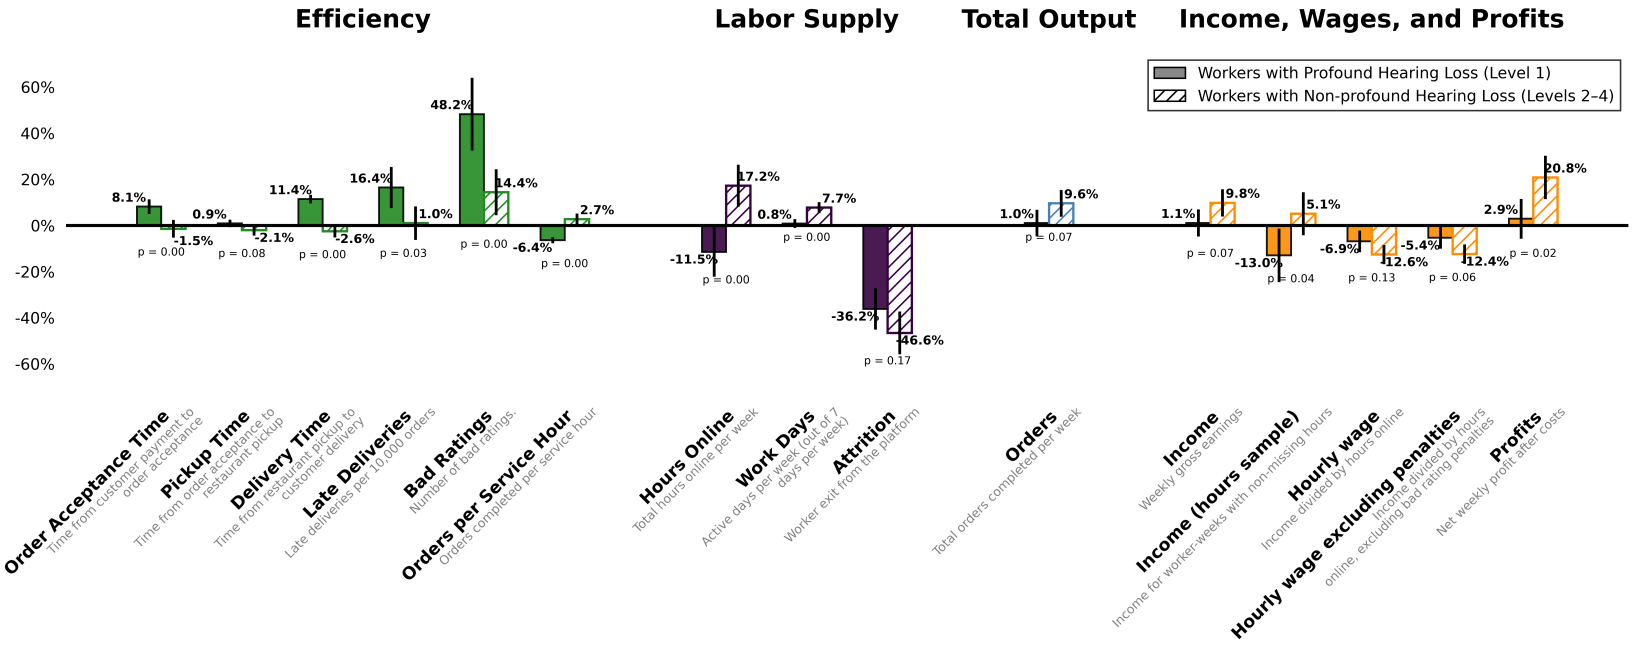
Figure A.3: Pre-AI Differences Between DHH and Non-disabled Workers: Gig and Team Workers Separately



Notes: This figure reproduces the analysis in Figure 1, but splits the sample by gig workers and team workers. As discussed in the main text, gig workers operate individually, exercise full discretion over their working hours, and decide which orders to accept. We plot pre-AI differences for DHH gig workers vs. non-disabled gig workers, as well as pre-AI differences for DHH team workers vs. non-disabled team workers. Each bar is estimated from a separate regression in the corresponding subsample (gig or team workers). The p-values shown are from tests of equality between gig and team workers. For example, looking at orders, among gig workers, DHH workers do 0.7% more orders than non-disabled workers, whereas among team workers, DHH workers do 8.7% more orders than non-disabled workers, and the difference is not statistically significant ($p = 0.13$).

The p-values are based on pooled regressions with interaction terms that test equality of the underlying OLS coefficients across groups. For outcomes estimated using IHS transformation (orders, delivery time, hours online, income), the tests are conducted on IHS coefficients; for outcomes estimated in levels (late deliveries, bad ratings, and attrition), the tests are conducted on level coefficients. Because the plotted coefficients are scaled by group-specific pre-AI means, the p-values need not correspond exactly to visual differences in the bars.

Figure A.4: Pre-AI Differences Between DHH and Non-disabled Workers: Splitting by Severity of Disability



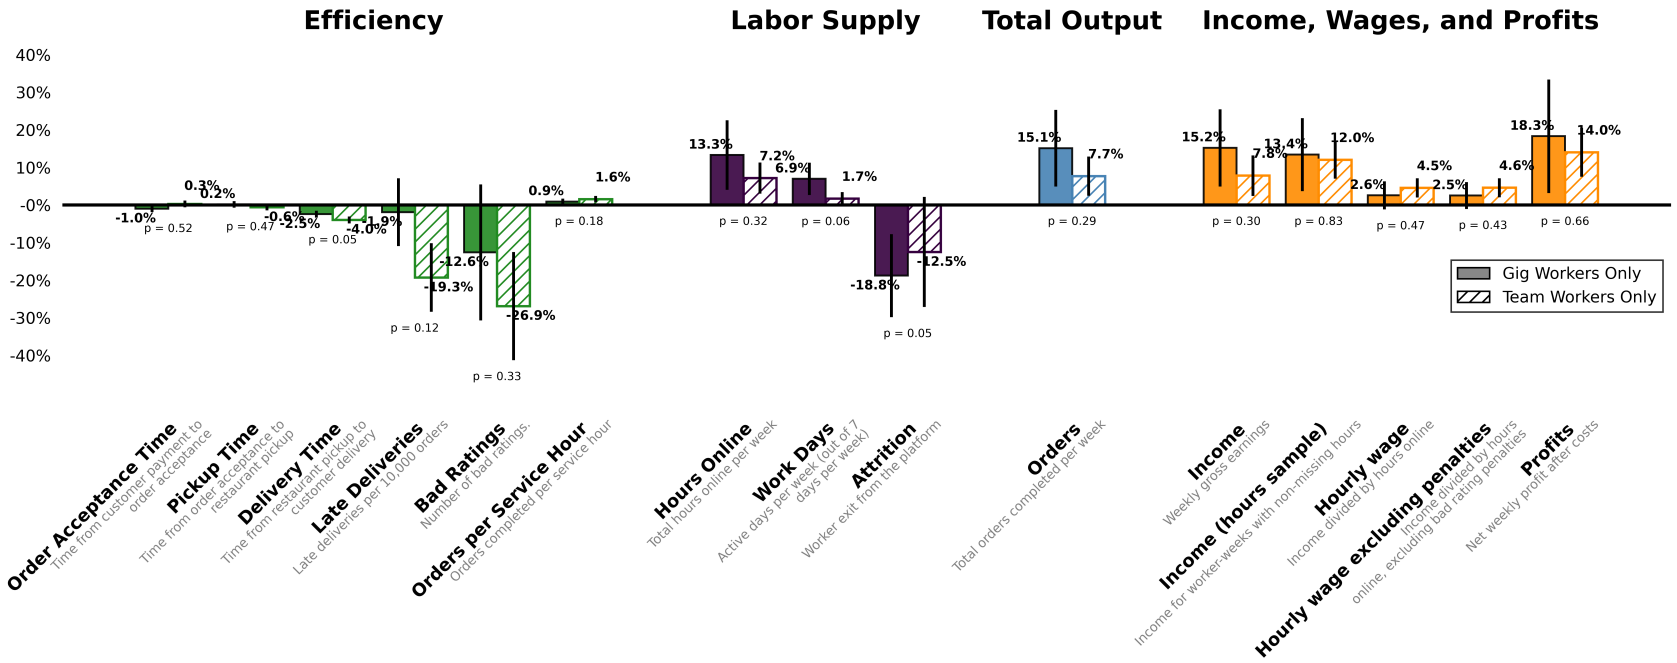
Notes: This figure reproduces the analysis in Figure 1, but splits DHH differences by severity of disability. As discussed in Section 1 of the main text, profound hearing loss (Level 1 in the Chinese system) is loss of more than 90 decibels. Non-profound hearing loss is loss of 41-90 decibels. The p-values shown are from tests of equality between profound and non-profound HL workers. We divide the raw coefficients by the pre-AI DHH mean within each severity level to obtain the percentages shown.

Table A.5: Pre-AI Differences Between DHH and Non-disabled Workers

	Baseline (1)	Gig Workers (2)	Team Workers (3)	Control for Avg. Distance from Restaurant to Customer (4)
Panel A: Efficiency				
Order Acceptance Time	0.034**	0.081***	-0.020	0.032**
	(0.016)	(0.030)	(0.014)	(0.015)
Pickup Time	-0.006	-0.014	-0.001	-0.007
	(0.009)	(0.013)	(0.013)	(0.009)
Delivery Time	0.045***	0.022	0.055***	0.041***
	(0.011)	(0.014)	(0.015)	(0.010)
Late Deliveries	0.087**	0.107**	0.095*	0.087**
	(0.035)	(0.048)	(0.053)	(0.035)
Bad Ratings	0.313***	0.310***	0.351***	0.313***
	(0.058)	(0.081)	(0.083)	(0.058)
Orders per Service Hour	-0.019**	-0.017**	-0.015	-0.017**
	(0.009)	(0.008)	(0.015)	(0.009)
Panel B: Labor Supply				
Hours Online	0.049	0.045	0.060*	0.049
	(0.044)	(0.049)	(0.034)	(0.044)
Work Days	0.043***	0.024	0.051***	0.043***
	(0.010)	(0.019)	(0.008)	(0.010)
Attrition	-0.414***	-0.284***	-0.630***	-0.414***
	(0.041)	(0.047)	(0.071)	(0.041)
Panel C: Total Output				
Orders	0.053**	0.007	0.087***	0.053**
	(0.026)	(0.048)	(0.024)	(0.026)
Panel D: Income, Wages, and Profits				
Income	0.054**	0.007	0.089***	0.054**
	(0.026)	(0.048)	(0.026)	(0.026)
Hourly Wage	-0.102***	-0.115***	0.010	-0.088***
	(0.020)	(0.022)	(0.020)	(0.019)
Hourly Wage Excluding Penalties	-0.094***	-0.107***	0.012	-0.080***
	(0.020)	(0.022)	(0.020)	(0.019)
Profits	0.118***	0.098	0.112***	0.118***
	(0.040)	(0.073)	(0.032)	(0.040)
Maximum worker-weeks in column	180,375	77,016	103,359	180,375
Workers	16,912	10,019	6,893	16,912

Notes: Standard errors clustered by worker in parentheses. Each coefficient here is from a separate regression. Each regression follows the procedure in Section 2.1. For all outcomes besides late deliveries, bad ratings, and attrition, we use an inverse-hyperbolic sine transform on y. For late deliveries, bad ratings, and attrition, the coefficients shown here are the OLS regression coefficients divided by the pre-AI DHH mean, so that the coefficients here can be interpreted as percentage differences. Significance levels: * p<0.1, ** p<0.05, *** p<0.01. Column 1 is the baseline specification using the full panel. Column 2 restricts the sample to gig workers only. Column 3 restricts the sample to team workers only. Column 4 uses the same specification as Column 1 but additionally controls for the average distance from restaurant to customer, where average distance is measured at the worker-week level.

Figure A.5: Impact of the AI Tool on DHH Worker Outcomes: Gig and Team Workers Separately



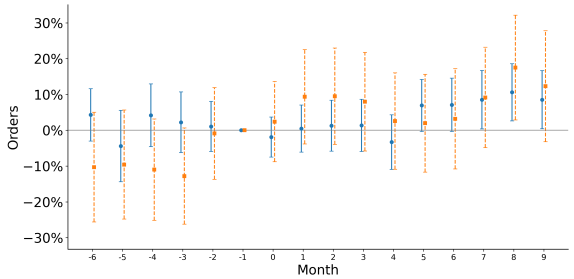
Notes: This figure reproduces the analysis in Figure 3, focusing exclusively on gig or team workers. As discussed in the main text, gig workers operate individually, exercise full discretion over their working hours, and decide which orders to accept. The p-values shown are from tests of equality between gig and team workers in the impact of the AI tool. For example, the AI tool raises the orders of gig workers by 15.1%, raises the orders of team workers by 7.7%, and the difference is statistically insignificant ($p = 0.29$).

Table A.6: The Impact of the AI Tool on the Outcomes of DHH Workers: Extended Results

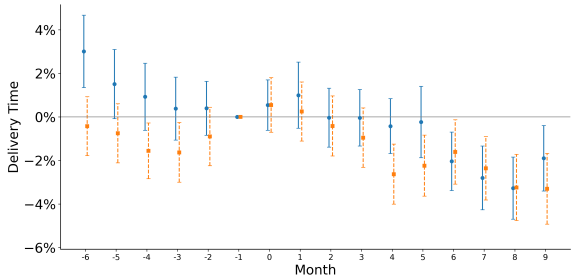
	Baseline (1)	Gig Workers (2)	Team Workers (3)	Exclude First 4 Months Post-AI (4)	Workers Present Both Pre-AI & Post-AI (5)	Restrict to Observations with Hours Data (6)	Control for Avg. Distance from Restaurant to Customer (7)
Panel A: Efficiency							
Order Acceptance Time	-0.008**	-0.010*	0.003	-0.010*	-0.006	-0.001	-0.009**
	(0.004)	(0.005)	(0.005)	(0.006)	(0.004)	(0.004)	(0.004)
Pickup Time	-0.004	0.002	-0.006	-0.005	-0.006	0.001	-0.005
	(0.004)	(0.005)	(0.005)	(0.005)	(0.004)	(0.004)	(0.004)
Delivery Time	-0.035***	-0.025***	-0.040***	-0.051***	-0.035***	-0.026***	-0.037***
	(0.004)	(0.006)	(0.006)	(0.006)	(0.004)	(0.005)	(0.004)
Late Deliveries	-0.112***	-0.019	-0.193***	-0.148***	-0.135***	-0.011	-0.112***
	(0.038)	(0.055)	(0.055)	(0.044)	(0.045)	(0.036)	(0.038)
Bad Ratings	-0.218***	-0.126	-0.269***	-0.252***	-0.197***	-0.086	-0.218***
	(0.068)	(0.110)	(0.088)	(0.083)	(0.063)	(0.073)	(0.068)
Orders per Service Hour	0.016***	0.009*	0.016***	0.021***	0.015***	0.009**	0.017***
	(0.004)	(0.005)	(0.005)	(0.005)	(0.004)	(0.004)	(0.004)
Panel B: Labor Supply							
Hours Online	0.106***	0.133**	0.072***	0.136***	0.100***	0.106***	0.106***
	(0.034)	(0.056)	(0.025)	(0.046)	(0.035)	(0.034)	(0.034)
Work Days	0.035***	0.069***	0.017	0.040***	0.033***	0.051***	0.035***
	(0.012)	(0.026)	(0.011)	(0.015)	(0.012)	(0.015)	(0.012)
Attrition	-0.210***	-0.188***	-0.125	-0.193***	Not app-	-0.119**	-0.210***
	(0.052)	(0.067)	(0.089)	(0.054)	licable	(0.053)	(0.052)
Panel C: Total Output							
Orders	0.105***	0.151**	0.077**	0.129***	0.101***	0.128***	0.105***
	(0.031)	(0.062)	(0.032)	(0.038)	(0.031)	(0.036)	(0.031)
Panel D: Income, Wages, and Profits							
Income	0.106***	0.152**	0.078**	0.132***	0.102***	0.130***	0.106***
	(0.031)	(0.062)	(0.033)	(0.039)	(0.032)	(0.036)	(0.031)
Hourly Wage	0.036**	0.026	0.046***	0.027	0.034**	0.036**	0.036**
	(0.014)	(0.023)	(0.016)	(0.019)	(0.015)	(0.014)	(0.014)
Hourly Wage Excluding Penalties	0.036***	0.025	0.046***	0.033*	0.035**	0.036***	0.036***
	(0.014)	(0.022)	(0.015)	(0.019)	(0.014)	(0.014)	(0.014)
Profits	0.168***	0.183**	0.140***	0.244***	0.184***	0.200***	0.168***
	(0.043)	(0.092)	(0.039)	(0.051)	(0.040)	(0.058)	(0.043)
Maximum worker-weeks in column	480,070	207,084	272,986	427,612	184,450	343,733	480,070
Workers	43,605	27,777	15,828	40,821	3,397	34,965	43,605

Notes: This table provides extended results on the impact of the AI tool on DHH workers. Each coefficient reported is from a separate DiD regression. Standard errors clustered by worker in parentheses. Significance levels: * p<0.1, ** p<0.05, *** p<0.01. N varies by outcome; see Table 2 for outcome-specific observations. Column 1 is the baseline specification that appears in Figure 3. Column 2 restricts the sample to gig workers only. Column 3 restricts the sample to team workers only. Column 4 excludes the first four months of data after the introduction of the AI tool. Column 5 restricts the sample to workers who have at least one week of data both pre- and post-AI; we do not analyze attrition in column 5 because we are analyzing a sample that is present both before and after the AI. Column 6 restricts to observations with hours data. Column 7 adds a control for average distance between restaurant and customer, where average distance is measured at the worker-week level.

