Dipartimento di Economia Marco Biagi

UNIMORE UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

DEMB Working Paper Series

N. 226

Task search and labour supply in the platform economy

Michele Cantarella[1], Chiara Strozzi[2]

October 2023

[1] MT School for Advanced Studies Lucca and University of Helsinki
Email: michele.cantarella@imtlucca.it
[2] University of Modena and Reggio Emilia and IZA
Email: chiara.strozzi@unimore.it

ISSN: 2281-440X online

Task search and labour supply in the platform economy*

Michele Cantarella[†] [1] and Chiara Strozzi[2]

[1]IMT School for Advanced Studies Lucca and University of Helsinki, email: michele.cantarella@imtlucca.it

[2]University of Modena and Reggio Emilia and IZA, email: chiara.strozzi@unimore.it

Abstract

In this paper, we study the wage elasticity of labour supply of platform workers by exploiting uncertainty in task search. Using a survey of on-location and online platform workers in Europe, we show that wage reductions due to task search are inversely related to increases in labour supply and that changes in earnings net of task search are also inversely related to labour supply. Our estimated backwards-bending labour supply curves are valid for all platform workers and are robust to a number of misspecification and endogeneity issues.

Keywords: *Platform economy, Labour supply, Job search, Reference Dependence, Monopsony*

JEL codes: J22, J33, J42, D83

*The authors would like to thank the participants of the UniMoRe Seminar Series on Political Economy, the 2023 Royal Economic Society Annual Conference, the 43rd Annual Meeting of the Finnish Economic Association, the 2022 AASLE conference, the Helsinki GSE Labor and Public Economics seminar series, the Labour Institute for Economic Research seminar series, the 2022 ECEFG Competition Policy Workshop, and the 2022 AIEL conference, along with members of the Digital Economy Unit at the JRC Growth and Innovation directorate, for their useful comments and suggestions. The data used in this analysis was collected by PPMI for the report "Study to support the impact assessment of an EU Initiative on improving the working conditions of platform workers" issued by DG EMPL - European Commission (No VC/2021/0093) (Barcevičius et al., 2021). The authors of this paper collaborated on the report in the context of a service request. The views expressed in the paper are the authors' own and do not necessarily reflect those of DG EMPL or PPMI.

†Corresponding author.

1 Introduction

Workers earning on a piece-rate basis are among the categories of workers who are the most exposed to income uncertainty. This group of workers comprises the self-employed (Parker et al., 2005) as well as any worker whose pay is related to current output, from high-skill professionals (Hart, 2008; Hart and Roberts, 2012) to taxi drivers (Camerer et al., 1997). Income uncertainty arises because, in piecework, the hours of paid work can vary depending on demand conditions.

The recent expansion of the platform economy extends income uncertainty to a vast new pool of workers who generate income by selling their services through online platforms (International Labour Office, 2021; Barcevičius et al., 2021). In most cases, these people are not directly employed by the platforms but are part of a pool of autonomous workers contracted on a task-by-task basis. Platforms act as matchmakers between clients and workers, who are paid a fixed reward on task completion.

Platform workers can offer their services either online or in person. The former are generally defined as online platform workers, including people working in low-skill micro-task platforms (i.e., Amazon Mechanical Turks, Crowdflower, etc.) and medium/high-skill freelancers (i.e., Upwork, Fiverr, etc.). The latter are better known as on-location (or on-demand) platform workers, usually providing services in person (i.e., Foodora riders, Uber drivers, and Taskrabbit handymen).[1]

Most of the people selling services through platforms are not employed by their platforms but rather are contracted as autonomous workers. Pay levels are usually set by the platform or are left to clients to decide. Workers are usually paid on task completion with a fixed reward, and entry requirements are often minimal.[2] As a matter of fact, once they join the platform, workers can work as long as they like, as long as there is demand for the tasks they

[1] Platform workers are a large subset of people generating income through online platforms, which also includes e.g. people selling goods on Ebay, or renting apartments through Airbnb. The distinction can be subtle, but platform workers are defined as such because they sell their own job services through online platforms and do not sell goods or accrue rents. For a comprehensive overview of online platforms and people finding work through them, see Berg et al. (2018).

[2] For example, Uber drivers might need to obtain a licence or show that they have access to a vehicle before entering a platform. See De Stefano and Aloisi (2018) for an overview of the contractual framework of platforms and their entry requirements.

perform.[3] Moreover, since most platform workers are self-employed, platforms do not have to sustain hiring/dismissal costs and/or provide fixed-hours contracts (Dube et al., 2020) while benefiting from perfect complementarity between labour and capital (as workers provide both their work and capital). These features make online platforms the ideal candidate for studying workers' behaviour in contexts with virtually no labour supply restrictions.

Labour market frictions can still emerge from demand-supply mismatches directly affecting the job search and match-making processes. Indeed, while algorithms can greatly improve the quality and speed of matching (Horton, 2017), these benefits are usually enjoyed by platforms and clients only, with workers facing high levels of unpaid work in the form of job search (Bogliacino et al., 2020), all facts which underline the monopsonistic nature of these markets (Dube et al., 2020; Kingsley et al., 2015).

The job search effort in online platforms is usually made at the expense of leisure.[4] As search inflates the hours of work but leaves the pay-out unchanged, workers are exposed to a source of uncertainty in the hourly pay rate through search shocks determining how long workers need to search for an extra hour of paid work. As a result, platform workers often end up spending more time than expected on the platform because of the time spent searching for available tasks, with workers working longer hours for lower pay and less when the pay is higher (Cantarella and Strozzi, 2021; Berg et al., 2018).

This stylised fact alone would suggest a backwards-bending labour supply curve at the intensive margin, thus implying a negative relationship between pay and how much time people devote to work in a given time span. This behaviour, which has also been observed among workers who would not be conventionally considered platform workers (such as Youtube creators in Barbos and Kaisen, 2022), could be explained by the possibility that workers would be working towards pre-defined earnings targets (Kukavica et al., 2022; Horton and Chilton,

[3] While platforms, either directly or indirectly, often exercise a significant degree of control over the activities of people finding work through them, workers usually retain extensive autonomy in setting their work schedules.

[4] We define leisure as any moment a worker can spend without being interrupted by work on short notice. We do not treat search time during platform work as leisure because, during this time, workers are still "on-call" and make themselves available to accept tasks. Often, tasks are assigned automatically or on a first-come, first-served basis. If a platform worker does not accept a task offer quickly, someone else will take the task. The implication is also that workers who want to make themselves available for work cannot plan to devote this search time to other activities.

2010), suggesting reference-dependent preferences. However, evidence on the wage elasticity of supply using data from the workers' side remains scarce and mixed, and it is also not supported by the same estimates produced with data from the platform side (Dube et al., 2020; Duch-Brown et al., 2022), which point at positive but inelastic (smaller than 1) wage elasticities of labour supply, at least for what concerns a subset of people working on online labour platforms (i.e. micro-task workers).[5]

Experimental trials focusing on settings comparable to platform economy jobs suggest similar conclusions. Bouhlel et al. (2022), recreating a sequential task setting, find that previous payoff losses lead to increases in the search effort for the next task, suggesting reference dependence. Orland and Rostam-Afschar (2021), recreating flexible arrangements, find that people can respond to uncertainty in wage by saving time or reallocating work to another work shift, but that the two actions are not perfect substitutes and depend on idiosyncratic loss aversion.

It is important to point out that the ambiguity about the slope of the labour supply curve can go beyond platform work. Indeed, negative wage elasticity estimates have also been observed among other groups of workers, especially the self-employed. Some have attributed these estimates to target-earning behaviour (Wales, 1973; Camerer et al., 1997; Martin, 2017). In contrast, an alternative interpretation suggests that negative estimates could result from wage uncertainty and that backwards-bending labour supply might be structural to self-employment (Parker et al., 2005). However, while the target earning hypothesis can be criticised on the basis of endogeneity issues such as division bias (Stafford, 2015; Farber, 2005), experimental studies have shown that backwards-bending estimates are plausible and occur among loss-averse individuals (as in Fehr and Goette, 2007). Recent studies (such as Zubrickas, 2023) have tried to reconcile this evidence by arguing that the heterogeneity in these estimates might result from heterogeneous responses to absolute and relative changes in the wage rate.

This debate is still ongoing in part because of the rise of the platform economy itself, which

[5]While in the case of self-employed workers "earnings elasticity" would be a more appropriate term than "wage elasticity", in this paper we will use the terms "wage", "earnings" and "pay rate" interchangeably in order to not depart from the nomenclature of the literature on self-employed labour supply.

has offered many new case studies on piece-rate workers to focus on. However, most studies on the labour supply of platform workers usually disregard uncertainty and make the implicit or explicit assumption that workers can forecast wage shocks at the beginning of each hour of work (as is the case for Chen et al., 2019). While some studies have tackled uncertainty in compensation in platforms (Butschek et al., 2022), the behavioural effects of uncertainty in task search remain mostly unexplored.

In this paper, we exploit variation in task search to study the behaviour of online platform workers to study how much uncertainty and idiosyncratic expectations affect workers' behaviour. Our empirical analysis takes advantage of a recent online survey on European platform workers conducted in 2021, which captures precise information on key variables for our study: paid hours of work, desired labour supply, earnings, and time devoted to search for the available tasks. We estimate two wage elasticity parameters. The first one is the (end-of-the-week) *actual wage elasticity*, which reveals how workers react to variations in wages caused by task search shocks during the working week. This elasticity captures uncertainty in earnings. The second one is the (start-of-week) *frictionless wage elasticity*, which reveals how much hourly earnings affect workers' supply *before* task search shocks are revealed, capturing heterogeneity in the expected level of earnings. Studying these elasticities independently allows us to settle the issue of the wage elasticity of platform workers and evaluate the role of reference dependence. Exploiting variation in search and hourly earnings in a novel empirical setting, we find that wage elasticity parameters are negative and inelastic and that both play a central role in explaining variation in labour supply. Our results are accompanied by a set of rigorous robustness checks that suggest that our estimates are robust to all sources of heterogeneity and endogeneity in earnings and in the search shock.

The contribution of our paper is threefold. First, our findings provide evidence on the labour supply elasticity of platform workers using supply-side data in an effort to better understand how task search uncertainty affects labour supply. Indeed, while most of the reference literature has focused on the monopsony power of digital labour platforms (Dube et al., 2020; Kingsley et al., 2015) and on the organisational implications of piecework (Lehdonvirta, 2018;

Alkhatib et al., 2017; Davis and Hoyt, 2020), so far no explicit attempt to analyse the role of task search - and piecework, in a more general sense - has been made. Our focus offers a novel insight into the nature of the platform economy that could be easily extended to other contexts, as our work contributes to the aforementioned debate on the labour supply of the self-employed workers. In this context, we also shed new light on the literature on the economics of piecework (Hart, 2008), on unpaid overtime (Bell and Hart, 1999) and on uncertainty in self-employment (Parker et al., 2005).