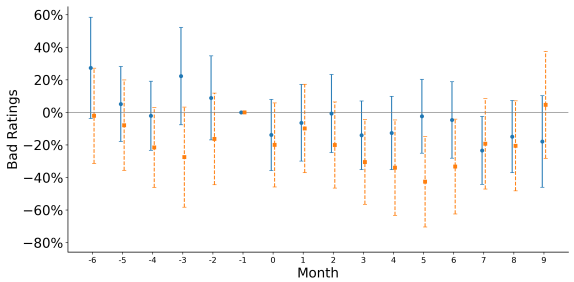
Figure A.6: Event Study Analysis: Splitting by Gig vs. Team Workers



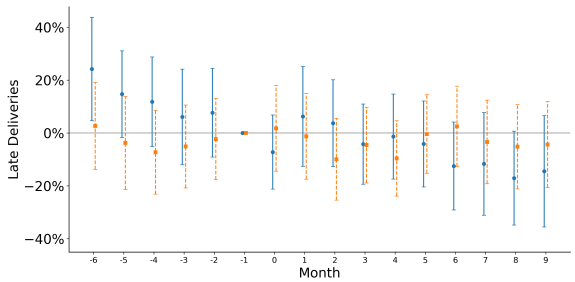
(a) Number of Orders Completed



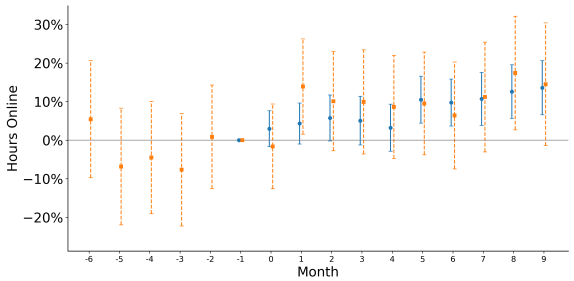
(b) Delivery Time



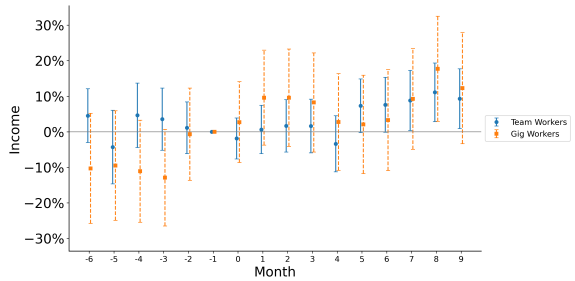
(c) Bad Ratings



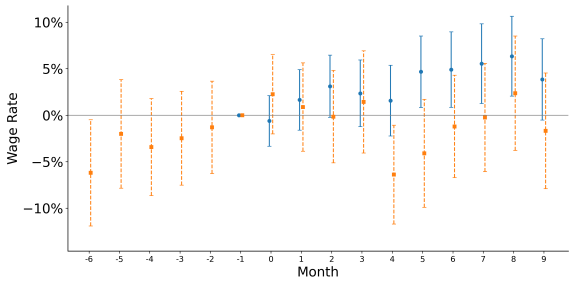
(d) Late Deliveries



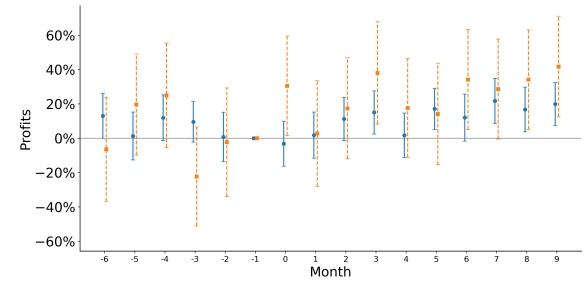
(e) Hours Worked



(f) Income



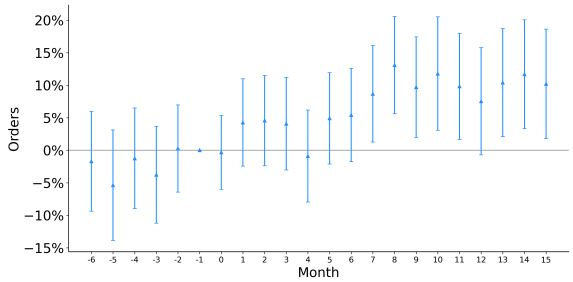
(g) Hourly Wage



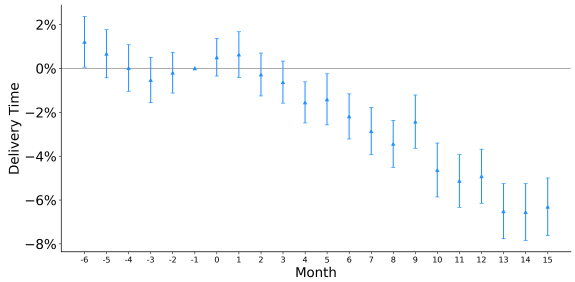
(h) Profit

Notes: This figure replicates the event studies in Figure 4 and Figure 5, but conducts the analysis separately for team workers and gig workers. Dots show the event-study coefficients, and whiskers show 95% confidence intervals clustered at the worker level. The blue dots are for team workers, and the orange squares are for gig workers. Hours worked data are missing for team workers prior to 1 month before the AI tool introduction.

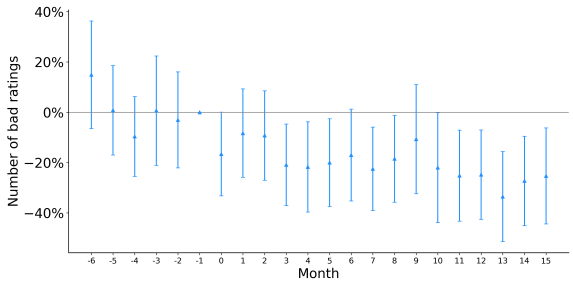
Figure A.7: Event Study Analysis: Robustness to Extended Post-Period



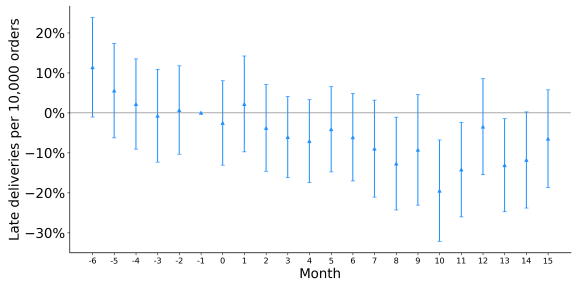
(a) Number of Orders Completed



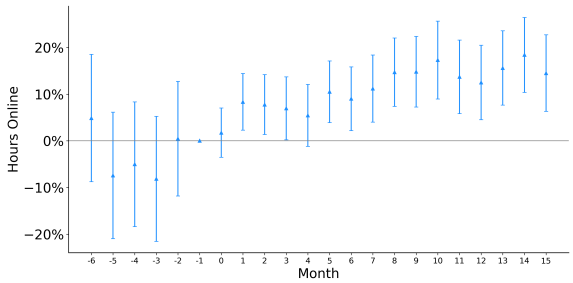
(b) Delivery Time



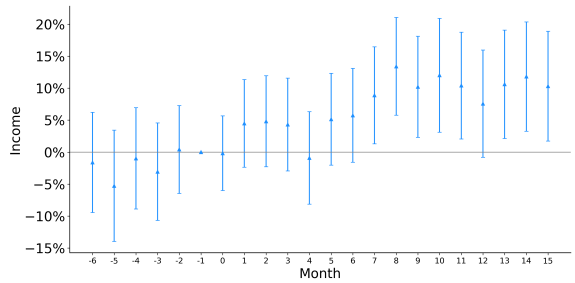
(c) Bad Ratings



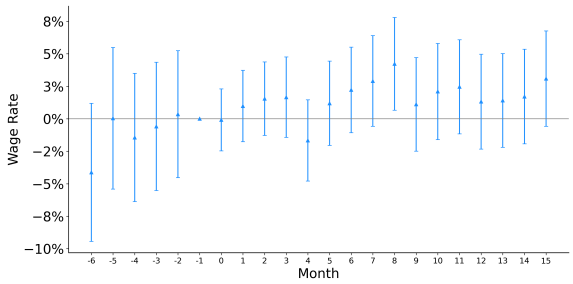
(d) Late Deliveries



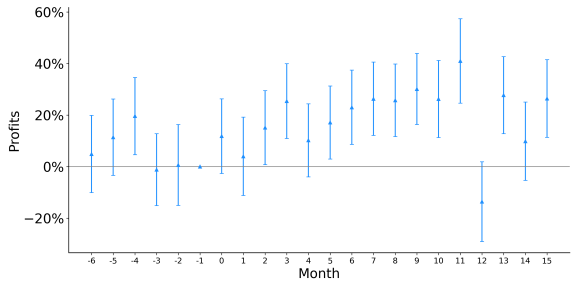
(e) Hours Worked



(f) Income



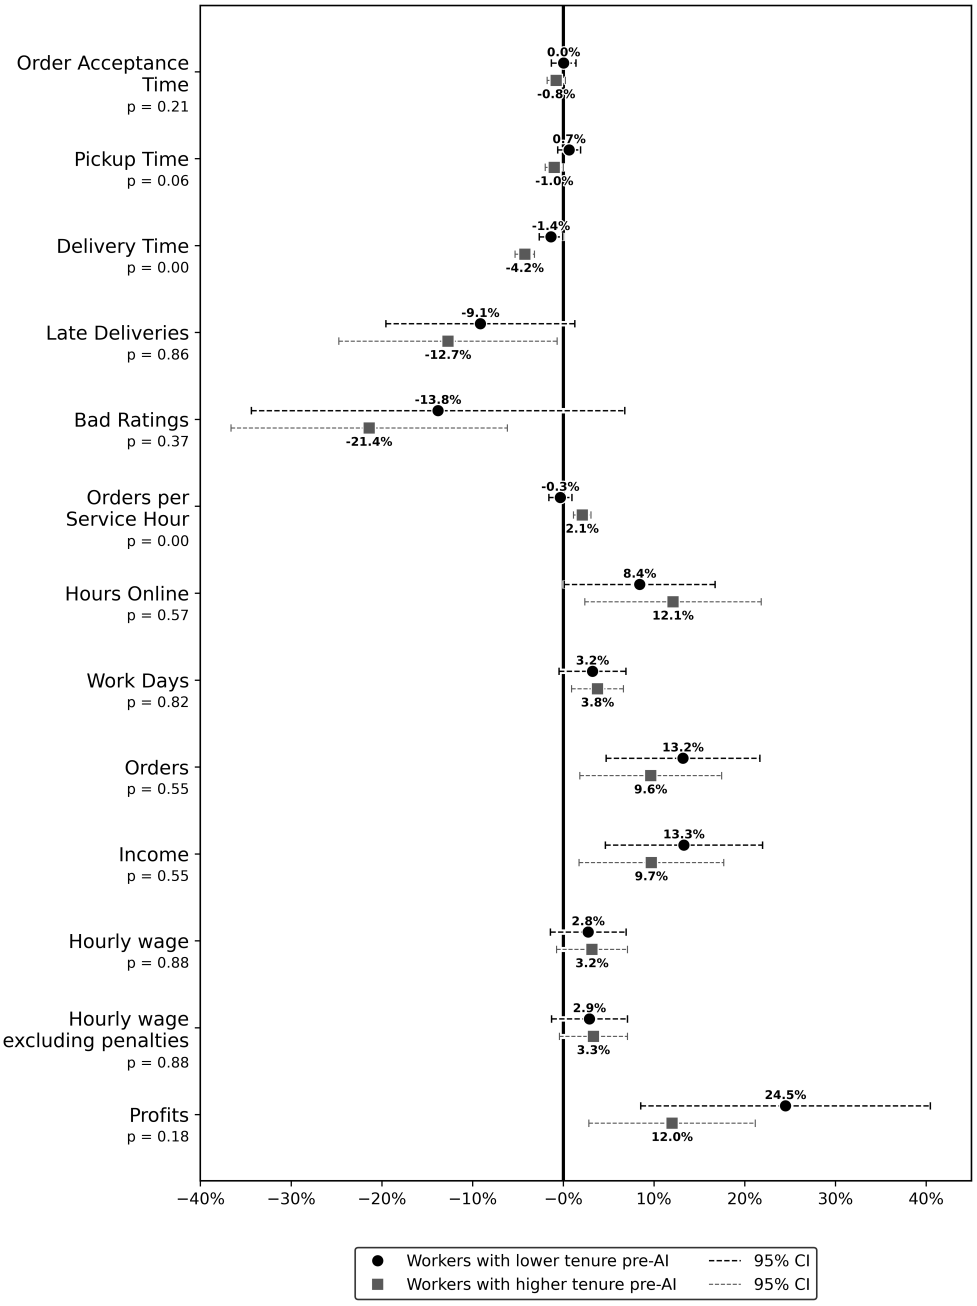
(g) Hourly Wage



(h) Profit

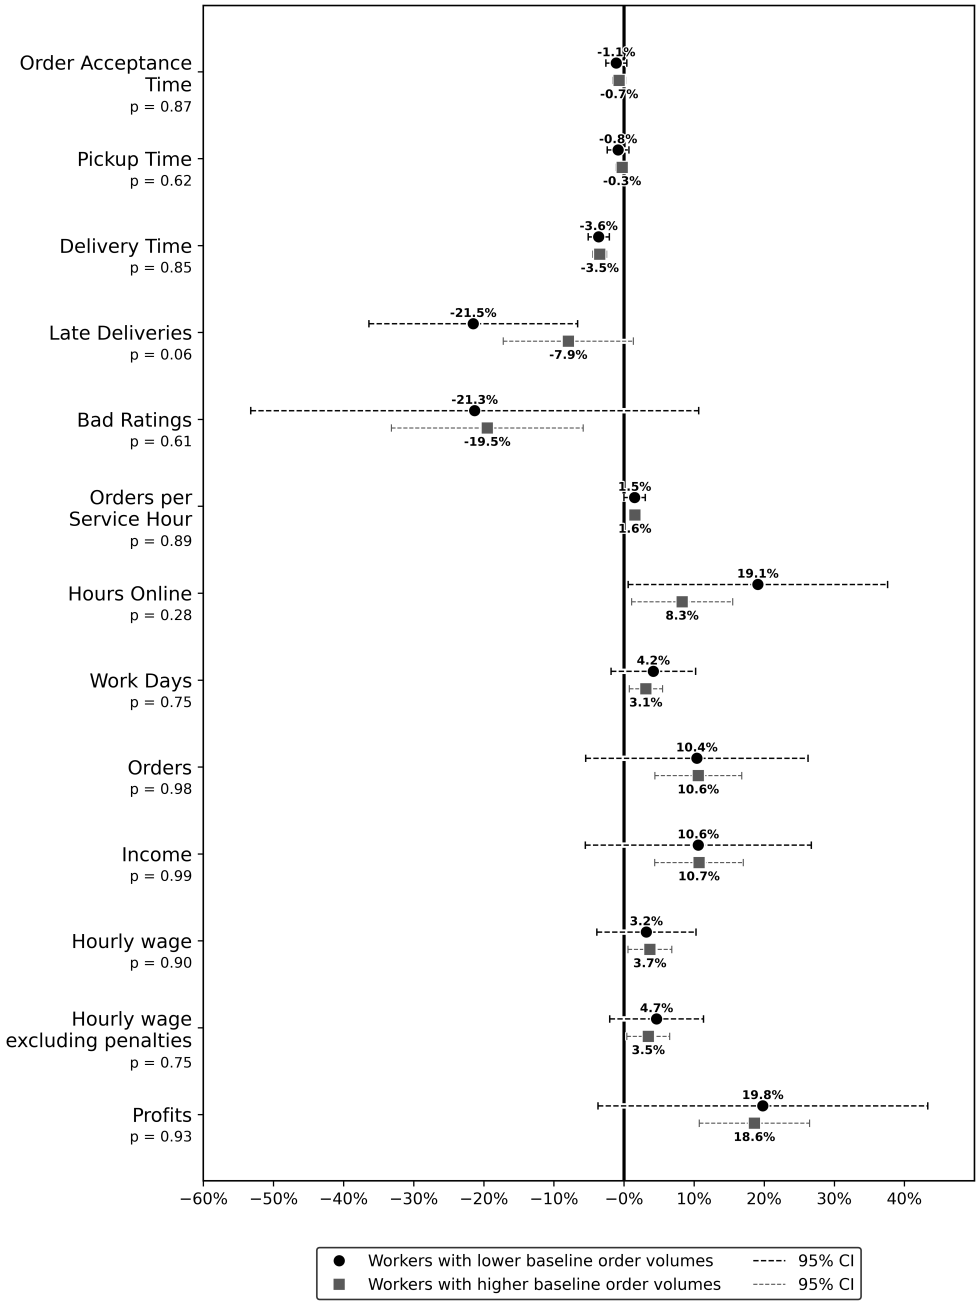
Notes: This figure replicates the event studies in Figure 4 and Figure 5 by extending the analysis through month 15 post-adoption (months 0–15). Dots show the event study coefficients, and whiskers show the 95% confidence intervals accounting for clustering by driver.