Secondly, we develop an intuitive yet novel method for estimating labour supply elasticities for all types of platform workers, which exploits the variation in actual and desired hours of work conditional on the intensity of uncertainty in task search. Our empirical approach belongs to the family of difference-in-differences methods and aims at absorbing away individual fixed effects by focusing on the comparison between actual and desired hours of work. Studying labour supply by comparing actual and desired hours of work is not new to the literature on "traditional" labour markets (see, for example, Euwals and van Soest, 1999; Stewart and Swaffield, 1997). However, in these contexts, demand shocks do not enter labour supply as search shocks but instead affect workers' employment status and contract type. This is not the case for platform work, as job search frictions are almost the only difference between desired and actual working hours, making platforms the ideal candidate for this kind of exercise. Similarly, demand-side search frictions (i.e. how long clients wait for the fulfilment of the task they ask for) have been exploited to study demand-side utility surplus in on-demand labour platforms such as Uber (Cohen et al., 2016; Lam et al., 2021) and labour supply in online labour platforms such as Amazon Mechanical Turks (Dube et al., 2020). Our work contributes to this literature by studying these frictions on the labour supply side, focusing on how long workers (instead of clients) wait.

Finally, our paper contributes to the ongoing policy discussion on online labour platforms. These platforms have recently garnered the attention of policy-makers across the world, with national courts and legislators working towards reassessing the employment status of platform workers (for an overview, see De Stefano et al., 2021). In Europe, EU institutions initiated this process in 2016 with the adoption of a European Agenda for the Collaborative Economy[6] as a part of an ongoing initiative to improve the working conditions of platform economy workers, recently culminating into the proposal of a EU Directive pushing towards reclassification of many autonomous platform workers into paid employees.[7] Advancing the theoretical framework and empirical evidence on the economics of the online labour markets is then paramount to better inform policy-makers and understand how these regulations can affect work in these platforms.

The paper is structured as follows. Section 2 discusses the economic implications of piecework from a theoretical perspective. Section 3 describes our data sources, while Section 4 presents our empirical approach and Section 5 discusses the conditions under which the estimates are valid and consistent. Results are discussed in Section 6, while Section 7 concludes. The Appendix offers additional statistics and robustness checks, along with Monte Carlo simulations testing the validity of our approach in a controlled setting.

2 Theoretical framework

In this section, we present a model of labour supply in the presence of job search and uncertainty. The model's main purpose is to show how the introduction of uncertainty in task search leads to difficulties in reaching an analytical solution to supply optimisation, motivating our empirical approach in Section 4.[8] The model we adopt draws from two main theoretical contributions: Arellano and Meghir (1992) and Parker et al. (2005). We rely on the former to specify labour supply in the presence of task search and refer to the latter to model uncertainty in task search. While we borrow from both sources, we also depart from each one by using more specific assumptions that could better fit the functioning of online platform labour

[6] Communication From the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions; A European agenda for the collaborative economy; COM/2016/0356 final.

[7] Proposal for a Directive of the European Parliament and of the Council on improving working conditions in platform work; COM/2021/762 final.

[8] The present section is not essential for understanding the empirical findings and implications of the research. Therefore, readers more interested in the practical aspects of the study can skip to the next section, where the research design and methods are described.

markets. Following these approaches to show how search shocks reduce to multiplicative wage shocks with the same implications for supply optimisation studied in Parker et al. (2005) is also valuable.

In our theoretical model, we assume that job offers (i.e., available tasks) appear directly in the platform (or are sent to the worker) as soon as they are available. Each task (job) is compensated by an advertised reward, which is known to the worker. In the reference period t, define as $\bar{w}$ average pay-out for the tasks that worker i can potentially perform.[9] This pay-out is determined by the demand-side characteristics of a task, including its on-demand or online nature, its difficulty, and any other requirement attached to it.

Platform workers devote time to searching for available tasks. This "search effort" is conducted at the expense of leisure (l) only, as paid work hours (h) can only result from search: with no search, workers find zero hours of paid work.

As in Arellano and Meghir (1992), search takes place by devoting leisure time to this activity, and as such, it would only be a source of disutility for the worker if job opportunities were not revealed through search.[10] In our framework, search generates additional hours of paid work in the same reference period. The time workers spend searching for a job (i.e., waiting for new job offers by the platform) is S, which is the number of hours workers spend looking for available tasks on the platform. As paid work is revealed through search, search is the only choice variable available to workers.

Paid work hours are then a function of the pay-out, idiosyncratic ability, the number of workers available to platforms and the search effort, $H(\bar{w}, a, L, S) = h$. This function is separable so that, at the end of the reference period:[11]

$$h = H_a(\bar{w}, a)H_s(\bar{w}, L, S). \tag{1}$$

[9] We henceforth omit the individual and time indices.

[10] Similarly, our model ties with the on-the-job leisure labour supply model of Dickinson (1999) inasmuch as we separate productive hours from total work hours. In our model, however, search effort is not considered as on-the-job leisure but rather as a source of additional hours of work. On-the-job leisure is not an issue for our empirical model because, in the survey, workers were asked specifically how much time they spend on the platform *searching* for tasks.

[11] We define as reference period the time that occurs between the start and the end of a working session. In our empirical setting, this lasts a week.

The first term refers to a wage/efficiency-specific shock, revealing the average amount of time it would take the worker to perform a task of value $\bar{w}$. It is a function of the average payout and the factor a, which can include ability, location, and any other factors contributing to how quickly a worker can complete a task. The second term is the search function, revealing how many tasks of average pay-out $\bar{w}$ are found conditional on search. Note that, as there is no substitution between labour and capital, capital is omitted from the labour demand function, and since L is also fixed in the reference period, the search function reduces to $H_s(\bar{w}, S)$.

We assume the search function takes a simple linear form, assuming no increasing or decreasing returns to the search or intercept terms. We further separate the function and model it as $H_s(\bar{w}, S) = H_w(\bar{w})S$. The term $H_w(\bar{w})$ can be treated as an idiosyncratic search shock $\rho = H_w(\bar{w})$ capturing labour demand by revealing how many potential tasks of value $\bar{w}$, on average, are turned into actual jobs for an extra unit of search.

This functional form is based on the intuition that the worker does not know the true functional form of $H_s(\bar{w}, S)$, so the idiosyncratic functional form that each worker experiences after searching can be approximated with a single parameter. The true entity of the final idiosyncratic shock ρ is unknown to workers, as the demand parameters and concurrent aggregate supply are also unknown. Workers can form expectations on the shock so that the realised shock will then equal its expectation plus a random stochastic term $\rho = E[\rho] + \theta$.

This specification comes with two important simplifications. First, we have assumed that the search shock is linear in search. Secondly, expectations on the shock are assumed, for now, to be uncorrelated with individual characteristics or the wage. These are naive simplifications which are only made to make our theoretical model concise and tractable. As discussed in Appendix B, a non-linear form implies that the search shock is also endogenous to the wage. Similar arguments are made when the expectations (and the ability to manipulate the shock) are endogenous to unobserved individual characteristics, as discussed in Section 5. Our empirical model (Section 4) lifts all these assumptions.

The *actual* number of hours spent on the platform adjusts the paid hours supply for search, and is defined as $h^A = h + S$, so that the total time endowment is $T = l + h + S = l + h^A$, where l is leisure. This means that the search effort is subjected to the following constraints:

$$S \leq T(1 + \rho H_a(\bar{w}, a))^{-1}; \quad S \geq 0. \tag{2}$$

The function $h = H_a(\bar{w}, a)H_s(\bar{w}, S)$ has the advantage of separating the number of tasks found to the time it took the worker to perform them and allows us to see how efficiency and search can affect the pay rate in two separate ways. To see how begin with the hourly *frictionless* rate of pay w. This is obtained by dividing platform income in the reference period by the total amount of paid hours, so that

$$w = \bar{w}H_s(\bar{w}, S)(H_a(\bar{w}, a)H_s(\bar{w}, S))^{-1} = \bar{w}H_a(\bar{w}, a)^{-1}. \tag{3}$$

The rate of pay corrected for unpaid hours is the hourly *actual* rate of compensation, and equals

$$w^A = \bar{w}H_s(\bar{w}, S)(H_a(\bar{w}, a)H_s(\bar{w}, S) + S)^{-1} \tag{4}$$

Unless the worker decides not to work, the search effort is always non-zero, so the hourly actual compensation will always be lower than the hourly frictionless compensation.

An important caveat is that if search is separable into an idiosyncratic shock and an effort component, then the rate of change between actual and nominal salary is independent to the search effort, and only depends on the search shock. The simple proof is detailed in Appendix B, and has important implications for our empirical strategy, as it justifies our approach of separating search unto its two fundamental demand shock and effort components.

Following from Parker et al. (2005), we assume, for simplicity, concave and separable utility. We assume that the leisure disutility from the hours of actual work equals the disutility from the hours paid work and search, so that $U_h(h^A) = U_h(h) + U_s(S)$. The optimisation problem for an individual at period t then follows $\max_{C,S}\{U(C, h^A)\}$. The expected utility is expressed as:

$$U(C,h^A) = \int_{-\infty}^{+\infty} U_c(C)\,dF(\theta) + U_h(T-l)$$
$$= \int_{-\infty}^{+\infty} U_c(\bar{w}(E[\rho]+\theta)S+\mu)\,dF(\theta) + U_h(T-l) \quad (5)$$

where C is consumption, which equals

$$C \equiv wh + \mu = \bar{w}(E[\rho]+\theta)S + \mu \quad (6)$$

where μ is a measure of other income which reflects net dissaving at the end of the period t.

An important feature of this model is that workers are always in control of how much leisure they are sacrificing. Uncertainty in search enters utility expectations in the left term of the equation in the form of a multiplicative shock on wages. The implication of this source of uncertainty in the budget constraint is immediately evident in equation (6) with reference to the work of Parker et al. (2005), as the authors have shown that, under multiplicative shocks such as the one our model reduces to, there is no solution for labour supply for an increase in uncertainty, even under separable utility. This can be shown from the first order condition for labour supply:

$$\bar{w}(E[\rho]+\theta)U_c^{'}(\cdot) + U_h^{'}(\cdot) = 0 \quad (7)$$

This means that, holding leisure constant, while workers can optimise labour supply through search, so that $S^* = S^*(\bar{w}E[\rho],\mu)$, there is no effective solution to the optimisation problem for an increase in ρ through θ. In other words, we cannot determine how platform workers respond to an increase in the search shock. Also, without holding leisure constant, idiosyncratic efficiency and, in turn, frictionless wages will enter the optimal search function.

The pay-out might be known for platform workers with control over the pay-out and for those for whom the platform controls the pay-out but has little time variation. However, for some other workers, uncertainty in pay-out could also be a factor. Further extensions of the model can either treat uncertainty in pay separately – so that $\bar{w} = (E[\bar{w}]+\psi)$ – or integrate this uncertainty term within θ. The implications of an increase in uncertainty remain, generally, the same.[12] What matters for our purposes is that, even in a highly stylised scenario in which the payout is certain, the model still reduces to a labour supply model with uncertainty in pay.

An alternative approach would require modelling the expectation formation process task-by-task in a manner that is not dissimilar to the one that Stenborg Petterson (2022) follows when modelling taxi drivers' reference points with a Markov process. However, solving the model for a change in the search shock before the worker experiences it remains effectively impossible.

An important implication of the utility maximisation problem is that optimal search will depend, *ceteris paribus*, on the frictionless hourly earnings (and, by extension, efficiency) and the expectations of the actual salary. Ultimately, these expectations remain unobserved, and their effect on supply remains ambiguous. This is a matter that can only be settled empirically.

3 Data

Our main source of information on platform workers comes from an online survey on platform work conducted by PPMI in June 2021 (for more information, see Barcevičius et al., 2021, and its online annex). We will refer to this survey as the PPMI survey from now on.

This survey was conducted in the context of a background report for the impact assessment of the aforementioned EU directive proposal for improving the conditions of platform workers.[13] The survey comprises a total of 10,938 respondents, sampled from a population of working age (16-74 y.o.) internet users from 9 EU countries (Denmark, France, Germany, Italy, Lithuania, the Netherlands, Poland, Romania and Spain).

The sampling frame implies that survey respondents are not necessarily active in online

[12] The implications of wage uncertainty for our empirical model are discussed later in Section 4.
[13] Ibid. 7.

platforms: people in "traditional" employment, along with the unemployed, are also sampled. Out of all respondents, 2,440 have produced income from online or on-location platforms at least once, while 1,722 have been active on platforms in the last 6 months. In relative terms, this is not a small sample considering (i) that other, more conventional, labour surveys usually understate platform work (especially when it is the respondent's secondary job) and fail to capture many of its dimensions (Bracha and Burke, 2021), and (ii) that our econometric approach is designed specifically to overcome sample size limitations.

The survey captures information on demographics, employment history and use of online platforms for all respondents. Full summary statistics for all individual-level variables used in the analysis are shown in Table 1. Occupation groups are omitted from the table for brevity, but are also used in our analysis and capture the current (or last) non-platform occupation using ISCO, and NACE one-digit classifications.

The survey includes an ad-hoc module for platform workers, containing information on remuneration, experience, hours of work, and working conditions. Among these, workers are asked to report which platform they work on, how many years they have worked in platforms, how regularly they have worked in the platform over the last 6 months, and whether the platform has any control over remuneration and/or working hours. Depending on the platform and the level of control over remuneration, we reclassify workers into four platform types: (i) online workers with control over pay, (ii) online workers with no control over pay, (iii) on-location workers with control over pay, (iv) on-location workers with no control over pay. A breakdown of the summary statistics by type of platform is available in Appendix C, Table C.2.

Looking at labour supply, platform workers are first asked, in the reference week, how much did they earn from platforms. Then, they are asked how many hours did they spend *searching or waiting for tasks/ work assignments*. The answer to this question will define the variable we will use for the job search effort.

Then workers are asked, again in the same week, how many hours they spent *implementing paid tasks/ work assignments*, yielding our paid hours variable, which returns actual hours after summing it with search. Immediately after this question, workers are asked to report *in an ideal situation* how many hours per week they would have preferred to work implementing paid tasks/ work assignments via online platforms. This is our desired hours variable.

The relationship between these labour supply variables in our sample is illustrated in Figure 1. The figure shows that the difference between paid and actual hours can be large. Notably, the tail of actual hours is thicker than the one of desired hours, reminding us of the fact that many of these workers spend more time on the platform than they wish. Still, several endogenous factors might contribute to these supply outcomes, so it is difficult to draw conclusions from this figure alone. In the next section, we develop an empirical methodology that can address these potential issues.

Table 1: Summary statistics

	(1) Mean	(2) sd	(3) p10	(4) p90
Paid hours of work	13.237	12.654	2	30
Actual hours of work	22.524	19.571	5	47
Desired hours of work	19.991	15.142	4	40
Frictionless hourly wage	32.577	177.361	0.681	46.154
Actual hourly wage	14.256	61.397	0.372	23.077
Gender: Female	0.408	0.492	0	1
Married or living with a partner	0.634	0.482	0	1
Age	35.052	12.381	20	53
Foreign born	1.921	0.270	2	2
Household size	3.060	1.215	1	5
Number of children	0.679	0.897	0	2
Education: Low (ISCED 1-2)	0.043	0.202	0	0
Education: Medium (ISCED 3-4)	0.307	0.461	0	1
Education: High (ISCED 5-8)	0.651	0.477	0	1
Non-platform income	2972.400	7332.163	50	4500
Employed in trad. job	0.476	0.500	0	1
Partner's income: higher	0.556	0.497	0	1
Partner's income: similar	0.246	0.431	0	1
Partner's income: lower	0.198	0.399	0	1
Experience in platforms (yrs.)	4.661	4.243	2	11
Regular platform worker	0.757	0.429	0	1
Platform's control over working hours, None	0.562	0.496	0	1
Observations		1595		

Notes: Mean, Standard Deviation and bottom/top deciles for individual-level descriptors. Sample of all workers having supplied at least 1 hour of paid platform work in the reference week. Partner's income statistics only estimated for the subsample of workers who are married/ living with a partner.

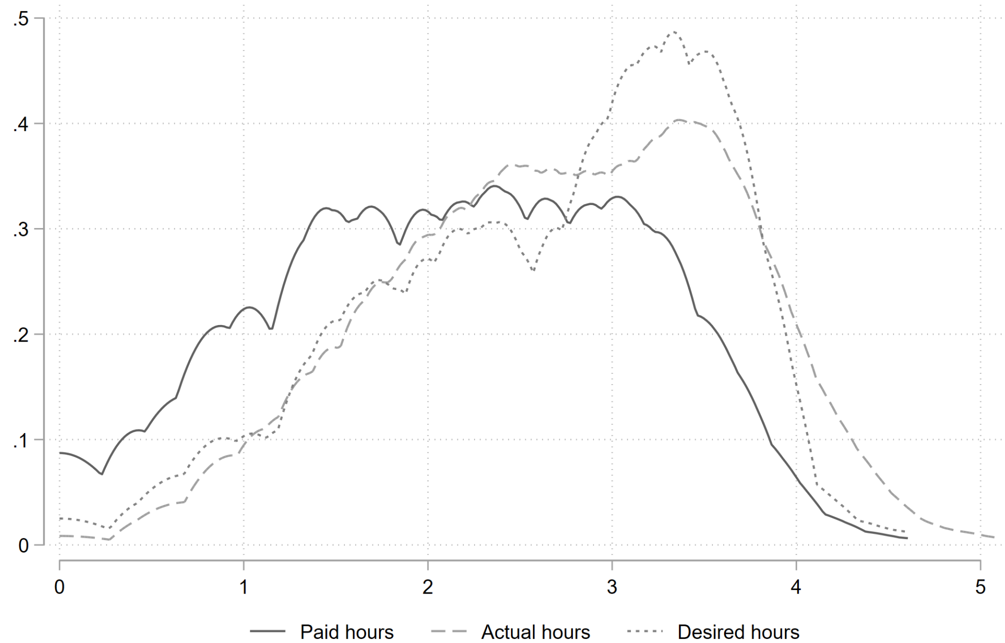

FIGURE 1: POTENTIAL LABOUR SUPPLY IN ONLINE LABOUR PLATFORMS

Notes: Kernel density plot (Epanechnikov smoothing) from a sample of workers having supplied at least 1 hour of paid platform work in the reference week. All values expressed in natural logarithms.

4 Empirical specification

In this section, we develop an empirical specification to estimate the elasticity of labour supply for workers participating in the platform economy.

Our estimation strategy exploits the piecework modality of work to estimate the elasticity of labour supply among workers from various types of platforms. As we have information on the paid hours of work h, the actual hours of work h^A and the desired hours of work h^D, our strategy is based on the intuition that, for each individual, these are all equilibrium outcomes to the same labour supply function conditional on the need to search for tasks and its underlying uncertainty.

Starting with labour supply in the absence of search, workers express their desired supply conditional on their salary w. This salary corresponds to the rate w, inclusive of idiosyncratic efficiency factors. Since no task search takes place in the desired equation, this is simultaneously the actual and frictionless rate. Omitting the individual subscript i, the *desired* labour supply function with no task search is expressed as:

$$ln(h^D) = \alpha_1 ln(w) + U^{'}\beta + X^{'}\delta + \eta_1 \tag{8}$$

where α_1 is the elasticity of (desired) labour supply to the rate of pay w. As in traditional labour supply functions, U is a vector of controls for any other source of income, including unearned income and income from other non-platform occupations, if available. X is a vector of observed and unobserved individual characteristics affecting labour supply participation. These can include both observed and unobserved preferences and all factors affecting efficiency in task completion, including ability and location characteristics.

It is important to underline that this is the supply of work that workers would provide *if* the frictionless wage w were also to be the actual hourly rate of pay. We later discuss in Section 5 the conditions under which this assumption stands.

In the presence of search, workers supply h^A actual hours of work to the platform. The increase in working hours caused by the unpaid search lowers the actual (and final) rate of pay from w to w^A. The *actual* labour supply equation is expressed as:

$$ln(h^A) = \alpha_1 ln(w^A) + \alpha_2 ln(w) + U^{'}\beta + X^{'}\delta + P^{'}\zeta + \eta_2 \tag{9}$$

where the parameters α_1 and α_2 capture the wage elasticity of supply under perfect and imperfect information, respectively. α_1 is the wage elasticity of supply given the actual salary, which eventually is revealed to the worker. Henceforth, we refer to this parameter as *actual wage elasticity of supply*. This parameter would sufficiently sketch out labour supply if workers had full information on the search shock. In that case, they would know the final value of w^A, and there would be no need to control for the frictionless wage w.

However, as we wish to test for the presence of imperfect information, and as we believe that these imperfections can affect the estimates of α_1, we then let frictionless salary re-appear in the equation, tying it to a different parameter, α_2. The parameter reveals, on average, how the overestimation of the final (actual) salary affects workers' supply schedules, capturing the effect of imperfect information. From this perspective, the parameter equals α_1 times the weight that is given to the frictionless salary, which is a function of expectations $E[\rho]$, so that $\alpha_2 = \alpha_1 D(E[\rho])$. The further these expectations are, on average, from producing the final actual salary w^A, the stronger the effect of α_2 will be. Henceforth, we refer to this parameter as *frictionless wage elasticity of supply.* The model offers an empirical implementation of the theoretical model discussed in Section 2, but by distinguishing between the actual and frictionless wage, we are now agnostic to the accuracy of workers' expectations.

If α_2 equals zero, then the actual supply equation reduces to a standard supply equation. Under imperfect information, the magnitude of α_2 grows as expectations deviate more and more from the final actual wage; α_2 will capture the *average effect* of these expectations, while idiosyncratic variation in expectations uncorrelated with the level of wage will affect α_2 through a disturbance term.[14]

It should be underlined that while the w term appears in both equations 8 and 9, it is associated with two different parameters in each equation. The point is that the terms associated with α_1 denote the actual - and final - wage associated with a given supply schedule. In the desired scenario, the absence of a search shock makes it so that the frictionless wage w is also the actual wage. This is not the case in the actual supply equation.

Demand-side platform factors outside the worker's control can also intervene to affect the amount of work available, so the P term is included to represent different types of platforms. As mentioned earlier in Section 3, we classify these platforms based on the on-location or online nature of the work performed, the level of control the platform exercises over pay, and the interaction between these two variables. This is not only motivated by our review of the literature but also by the intuition that the services sold offline and online will also vary significantly in nature and feature different demand elasticities. Platforms can also exploit their monopsony powers when they can exercise control over pay, ultimately affecting the total amount of work available. We then control for these characteristics in the P vector as these can affect the equilibrium levels of supply. Unobserved supply-side factors determining

[14]However, this disturbance, if unobserved and large enough, can have an effect on the consistency of our estimates; please refer to the discussion in Section 5 and the Monte Carlo simulations in Appendix D.1 for further details.

access into each platform type are to be included in the X vector instead.