Figure A.8: Heterogeneity in Effects of AI Tool by Pre-AI Tenure



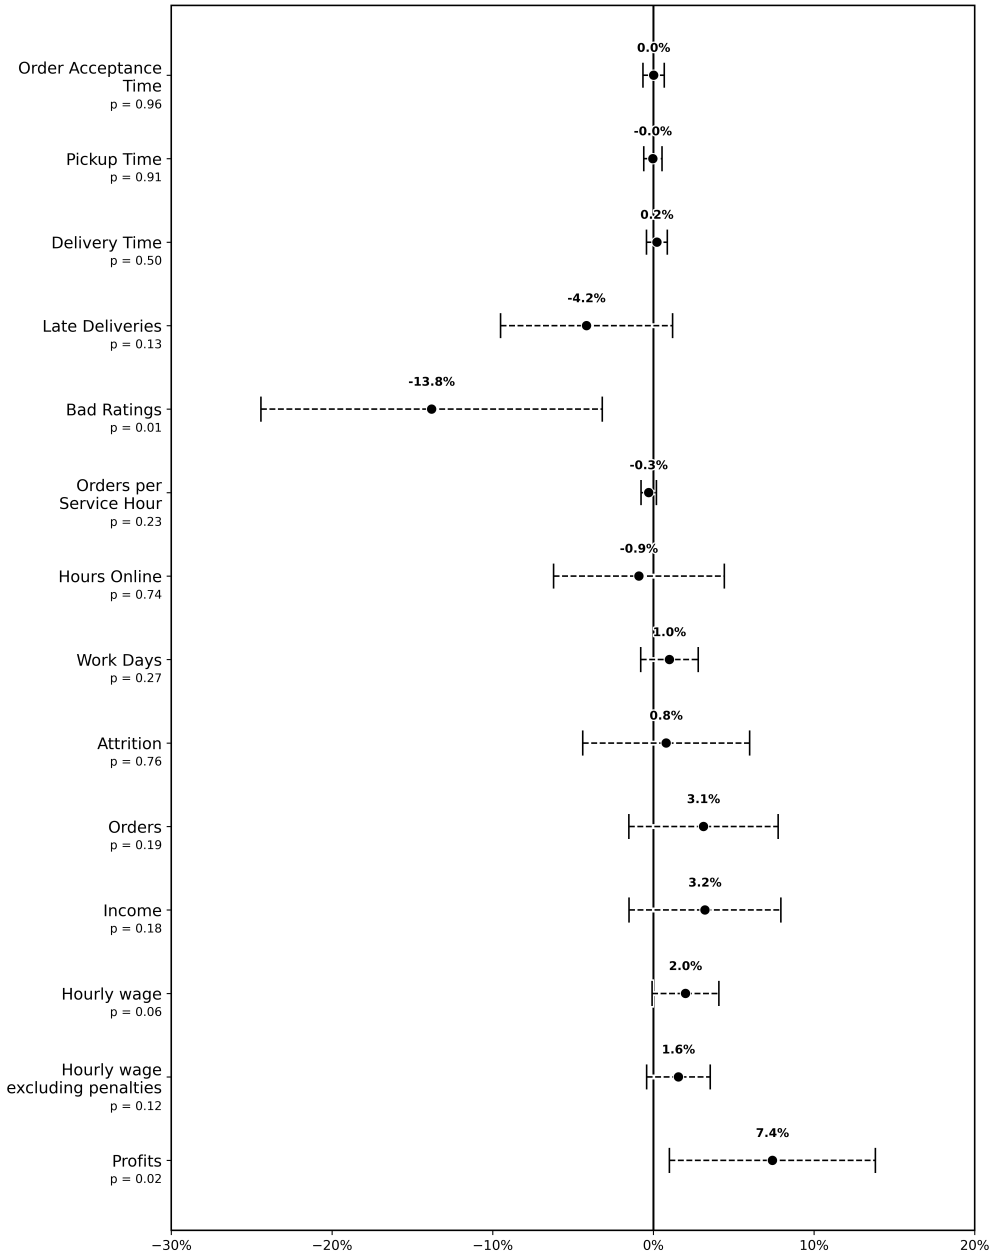
Notes: This figure presents heterogeneous treatment effects of the AI tool by pre-AI worker tenure. For each outcome, we estimate our main DiD model separately for workers within that group. The p-values shown are tests for equality of treatment effects between groups and are estimated using a standard triple-interaction model (including all three double interactions). The p-value is a test of whether the coefficient on $DHH_i \times Post_t \times z_i$ equals zero, where z_i is the heterogeneity characteristic of interest, in this case, pre-AI tenure. Pre-AI tenure is categorized as Low or High using a median split across workers, with tenure measured as of the date of AI implementation, t_0. The sample consists of workers who are present at t_0. There are observations from 482 DHH workers and 2,995 non-disabled workers.

Figure A.9: Heterogeneity in Effects of AI Tool by Pre-AI Productivity



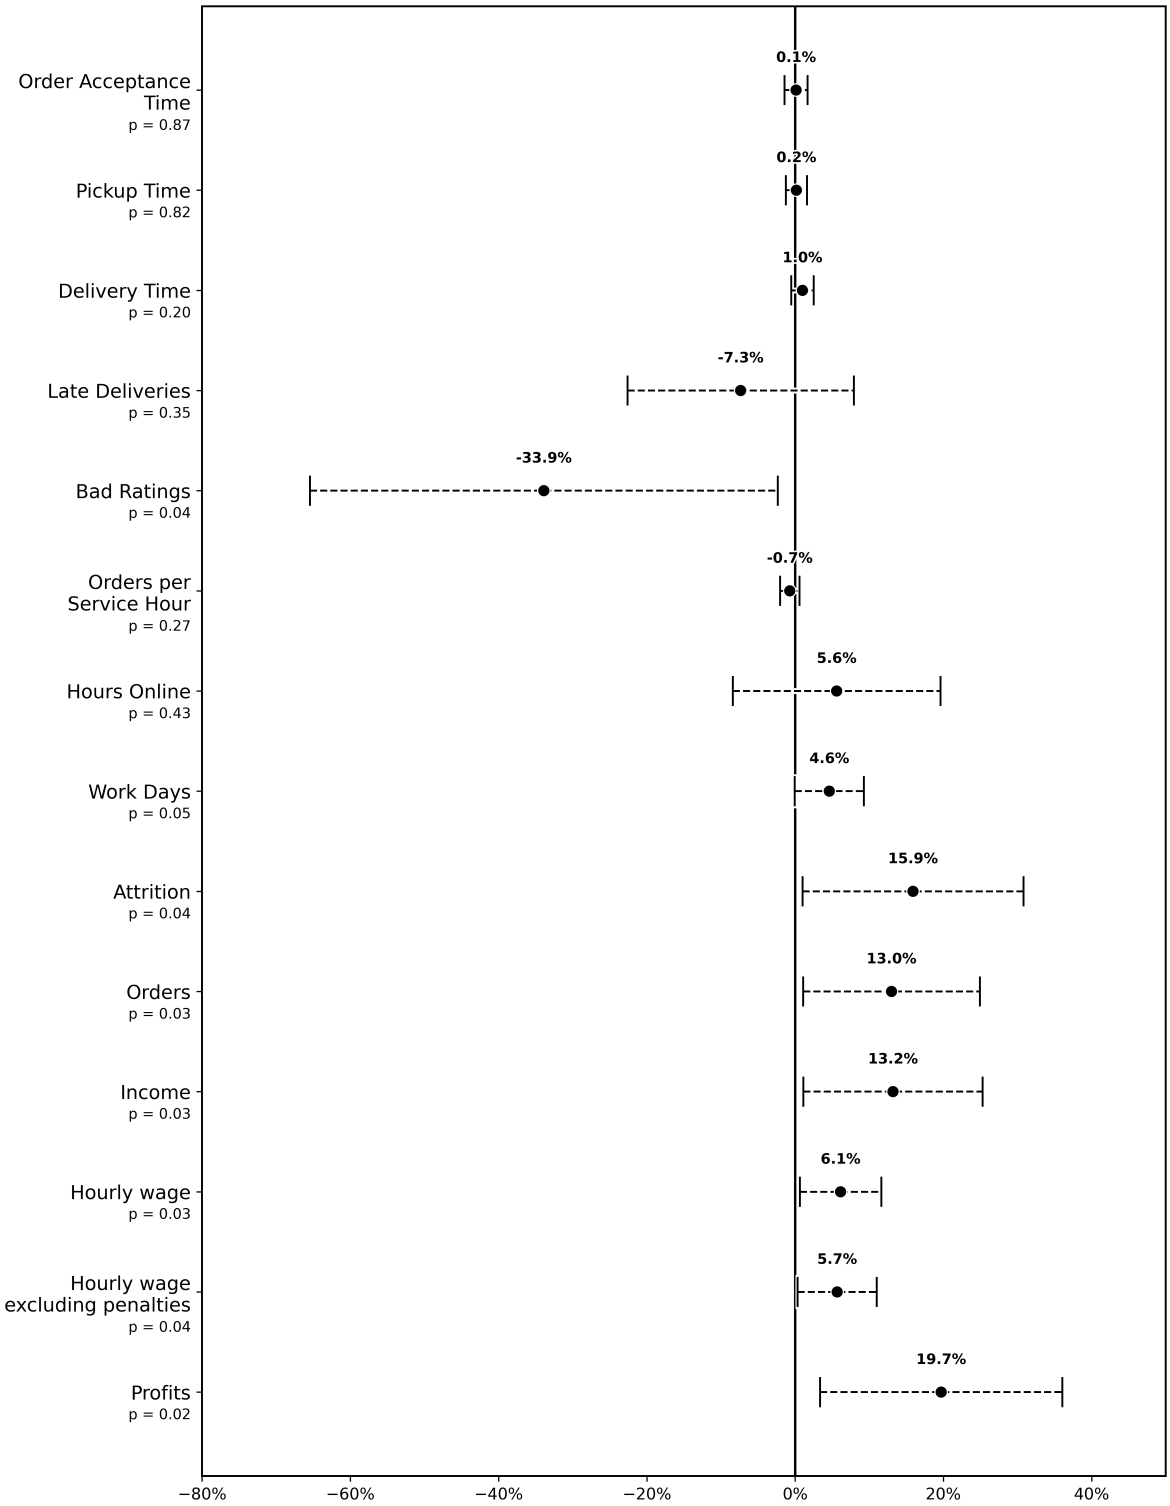
Notes: This figure is similar to Figure A.8, but splits by pre-AI productivity. To construct the productivity measure, we first estimate worker fixed effects from a regression of the inverse hyperbolic sine of weekly orders on worker and year-week fixed effects, using only pre-AI observations. We then regress these estimated worker fixed effects on an indicator for team (versus gig) configuration, so that the resulting residuals capture productivity variation net of systematic differences between team and gig workers. Workers are classified as high- (low-) productivity if their residual falls above (below) the sample median. The sample consists of worker-weeks from 1,515 DHH and 15,397 non-disabled workers.

Figure A.10: Robustness on Heterogeneity in the Effects of AI Tool by Disability Severity: Dosage Response Model



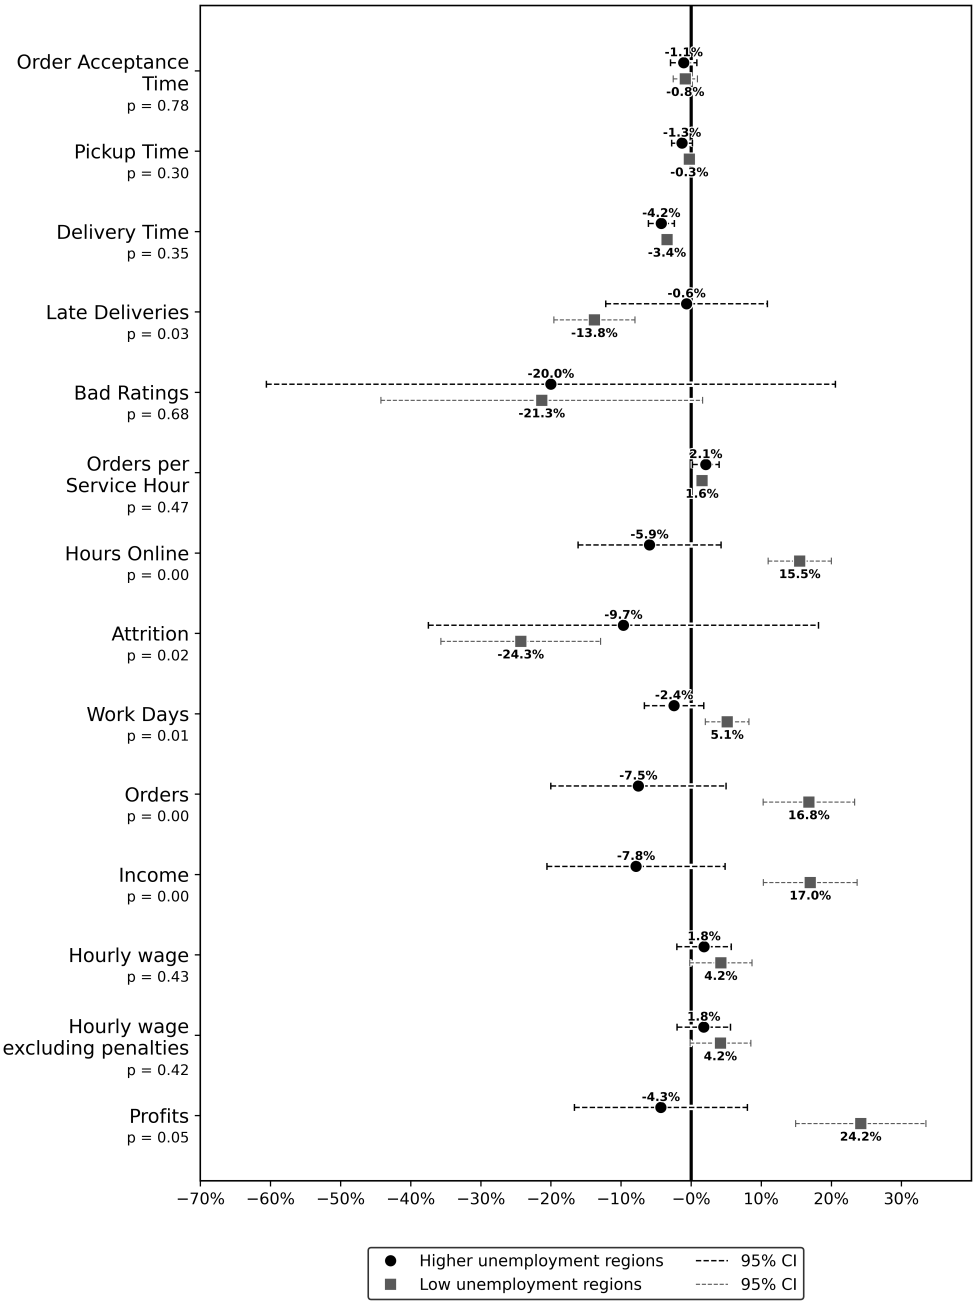
Notes: This figure is similar to Figure 9 in the main text. The difference is that instead of splitting up DHH workers by profound and non-profound hearing loss, we estimate a dosage response model. Each dot shows the dosage-response coefficient; the p-value tests whether it equals zero. The original Level 1 (most severe, $>$ 90 dB hearing loss) to Level 4 (least severe, 41–60 dB) is transformed using Severity $= 5 -$ Level. The figure plots the coefficient on the triple interaction of $DHH_i \times Post_t \times \text{Severity}_i$. Whiskers indicate 95% confidence intervals clustered by driver.

Figure A.11: Robustness on Heterogeneity By Disability Severity: Control for Other Interactions



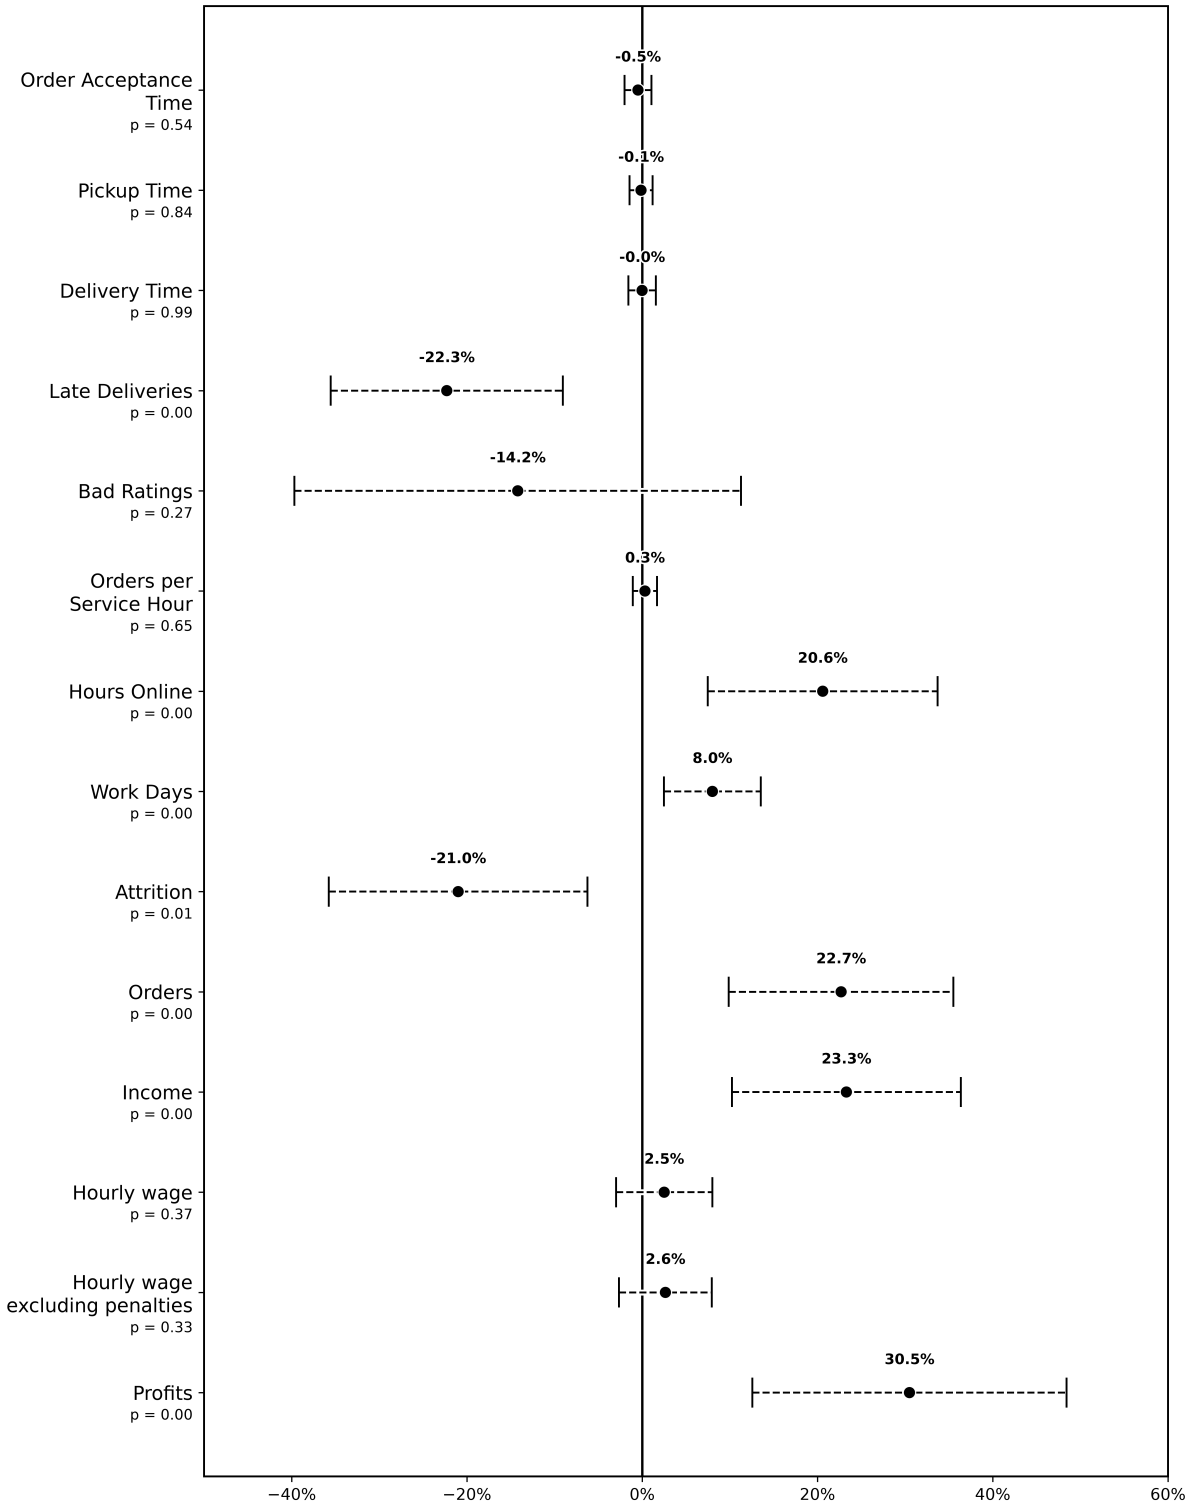
Notes: This figure presents a robustness check on heterogeneity by disability severity. It presents analyses similar to those in Figure 9, but also controls for triple interactions, $DHH \times Post \times Female$ and $DHH \times Post \times Age$, as well as relevant double interactions. This helps address whether our heterogeneity by severity is standing in for heterogeneity by demographics. The coefficients plotted here are the triple interactions $DHH_i \times Post_t \times \text{Profound}_i$. A coefficient different from 0 indicates a differential effect by disability severity.

Figure A.12: Heterogeneity in the Effects of the AI Tool by Regional Unemployment Level



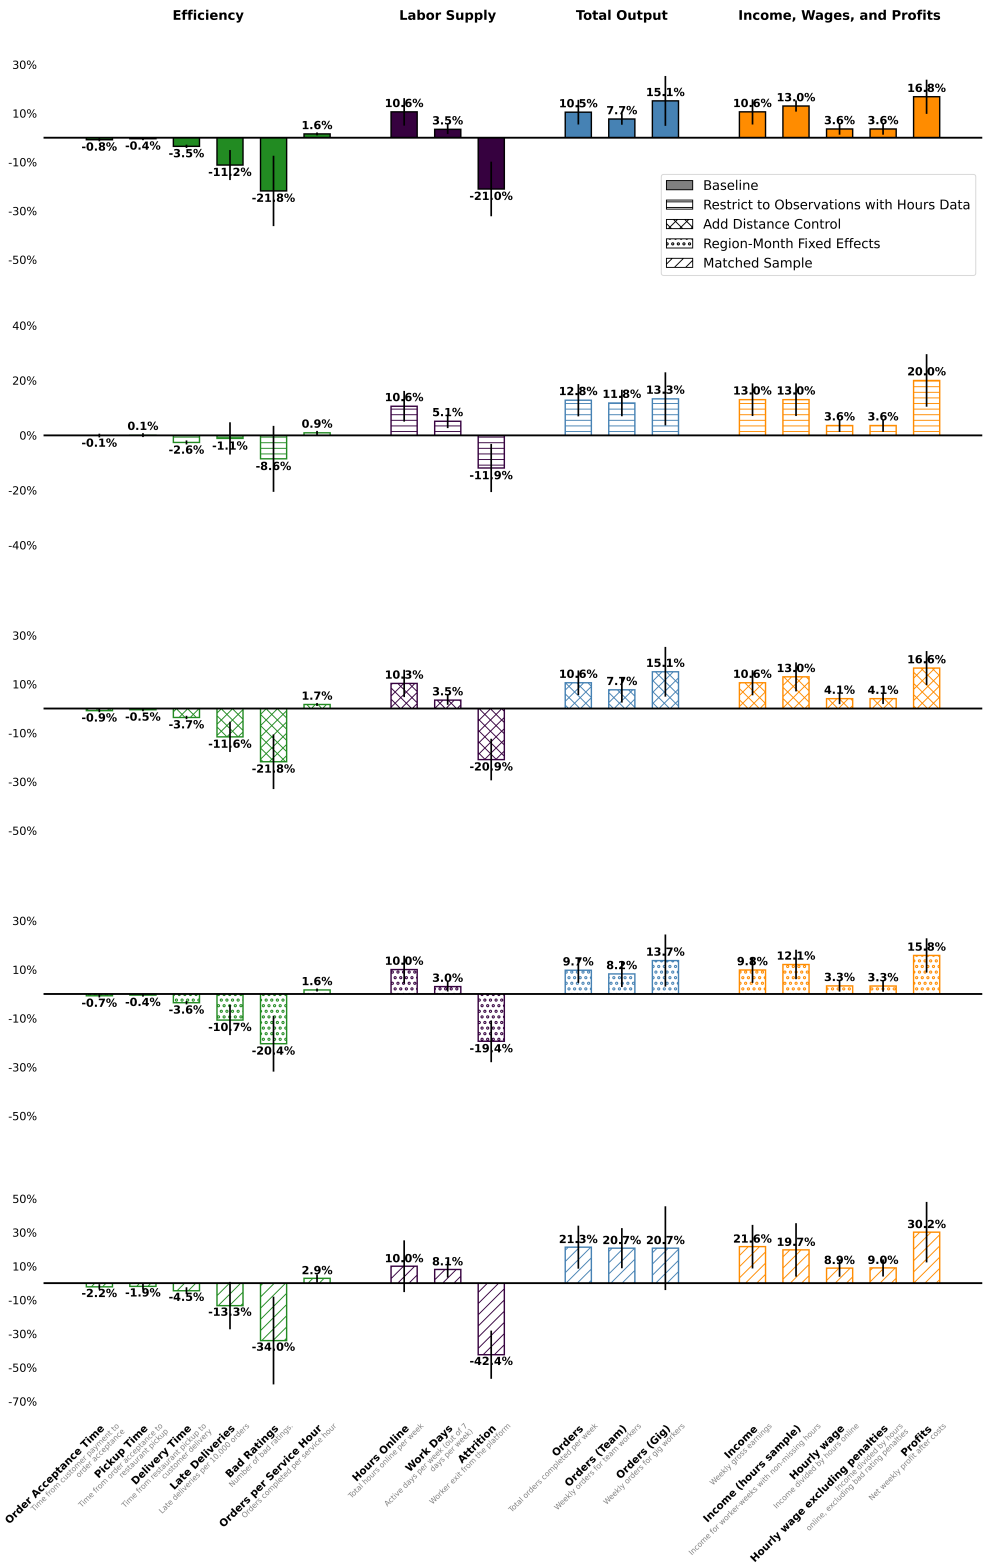
Notes: This figure presents heterogeneous treatment effects by regional unemployment level. It is similar to Figure A.8, but instead splits the sample by regional unemployment rate. Regions are divided into high- and low-unemployment regions using the median unemployment rate as the cutoff. The sample includes 31 provincial-level administrative regions ("regions") in mainland China. Unlike for other heterogeneity analyses, standard errors are block-bootstrap clustered by region (with 500 replications), reflecting that we study a characteristic that varies by region. We employ the block-bootstrap cluster since the number of clusters at 31 is relatively small. The sample consists of worker-weeks from 3,982 DHH and 36,719 non-disabled workers. The sample size is smaller than in our main analyses because region is missing for some workers. Information on unemployment is from the 2021 China Labour Statistical Yearbook. Low-unemployment regions are Qinghai, Xinjiang, Shanxi, Hunan, Jiangsu, Anhui, Guangdong, Guangxi, Zhejiang, Tibet, Shanghai, Jiangxi, Shandong, Chongqing, and Hubei. High-unemployment regions are Hebei, Hainan, Beijing, Heilongjiang, Jilin, Fujian, Henan, Gansu, Shaanxi, Sichuan, Tianjin, Inner Mongolia, Yunnan, Ningxia, Liaoning, and Guizhou.

Figure A.13: Robustness on Heterogeneity in the Effects of the AI Tool by Regional Unemployment Level: Control for Other Interactions



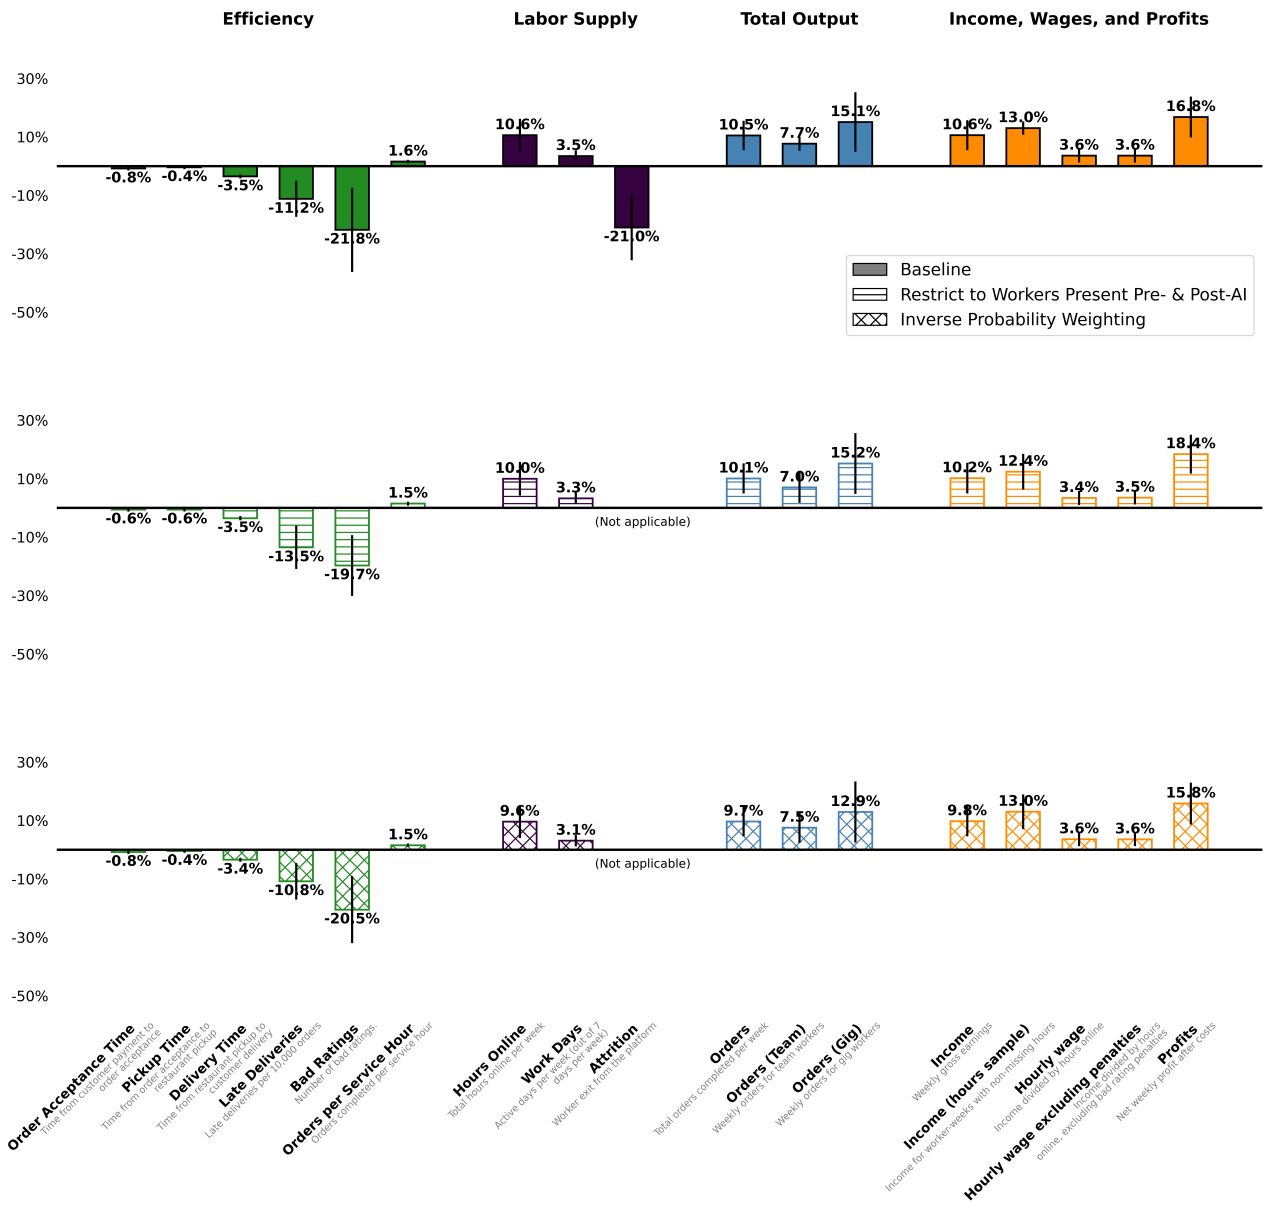
Notes: This figure presents a robustness check on heterogeneity by regional unemployment. It presents analyses similar to those in Figure A.12, but also controls for triple interactions, $DHH \times Post \times Female$ and $DHH \times Post \times Age$, as well as relevant double interactions. This helps address whether our heterogeneity by regional unemployment is standing in for heterogeneity by demographics. The coefficients plotted here are the triple interactions $DHH_i \times Post_t \times$ Lower Unemployment Region$_i$, where Low-Unemployment Region$_i$ is a dummy for being in a lower unemployment region. A coefficient different from 0 indicates a differential effect by regional unemployment level.

Figure A.14: Figure 3 Robustness: Checks on Hours Sample, Matched Sample, and Controls



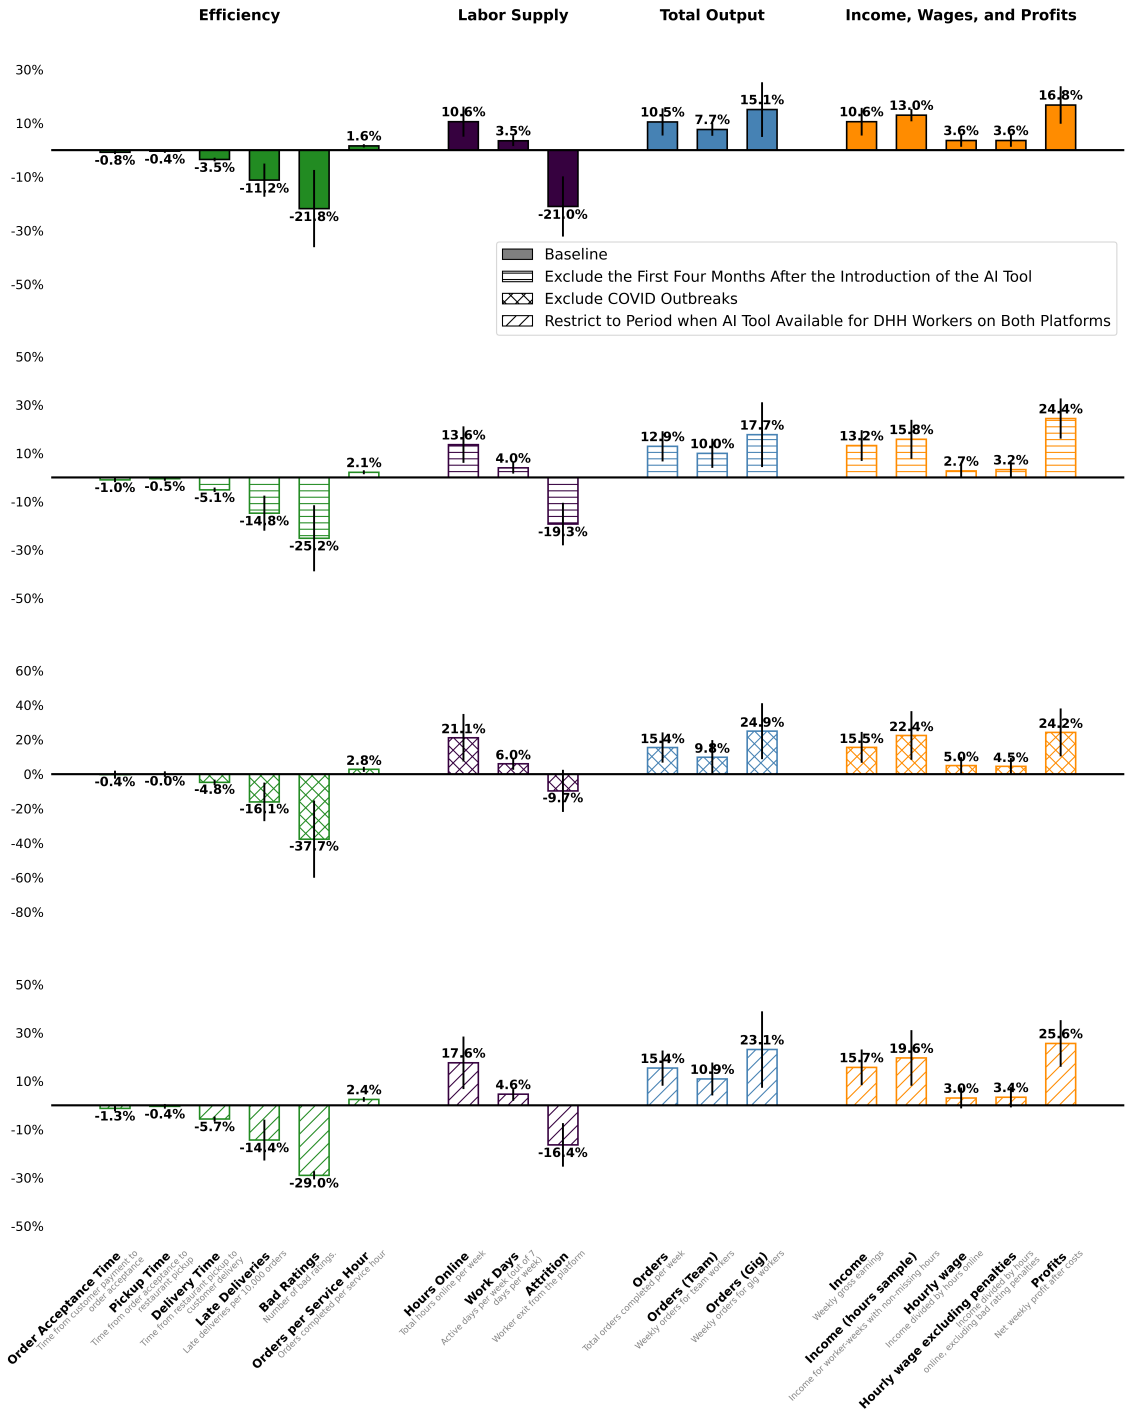
Notes: The top panel replicates the baseline analyses in Figure 3. The 2nd panel restricts the sample to observations with hours data, since hours are missing for many worker-weeks (see Figure A.2); the resulting sample contains 343,733 worker-week observations from 34,965 workers. The 3rd panel adds a worker-week-level control for the average restaurant-to-customer distance per order. This variable proxies for order difficulty, since longer-distance orders may require more travel time, involve greater delivery-time uncertainty, and take workers farther from dense restaurant areas or into less familiar delivery locations. This specification uses the same sample as the baseline. The 4th panel adds region-month dummies. The 5th panel performs the DiD analysis using exact matching, matching each DHH worker to one non-disabled worker based on entry year, work model (gig or team), city, 5-year age bin, and gender, leaving 2,354 DHH and 2,354 non-disabled workers.