Equations (8) and (9) can be difficult to estimate because of the endogeneity of wage caused by the unobserved part of the X vector. In other words, in equation (8), the salary w attached to α_1 still depends on unobserved ability, which could be endogenous to the hours of work. For example, less-efficient workers can take longer to complete certain tasks, while other platform workers can manipulate their frictionless salary, some by setting the price of their services if in power to do so, others by searching in low-supply hours. Equation (9), instead, should consistently estimate the α_1 parameter by controlling the way these factors affect w^A through w, but w^A might also be endogenous if workers with higher skills also have better knowledge of the search shock or might be able to manipulate it (for example, by searching during specific hours or from multiple platforms), enabling them to raise their actual salary. At the same time, the α_2 parameter associated with w will still suffer from the same endogeneity problems.

A solution would be to find a way to instrument the wage terms (see Blundell and Macurdy, 1999, for an overview), but valid instruments are not always available and are known to underperform in small samples. This is further complicated by the fact that we should instrument simultaneously for both wage terms. We can, however, exploit information on desired and actual labour supply to retrieve these parameters.

In this context, the solution to this problem lies in the fact that the two equations reveal how many hours workers would wish to supply as only the idiosyncratic search shock and its expectations change, because the frictionless wage in the desired equation is also, in this scenario, the final and actual wage. This leads to the system of equations:

$$\begin{cases} ln(h^D) = \alpha_1 ln(w) + U^{'}\beta + X^{'}\delta + \eta_1 \\ ln(h^A) = \alpha_1 ln(w^A) + \alpha_2 ln(w) + U^{'}\beta + X^{'}\delta + P^{'}\zeta + \eta_2 \end{cases} \tag{10}$$

The equations in the system (10) share a subset of parameters by definition, as they both model two optimal combinations of work and leisure under different levels of worker's utility deriving from different combinations of leisure, work and search.

Fundamentally, demand-side restrictions leading to the change in final salary from its original frictionless value are the only differences between the two equations. This is motivated by the intuition that desired and actual hours of work are the product of the same supply function and that the presence and intensity of search is effectively the only major source of demand-side restrictions to the supply of work, as any other restriction to the hours of work is either absent or observable because the autonomous nature of work makes workers generally free to supply as much work as they wish.

If this condition is satisfied (as discussed later in Section 5), then the two equations can be differenced to study the change in wage across the two states and cancel out the unobserved term. Taking differences from the equations in system (10) and simplifying, we reach:

$$ln(h^A/h^D) = \alpha_1 ln(w^A/w) + \alpha_2 ln(w) + P^{'}\zeta + \eta_3 \tag{11}$$

obtaining first-difference estimators of the wage elasticity of labour supply. Taking differences between the two equations, we now focus on how *changes* in wage affect *changes* in supply, keeping all idiosyncratic workers' characteristics fixed as the time-invariant $X^{'}\delta$ term is now cancelled out. First-differencing allows us to obtain conditional independence for the *level* of salary (w), which is no longer endogenous to the outcome.[15] Most importantly, variations in its *change* (w^A/w) can now also be studied net of w. As a result, the final search shock is incorporated in w^A/w, while initial expectations are incorporated by w, with the correlation between the two being now uncorrelated with the outcome.[16] In fact, in contrast with equation (9), the w term can now be removed from the equation without the α_1 parameter being affected.

For simplicity, we henceforth refer to this model as the *differenced* specification. Interpretation of the terms in the equation remains straightforward, as the parameters α_1 (wage

[15] This can also be shown via Monte Carlo simulation, as discussed in Appendix D.

[16] Effectively, specification (11) offers difference-in-differences estimates for labour supply by comparing outcomes across two different states depending on the intensity of the search shock. In our case, treatment is continuously defined so that deviations from the desired supply for "controls" experiencing low frictions are compared with outcomes for "treatments" experiencing high frictions. In the difference-in-differences jargon, the "common trends" assumption is satisfied if, under the same search shock (or in the absence of it), the change in actual and desired supplied hours of work would be identical for all workers. From another perspective, equation (11) estimates quasi-Frischian elasticities of supply by treating the desired status as a pre-treatment scenario and absorbing away the individual-fixed effects via differencing.

elasticity) and α_2 (frictionless wage elasticity) can be interpreted as they appear in the system (10).[17] The main difference is that α_1 can now also be interpreted as the elasticity of supply to the change from frictionless to actual wage.

This setting offers two other advantages. Firstly, it allows for endogeneity in the search shock and its expectation because the components of X that might have affected w^A/w are also cancelled out. In other words, while endogeneity in the search shock is unlikely, the endogenous component of the search shock can also be cancelled out by differencing; see Section 5 for further discussion. A more intuitive argumentation suggests that the ability to manipulate the search shock precedes the final shock experienced in the reference week.

Secondly, note that this specification addresses division bias in the wage elasticity term, which is a common problem in labour supply models computing wages from income and hours of work when there is measurement error in the latter (Farber, 2005; Stafford, 2015). In our model, the actual wage is measured through variations in the hours of work, keeping income fixed. As there is no reason to assume that measurement error should differ for search and paid hours, then w/w^A will not depend on the misspecification of the hours worked, unlike other labour supply models.

Finally, while the model will absorb individual characteristics, heterogeneities in labour supply between platform types can still be driven by self-selection and be studied by adding interactions between the elasticity term α_1 and the platform types included in vector P. Specifically, we are interested in evaluating how these elasticity parameters can change depending on whether services are offered online, whether workers have control over their pay, or whether workers have access to a second "traditional" job.[18]

As a final note, it should be noted that a number of conditions have to hold for this model to work correctly. These are discussed in detail, complete with a set of robustness checks, in Section 5.

[17] An interesting caveat of this first-difference approach is that knowing that w^A/w equals h/h^A, the equation can be rewritten as $ln(h^A/h^D) = \alpha_1 ln(h/h^A) + \alpha_2 ln(w) + P^{'}\zeta + \eta_3$. The parameter α_1 is still consistently estimated.

[18] This second check is offered in Appendix A, Table A.1.

5 Conditions for estimation validity

We here discuss the conditions under which the estimated parameters converge to the true elasticities. A number of factors can threaten the validity and consistency of the proposed models.

In general, for the model to be valid and consistent, we should be able to treat desired (8) and actual (9) supply equations as labour supply equations expressed in two different points in time, with demand-side constraints being the only differences between the two. In this way, the final differenced equation (11) can be treated similarly to an empirical Frisch labour supply function so that wages can be correlated with the idiosyncratic term as long as the variation in wage is uncorrelated with the error term. Additionally, in contrast with standard panel data approaches to the estimation of Hicksian labour supply, our approach is unaffected by division bias, which is caused by measurement error in the hours of work when the wage is obtained by dividing income by the former. This is because, as long as measurement error in hours of search and paid work is comparable, the variation between the frictionless and actual wage cannot be driven by measurement error.[19]

For our model to reduce to a Frisch supply function, a number of conditions need to hold. These can be summarised into two classes, which we discuss in detail in this Section: (i) independence of the search shock and (ii) differential measurement error in the supply functions.

Endogeneity in the search shock poses the main threat to the validity of both α_1 and α_2 elasticity estimates, and we identify two forms of endogeneity that our model can encounter. Under "weak" endogeneity, the shock is correlated with unobserved ability that affects both supply equations, while under "strong" endogeneity, the shock is correlated with unobserved ability that affects only one of the two supply equations. On the other hand, measurement error will affect not only the magnitude of α_2, but also the consistency of both elasticity estimates if measurement error features an idiosyncratic component. This variability could come from variations in expectations or other sources of variation in the wage, which we group

[19]There are, however, other sources of "differential" measurement error that can affect our estimates, as we discuss later in this section.

under the umbrella of "differential measurement error".

The following sections provide a set of rigorous and exhaustive tests to deal with the above potential issues in model validity All tests presented in this section are reproduced in Section 6. Appendix A offers additional robustness checks, including specifications with heterogenous effects conditional on access to traditional employment. Monte Carlo simulations are offered in Appendix D to provide an intuitive representation of the properties of our models and their robustness to these conditions.

5.1 Independence of the search shock

First, we turn to the independence of the search shock ρ. This issue is tightly related to the functional form of the search function.

We assumed earlier in Section 2 that the search shock is linear. This might not necessarily be the case in the empirical context. We have discussed that changes in information about the search shock affect the search effort through changes in wage expectations. However, if the search shock is non-linear, then the value of w^A will also change under different levels of search.[20] We know from our theoretical model in Section 2 that the search effort is the choice variable under uncertainty and, as such, is endogenous to the frictionless wage. Under a non-linear search shock, workers with higher w will adjust their search effort and experience a different shock.

A related problem is tied to the possibility that workers might endogenously affect the search shock or, equivalently, that access to information on the search shock is endogenous. Some workers might search during peak demand hours or, in some instances, be allowed to search over multiple platforms, and this behaviour might correlate with unobserved individual ability. This would suggest that the search shock can also be endogenous to individual characteristics X, which could, in turn, influence w already.

This discussion suggests that both (i) w and (ii) w^A/w can be endogenous due to the search shock. The first is a case of "weak" endogeneity and affects only the actual supply model (9). In this case, the same unobserved factors from X that affect w also affect the

[20]See Appendix B for a discussion.

shock ρ. In the presence of endogeneity in the search shock, α_1 will be incorrectly estimated in the actual supply model. These factors, however, are absorbed through first-differencing: in such a scenario, estimates from the differenced model can still be deemed reliable because this is the only model that controls for both the frictionless salary w and the rate of change in salary w^A/w, while also cancelling out X.[21]

The second is a case of "strong" endogeneity, which would affect the differenced model (11). In this case, the threat to identification would come from the possibility that unobserved ability factors affected h^A and h^D differently (so that they are not cancelled out) while also affecting the shock. Should this be the case, then α_1 and α_2 will also be incorrectly estimated in the differenced model too. Under "strong" endogeneity, the differenced equation updates to:

$$ln(h^A/h^D) = \alpha_1 ln(w^A/w) + \alpha_2 ln(w) + (X^A - X)^{'}\delta + P^{'}\zeta + \eta_3 \tag{12}$$

Where $(X^A - X)^{'}$ is the residual subset of individual characteristics whose effect on the difference outcome is non-zero, i.e. the characteristics that affect supply only in the actual equation. $(X^A - X)^{'}$ can only be correlated with w^A/w, regardless of whether X is already be correlated with X. As we discuss in the next subsection, this issue can be related to differential measurement error when the error is also related to the search shock.

This strong endogeneity scenario is, in general, unlikely. There is very little workers can do to influence the task search process. The strategies workers can use to improve search times are not particularly hard to master, and it is unlikely that any individual ability that does not already affect the frictionless wage would help workers reduce the search shock they will experience. Still, it is worth testing for the presence of strong endogeneity to dispel all doubts.

For an exhaustive test of weak independence, estimates for α_1 can be validated by performing specifications tests on actual wage elasticity estimates, which should be consistent between the actual (9) and differenced supply (11) models. Each of the two models relies on different assumptions, so if both models produce the same results, we can safely argue that

[21]The stability of the differenced model can also be appreciated in our simulation. See Appendix D.2 for further reference.

both models are correct. A large enough change in the α_1 coefficient between the differenced and the actual supply model would suggest that the shock is not unconditionally exogenous. A similar test can be performed by removing the frictionless wage term from the differenced supply model and then measuring the change in coefficients from the differenced model (11). This second test will indicate whether the w^A/w term is actually rendered independent from w via differencing.