Figure A.15: Robustness of Figure 3: Robustness Checks on Differential Attrition



Notes: The top panel replicates the baseline analyses in Figure 3. The middle panel restricts the sample to drivers observed for at least one week both before and after the introduction of the AI tool; this restricted sample includes 184,450 worker-week observations from 3,397 workers. The bottom panel estimates the main DiD specification using inverse probability weights to account for differential attrition. Specifically, we first estimate each worker's probability of remaining in the sample using a logit model as a function of time-invariant worker characteristics, namely city-of-work fixed effects, age fixed effects, gender, and team status (gig vs. team), as well as year-week fixed effects. We then weight observations by the inverse of the predicted probability. The inverse probability weighting specification uses the baseline sample of 480,070 worker-week observations from 43,605 workers. Attrition is not analyzed in the middle or bottom panels, as we use strategies there to address concerns about differential attrition, and analyzing attrition there would be circular.

Figure A.16: Robustness of Figure 3: Robustness Checks Related to Longer-run Effects and External Validity



Notes: The top panel replicates the baseline analyses in Figure 3. The second panel excludes the first four months after the introduction of the AI tool, yielding 427,612 worker-week observations from 40,821 workers. The third panel excludes worker-week observations from regions-weeks affected by COVID outbreaks, leaving 234,004 worker-week observations from 28,303 workers. Under China's "Dynamic Zero-COVID" policy, restrictions were localized rather than synchronized nationwide, so their timing and stringency varied across regions and within regions. Oftentimes, only part of a region would have a COVID outbreak, so we are being conservative in excluding region-weeks with any outbreak. We identify affected region-weeks using official government announcements and verify the timing with contemporaneous news reports. The bottom panel excludes the period when the AI tool was available on our platform but not yet on the other major food delivery platform. Since China's food delivery market is largely dominated by these two platforms, this restriction addresses the possibility that temporary cross-platform spillovers, such as worker or customer reallocation, affect our estimates. This restricted sample includes 391,839 worker-week observations from 38,164 workers.

Table A.7: Bad Ratings: Pre-AI Differences, DiD Estimates, and DiD Robustness Checks

	(1)	(2)	(3)
Panel A: Pre-AI Means			
	Non-disabled Workers	DHH Workers	
Mean	0.10	0.15	
Panel B: Pre-AI Difference, Raw and Scaled			
	OLS, raw coefficient	OLS coefficient divided by non-disabled mean	OLS coefficient divided by DHH mean
DHH	0.047***	0.470***	0.313***
	(0.009)	(0.087)	(0.058)
Observations	180,375	180,375	180,375
Panel C: DiD Estimate, Raw and Scaled			
	OLS, raw coefficient	OLS coefficient divided by DHH mean	
$DHH \times Post$	−0.032***	−0.218***	
	(0.010)	(0.068)	
Observations	480,070	480,070	
Panel D: DiD Robustness, Dependent Variable = Bad Ratings per Week			
	OLS; coefficient divided by pre-AI DHH mean	Poisson	OLS; coefficient divided by pre-AI DHH mean, controlling for completed orders
$DHH \times Post$	−0.218***	−0.185***	−0.239***
	(0.068)	(0.071)	(0.068)
Observations	480,070	480,070	480,070
Panel E: DiD Robustness, Dependent Variable = Bad Ratings per 10,000 Orders			
	OLS	OLS, order-weighted	
$DHH \times Post$	−2.350**	−2.197***	
	(1.024)	(0.568)	
Observations	480,070	480,070	

Notes: Panel A reports pre-AI means for non-disabled and DHH workers. Panel B reports the regression-adjusted pre-AI difference (DHH vs. non-disabled) from estimating equation (2): the raw OLS coefficient in levels, the coefficient divided by the non-disabled pre-AI mean, and the coefficient divided by the DHH pre-AI mean. Panel C reports the baseline DiD estimate from equation (3): the raw OLS coefficient and the coefficient divided by the DHH pre-AI mean. Panels D and E examine robustness of the DiD results to alternative estimation methods, functional forms, and control variables. Each DiD coefficient is from a difference-in-differences specification using the same fixed effects as in the main text. Standard errors clustered by worker in parentheses. Significance levels: * p<0.1, ** p<0.05, *** p<0.01.

Table A.8: Late Deliveries: Pre-AI Differences, DiD Estimates, and DiD Robustness Checks

	(1)	(2)	(3)
Panel A: Pre-AI Means			
	Non-disabled Workers	DHH Workers	
Mean	4.79	5.11	
Panel B: Pre-AI Difference, Raw and Scaled			
	OLS, raw coefficient	OLS coefficient divided by non-disabled mean	OLS coefficient divided by DHH mean
DHH	0.497*	0.104**	0.097**
	(0.275)	(0.057)	(0.054)
Observations	180,375	180,375	180,375
Panel C: DiD Estimate, Raw and Scaled			
	OLS, raw coefficient	OLS coefficient divided by DHH mean	
$DHH \times Post$	−0.853*	−0.167*	
	(0.497)	(0.097)	
Observations	480,070	480,070	
Panel D: DiD Robustness, Dependent Variable = Late Deliveries per Week			
	OLS; coefficient divided by pre-AI DHH mean	Poisson	OLS; coefficient divided by pre-AI DHH mean, controlling for completed orders
$DHH \times Post$	−0.167*	−0.164*	−0.255**
	(0.097)	(0.094)	(0.103)
Observations	480,070	480,070	480,070
Panel E: DiD Robustness, Dependent Variable = Late Deliveries per 10,000 Orders			
	OLS	OLS, order-weighted	
$DHH \times Post$	−51.15***	−40.60**	
	(17.09)	(17.28)	
Observations	480,070	480,070	

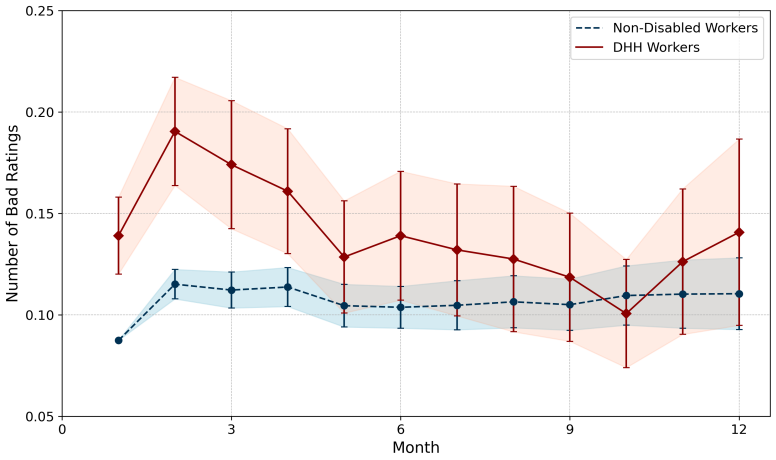
Notes: Panel A reports pre-AI means for non-disabled and DHH workers. Panel B reports the regression-adjusted pre-AI difference (DHH vs. non-disabled) from estimating equation (2): the raw OLS coefficient in levels, the coefficient divided by the non-disabled pre-AI mean, and the coefficient divided by the DHH pre-AI mean. Panel C reports the baseline DiD estimate from equation (3): the raw OLS coefficient and the coefficient divided by the DHH pre-AI mean. Panels D and E examine robustness of the DiD results to alternative estimation methods, functional forms, and control variables. Each DiD coefficient is from a difference-in-differences specification using the same fixed effects as in the main text. Standard errors clustered by worker in parentheses. Significance levels: * p<0.1, ** p<0.05, *** p<0.01.

Table A.9: Attrition: Pre-AI Differences and DiD Estimates

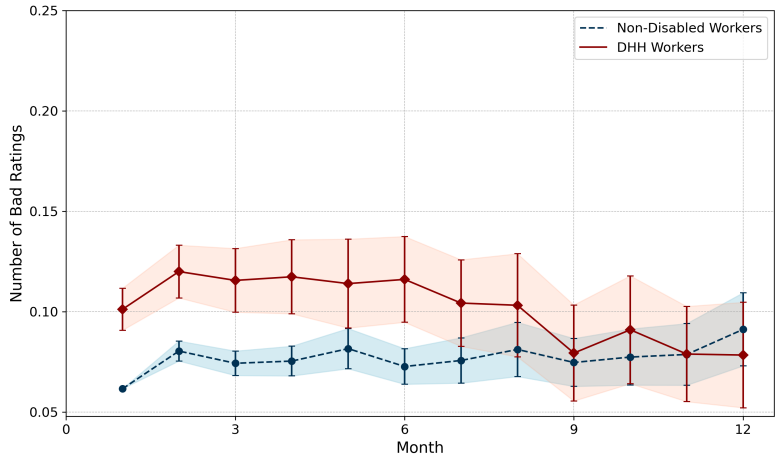
	(1)	(2)	(3)
Panel A: Pre-AI Means			
	Non-disabled Workers	DHH Workers	
Mean	0.076	0.056	
Panel B: Pre-AI Difference, Raw and Scaled			
	OLS, raw coefficient	OLS coefficient divided by non-disabled mean	OLS coefficient divided by DHH mean
DHH	-0.023^{***}	-0.301^{***}	-0.414^{***}
	(0.002)	(0.030)	(0.041)
Observations	180,375	180,375	180,375
Panel C: DiD Estimate, Raw and Scaled			
	OLS, raw coefficient	OLS coefficient divided by DHH mean	
$DHH \times Post$	-0.012^{***}	-0.210^{***}	
	(0.003)	(0.052)	
Observations	456,192	456,192	

Notes: Panel A reports pre-AI means for non-disabled and DHH workers. Panel B reports the regression-adjusted pre-AI difference (DHH vs. non-disabled) from estimating equation (2): the raw OLS coefficient in levels, the coefficient divided by the non-disabled pre-AI mean, and the coefficient divided by the DHH pre-AI mean. Panel C reports the baseline DiD estimate from equation (3): the raw OLS coefficient and the coefficient divided by the DHH pre-AI mean. Panels A and B are based on the pre-AI sample of 180,375 worker-weeks, whereas Panel C is based on the full attrition sample of 456,192 worker-weeks (the attrition outcome excludes the final weeks of 2023; see the notes to Table 2). Significance levels: * p<0.1, ** p<0.05, *** p<0.01.

Figure A.17: Robustness of Figure 7: Tenure-Productivity Curves in Bad Ratings, Before and After the AI Tool (Count of Bad Ratings instead of Ratings per 10,000 Orders)



(a) Pre-AI tool



(b) Post-AI tool

Notes: This figure is similar to Figure 7, but the outcome is number of bad ratings. Error bars show 95 percent confidence intervals accounting for clustering by worker. The samples in each panel are the same as in Figure 7.

B Prior Evidence on Productivity of Disabled Workers

This appendix reviews prior evidence on the workplace performance of disabled relative to non-disabled workers, and discusses why comparisons based on measured individual productivity at scale have not previously been feasible.

What existing reviews conclude. Systematic reviews find little quantitative evidence. Lindsay et al. (2018) screen over 6,000 studies on the benefits of hiring people with disabilities and retain 39 for review; only eight report evidence related to productivity, and these rely on self-reported or perception-based measures. A more recent review reaches similar conclusions (Saha et al., 2023). Researchers across fields have noted the gap directly. Narayanan and Terris (2020) observe that the existing literature does not study the productivity implications of employing disabled workers using microdata. In vocational rehabilitation, Cimera (2002, 2006) proposes an accounting framework for measuring the employer costs and benefits of hiring workers with disabilities—closely anticipating the profitability question we study—but notes that no study had carried out such an analysis.

The closest existing studies. Prior evidence falls into four categories.

Employer surveys and ratings. Graffam et al. (2002) survey 643 Australian employers, each rating one employee with a disability against a hypothetical "average" employee, finding lower ratings on speed and quality of work but better attendance and lower recruitment and compensation costs.[1] A related literature studies employer attitudes and hiring intentions—rather than ratings of realized performance—typically finding that employers associate hiring workers with disabilities with reliability and corporate social responsibility, while also citing concerns about skill and cost (Piramanayagam and Seal, 2021; Dean et al., 2022). Relatedly, studies of customer responses to disabled frontline employees in retail and hospitality settings generally find favorable or neutral reactions, moderated by service failures (Rosenbaum et al., 2017; Madera et al., 2020).

Small-sample studies. Hernandez and McDonald (2010) compare 95 workers with disabilities to 219 workers without disabilities across 13 US firms using supervisor performance ratings, tenure, and absence records, finding similar performance ratings and longer average tenure for workers with disabilities; Brickey and Campbell (1981) document longer tenure among 17 fast-food workers with intellectual disabilities at one firm; Seva (2020) conducts a two-day time study of four workers with autism at one Philippine firm; and Kaletta et al. (2012), an internal company report, studies trucking outcomes for 430 drivers (131 with a disclosed disability, including 21 who are deaf or hard of hearing), finding fewer safety incidents among drivers with disabilities.

Team-level operational data. Using daily production data from a single work-integration social enterprise (a nonprofit whose mission is employing disabled workers), Narayanan and Terris (2020) and Zhang et al. (2023) relate the mix of disability types within apparel-manufacturing teams to team output; these studies do not measure individual productivity and do not compare disabled with non-disabled workers (the share of non-disabled workers on a team enters only as a control variable).

[1] Various websites refer to internal DuPont surveys from 1990 indicating broadly similar supervisor ratings for disabled and non-disabled workers, but we have been unable to locate the underlying report.

Non-productivity outcomes. Large-scale evidence exists for outcomes other than productivity: occupational injuries measured in household survey data (Zwerling et al., 1997), wage gaps—from which productivity differences are inferred rather than measured (Malo and Pagán, 2012; Liao et al., 2024; Forth and Jones, 2025)—and firm-level evidence on disability-quota compliance at publicly listed Chinese firms, based on aggregate rather than individual-level productivity data (Jing et al., 2022).

Why measured comparisons have been infeasible. Three data constraints reinforce one another. First, mainstream employers rarely record disability status in a form that can be linked to individual output: disclosure is voluntary, and information about disabilities and accommodations is typically confidential. Second, the settings that do systematically record disability status—disability insurance programs, quota-compliance records, and sheltered or supported employment—either do not measure output, lack non-disabled workers performing the same job, or operate under non-market conditions. Third, individual output itself is measured in only a minority of jobs, mainly piece-rate manufacturing and logistics settings. Our setting resolves all three constraints simultaneously: disability status is verified by government-issued certificate at registration, the Platform's app records each worker's output and speed at each stage of the delivery process, and open hiring places disabled and non-disabled workers in the same job under identical conditions. This complements the discussion of the value of granular administrative data in the Conclusion.

Relation to our findings. Although the earlier evidence is perception-based or small in scale, its direction anticipates our administrative findings: workers with disabilities are consistently described as somewhat slower on average but more reliable and less costly to retain (Graffam et al., 2002; Hernandez and McDonald, 2010). Our contribution is to establish these patterns—and their consequences for wages and profitability—with measured individual productivity at scale, and to evaluate an AI-enabled accommodation that this earlier literature could not have anticipated.

C Details on the AI Tool: Intelligent Outbound Call

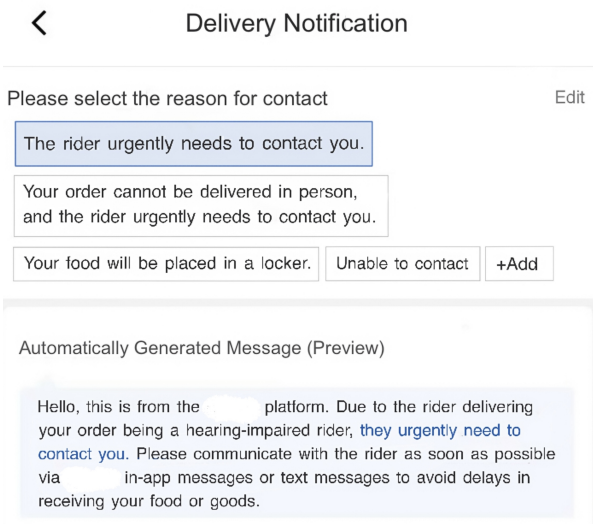
The IOC feature is integrated into the delivery app. Workers can select from predefined default prompts or input customized prompts relevant to specific situations. These text messages are converted into audio through an advanced TTS system and sent to the customer via a cellphone call that appears to originate from the worker's phone number.

Panel (a) of Figure C.1 displays the default prompt menu provided by the app. Workers can quickly select an appropriate option, preview the automatically generated message, and initiate the automated call instantly.

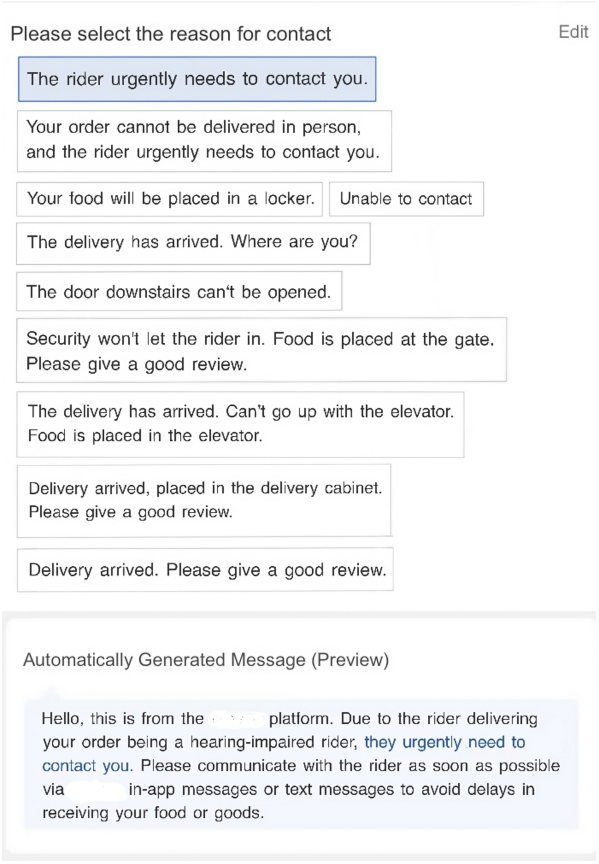
For situations beyond standard prompts, workers have the flexibility to create and store personalized messages. Panel (b) of Figure C.1 illustrates the menu with customized additions, showing prompts a worker has added alongside the defaults to address real-world delivery challenges.

Appendix D provides background on the TTS technology underlying the IOC, and on why AI-based speech synthesis is central to the tool's effectiveness.

Figure C.1: Screenshot of the Intelligent Outbound Call Feature



(a) Default Prompt Menu



(b) Menu with Customized Additions

Notes: Panel (a) shows the default prompt menu for the Intelligent Outbound Call feature. Panel (b) shows the menu with customized additions from one of the DHH workers we worked with, illustrating prompts they added alongside the defaults. The request for "Please give a good review" is likely a polite way of asking to avoid a bad rating.

D Text-to-Speech Technology and the Role of AI

This appendix provides background on text-to-speech (TTS) technology to clarify the sense in which the IOC is an AI-based tool. See Tan et al. (2021), Taylor (2009), and Vapi (2025) for further general discussion and chronology.

Six decades of intelligible but robotic speech. Synthetic speech long predates modern AI. Mechanical speaking machines date to the 18th century, and Bell Labs demonstrated the first electronic speech synthesizer, the VODER, in 1939 (Klatt, 1987). From the 1970s onward, rule-based "formant" synthesizers such as MITalk and DECtalk—the latter famously providing Stephen Hawking's voice—made TTS practical for accessibility applications (Klatt, 1987). In the 1990s and 2000s, two approaches dominated: concatenative synthesis, which stitches together fragments of recorded human speech (Hunt and Black, 1996), and statistical parametric synthesis (Zen et al., 2009). Each had characteristic artifacts: audible "joins" between fragments in the former, and muffled or buzzy speech in the latter (Tan et al., 2021). Thus, across four decades of TTS systems, synthetic speech remained intelligible but unmistakably machine-like—a listener could nearly always tell within moments that the voice was not human.

The neural TTS breakthrough. Deep learning changed this discontinuously. Rather than following hand-coded rules or splicing recorded fragments, neural TTS systems infer from large corpora of recorded speech what natural speech sounds like—its prosody, intonation, and pauses—an instance of prediction, the core function of modern AI (Agrawal et al., 2018). In 2016, researchers at Google DeepMind introduced WaveNet, a deep neural network that generates raw audio waveforms, producing speech that human listeners rated as "significantly more natural sounding than the best parametric and concatenative systems for both English and Mandarin" (van den Oord et al., 2016)—the latter being the language relevant in our setting. In 2017–2018, the Tacotron systems replaced the traditional multi-stage synthesis pipeline with neural networks trained end-to-end (Wang et al., 2017). Tacotron 2 achieved a mean opinion score (MOS) for naturalness of 4.53, compared to 4.58 for professionally recorded human speech (Shen et al., 2018); that is, raters scored the synthetic voice within a hair of a professional human recording. Distillation methods soon made these models fast enough for large-scale commercial use (van den Oord et al., 2018), and by around 2020, neural TTS was deployed in commercial voice assistants and offered at very low marginal cost per call by major cloud providers, including, for Mandarin, by China's leading technology companies. Progress has continued since, with recent systems achieving formal "human parity" in naturalness relative to human recordings (Chen et al., 2024).

Why AI-driven naturalness matters for the IOC. An IOC call succeeds only if the customer, upon picking up, keeps listening to an automated call and then acts on its message. Robotic-sounding calls fare poorly on the first step: once customers hear a robotic voice, it triggers the heuristics they use to screen out telemarketing and scam robocalls, both widespread in China, and they routinely hang up. Naturalness also matters for the second step. The IOC is strictly one-way—the customer cannot talk back, and replies, if needed, are sent by in-app message—so the call is not a conversation but a single spoken message that must, on its own, come across as clear and credible enough to act on (Shah et al., 2023; Arcot,

2025). Our fieldwork (Appendix H) points the same way: the DHH workers we accompanied used the IOC on nearly every order, the calls were generally effective at prompting customers to act (e.g., opening access-controlled doors), and after a call the worker could typically leave the order at the designated spot with no further interaction. By contrast, when a coauthor simulated the pre-AI experience—delivering orders while wearing earplugs and relying on in-app messages—customers often ignored the messages or replied only after several minutes, and direct phone calls placed without speaking left customers audibly confused and irritated.

The timing of the IOC also reflects this technological trajectory. The tool was developed and launched in mid-2022, roughly two years after near-human neural TTS had become cheap and widely available. We cannot directly observe how a version of the IOC built on pre-neural TTS would have performed. However, to our knowledge, no such tool was deployed on delivery platforms in the pre-neural era, and both major Chinese food-delivery platforms introduced TTS calling tools for DHH workers only in 2022, after neural TTS had matured.

E Additional Discussion and Analysis

E.1 Impact of AI Tool on Good Ratings of DHH Workers

This section examines how the AI tool affects the number of good ratings received by DHH workers. As discussed in the main text, our main text analyses focus on bad ratings because bad ratings are much more important to the Platform and have monetary consequences for workers. By contrast, based on our interviews with the Platform, good ratings do not generate direct practical benefits for workers. We apply the same empirical framework used in the main text: column (1) reports the pre-AI difference specification from equation (2), and column (2) reports the overall treatment effect using the difference-in-differences specification from equation (3). Table E.1 presents the results.

Column (1) shows that, prior to the introduction of the AI tool, DHH workers receive 16.5% more good ratings than non-disabled workers in the pre-AI period. In comparison, our main results on bad ratings show that DHH workers also received 31.3% more bad ratings before the AI tool was introduced. Together, these findings suggest that DHH workers receive more customer feedback of both types, which likely partly reflects that they complete more orders. Column (2) shows that the AI tool did not have a statistically significant effect on good ratings for DHH workers relative to non-DHH workers.

Taken together, these results suggest that while DHH workers were already receiving more good ratings prior to the AI tool, the tool itself did not further improve this dimension of customer feedback. This contrasts with our findings for bad ratings in the main text, where the AI tool meaningfully reduced negative feedback. One interpretation is that the AI tool mainly operates by reducing downside communication failures, rather than shifting the upper tail of customer satisfaction. Importantly, the null effect of the AI tool on good ratings (column 2) suggests that customer sympathy for DHH workers is not driving our main results: if sympathy were the mechanism, we would expect good ratings to improve alongside the reduction in bad ratings, which we do not find.

Table E.1: Pre-AI Differences and Impact of AI Tool on Good Ratings of DHH Workers

	(1)	(2)
	Pre-AI Difference	DiD Result
DHH_i	0.165***	
	(0.038)	
$DHH_i \times Post_t$		0.006
		(0.051)
Observations	180,375	480,070

Notes: Column (1) reports pre-AI differences in good ratings between DHH and non-disabled workers. Column (2) reports DiD estimates of the impact of the AI tool on good ratings for DHH workers relative to non-disabled workers. As for bad ratings, we estimate non-transformed OLS models and divide by the pre-AI DHH mean to calculate percentage differences. Standard errors clustered by worker in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

E.2 Impact of AI Tool on Entry of DHH Workers

This section examines how the introduction of the AI tool affects the entry of DHH workers onto the Platform. Our analysis in the main text is based on workers who join the Platform and complete at least one order. However, if the AI tool makes work on the Platform more attractive for DHH workers, it could also spur more DHH workers to join.

We cannot observe the set of DHH workers potentially interested in working on the Platform. Instead, we collapse the data by disability type (DHH or non-disabled) and year-week. We have 3 years of data, so 3*52 = 156 year-weeks, and there are two disability types, so we have 312 observations. We then estimate a model very similar to equation (3):

$$Y_{gt} = \alpha + \gamma \cdot DHH_g \times Post_t + \delta \cdot DHH_g + \xi_t + \varepsilon_{gt}. \tag{E.1}$$

where g denotes group (i.e., DHH or non-disabled) and t denotes time in year-weeks; Y_{gt} is the inverse hyperbolic sine of the number of new entrants on the Platform in group g during year-week t; and DHH_g is a dummy for being DHH. We are interested in estimates of γ.

Table E.2 presents the regression results. The AI tool increases the number of DHH registrations on the Platform by 27.5%.

Benchmarking the entry response. To benchmark this magnitude, we relate the entry response to the tool's effects on worker pay. The AI tool raises DHH weekly income by about 11% (Figure 3), so the entry response implies an elasticity of entry with respect to income of roughly 2.5. This is broadly in line with quasi-experimental estimates of firm-level labor supply and recruitment elasticities, which generally fall between about 1.5 and 5: Sokolova and Sorensen (2021) report a median direct estimate of 1.4 in a meta-analysis; Dal Bo et al. (2013) estimate an experimental application elasticity of 1.8; Bassier et al. (2022)

Table E.2: Impact of AI Tool on Entry of DHH Workers: DiD Results

	(1)
$DHH_g \times Post_t$	0.275**
	(0.110)
Observations	312

Notes: An observation is a year-week × disability type (i.e., DHH or non-disabled). Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

estimate separation-based elasticities around 4; Emanuel and Harrington (2024) estimate recruitment elasticities of 3.2–4.2; and Azar et al. (2022) estimate a firm-level application elasticity of 4.8. Relative to the tool's 3.6% effect on hourly wages alone, the implied elasticity is larger (roughly 7.6), suggestive that DHH workers value the tool beyond its effect on pay, e.g., due to reduced communication stress.

Caveats. Several caveats apply to the entry analysis. First, unlike our main results, this is a group-by-week design without worker fixed effects; identification requires that entry flows of DHH and non-disabled workers would have trended similarly absent the tool. Second, our data include the universe of DHH entrants but a random sample of non-disabled entrants. Because the sample was drawn uniformly over the sample period, sampled counts are proportional to true entry, and the proportionality constant is absorbed by the group effect in equation (E.1), leaving γ unaffected. Third, in line with our sample construction, entry is measured as registering and completing at least one week of work. Fourth, the elasticity benchmarks above are estimated from wage variation holding job attributes fixed, whereas the AI tool changed pay, amenities, and awareness of the Platform's accommodations together; the implied elasticities should thus be viewed as a summary lens rather than as structural estimates.

E.3 Mediation Analysis

We run a mediation analysis to examine whether the treatment's effect on order volume operates through algorithmic reallocation. As discussed in the main text, the Platform's algorithm looks favorably upon drivers with fewer bad ratings and fewer late deliveries. If this is the primary channel, the AI tool would raise orders mainly by improving late deliveries and bad ratings, which the algorithm then rewards—with observed gains in order volume reflecting task reallocation rather than a direct increase in worker capacity. We examine two potential mediating channels: late deliveries and bad ratings.[2] For each channel, we estimate (i) the total effect of the AI tool on orders, (ii) the effect of the AI tool on the mediator,

[2]We analyze late deliveries and bad ratings each in terms of a ratio per 10,000 orders, doing so to avoid a "reverse causality" issue where someone has lots of bad ratings due to working more orders in a week.

and (iii) the effect of the AI tool on orders when controlling for the mediator. Results are reported in Table E.3.

Table E.3: Mediation Analysis: Is the Effect on Orders Driven by Algorithmic Reallocation?

Dep. Var.:	Orders	Late Deliveries	Orders	Bad Ratings	Orders
	(1)	(2)	(3)	(4)	(5)
$DHH_i \times Post_t$	0.105^{***}	-51.15^{***}	0.103^{***}	-2.350^{**}	0.105^{***}
	(0.0307)	(17.09)	(0.0307)	(1.024)	(0.0307)
Late Deliveries (per 10k orders)			$-3.24 \times 10^{-5***}$		
			(4.30×10^{-6})		
Bad Ratings (per 10k orders)					$-8.68 \times 10^{-5***}$
					(3.34×10^{-5})
Observations	480,070	480,070	480,070	480,070	480,070

Notes: Standard errors clustered by worker in parentheses. Significance levels: * p<0.1, ** p<0.05, *** p<0.01. Column 1 reports the baseline DiD treatment effect on orders following equation (3). Columns 2 and 4 use late deliveries and bad ratings, respectively, as the dependent variable with the same controls. Columns 3 and 5 re-estimate the baseline orders regression adding late deliveries and bad ratings, respectively, as additional controls. Late deliveries and bad ratings are measured per 10,000 orders.

Column 1 of Table E.3 gives the baseline total effect: the treatment increases the IHS-transformed number of orders by 0.105. Column 2 shows that the treatment reduces late deliveries by 51.15 per 10,000 orders, and column 4 shows that it reduces bad ratings by 2.350 per 10,000 orders. Columns 3 and 5 re-estimate the baseline orders regression controlling for each mediator in turn. The treatment coefficient attenuates only slightly to 0.103 when controlling for late deliveries (column 3), implying that the late delivery channel accounts for about 1.6% of the total effect. The coefficient is essentially unchanged at 0.105 when controlling for bad ratings (column 5), implying that the bad ratings channel accounts for less than 0.2% of the total effect.

Consistent with the algorithm penalizing late deliveries and bad ratings, both mediators are negatively associated with orders (columns 3 and 5). However, controlling for either of these does little to diminish the treatment effect of the AI tool on orders. The dominant mechanism linking the AI tool to higher order volume therefore lies outside of bad ratings and late deliveries. This supports that algorithmic reallocation is not an important mechanism for our treatment effect on orders.

E.4 How Much Do Workers Gain? Welfare Calculations Based on the Conceptual Framework

The DiD estimates show that the AI tool raises DHH workers' weekly income by about 11% (Figure 3). This section uses the conceptual framework from Section 1.1 to ask how

much of that income gain constitutes a *welfare* gain for workers. The distinction matters because roughly two-thirds of the income gain reflects increased hours, and hours have an opportunity cost: time reallocated to the Platform was previously spent on other valued activities.

A sufficient-statistics result. In the framework, a worker chooses the share of work time on the Platform, h, to maximize $wzh + R(1-h)$, where $R(\cdot)$ is a concave return to off-Platform time; the functional form in Section 1.1 is a special case, but nothing below depends on it. Let $\hat{\gamma}_w$ and $\hat{\gamma}_h$ denote the DiD estimates of the AI tool's effects on hourly wages and on weekly hours (both approximately log points). By the envelope theorem, $dV/d(wz) = h$, where V is the worker's indirect utility. To second order, the welfare gain as a share of pre-AI weekly earnings E_0 is given by equation (4) of the main text, restated here for convenience:

$$\frac{\Delta V}{E_0} = \underbrace{\left(e^{\hat{\gamma}_w} - 1\right)}_{\text{inframarginal wage gain}} + \underbrace{\tfrac{1}{2}\,\hat{\gamma}_h\,\hat{\gamma}_w}_{\text{re-optimization gain}} .$$

The interpretation of the two terms is given in the main text. Notably, equation (4) requires no knowledge of the outside-option parameters r and β, of the worker's time endowment, or of how off-Platform time is spent. It requires only that workers freely choose their hours and that the piece rate w is unchanged by the tool.[3]

Results. Panel A of Table 4 applies equation (4). In the baseline, workers gain the equivalent of 3.9% of pre-AI weekly earnings (SE 1.5), of which 3.7 percentage points is the inframarginal wage gain and 0.2 percentage points the re-optimization term. Thus, while weekly income rises by roughly 11%, worker surplus rises by roughly 4%: most of the hours-driven income gain is compensated reallocation of time rather than surplus. Results are nearly identical using the panel of workers present both pre- and post-AI (3.6%), indicating composition changes are not driving the calculation.