Moving to strong independence, as a first test, we can check for changes in the coefficients of interest (α_1 and α_2) in the differenced model before and after all other controls have been removed. After removing the observable set of controls, a change in the coefficients could suggest that components of the X term, which might affect only the search shock and the outcome, have not been absorbed correctly. We want the elasticity coefficients to remain statistically unchanged for the test to succeed.

Our theoretical considerations suggest a second test. Workers are assumed to optimise search depending conditional on w only (which already includes skills and efficiency). Hence, after controlling for w, we expect to see no effect on the search effort from other individual characteristics. Regressing search effort on the full set of observable controls, we expect the search effort to depend on no other individual characteristics other than the frictionless wage. If other skill factors that might affect the frictionless wage also affect the search effort, then we cannot rule out that these factors also affect the actual wage and the shock.

A final test for "strong" independence is tied to the tests for differential measurement error. Being violated only when a subset of predictors affects one of the two equations, strong independence is strictly related to differential measurement error, suggesting that tests for the latter, if passed, will also dispel "strong" endogeneity concerns.

5.2 Differential measurement error

The second condition relates to differential measurement error, which occurs when one of the two labour supply equations (i.e. desired labour supply and actual labour supply) suffers from sources of measurement errors that do not already affect the other. If this form of error is also correlated with the search shock,[22] then the "strong" independence assumption will also fail, as discussed above.

There are two forms of differential measurement error that can affect our estimates. The first involves the X term. Differential measurement error implies that individual characteristics might contribute differently to desired supply schedules as they do to actual supply schedules. If these characteristics are observed, it is sufficient to control for them in the differenced equation, but if they are not, the error will be observed by the term η_3. The idea that some characteristics which do not affect the desired supply might affect the actual supply is not unreasonable, as discussed in the subsection above.

The second form concerns variation in the frictionless wage term w. This can arise from a number of factors, which are all related to variation in w during the estimation window. In our design, *ex-ante* workers do not know the entity of the search shock but make expectations based on the frictionless rate when realising their actual supply, *ex-post* they express desired supply keeping this same frictionless rate fixed. The implication is that the frictionless rate w is assumed to be the same *before* and *after* the working week. Allowing w to change, the desired hours equation will be based on the end-of-the-week frictionless wage w_1, while the actual hours equation will be based on the start-of-the-week frictionless wage w_0.

Both w_0 and w_1 could be unobserved and be affected by several sources of variation. A first source of measurement error could arise from the idiosyncratic variation in the actual wage w_0 expectations among individuals. Variations in idiosyncratic ability and location factors might cause further variation in w_0 in the actual supply equations. Similar arguments apply to w_1, as workers might overstate their ability or factor demand-side constraints in their frictionless wage.

Including all forms of differential measurement error, the differenced equation updates to:

$$ln(h^A/h^D) = \alpha_1 ln(w^A/w_1) + \alpha_2 ln(w_0) + (X^A - X)^{'}\delta + P^{'}\zeta + \eta_3 \tag{13}$$

Measurement error has different implications on the estimation depending on its nature.

[22] For example, by helping reduce the search effort.

Error in X, if tied to an unobserved component, will severely affect the consistency of our estimates and will also pose as a source of bias if correlated with the search shock ρ.

Error in w will have different effects depending on whether the error features an idiosyncratic component. If it does not, then the error will simply be absorbed by α_2, moving the estimate in a given direction. In this case, differences between $w_{0,1}$ and w are to be considered systematic, and only our interpretation of α_2 would need to be updated, as the frictionless wage elasticity will now incorporate the time-variant wage shocks that are shared across all workers. The effect on our interpretation of the α_2 term is minor, as in this case the variation in wage would simply features a systematic component shared by all workers. This could be entirely attributed to expectation biases already absorbed by α_2, as it is unlikely for the wage to be systematically higher (or lower) in one of the two equations.

The implications of an idiosyncratic error term are more nuanced. Recall that the α_2 parameter already results from a weighted function of wage expectations and the actual wage elasticity, such that $\alpha_2 = \alpha_1 D(E[\rho])$. If this wage also features an idiosyncratic component, then this variation will be absorbed by α_2, with the estimates growing increasingly inconsistent as the variance of the idiosyncratic component grows.[23] The effects of these idiosyncratic disturbances are illustrated in detail through Monte Carlo simulation in Appendix D.1.

We could argue that some of this measurement error is, by a large part, negligible. Variation in expectations, skills, and location factors is unlikely to be high enough to influence w_1 during the reference week. Similarly, we could argue that workers do take into account their idiosyncratic ability but not platform-side constraints when expressing their desired supply, which we define as assumptions of *"Self-consciousness"* and *"No devil's advocate"*, respectively.[24] However, the error *between* workers must also be negligible for the model to be unaffected. This might be difficult to defend when it comes to expectations. This suggests

[23] All sources of measurement error in w will affect our elasticity estimates in this way.

[24] According to the former, workers should consider how efficiently they can perform a task when expressing desired labour supply. The intuition is that, if offered unlimited tasks at a fixed pay-out, workers will never assume they can complete them immediately but will take into account idiosyncratic factors affecting how efficiently they can perform them. According to the latter, when expressing desired hours of work in ideal conditions, sources of disutility arising from the demand side should not be considered. In other words, workers should not express their preference for hours of work based on the search shock they experienced in the reference period because, in ideal conditions, the search effort should be zero as it comes at a net loss of leisure.

that some level of idiosyncratic error will be inevitable. There are a few ways to test for the presence of measurement error, but if w is assumed to be endogenous, it will not be possible to cross-check for validity across models as we did for α_1.

The most exhaustive test for this idiosyncratic error involves the introduction of an interaction term between w^A/w and w in the first-differenced equation. The intuition comes from the fact that w^A/w should capture variation in wage net of w. If w varies between one of the two terms (so that, for example, we have $w^A/w= w^A/w_1$, or $w = w_0$), then the observed frictionless wage term will absorb part of this variation, introducing spurious correlation with w^A/w on the estimation phase.

It follows that this residual variation can be assessed by adding an interaction term between w^A/w and w, whose parameter we define as α_3. If the "true" unobserved w differs significantly between the two wage terms, the term will have a non-zero coefficient, and all elasticity estimates will change after the introduction of the interaction term. If this variation is, instead, negligible, then the parameter α_3 should be approximately zero. An interesting caveat of this relationship is that the interaction term can, by chance, be zero even if idiosyncratic variation in measurement error is high, simply because it is more likely to be zero when variation in α_2 is low enough. But as long as this term is zero, α_1 and α_2 would remain identified. This is demonstrated by our Monte Carlo simulations Appendix D.1, Figure D.4: note how when α_3 is approximately zero, then α_1 and α_2 are close to their "real" predetermined values.

Measurement error in X will also cause similar issues and also be detectable through this test. Measurement error in w, net of idiosyncratic variation, will instead be absorbed by α_2. These conditions are reproduced via Monte Carlo simulation for demonstration purposes in Appendix D.2, Table D.3.

The test above will check for the presence of error but not the level of idiosyncratic variability in wages. A second test, detailed in Appendix D.1, can provide an upper bound for the variability of w. In short, this can be retrieved simply by looking at the standard error of the estimate of the elasticity of the interaction term, which grows linearly with the variation in w. As a rule of thumb, if the standard error is too high, then the level of variation in idiosyncratic expectations is probably too high. Our discussion suggests that, while the model can tolerate some unobserved idiosyncratic variation in wage, too much variation will increase the variability of the results and can only be accommodated by increasing the sample size.[25]

Other checks can also be performed to test the other conditions more directly. For example, the *"No devil's advocate"* condition can be tested by adding platform-side controls in the desired hours equation (8) and testing if the estimated elasticities are affected by the inclusion of these controls; they should if workers take into account platform-side constraints when expressing their desired supply. Similar tests can be conducted on the *"Self-consciousness"* condition too. These are discussed in detail in Appendix A.

6 Results

Our labour supply estimates are shown in Table 2 and 3. In Table 2, we offer average elasticities for all platforms. Platform-specific elasticities (*heterogeneous* elasticities, henceforth) are shown in Table 3. Platform types are captured by the online or on-location nature of the services offered (*On-location worker*) and the degree of control over pay that the worker enjoys (*Control over pay*), and the interactions between the two. All hours of work and wage variables are expressed as natural logarithms.

Our control variables are detailed below. These are split into sets of *individual controls*, *occupation fixed effects*, *platform controls* and *country fixed effects*.

Our set of *individual controls* (X) includes three subsets of variables. The first subset includes variables which are traditionally included in most labour supply models: age and its squared term, education (ISCED), foreign nationality, marital status, gender, the interaction between marital status and gender, partner's income (equal, higher or lower than respondent), household size, and the number of dependent children.

The second subset of individual characteristics includes controls for workers' idiosyncratic experiences on platforms, which can potentially generate non-linearities in search. Specifically,

[25] This will be the case, for example, for some of our estimates for the heterogeneous elasticities models, which do not benefit from large enough subsample sizes.

Table 2: Labour supply elasticity estimates

	Hours, desired (ln)		Hours, actual (ln)		Hours, first-differenced (ln)					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
α_1	-0.083***	-0.095***	-0.141***	-0.220***	-0.200***	-0.230***	-0.234***	-0.219***	-0.238***	-0.233***
	(0.014)	(0.014)	(0.016)	(0.045)	(0.045)	(0.045)	(0.045)	(0.044)	(0.042)	(0.057)
α_2				0.081		-0.043**	-0.037**	-0.030*	-0.020	-0.042*
				(0.042)		(0.013)	(0.012)	(0.013)	(0.012)	(0.017)
On-location worker		-0.072	0.115	0.108	0.165**	0.186***	0.206***			0.186***
		(0.081)	(0.070)	(0.075)	(0.052)	(0.053)	(0.053)			(0.053)
Control over pay		0.003	0.103	0.082	0.053	0.061	0.055			0.062
		(0.131)	(0.111)	(0.112)	(0.105)	(0.104)	(0.101)			(0.104)
On-location worker × Control over pay		-0.020	-0.269*	-0.258*	-0.215*	-0.234*	-0.266*			-0.234*
		(0.122)	(0.108)	(0.109)	(0.105)	(0.105)	(0.104)			(0.105)
Experience in platforms (yrs.)	-0.019**	-0.016*	-0.010	-0.012	0.002	0.003		0.005		0.003
	(0.007)	(0.007)	(0.008)	(0.008)	(0.006)	(0.006)		(0.006)		(0.006)
Regular platform worker	0.155**	0.135*	0.131*	0.112*	-0.031	-0.028		-0.062		-0.028
	(0.057)	(0.057)	(0.053)	(0.054)	(0.050)	(0.050)		(0.044)		(0.050)
α_3										0.001
										(0.019)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Last occupation FE	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes	No	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes	No	Yes
Platform controls	No	Yes	Yes	Yes	Yes	Yes	Yes	No	No	Yes
Adjusted R-Squared	0.131	0.178	0.179	0.180	0.148	0.155	0.112	0.101	0.040	0.155
Observations	1590	1590	1584	1584	1580	1580	1581	1580	1581	1580

Notes: SE clustered by occupation/sector clusters in parentheses. All wage and labour supply variables are expressed in natural logarithms.