Accounting for amenities. Panel A treats the entire hours response as movement along a stable labor supply curve to the Platform. However, the implied hours elasticity of 2.9 is high relative to existing estimates. Chen et al. (2019) estimate driver-level labor supply elasticities for Uber drivers and find that the vast majority lie between one and two. Using experimental week-long reductions in Uber's commission, Angrist et al. (2021) estimate an intertemporal substitution elasticity of about 1.2, in line with experimental estimates for Uber drivers in Caldwell and Oehlsen (2022) and for other workers with flexible hours (Fehr and Goette, 2007; Mas and Pallais, 2017). Because these estimates are identified from transitory wage variation, while the AI tool raises the return to Platform work persistently—and responses to transitory wage changes weakly exceed responses to persistent ones—elasticities of one to two, if anything, overstate the relevant elasticity, implying that the amenity values below are, if anything, understated. A natural interpretation, supported by our qualitative evidence (Appendix H.1), is that the AI tool also raised the non-pecuniary

[3]Consistency with the functional form in Section 1.1: under the CES outside option, the elasticity of Platform hours with respect to the wage equals $\beta(1-h)/h$ at the worker's optimum. The observed elasticity $\hat{\gamma}_h/\hat{\gamma}_w \approx 2.9$ therefore pins down only combinations of β and h; β itself is not needed for the welfare calculation.

value of Platform work—reducing stress and difficult customer interactions.[4] Suppose the true hours elasticity is η and the tool raises the effective return to Platform time (wage plus amenity value a) rather than the wage alone. The observed hours response then reveals $\Delta \ln(wz + a) = \hat{\gamma}_h/\eta$, and the welfare gain becomes $(e^{\hat{\gamma}_h/\eta} - 1) + \frac{1}{2}\hat{\gamma}_h(\hat{\gamma}_h/\eta)$. Panel B reports results for $\eta \in \{1, 1.5, 2\}$, spanning the range of estimates above: the implied amenity is worth 1.8–7.5% of the hourly wage, and total worker surplus is 5.7–11.7% of pre-AI earnings. Under the amenity interpretation, Panel A is a lower bound on worker gains. The main force that could push in the opposite direction—workers targeting income and supplying hours beyond the point where the marginal hour is worth its opportunity cost—is inconsistent with our heterogeneity results by disability severity, as discussed in Section 3.6.

Heterogeneity by severity. Because equation (4) applies group by group, welfare gains by disability severity follow directly from the severity-specific wage and hours estimates (Figure 9): 7.4% of earnings for workers with profound hearing loss versus 1.3% for those with non-profound hearing loss (Panel C). The surplus from the AI tool is concentrated among the workers facing the steepest communication barriers.

Entry and retention margins. The calculations above apply to inframarginal (incumbent) workers. As shown in Appendix E.2, the AI tool also increased DHH entry onto the Platform by 27.5%. By revealed preference, each additional entrant gains between zero and the incumbent gain, and the same logic applies to marginal stayers.[5] Aggregate worker surplus therefore scales with the (larger) post-AI DHH workforce; the per-worker figures in Table 4 are conservative for incumbents and an upper bound for entrants.

Discussion and caveats. Taken together, these calculations imply that DHH workers gain the equivalent of roughly 4–12% of weekly earnings from the AI tool, alongside the Platform's 17% gain in per-worker profits: the benefits of the tool are shared between the Platform and its workers, with worker gains concentrated among the profoundly deaf. Three caveats apply. First, the standard errors treat the wage and hours DiD estimates as independent; a joint bootstrap on the microdata would account for their covariance. Second, for team workers, the accelerating piece-rate schedule implies part of the wage gain reflects movement along the pay schedule rather than higher z. Third, the welfare gains are weekly flows during our sample period and do not incorporate longer-run effects.

F Further Information on Hearing Loss Levels

Below is an English translation of Section 5.3 of GB/T 26341-2010, done by ChatGPT on 1/25/26. Note that drivers with HL of 40 dB or less are not issued a certificate of disability.

- *5.3.2 Hearing Disability Level I. The structure and function of the auditory system are extremely severely impaired. The average hearing loss of the better ear is greater than*

[4] In principle, an excess hours response could also reflect relaxed order rationing tied to pre-AI performance metrics. However, the mediation analysis in Appendix E.3 indicates that algorithmic reallocation is not an important mechanism, so we focus on the amenity interpretation. That evidence also supports the assumption underlying Panel A that workers freely choose their hours.

[5] If potential entrants' outside options are locally uniformly distributed near the entry margin, the average entrant gains approximately half the incumbent gain.

90 dB HL. Speech communication cannot be carried out relying on hearing. There are extremely severe limitations in activities such as understanding and communication, and extremely serious barriers exist in participation in social life.

- *5.3.3 Hearing Disability Level II. The structure and function of the auditory system are severely impaired. The average hearing loss of the better ear is between 81 and 90 dB HL. There are severe limitations in activities such as understanding and communication, and serious barriers exist in participation in social life.*

- *5.3.4 Hearing Disability Level III. The structure and function of the auditory system are moderately to severely impaired. The average hearing loss of the better ear is between 61 and 80 dB HL. There are moderate limitations in activities such as understanding and communication, and moderate barriers exist in participation in social life.*

- *5.3.5 Hearing Disability Level IV. The structure and function of the auditory system are moderately impaired. The average hearing loss of the better ear is between 41 and 60 dB HL. There are mild limitations in activities such as understanding and communication, and mild barriers exist in participation in social life.*

G Potential Nondisclosure of DHH Status

This Appendix discusses the issue of potential nondisclosure of DHH status by workers to the Platform, which determines whether a worker is classified as DHH in our data.

Why workers disclose. As noted in Section 1 of the main text, based on our qualitative work and understanding of the Platform, essentially all certified DHH workers disclose their disability to the Platform, for two main reasons.

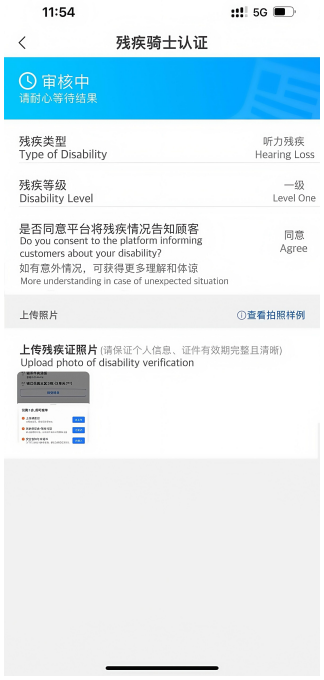
First, it is difficult to hide, especially for team workers. In the team work model, drivers have managers, and a worker's manager would need to know if they are DHH. Drivers in both work models (gig and team) also generally receive calls from customer service or customers, and the Platform must decide whether those calls should go through.

Second, DHH workers have a strong personal incentive to disclose their disability so that the Platform can accommodate them. People with disabilities in China are accustomed to disclosing their disability to employers and the government in order to receive potential accommodations. For example, if DHH workers do not disclose their status to the Platform, customers cannot be informed of their disability and may be less understanding in unexpected situations. Figure G.1 shows the screen from the registration process where workers can disclose their DHH status.

Comparison of severity between our data and the broader population. Table G.1 compares the distribution of hearing loss severity in our Platform sample to that among registered disability certificate holders in China. Workers with profound hearing loss (Level 1) account for 54% of DHH workers on the Platform, compared to 26% among certificate holders overall in 2012.[6]

[6]The 2012 data are from the *China Statistical Yearbook on the Work for Persons with Disabilities*. In more recent publications, severity levels are reported for all disability types combined rather than separately

Figure G.1: Screen During Registration Process where Workers Indicate Disability Status



Notes: This figure shows the original text, as well as text translated into English.

Table G.1: Distribution of Hearing Loss Severity

Source	Level 1	Level 2	Level 3	Level 4
Registered certificate holders (2012)	26%	24%	23%	26%
Our Platform sample	54%	16%	14%	16%

Notes: This table compares disability severity prevalence between the overall Chinese population and our data, which covers new DHH workers on the Platform between 2021-2023.

In principle, an analyst might try to use quantitative tests to detect nondisclosure of DHH status on the Platform, but this is extremely difficult. One might compare the share of DHH workers in our data to population shares, overall or by severity. However, such comparisons also reflect selection of DHH workers onto the Platform. While the share of drivers with profound HL is higher in our sample than in the population, we believe this primarily reflects selection based on outside options. As in many countries, DHH workers face substantial discrimination and barriers to employment, particularly those with profound HL, leading many to gravitate toward gig work.

Implications of potential nondisclosure. How would any nondisclosure of DHH status affect interpretation of our results? Any nondisclosure that is fixed over time (pre-

for hearing loss.

sumably the most natural kind) would not affect the main DiD results on the impact of the AI tool, since these specifications include worker fixed effects. Moreover, undisclosed DHH workers would not receive the AI tool—which was rolled out only to workers who disclosed their DHH status—and would thus remain in the control group throughout, creating no treatment misclassification. It would only affect the pre-AI comparisons, potentially making DHH workers appear less productive if higher-performing DHH workers were less likely to disclose their status.[7]

A different issue would arise if the AI tool itself affected disclosure of DHH status. While this is theoretically possible, none of our qualitative data or conversations with the Platform suggest that the AI tool made workers more likely to disclose their disability.[8]

H Qualitative Data

H.1 Interviews

We sought to better understand how DHH workers interact with the AI tool, how frequently they use it, and, more broadly, their daily routines, work practices, and the challenges they encounter. Thus, we conducted a group interview in summer 2024 with seven DHH workers, all team workers from a station in Hangzhou. Workers were recruited through the station manager, who asked for volunteers to discuss the experiences of DHH workers in food delivery. To facilitate communication with DHH workers, we hired a translator proficient in sign language. Workers were seated around a long table. Two coauthors took notes, along with representatives of the Platform. Workers received compensation from the Platform, around 300-400 CNY for their time. The group interview lasted around 1.5 hours.

Table H.1 summarizes the demographics of the workers interviewed. As shown in Table 1, most workers on the Platform, both DHH and non-disabled, are men. Our interview sample also consists entirely of men, but it is somewhat younger than the overall DHH worker population on the Platform. We therefore view it as broadly similar in gender composition, while noting that it skews younger in age. The DHH workers interviewed varied in severity — some had profound hearing loss and some had non-profound hearing loss — though we did not record the exact severity level for each interviewee. These workers had disclosed their disability to the Platform.

None of the 7 drivers wore hearing aids to the interview. This is consistent with our general observation that drivers on the Platform do not typically wear hearing aids while working. This is for multiple reasons, including: (1) hearing aids being costly, (2) hearing aids amplifying background noise (e.g., wind and traffic) and being uncomfortable while working, and (3) hearing aids being less helpful for some drivers with profound hearing loss.

AI tool usage and user experience. During the interviews, a representative of the Platform asked if workers were aware of the AI-based IOC tool and did they use it. All

[7] That is, it would work against finding that DHH do well for the Platform in terms of output and profits.

[8] In our data, DHH status is measured as of a person's last record in the dataset. If someone disclosed their DHH status as a result of the AI tool, they would be classified as DHH for the full period of our data. Our effects are identified from within-person changes in outcomes for DHH workers.

Table H.1: Summary of Interviewee Demographics

Age	Hometown	Tenure on Platform	Marital Status	Gender
22	Fujian Province	1+ years	Single	Male
26	Chongqing Municipality	3 years	Single	Male
27	Chongqing Municipality	4 years	Single	Male
24	Henan Province	3 months	Single	Male
25	Henan Province	3+ months	Single	Male
24	Henan Province	1+ years	Single	Male
23	Henan Province	1+ years	Single	Male

workers answered Yes to both questions. DHH workers were also asked what they thought about the tool, and they reported that the tool is useful for them.

DHH workers reported that the AI tool was especially useful for questions of where to leave the food, how to get into the building, and whether to knock on the customer's door (e.g., some customers could penalize drivers for unwanted knocking with a bad rating). This reduced misunderstandings and incidents, and it also reduced time waiting around for a customer's response. Workers are often quite anxious about bad ratings and late deliveries, and helping with these outcomes makes workers feel more comfortable and reduce stress related to these outcomes.

While some drivers interviewed had joined the Platform after the introduction of the AI tool, two drivers had started pre-AI. These drivers reported that there were significant bottlenecks pre-AI, such as not being able to reach out to customers upon delivery.

Outside job opportunities . The interviews suggested to us that DHH workers often lack strong outside options. When we asked the DHH workers interviewed about what other jobs they could do, they mentioned jobs such as assembly line work and food service work. Some explicitly said that there were not many other jobs available to them. A few also mentioned that they hoped to save money from food-delivery work so that they could eventually open a small business, such as a grocery store, in their hometown.

In addition to asking about AI usage and outside options, we list five other main questions we asked, as well as summarize the answers of the 7 DHH worker respondents. These are based on the notes taken by coauthors.

Q: What specific challenges have you encountered during deliveries?
A: Workers highlighted several challenges:

- One worker described an incident where a customer interaction required mediation by the station manager due to communication barriers. Despite the manager's assistance, the worker received a negative review, significantly affecting his morale and leading him to consider leaving the job.

- Workers reported examples of situations where customers left bad ratings after communication issues, citing reasons such as "not calling" or "not speaking clearly."
- Another worker expressed frustration with having multiple deliveries classified as significantly late deliveries (i.e., beyond the expected time windows). Such labels negatively impacted his performance record.
- A worker mentioned taking proactive measures during deliveries to reduce confusion, such as sending a pre-written explanatory message to customers with the station's contact information to address issues caused by delays or miscommunications.

Q: During pickup, do you encounter communication barriers?

A: In general, DHH workers felt much better about pickup relative to delivery. However, workers reported experiencing occasional communication barriers during pickups. Four respondents described needing to repeatedly remind restaurant staff—often two to three times—due to communication difficulties, while non-disabled riders typically resolve issues swiftly by shouting a single instruction that kitchen staff immediately understand. To remind kitchen staff about orders they needed to get, some DHH workers would bang on a counter and show their phones to the workers. Some workers perceived that non-disabled riders tend to receive meals faster because their vocal presence exerts greater urgency on kitchen staff.[9]

Q: Have you experienced forced stops during deliveries?

A: Four workers reported that switching between Platform apps, such as moving from navigation to the order interface, frequently results in forced stops. One rider specifically noted that managing two phones simultaneously—one dedicated to navigation and another to orders—necessitated frequent stops, causing significant disruptions.

Q: Do you receive advance notifications about potentially overdue orders?

A: All workers acknowledged receiving in-app reminders, but due to the nature and pace of their work, they rarely had sufficient time to read or react to these notifications.

Q: Do you have access to rest facilities or water during work?

A: Workers reported using specialized restroom-locating maps to find nearby public facilities for quick rest and water. They seldom return to the station for breaks.

H.2 Fieldwork

To deepen our understanding of challenges faced by DHH delivery workers, we undertook immersive fieldwork alongside them. Since the field notes cover some topics not relevant to the present research (e.g., how workers swap electric motorcycle batteries, discounts restaurants offer to delivery workers), for brevity, we omit some of these from the Appendix. We conducted two practical explorations:

1. One coauthor spent several days directly observing 3 DHH workers by accompanying them full-time, seated on the back seat of an electric motorcycle during their delivery

[9] While workers perceived communication barriers during pickup, this did not translate into a statistically significant pickup-time difference in our quantitative analysis (see Figure 1).

shifts. This hands-on approach enabled us to closely observe and document DHH workers' AI use, daily routines, and difficulties they encounter. The first was a pilot, during which we worked for around an hour as a starting point, whereas the other two sessions lasted about 11-12 hours each.

2. To gain first-hand insight into the communication challenges faced by DHH workers, a different coauthor conducted his delivery work wearing earplugs to mimic hearing impairment (see Panel (b) of Figure H.1), at the invitation and encouragement of the station manager.

While waiting for orders during off-peak hours, workers would often gather in certain areas, such as a specific place in the plaza. Among the DHH workers we met there, both gig workers and team workers had disclosed their disability status to the Platform.

In the text that follows, the first person "I" refers to the perspective of a coauthor instead of a subject. All photographs were taken by the author group. Images were converted to black-and-white to anonymize the color of the Platform's uniforms, and all Platform-identifying information was removed to preserve anonymity. To protect their anonymity, worker faces are obscured, whereas coauthor faces are retained.

Fieldwork Records, July 1, 2024, Hangzhou

Subject: Male, 26 years old, 3 years of employment, non-profound hearing loss
Duration: Around 1 hour (3:00 PM to 4:00 PM); 3 orders delivered
Observations:

1. During the pilot, we observed that upon arrival at, or near, the delivery destination, the worker usually used the AI tool to place a call to the customer. Sometimes the screen displayed a notification that the customer did not answer the call; in those situations, the worker repeatedly used the function "initiate IOC" to send the message again.

2. In cases where there were multiple orders to be delivered to one place, the worker could use the AI tool to simultaneously send out IOCs for multiple orders.

3. There was one case in which we noticed that the worker paid particular attention to traffic safety. During a delivery, we were about to enter a small tunnel. The worker stopped at the entrance, honked the horn first, even though he himself could not hear it, and then entered the tunnel cautiously.

Fieldwork Records, July 2, 2024, Hangzhou

Subject: Male, 27 years old, 3 years of employment, non-profound hearing loss
Duration: 10 hours 50 minutes (10:00 AM – 8:50 PM), 33 orders delivered
Observations:

1. In total, the worker delivered around 33 orders during the hours observed, and I observed that he used the AI tool for all deliveries. He was quite familiar with the tool

Figure H.1: Immersive Fieldwork with Delivery Workers



(a) Coauthor accompanying a DHH worker during deliveries



(b) Coauthor delivering orders with earplugs to simulate hearing impairment

and used it smoothly. I also briefly discussed his experience with the AI tool with him, communicating through a chat app similar to WhatsApp. The worker explained that communicating with customers at the final delivery stage had become much more effective. In most cases, after using the intelligent outbound calling feature, he only needed to drop off the food at the designated location, without further interaction or communication.

2. Delivery workers tend to gather at a common meeting point after completing their orders. For instance, we were near Hangzhou 303 Life Plaza, where many workers frequently returned. Workers tend to avoid orders exceeding 4 kilometers, and after delivering such "distant orders," they usually navigate back to common gathering points.

3. A common tactic I noticed was workers pressing multiple elevator buttons when making deliveries. For instance, if a worker had to deliver to the 20th floor, he would also press buttons for the 21st and 22nd floors. This strategy allows them to reduce waiting time on the way down.

4. Over the course of a day, I noticed that many of the deliveries were from the same restaurants or to the same apartment complexes. The worker often made repeated trips to the same locations throughout the day.

5. During the lunchtime peak hours, the worker was handling six orders simultaneously, which led to a delay in one of the deliveries. After this delay, the number of orders assigned to them decreased significantly. One reason for the delay was that the second order was held up due to the restaurant's slow preparation and packaging process. The authors confirmed with the platform that no additional time is granted for subsequent orders the worker is simultaneously carrying in such situations. Another cause was the worker's decision to wait for the delayed order instead of collecting food for the other orders from nearby restaurants in the same mall. The worker could have optimized their pickup sequence to avoid unnecessary waiting.

6. The top-performing worker at this station managed to complete approximately 2,823 orders in a month, earning around 20,000 CNY ($2,800 USD). However, this required working over 16 hours a day. At times, the worker would fall asleep on their electric motorcycle between orders, waking up only when the next one was assigned. Occasionally, they would even fall asleep while waiting at traffic lights.

Fieldwork Records, July 7, 2024, Hangzhou

Subject: Male, 24 years old, 1 year of employment, profound hearing loss
Duration: 12 hours 20 minutes (7:45AM – 8:05PM), 75 orders delivered
Observations:

1. Among all the orders delivered, I observed that the worker used the AI tool for most orders. For some orders, customers had left specific remarks when placing the order, such as "no need to contact me, just leave the order at the front door." (around 2 -

5% based on my observation). In these cases, the worker did not use the IOC feature. Instead, he simply took a photo of the food upon delivery and uploaded it to the app as proof of delivery.

2. At 7:45 AM, we started the day. The worker provided me with a full set of sun protection gear, including a sun-protection suit and a neck gaiter. The temperature that day ranged from 38-40C. Figure H.3 shows the equipment required for working outdoors.

3. Around 9 AM, we received an order for a large watermelon. However, the mall where the order was placed didn't open until 10 AM, so we had to pick it up from a locker in the underground parking lot. The security guard wouldn't allow us to ride the motorcycle down to the parking lot, so we had to run from the garage entrance to the basement, grab the watermelon, and run back up to avoid a late delivery.

4. Later in the morning, we received an order for McDonald's. The Platform indicated that the food should be picked up from a collection locker. When we entered the store, a staff member in the kitchen shouted, "The food is on the counter," but the worker didn't hear and went straight to the locker. After unsuccessfully entering the pick-up code several times, the worker finally went to the counter to communicate with the staff and retrieve the meal.

5. The intelligent outbound call is frequently used when workers request customers to open doors, and it is generally very effective. However, there was one exception during our deliveries. In this case, when we arrived at the destination, the worker followed the customer's instructions to contact the apartment complex management via the intercom. After the system connected, no one spoke on the other side. Although I heard the door unlocking sound, the door did not automatically pop open, and there was no visible indication that it had unlocked. The worker continued waiting by the intercom until I eventually walked over and pulled the door open.

6. One order was delivered to an urban village [chengzhongcun in Mandarin], but the address was inaccurate. Upon arriving, the worker used the intelligent outbound call to contact the customer, who messaged back, "I'm upstairs." However, it was unclear what the customer meant by "upstairs." I then heard the customer shouting from inside a factory behind us, but by that time, the worker had already started riding in the opposite direction. I stopped the worker, informing him that the customer was inside the factory, and we managed to find the customer. In such situations, a worker with a hearing disability could waste significant time locating the exact address. Later that afternoon, we encountered a similar issue with an order delivered to a construction site.

7. There was another incident where a restaurant incorrectly packed an order. The customer had ordered noodles and fresh tomatoes, but the restaurant forgot to include the tomatoes. The restaurant called the worker during delivery; as the worker could not hear the caller, the phone was passed to the accompanying coauthor to handle the conversation. We were already halfway to the destination, and returning would have

delayed the delivery, so I explained to the restaurant that we were near the destination and couldn't return, and suggested they contact the customer directly to explain the situation.

8. In general, when a situation requires in-depth communication with restaurants or customers, a worker with a hearing disability still faces difficulties and takes significantly more time to resolve the issue compared to a non-disabled worker. While these types of problems are difficult to address, they occur rarely. In most cases, the intelligent outbound calls are sufficient to assist workers.

9. The worker I worked with mentioned, "I typically start running orders at 7 AM and take a break around 3 PM. Then I go back online from 4 PM until 8 or 9 PM." My observations and experiences working alongside different full-time workers were similar. According to them, to earn a decent income, they need to work 10-13 hours a day.

10. Below is a summary of orders that I found particularly difficult for workers to deliver:

    - Some residential complexes have very slow elevators (there are many residents, but very few elevators).
    - Some older complexes do not have elevators.
    - Security guards at certain complexes do not allow workers to enter with electric motorcycles.
    - Apartments in large residential complexes can be difficult to navigate.
    - Delivering to urban villages is particularly challenging, as the locations are difficult to find and often require climbing stairs without elevators.

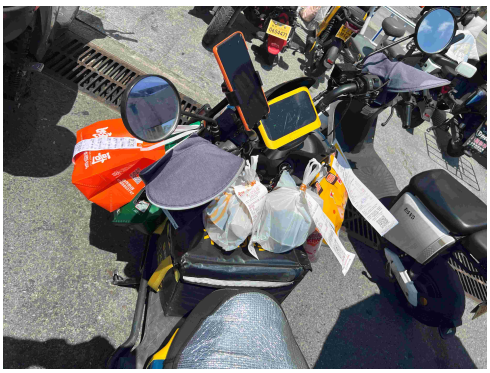
Fieldwork Records, July 2024, Hangzhou

Subject: None. Coauthor delivering by himself wearing earplugs
Duration: Part-time, multiple days; 3–5 orders per session
Observations: For the coauthor who delivered orders while wearing earplugs, he worked part-time over multiple days and delivered 3–5 orders each time entirely on his own. He used the AI tool for all orders and found it highly effective.

Because customers expect to receive phone calls upon delivery, the inability to communicate by phone can create significant difficulties for DHH workers. Thus, to better understand this challenge, on some early orders the coauthor also tried both in-app messaging and direct phone calls without speaking, intentionally simulating the communication constraints faced by DHH workers before the AI implementation. In-app messages were often ignored by customers or answered only after a delay, usually several minutes, which was particularly frustrating because the worker needed to continue delivering other orders. Direct phone calls without speaking were also uncomfortable for the coauthor. Customers would often say "hello?" multiple times and, based on their tone, appeared to become increasingly impatient or irritated. After AI implementation, the direct in-app phone call feature was replaced by the AI tool for DHH workers; thus, the coauthor used direct phone calls only to understand the communication environment faced by DHH workers before the AI tool became available.

In this context, the IOC feature was very helpful and easy to use at the delivery stage. As a new worker, the coauthor was unfamiliar with many delivery destinations and often arrived close to the expected arrival time. As a result, saving even one minute, or sometimes just a few seconds, by using the AI tool instead of waiting for in-app message responses or making direct phone calls significantly reduced stress and helped prevent late deliveries.

Figure H.2: Concurrent Orders During Peak Hours



Notes: This figure illustrates multiple orders being simultaneously handled by a delivery worker during peak hours. Experienced workers typically manage 6-10 concurrent orders during peak hours. The maximum number of orders a rider can handle simultaneously is set by the Platform's algorithm, considering factors like tenure and past performance. Further details on the assignment algorithm are discussed in Section 1.

Figure H.3: Essential Equipment for Outdoor Work



Notes: Figure H.3 illustrates the necessary protective gear for outdoor delivery work, including a sun-protection suit, helmet, and neck gaiter. Typical summer temperatures in the focal city range from 38C-40C. Complete sun-protection equipment is essential, and workers must remain well-hydrated to avoid heat-related illnesses.

Appendix References

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb, *Prediction Machines: The Simple Economics of Artificial Intelligence*, Boston, MA: Harvard Business Review Press, 2018.

Angrist, Joshua D., Sydnee Caldwell, and Jonathan V. Hall, "Uber versus Taxi: A Driver's Eye View," *American Economic Journal: Applied Economics*, 2021, *13* (3), 272–308.

Arcot, Siva Venkatesh, "Comparative Analysis of AI-powered Outbound Dialer Campaigns vs. Legacy Outbound Dialer Campaigns (Contact Center)," *Asian Journal of Research in Computer Science*, Feb. 2025, *18* (3), 73–84.

Azar, José, Steven Berry, and Ioana Marinescu, "Estimating Labor Market Power," Working Paper 30365, National Bureau of Economic Research 2022.

Bassier, Ihsaan, Arindrajit Dube, and Suresh Naidu, "Monopsony in Movers: The Elasticity of Labor Supply to Firm Wage Policies," *Journal of Human Resources*, 2022, *57* (S), S50–S86.

Brickey, M. and K. Campbell, "Fast Food Employment for Moderately and Mildly Retarded Adults: The McDonald's Project," *Mental Retardation*, 1981, *19* (3), 113–116.

Caldwell, Sydnee and Emily Oehlsen, "Gender Differences in Labor Supply: Experimental Evidence from the Gig Economy," *University of California, Berkeley Working Paper*, 2022.

Chen, M. Keith, Peter E. Rossi, Judith A Chevalier, and Emily Oehlsen, "The Value of Flexible Work: Evidence from Uber Drivers," *Journal of Political Economy*, 2019, *127* (6), 2735–2794.

Chen, Sanyuan, Shujie Liu, Long Zhou, Yanqing Liu, Xu Tan, Jinyu Li, Sheng Zhao, Yao Qian, and Furu Wei, "VALL-E 2: Neural Codec Language Models are Human Parity Zero-Shot Text to Speech Synthesizers," *arXiv preprint arXiv:2406.05370*, 2024.

Cimera, Robert E., "The Monetary Benefits and Costs of Hiring Supported Employees: A Primer," *Journal of Vocational Rehabilitation*, 2002, *17* (1), 23–32.

_ , "The Monetary Benefits and Costs of Hiring Supported Employees: Revisited," *Journal of Vocational Rehabilitation*, 2006, *24* (3), 137–144.

Dal Bo, Ernesto, Frederico Finan, and Martín A. Rossi, "Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service," *Quarterly Journal of Economics*, 2013, *128* (3), 1169–1218.

Dean, Evan E., Veerle Garrels, Hanne Marie Høybråten Sigstad, and Christian Wendelborg, "Employer Willingness to Hire People with Intellectual Disability: The Influence of Employer Experience, Perception of Work, and Social Responsibility," *Journal of Vocational Rehabilitation*, 2022, *57* (1), 85–95. doi:10.3233/JVR-221201.

Emanuel, Natalia and Emma Harrington, "The Payoffs of Higher Pay: Labor Supply and Productivity Responses to a Voluntary Firm Minimum Wage," Staff Report 1182, Federal Reserve Bank of New York 2024.

Fehr, Ernst and Lorenz Goette, "Do Workers Work More if Wages Are High? Evidence from a Randomized Field Experiment," *American Economic Review*, 2007, *97* (1), 298–317.

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, Megan Schouweiler, and Michael Westberry, "IPUMS CPS: Version 12.0 [dataset]," 2024. https://doi.org/10.18128/D030.V12.0.

Forth, John and Melanie K. Jones, "The Disability Pay Gap Within and Across Firms," IZA Discussion Paper 17679, Bonn, Germany February 2025.

Graffam, Joseph, Kaye Smith, Alison Shinkfield, and Udo Polzin, "Employer Benefits and Costs of Employing a Person with a Disability," *Journal of Vocational Rehabilitation*, 2002, *17* (4), 251–263.

Hernandez, Brigida and Katherine McDonald, "Exploring the Costs and Benefits of Workers with Disabilities," *Journal of Rehabilitation*, 2010, *76* (3), 15–23.

Hunt, Andrew J. and Alan W. Black, "Unit Selection in a Concatenative Speech Synthesis System Using a Large Speech Database," in "Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)" 1996, pp. 373–376.

Jing, Jiaqi, Xiaoqing Feng, Jianbo Song, and Boya Li, "Does the Inclusion of Disabled Employees Affect Firm Performance? Empirical Evidence from China," *Sustainability*, 2022, *14* (13), 7835.

Kaletta, J. P., D. J. Binks, and R. Robinson, "Creating an Inclusive Workplace: Integrating Employees with Disabilities into a Distribution Center Environment," *Professional Safety*, 2012, *57* (6), 62–71.

Klatt, Dennis H., "Review of Text-to-Speech Conversion for English," *Journal of the Acoustical Society of America*, 1987, *82* (3), 737–793.

Liao, Juan, Man Gao, Xiji Zhu, and Yu Yang, "Productivity Difference or Discrimination Effect?

Disability-Related Wage Gap in China," *Asia Pacific Journal of Human Resources*, 2024, *62* (2), e12400. doi:10.1111/1744-7941.12400.

Lindsay, Sally, Elaine Cagliostro, Mikhaela Albarico, Neda Mortaji, and Leora Karon, "A Systematic Review of the Benefits of Hiring People with Disabilities," *Journal of Occupational Rehabilitation*, 2018, *28* (4), 634–655.

Madera, Juan M., D. Christopher Taylor, and Nelson A. Barber, "Customer Service Evaluations of Employees with Disabilities: The Roles of Perceived Competence and Service Failure," *Cornell Hospitality Quarterly*, 2020, *61* (1), 5–18. doi:10.1177/1938965519882315.

Malo, Miguel Á. and Ricardo Pagán, "Wage Differentials and Disability Across Europe: Discrimination and/or Lower Productivity?," *International Labour Review*, 2012, *151* (1–2), 43–60.

Mas, Alexandre and Amanda Pallais, "Valuing alternative work arrangements," *American Economic Review*, 2017, *107* (12), 3722–59.

Narayanan, Sriram and Ed Terris, "Inclusive Manufacturing: The Impact of Disability Diversity on Productivity in a Work Integration Social Enterprise," *Manufacturing & Service Operations Management*, 2020, *22* (6), 1112–1130.

Piramanayagam, Senthilkumaran and Partho Pratim Seal, "Employers' Attitudes and Hiring Intentions towards Persons with Disabilities in Hotels," *Disability, CBR & Inclusive Development*, 2021, *31* (4), 116–133. doi:10.47985/dcidj.433.

Rosenbaum, Mark Scott, Rojan Baniya, and Tali Seger-Guttmann, "Customer Responses towards Disabled Frontline Employees," *International Journal of Retail & Distribution Management*, 2017, *45* (4), 385–403. doi:10.1108/IJRDM-08-2016-0133.

Saha, Sabhasachi, Abhimanyu Sharotry, Francis Méndez Mediavilla, Jesús Jiménez, and Eduardo Pérez, "Productivity and Reliability of People with a Disability: A Literature Review," in "Proceedings of the IISE Annual Conference & Expo 2023" 2023.

Seva, Rosemary R., "A Productivity Assessment of PWD Employees in a Philippine Company," *DLSU Business & Economics Review*, 2020, *29* (2), 93–103.

Shah, Shariq, Hossein Ghomeshi, Edlira Vakaj, Emmett Cooper, and Shereen Fouad, "A Review of Natural Language Processing in Contact Centre Automation," *Pattern Analysis and Applications*, 2023, *26* (3), 823–846.

Shen, Jonathan et al., "Natural TTS Synthesis by Conditioning WaveNet on Mel Spectrogram Predictions," in "Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)" 2018, pp. 4779–4783.

Sokolova, Anna and Todd Sorensen, "Monopsony in Labor Markets: A Meta-Analysis," *ILR Review*, 2021, *74* (1), 27–55.

Tan, Xu, Tao Qin, Frank Soong, and Tie-Yan Liu, "A Survey on Neural Speech Synthesis," *arXiv preprint arXiv:2106.15561*, 2021.

Taylor, Paul, *Text-to-Speech Synthesis*, Cambridge, UK: Cambridge University Press, 2009.

van den Oord, Aaron, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu, "WaveNet: A Generative Model for Raw Audio," *arXiv preprint arXiv:1609.03499*, 2016.

_ , Yazhe Li, Igor Babuschkin, Karen Simonyan, Oriol Vinyals, Koray Kavukcuoglu, George van den Driessche, Edward Lockhart, Luis C. Cobo, Florian Stimberg, Norman Casagrande, Dominik Grewe, Seb Noury, Sander Dieleman, Erich Elsen, Nal Kalchbrenner, Heiga Zen, Alex Graves, Helen King, Tom Walters, Dan Belov, and Demis Hassabis, "Parallel WaveNet: Fast High-Fidelity Speech Synthesis," in "Proceedings of the 35th International Conference on Machine Learning (ICML)" 2018, pp. 3918–3926.

Vapi, "A History of Text-to-Speech: From Mechanical Voices to AI Assistants," https://vapi.ai/blog/history-of-text-to-speech June 2025. Accessed July 10, 2026.

Wang, Yuxuan, RJ Skerry-Ryan, Daisy Stanton, Yonghui Wu, Ron J. Weiss, Navdeep Jaitly, Zongheng Yang, Ying Xiao, Zhifeng Chen, Samy Bengio, Quoc Le, Yannis Agiomyrgiannakis, Rob Clark, and Rif A. Saurous, "Tacotron: Towards End-to-End Speech Synthesis," in "Proceedings of Interspeech" 2017, pp. 4006–4010.

Zen, Heiga, Keiichi Tokuda, and Alan W. Black, "Statistical Parametric Speech Synthesis," *Speech*

Communication, 2009, *51* (11), 1039–1064.

Zhang, Ying, Sriram Narayanan, Tharo Soun, Kalyanmoy Deb, and Dustin Cole, "Maximizing Disability Diversity, Language Diversity, and Productivity: A Study in Apparel Manufacturing," *Production and Operations Management*, 2023, *32* (12), 3783–3800.

Zwerling, Craig, Paul S. Whitten, Charles S. Davis, and Nancy L. Sprince, "Occupational Injuries Among Workers with Disabilities: The National Health Interview Survey, 1985–1994," *JAMA*, 1997, *278* (24), 2163–2166.