*p<.05, **p<.01, ***p<.001

Table 3: Labour supply elasticity estimates, by platform type

	Hours, desired (ln)		Hours, actual (ln)		Hours, first-differenced (ln)				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
α_1 (Online worker × No control over pay)	-0.075***	-0.070**	-0.110***	-0.271***	-0.233***	-0.250***	-0.256***	-0.175**	-0.173**
	(0.020)	(0.023)	(0.024)	(0.054)	(0.065)	(0.064)	(0.064)	(0.055)	(0.053)
α_1 (Online worker × Control over pay)	-0.067***	-0.108***	-0.152***	-0.145	-0.167*	-0.200*	-0.216**	-0.261***	-0.311***
	(0.019)	(0.025)	(0.025)	(0.087)	(0.083)	(0.082)	(0.077)	(0.077)	(0.074)
α_1 (On-location worker × No control over pay)	-0.105***	-0.094**	-0.189***	-0.285*	-0.141	-0.218*	-0.210	-0.356***	-0.389***
	(0.027)	(0.035)	(0.031)	(0.118)	(0.111)	(0.109)	(0.106)	(0.092)	(0.092)
α_1 (On-location worker × Control over pay)	-0.121***	-0.170***	-0.175***	0.002	-0.170	-0.188	-0.191	-0.215	-0.271*
	(0.031)	(0.049)	(0.041)	(0.147)	(0.138)	(0.155)	(0.146)	(0.131)	(0.120)
α_2 (Online worker × No control over pay)				0.170**		-0.027	-0.017	-0.034	-0.020
				(0.054)		(0.017)	(0.016)	(0.017)	(0.016)
α_2 (Online worker × Control over pay)				-0.006		-0.044*	-0.036*	-0.027	-0.021
				(0.080)		(0.018)	(0.018)	(0.018)	(0.017)
α_2 (On-location worker × No control over pay)				0.094		-0.091***	-0.082**	-0.047	-0.032
				(0.115)		(0.026)	(0.026)	(0.025)	(0.023)
α_2 (On-location worker × Control over pay)				-0.184		-0.026	-0.045	-0.023	-0.031
				(0.139)		(0.043)	(0.037)	(0.037)	(0.034)
On-location worker		-0.038	0.206*	0.258*	0.222*	0.327**	0.355***		
		(0.119)	(0.093)	(0.120)	(0.099)	(0.102)	(0.101)		
Control over pay		0.062	0.148	0.239	0.094	0.123	0.113		
		(0.143)	(0.118)	(0.134)	(0.112)	(0.118)	(0.113)		
On-location worker × Control over pay		0.070	-0.329*	-0.244	-0.272	-0.400*	-0.379*		
		(0.184)	(0.139)	(0.188)	(0.167)	(0.183)	(0.179)		
Experience in platforms (yrs.)	-0.020**	-0.016*	-0.010	-0.012	0.002	0.003		0.005	
	(0.006)	(0.007)	(0.008)	(0.007)	(0.006)	(0.006)		(0.006)	
Regular platform worker	0.142*	0.133*	0.131*	0.106*	-0.034	-0.030		-0.051	
	(0.058)	(0.057)	(0.053)	(0.053)	(0.050)	(0.051)		(0.045)	
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Last occupation FE	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes	No
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes	No
Platform controls	No	Yes	Yes	Yes	Yes	Yes	Yes	No	No
Adjusted R-Squared	0.133	0.181	0.182	0.187	0.148	0.158	0.115	0.104	0.046
Observations	1590	1590	1584	1584	1580	1580	1581	1580	1581

Notes: SE clustered by occupation/sector clusters in parentheses. All wage and labour supply variables are expressed in natural logarithms.

*p<.05, **p<.01, ***p<.001

we account for the possibility that more experienced workers are better informed about optimal search schedules by controlling for years of experience in platforms and the regularity of the worker's activity in platforms over the last six months.[26]

The third and final subset controls for characteristics related to the respondents' access to "traditional" labour markets, which will inevitably affect the amount of time that can be allocated to platform work. Namely, we control for their employment status (i.e., whether they have another non-platform job) and for total income from other sources (in logarithm). This can include income from another job but also social security transfers.

Furthermore, we add intercepts for the interaction of the ISCO-08 occupation and NACE-Rev2 industry codes of the last non-platform occupation held by the respondent, treating the absence of prior labour market experience (i.e., the "ever-unemployed") as the baseline level (*Last occupation FE*). As discussed in Cantarella and Strozzi (2021), the inclusion of these occupation controls does not constitute a "bad control" situation, as these conditions relate to outcomes that often predate access to platforms and, as such, they can proxy for unobserved ability. We also cluster standard errors by each occupation-sector cell (for a total of 180 clusters). In this way, we implicitly account for the possibility that the error residual follows a similar structure for people of similar skill.

Demand-side platform characteristics are instead included in the vector of *Platform controls*, i.e. the P vector discussed earlier in Section 4. These include the generic platform types discussed above (*Online vs on-location; Control over pay*) but also a rich set of other minute characteristics of platform jobs which can capture demand-side variations in the search shock.

These include the presence of the following benefits provided by the platforms: guaranteed minimum monthly or weekly salary, private health insurance, paid holidays, paid sick leave, insurance against accidents at work, work-related training, pension contributions, and paid parental/ paternal/ maternal leave.[27]

[26] While we can control for the number of platforms the worker was active in at the time of the interview to account for the fact that some workers might be able to search from multiple platforms simultaneously, we prefer to exclude such control because it is clearly endogenous to the search effort. Nonetheless, we have tested all our specifications with this control included and found no difference in our results.

[27] All responses are categorical: Not applicable/ Don't know; No; Yes; I have it, but not because of the platform.

Organisational aspects of work in the platform are also included, specifically whether the platform: pays workers periodically rather than on task completion, sets working schedules and/or minimum work periods, guarantees at least a minimum amount of work every week or month, allocates clients and/ or work assignments, disciplines workers when they refuse clients or work assignments (e.g., account terminated, fewer assignments, etc.), provides tools, materials and/ or protective equipment, monitors the implementation of work assignments, and prevents workers from working via other similar platforms.[28]

Finally, intercepts for each country in the survey are added too (*Country FE*), to account for differences in platforms and contracts across the countries surveyed.

Looking at the main results, columns (1) and (2) from both Tables display wage elasticities for desired labour supply. Column (1) corresponds to equation (8), while column (2) adds platform-side controls. In all instances, the results seem to suggest a backwards-bending labour supply curve, with a percentage increase in the (frictionless) wage leading, on average, to a statistically significant ~ 0.08 percentage point reduction in desired hours. There are no differences in sign between types of platforms, and the largest difference in magnitude is 0.046 log points.

While the elasticity estimates are not reliable because the frictionless wage may be endogenous to the desired supply, the inclusion of platform-side controls in column (2) allows us to perform a first test for measurement error in desired hours. As we hypothesized, platform-side factors do not seem to affect workers' desired supply, pointing at a mere change of ~ 0.012 points in the average estimated wage elasticity. A Z-test between the coefficients suggests that we cannot reject the null hypothesis that the coefficients are the same (Z-test: $0.606 < 1.96$). Platform types also appear to have little effect on desired hours, as all platform type coefficients are statistically not significant. Table A.1 from Appendix A provides further checks in columns (1) and (2), controlling for actual salary and search. Similarly, negligible changes are also to be noted for the heterogenous elasticities model, as estimated elasticities show minimal changes after adding platform-side controls to the specification.

These findings provide preliminary evidence in support of our assumptions on differential

[28] All responses are categorical: Not applicable/ Don't know; No; Yes.

measurement error (and on the "No devil's advocate" condition specifically), suggesting that workers do not generally consider platform-side constraints when expressing their desired labour supply. In other words, the "ideal conditions" do imply zero search, suggesting that desired hours are revealed net of demand-side effects.

Actual hours are studied in columns (3) and (4) of both tables, corresponding to the specification from equation (9), with column (3) omitting the frictionless wage. Estimates for α_1 from column (3) are negative and inelastic, and the magnitude of the effect is stronger than the one estimated for the desired hours specifications.

However, column (3) results might still suffer from the same endogeneity issues affecting the desired hours equation. If the frictionless wage is not controlled for, the actual wage might be endogenous to the actual hours of work. Introducing the frictionless wage term in column (4) confirms our suspicions: the wage elasticity term, while still negative and inelastic, increases in magnitude, suggesting that a point percentage increase in (actual) wage leads, on average, to a 0.22 per cent reduction in the hours worked. This same pattern can be observed in Table 3 when looking at platform-specific elasticities, this time with some elasticities turning fully inelastic after introducing the frictionless wage term. In column (4) the estimates for α_2 are not statistically different from zero, pointing to the absence of a frictionless wage effect. However, these estimates for α_2 are far from final because the frictionless wage term might still be endogenous. Estimates for α_1 are more reliable but, as discussed earlier in Section 5, are only valid if w is time-invariant and if the search shock is uncorrelated with unobserved characteristics in X.

We then study the differenced equation (11) in columns (5) to (9) in Tables 2 and 3. The first important result is that wage elasticity coefficients are relatively unchanged from the ones estimated in the labour supply model with actual hours worked, thus supporting a backwards-bending supply along with the target-earning hypothesis.

The initial specification from column (5) omits the frictionless wage term, but the wage elasticity can still be validly estimated if the components of X affecting w have been absorbed out via differencing. This provides a first test for the general ability of our model to absorb the endogenous components of w. This seems to be the case: looking at average elasticities, a point percentage increase in wage would correspond to a 0.20 per cent reduction in the hours of work, significant at the 0.001 level. Adding back the control for the frictionless wage in column (6), the wage elasticity coefficient turns to -0.23 per cent. Our tests suggest that the two coefficients are not statistically different from each other (Z-test: $0.06 < 1.96$), and that the search shock is independent from w.

Platform-specific elasticities from Table 3 point at a backwards-bending supply for all workers, albeit statistically significant at the 0.001 level for the subset of online workers with no control over pay, which is also our sample's largest group of workers. The standard errors are larger due to the smaller size of each subsample of platforms, but, in general, there are no differences in a statistical sense between the coefficients from columns (5) and (6).

We can now test for "weak" endogeneity. Comparing the α_1 coefficients between the actual (column 4) and differenced (column 6) supply models provides the most exhaustive of our robustness checks, as discussed in Section 5. Since each model depends on a different set of assumptions, it is unlikely for both models to be invalid and, at the same time, to converge to the same estimate for wage elasticity. Our tests reveal that the coefficients are not statistically different from each other (Z-test: $0.157 < 1.96$), suggesting that the search shock is exogenous to individual characteristics X shared between the actual and desired equation.

Columns (7) to (9) test the robustness of the estimates from column (6) by removing sets of variables in a stepwise fashion, completing our set of tests for endogeneity by testing for its "strong" variant.

We begin in column (7) with the most important of these tests by removing all individual-level controls and last-occupation fixed effects corresponding exactly to the differenced supply model from equation (11), in which X was cancelled out.[29] The removal of all regressors has a very small effect on our elasticity estimates ($\alpha_1 = -0.234$, $\alpha_2 = -0.037$) with a minor influence over the overall predictive power of the model. Further tests reveal that we cannot reject the null hypothesis that both coefficients α_1 (Z-test: $0.063 < 1.96$) and α_2 (Z-test:

[29]Note that these include platform-specific experience terms, which are the most likely characteristics that can potentially affect the search shock but not the frictionless wage.

$0.339 < 1.96$) remain the same even after removing all other individual controls in the model. Similar conclusions can be drawn from the heterogeneous elasticity models from Table 3, as the differences between columns (6) and (7) remain minimal and statistically not significant. These results suggest that the search shock is independent of unobserved individual characteristics affecting only one of the two supply equations.

The cancelling of the X vector suggests that estimates for the frictionless wage elasticity α_2 from column (6) of both Tables 2 and 3 are net of unobserved individual characteristics affecting the salary. In column (6), we estimate that a point percentage increase in the frictionless wage leads to a 0.044% reduction in the final supply. The negative sign is consistent with the backwards-bending estimates of the wage elasticity of supply. This finding suggests that, on average, workers with higher frictionless salaries work less than other workers who have experienced the same search shock but that this behaviour is driven entirely by expectations. In other words, as workers with high frictionless wages adjust their supply for the search shock, they display the same wage elasticity as other workers, but initial expectations based on the frictionless wage make these workers lose several hours of work, during which these workers would have preferred to work in order to reach their earnings target. Further simplifying, we come to the conclusion that information gaps on the search shock make workers miss their earnings targets.

It is also apparent that experience in the platform ceases to be a factor in this information process: notably, the coefficients for years of experience in the platform and for regularity of activity in the platform turn statistically zero once we move to the differenced specifications.

We repeat these tests in columns (8) and (9), first by removing all platform-side controls and then all platform and individual controls. While we do not expect our estimates to be robust to the removal of platform-side controls,[30] we still note that our wage elasticity estimates remain fairly robust to the removal of these controls. The fact that the α_1 coefficient remains statistically unchanged even after nearly all controls have been removed (column 9, Table 2) suggests that labour supply elasticities might even be estimated semi-parametrically should within-reference week variation in wage be accounted for.

[30] Recall that, in the model (11), the P vector is not cancelled out.

Table 4: Search effort, individual determinants

	Search effort (ln)			
	(1)	(2)	(3)	(4)
w (ln)	-0.045**			
	(0.015)			
Non-platform income (ln)	0.005	-0.002	-0.002	-0.007
	(0.012)	(0.011)	(0.011)	(0.011)
Non-platform income (zero)	-0.023	-0.107	-0.112	-0.122
	(0.138)	(0.136)	(0.133)	(0.135)
Education: Medium (ISCED 3-4)	-0.081	-0.077	-0.070	-0.079
	(0.123)	(0.122)	(0.118)	(0.117)
Education: High (ISCED 5-8)	0.025	0.020	0.001	-0.023
	(0.129)	(0.128)	(0.123)	(0.120)
Gender: female	-0.012	0.012	0.011	-0.012
	(0.082)	(0.082)	(0.078)	(0.074)
Married or living with a partner	-0.045	-0.032	0.010	0.036
	(0.069)	(0.069)	(0.067)	(0.068)
Female × partner	0.040	0.016	-0.007	-0.012
	(0.109)	(0.109)	(0.105)	(0.106)
Partner's income: similar	-0.060	-0.056	-0.059	-0.039
	(0.074)	(0.074)	(0.072)	(0.070)
Partner's income: lower	-0.063	-0.067	-0.073	-0.144
	(0.076)	(0.076)	(0.077)	(0.076)
Age	0.007	0.009	0.012	0.005
	(0.014)	(0.014)	(0.013)	(0.013)
Foreign born	-0.017	-0.030	-0.050	-0.063
	(0.102)	(0.103)	(0.100)	(0.095)
Household size	0.004	0.004	0.002	0.007
	(0.027)	(0.027)	(0.027)	(0.028)
Number of children	0.023	0.015	0.016	0.032
	(0.032)	(0.032)	(0.031)	(0.031)
Regular platform worker	-0.043	-0.043	-0.049	-0.079
	(0.055)	(0.056)	(0.055)	(0.055)
Experience in platforms (yrs.)	-0.004	-0.005	-0.003	0.000
	(0.007)	(0.007)	(0.006)	(0.006)
Country FE	Yes	Yes	Yes	Yes
Last occupation FE	Yes	Yes	No	No
Platform controls	Yes	Yes	Yes	No
Adjusted R-Squared	0.090	0.084	0.067	0.014
Observations	1580	1580	1580	1580

Notes: SE clustered by occupation/sector clusters in parentheses. Joint test of significance (individual controls): Specification (1) $F(17, 179) = 0.57$, Prob $> F = 0.900$; Specification (2) $F(16, 179) = 0.60$, Prob $> F = 0.883$; Specification (3) $F(16, 179) = 0.59$, Prob $> F = 0.890$; Specification (4) $F(16, 179) = 1.36$, Prob $> F = 0.164$.

*p<.05, **p<.01, ***p<.001

The α_2 coefficient loses some of its magnitude (but the magnitude of the standard error is unchanged) after the removal of these controls, suggesting that the platform characteristics might play an important role in shaping workers' expectations or introduce further variability in the frictionless salary. Similar findings can be drawn from the heterogeneous elasticity estimates from Table 3. Here, the standard error is not low enough to warrant significance for many of the estimated elasticities. The smaller size of each platform-specific cell probably prevents us from absorbing the variability in w, explaining the larger variation across each platform-specific estimate relative to the results from Table 2. The estimates remain, nonetheless, far from inconclusive, as the actual and frictionless wage elasticities remain negative and higher than -1.

We conduct a final test for the "strong" exogeneity of the search shock. Earlier in Section 5, we mentioned that the search effort should only depend on the frictionless wage unless other unobserved ability factors were to affect the search shock. In table Table 4, column (1), we test whether the search effort is correlated with individual characteristics after controlling for the frictionless wage. Our results suggest that this is not the case: none of the individual-level coefficients are statistically significant, and a joint test of significance suggests that we cannot reject the hypothesis that these coefficients are jointly zero at the 5% significance level.

These results are sufficient to dispel any claims that workers' ability is endogenous to the search shock. Nonetheless, we test for joint significance of the individual controls vector again after removing the wage (column 2), occupation fixed effects (3), and platform controls (column 4) to find that these individual-level predictors have no effect on the search effort.[31]

Finally, we conduct a final test for the presence of measurement error in column (10) of Table 2 by introducing an interaction between the two wage terms, yielding the α_3 parameter. As mentioned earlier in Section 5, the introduction of the interaction term should leave our elasticity estimates unchanged from our main specification. Our statistical tests suggest that there is no difference in a statistical sense between the elasticity estimates from columns (6) and (10) (α_1, Z-test: $0.041 < 1.96$; α_2, Z-test: $0.046 < 1.96$). Furthermore, a non-zero

[31] Furthermore, these results suggest that these ability factors do not even indirectly affect search through the frictionless salary, reinforcing the results from Cantarella and Strozzi (2021) which suggested that the frictionless rate of pay in online platforms is independent of idiosyncratic ability.

interaction effect would suggest that there is residual variation in supply driven by the wage that is not explained by our two wage terms. However, the estimated interaction coefficient α_3 is statistically not significant and close to zero in magnitude (0.001), further reinforcing the validity of our estimates.

This test suggests that measurement error is negligible and that, in this scenario, it does not affect the validity and consistency of the estimates for both elasticities. While this could be serendipitous, the standard error of the estimated interaction parameter allows us to derive an upper bound for the between-equation idiosyncratic variance of the frictionless wage itself, which, as discussed, is absorbed by the α_2 parameter. As discussed in Appendix D.1, the underlying variance appears to be no larger than 20% of each worker's frictionless wage.

Appendix A offers some more minor checks. Among these, we test for other minor measurement error issues and heterogeneity in labour supply conditional on access to "traditional" labour markets.

7 Conclusions

In this paper, we have studied the wage elasticity of workers' labour supply in the platform economy using data from a recent survey on European platform workers. Adopting a novel approach that exploits information on labour supply outcomes under perfect and imperfect information on task search, we have estimated labour supply elasticities for platform workers. Our method is robust to several rigorous checks and produces wage elasticities that are consistent with the target-earning literature and that validate results from experimental trials in platform labour markets. This method can be extended to other contexts with some adjustments, freeing up degrees of freedom in small samples and proving particularly useful in contexts where longitudinal data is unavailable. Semi-parametric applications of the method can also produce idiosyncratic elasticities under specific circumstances.

Our results show that workers exhibit a backwards-bending labour supply curve regardless of platform type. In particular, we find that workers adjust to pay rate variation caused by task search - i.e. the *actual* wage - by increasing their supply by about 0.20 per cent for each percentage decrease in wage.

Our findings also indicate that earnings unadjusted for task search play a key role in explaining the residual variation in the labour supply of platform workers. For each percentage point increase in the rate of pay net of search - i.e. the *frictionless* wage - the supply is reduced by around 0.04 percentage points. As wage variation from the frictionless rate can only be negative, this behaviour suggests that workers initially undershoot their supply schedule based on their frictionless rate of pay and then try to readjust it during the working week for a net income loss.

Overall, these results suggest that both variation in job search and individual expectations play a central role in shaping the labour supply of platform workers and that backwards-bending preferences are relevant to both start-of-the-week scheduling and infra-week adjustments. These considerations suggest the presence of target-earning behaviour in the context of reference dependence, where previous payoff losses lead to increases in the search effort for the next task.

These results have important implications. Firstly, our findings support existing evidence on the monopsony power enojoyed by platforms. In the presence of negative and inelastic wage elasticities, this monopsony power is even greater than what suggested by current studies on platform data, where labour supply is found mainly inelastic. Furthermore, the persistence of negative estimates for search-unadjusted frictionless earnings suggests that platforms enjoy a further channel of monopsonistic power. In fact, platforms are not only able to take advantage of backwards-bending elasticities in the change in salary caused by search, but they also benefit from similar negative elasticities for changes in frictionless salary levels in general. This is all to the benefit of platforms and clients, with workers experiencing a net loss in utility and, possibly, job quality deterioration as they work harder for less.

Secondly, our results extend these considerations to the entirety of platform workers. Negative wage elasticities are found for workers on any platform, no matter if the service is provided online or in-person or if the worker has control over pay or not. The fact that these market inefficiencies persist even after accounting for differences in platform types suggests that this behaviour is structural to platform work and that, in line with the conclusions of behavioural theory on workers' behaviour, platform workers are loss averse on average. Whether this loss aversion character is endogenous to platform workers or if it is caused instead by the freelancing nature of these jobs is up for debate.

Finally, it is worth taking into account that search and demand-side shocks might generate similar behaviour in other labour markets through unpaid overtime, bonuses, or commissions. Along this line, our results could be interpreted as an invitation to the reassessment of monopsony power in more "traditional" labour markets as well. Future research could build on the methods we have developed in this paper, for example, by adapting them to contexts in which supply is constrained by fixed-hours schedules or by incorporating wage expectations in our econometric setting. Experimental trials can also explore the connection between stated choice supply preferences and actual supply behaviour, with an eye on the heterogeneities connected to wage expectations.

References

Alkhatib, A., Bernstein, M. S., and Levi, M. (2017). Examining crowd work and gig work through the historical lens of piecework. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, page 4599–4616, New York, NY, USA. Association for Computing Machinery.

Arellano, M. and Meghir, C. (1992). Female labour supply and on-the-job search: An empirical model estimated using complementary data sets. *The Review of Economic Studies*, 59(3):537–559.

Barbos, A. and Kaisen, J. (2022). An example of negative wage elasticity for youtube content creators. *Journal of Economic Behavior & Organization*, 203:382–400.

Barcevičius, E., Gineikytė-Kanclerė, V., Klimavičiūtė, L., and Martin, N. R. (2021). Study to support the impact assessment of an eu initiative to improve the working conditions in platform work. Final report to the European Commission, PPMI.

Bell, D. N. F. and Hart, R. A. (1999). Unpaid work. *Economica*, 66(262):271–290.

Berg, J., Furrer, M., Harmon, E., Rani, U., and Silberman, M. S. (2018). Digital labour platforms and the future of work: Towards decent work in the online world. International Labour Office.

Blundell, R. and Macurdy, T. (1999). Labor supply: A review of alternative approaches. In *Handbook of Labor Economics*, pages 1559–1695. Elsevier.

Bogliacino, F., Codagnone, C., Cirillo, V., and Guarascio, D. (2020). Quantity and quality of work in the platform economy. In *Handbook of Labor, Human Resources and Population Economics*, pages 1–28. Springer International Publishing.

Bouhlel, I., Chessa, M., Festré, A., and Guerci, E. (2022). When to stop searching in a highly uncertain world? a theoretical and experimental investigation of "two-way" sequential search tasks. *Journal of Economic Behavior & Organization*, 203:80–92.

Bracha, A. and Burke, M. A. (2021). How big is the gig? The extensive margin, the intensive margin, and the hidden margin. *Labour Economics*, 69:101974.

Butschek, S., González Amor, R., Kampkötter, P., and Sliwka, D. (2022). Motivating gig workers – evidence from a field experiment. *Labour Economics*, 75:102105.

Camerer, C., Babcock, L., Loewenstein, G., and Thaler, R. (1997). Labor supply of New York City cabdrivers: One day at a time. *The Quarterly Journal of Economics*, 112(2):407–441.

Cantarella, M. and Strozzi, C. (2021). Workers in the crowd: the labor market impact of the online platform economy. *Industrial and Corporate Change*, 30(6):1429–1458.

Chen, M. K., Rossi, P. E., Chevalier, J. A., and Oehlsen, E. (2019). The value of flexible work: Evidence from Uber drivers. *Journal of political economy*, 127(6):2735–2794.

Cohen, P., Hahn, R., Hall, J., Levitt, S., and Metcalfe, R. (2016). Using big data to estimate consumer surplus: The case of Uber. Working Paper 22627, National Bureau of Economic Research.

Davis, M. and Hoyt, E. (2020). A longitudinal study of piece rate and health: evidence and implications for workers in the us gig economy. *Public Health*, 180:1–9.

De Stefano, V. and Aloisi, A. (2018). *European legal framework for "digital labour platforms"*. European Commission. Joint Research Centre.

De Stefano, V., Durri, I., Stylogiannis, C., and Wouters, M. (2021). Platform work and the employment relationship. ILO Working paper 27, International Labour Organisation.

Dickinson, D. L. (1999). An experimental examination of labor supply and work intensities. *Journal of Labor Economics*, 17(4):638–670.

Dube, A., Jacobs, J., Naidu, S., and Suri, S. (2020). Monopsony in online labor markets. *American Economic Review: Insights*, 2(1):33–46.

Duch-Brown, N., Gomez-Herrera, E., Mueller-Langer, F., and Tolan, S. (2022). Market power and artificial intelligence work on online labour markets. *Research Policy*, 51(3):104446.

Euwals, R. and van Soest, A. (1999). Desired and actual labour supply of unmarried men and women in the Netherlands. *Labour Economics*, 6(1):95–118.

Farber, H. S. (2005). Is tomorrow another day? The labor supply of New York City cabdrivers. *Journal of Political Economy*, 113(1):46–82.

Fehr, E. and Goette, L. (2007). Do workers work more if wages are high? Evidence from a randomized field experiment. *American Economic Review*, 97(1):298–317.

Hart, R. A. (2008). Piece work pay and hourly pay over the cycle. *Labour Economics*, 15(5):1006–1022.

Hart, R. A. and Roberts, J. E. (2012). Real wage cyclicality and the Great Depression: evidence from British engineering and metal working firms. *Oxford Economic Papers*, 65(2):197–218.

Horton, J. J. (2017). The effects of algorithmic labor market recommendations: Evidence from a field experiment. *Journal of Labor Economics*, 35(2):345–385.

Horton, J. J. and Chilton, L. B. (2010). The labor economics of paid crowdsourcing. In *Proceedings of the 11th ACM Conference on Electronic Commerce*, EC '10, page 209–218, New York, NY, USA. Association for Computing Machinery.

International Labour Office (2021). *World employment and social outlook 2021*. International Labour Office, Switzerland.

Kingsley, S. C., Gray, M. L., and Suri, S. (2015). Accounting for market frictions and power asymmetries in online labor markets. *Policy & Internet*, 7(4):383–400.

Kukavica, A., McKenna, S., Shum, M., Chen, K., and Camerer, C. F. (2022). Beyond taxis: Reference-dependence in rideshare drivers' labor supply. Mimeo, California Institute of Technology.

Lam, C. T., Liu, M., and Hui, X. (2021). The geography of ridesharing: A case study on New York City. *Information Economics and Policy*, 57:100941.

Lehdonvirta, V. (2018). Flexibility in the gig economy: Managing time on three online piece-work platforms. *New Technology, Work and Employment*, 33(1):13–29.

Martin, V. (2017). When to quit: Narrow bracketing and reference dependence in taxi drivers. *Journal of Economic Behavior & Organization*, 144:166–187.

Orland, A. and Rostam-Afschar, D. (2021). Flexible work arrangements and precautionary behavior: Theory and experimental evidence. *Journal of Economic Behavior & Organization*, 191:442–481.

Parker, S. C., Belghitar, Y., and Barmby, T. (2005). Wage uncertainty and the labour supply of self-employed workers. *The Economic Journal*, 115(502):C190–C207.

Stafford, T. M. (2015). What do fishermen tell us that taxi drivers do not? An empirical investigation of labor supply. *Journal of Labor Economics*, 33(3):683–710.

Stenborg Petterson, M. (2022). Estimation of a latent reference point: Method and application to NYC taxi drivers. Mimeo, Department of Economics Brown University.

Stewart, M. B. and Swaffield, J. K. (1997). Constraints on the desired hours of work of British men. *The Economic Journal*, 107(441):520–535.

Wales, T. J. (1973). Estimation of a labor supply curve for self-employed business proprietors. *International Economic Review*, 14(1):69.

Zubrickas, R. (2023). The relative income effect and labor supply. *Journal of Economic Behavior & Organization*, 209:176–184.

Appendix

A Additional robustness checks

Table A.1 offers further robustness checks of our results from Table 2.

Columns (1) and (2) offer further tests to the *"Self-consciousness"* and *"No devil's advocate"* conditions. The former can be tested by removing all individual controls and checking for a change in the wage term. While this is far from an exhaustive test, a lack of change would indicate that individuals already take ability into account when expressing their desired labour supply. According to our results, this seems to be the case, as the elasticity coefficient is not particularly affected by the removal of individual controls (moving from -0.083 to -0.076, Z-test: $0.35 < 1.96$).

The latter is tested by adding task search as a regressor in the desired hours equation in column (2). Recall that search is a platform-side constraint that is not supposed to enter the desired hours equation. In case task search is correlated with the desired labour supply, a significant change in the wage term would indicate that desired supply schedules are endogenous to the wage. This is not the case, as the change in coefficient from the original estimate is not statistically different from zero (Z-test: $0.11 < 1.96$).

Turning to the differenced equation, we conduct one final check to study whether workers with access to other sources of income display different behaviour in the labour market. Individuals with no access to jobs other than platform work may be less risk-averse and display different elasticities. We test this in columns (3) and (4) of Table A.1 by introducing interactions between wage and employment status in non-platform jobs so that heterogeneous elasticities can be studied. However, the estimated interaction coefficients for α_1 and α_2 are statistically not significant, ruling out the hypothesis that workers without access to other occupations might act differently.[32]

[32] In any case, this is far from a definitive check because access to other occupations might be endogenous to the salary, even if the ever-unemployed status is controlled for within the last occupation FE vector.

Table A.1: Labour supply elasticity estimates

	Hours, desired (ln)		Hours, first-diff. (ln)	
	(1)	(2)	(3)	(4)
α_1	-0.076***	-0.081***	-0.187***	-0.185***
	(0.014)	(0.012)	(0.052)	(0.051)
α_2			-0.042**	-0.038*
			(0.013)	(0.016)
On-location worker	-0.100	-0.049	0.186***	0.186***
	(0.079)	(0.071)	(0.053)	(0.053)
Control over pay	0.091	-0.025	0.064	0.065
	(0.059)	(0.115)	(0.104)	(0.104)
On-location worker × Control over pay	0.002	-0.016	-0.233*	-0.234*
	(0.118)	(0.114)	(0.105)	(0.104)
Search (ln)		0.377***		
		(0.021)		
Regular platform worker		0.155**	-0.028	-0.027
		(0.052)	(0.050)	(0.050)
Experience in platforms (yrs.)		-0.014*	0.003	0.003
		(0.006)	(0.006)	(0.006)
α_1 (Employed in trad. job)			-0.094	-0.101
			(0.100)	(0.102)
α_2 (Employed in trad. job)				-0.010
				(0.022)
Employed in trad. job			-0.040	-0.028
			(0.070)	(0.077)
Country FE	Yes	Yes	Yes	Yes
Last occupation FE	No	Yes	Yes	Yes
Individual controls	No	Yes	Yes	Yes
Platform controls	No	Yes	Yes	Yes
Adjusted R-Squared	0.080	0.312	0.156	0.156
Observations	1591	1580	1580	1580

Notes: SE clustered by occupation/sector clusters in parentheses. All wage and labour supply variables are expressed in natural logarithms.

*p<.05, **p<.01, ***p<.001

B Non-linear search shock

Assuming a linear form for the search function, it is easy to show that the rate of change from the *frictionless* to *actual* wage is independent of search. Omitting the individual subscripts, we obtain this equation:

$$\frac{w^A}{w} = \frac{h}{h^A} = \frac{H_a(\bar{w}, a)\rho S}{H_a(\bar{w}, a)\rho S + S} \tag{14}$$

Taking the derivative of both sides of this equation with respect to S, the rate of change in w^A/w is zero:

$$\frac{dw^A}{dwS} = \frac{d}{dS}\left(\frac{H_a(\bar{w}, a)\rho S}{H_a(\bar{w}, a)\rho S + S}\right) = 0 \tag{15}$$

This shows that, under the assumption of a linear search shock, the search effort cannot determine the rate of change between actual and frictionless salary, but only by the search shock. This is, after all, already an implication of the nature of w and w^A.

The only consideration that applies is that the choice variable (the search effort) has to be independent of efficiency a for the function $H(W(a), S)$ to be separable. This assumption is reasonable and, at the same time, similar to assumptions which have been treated as standard in the related literature, such as Stenborg Petterson (2022). In other words, the decision to keep searching is independent of efficiency factors after holding the pay rate fixed.

We then turn to the non-linear specification. The claim that w^A is not related to S relies on the specific functional form assumption of ρS. Different specifications might change this result.

For example, assume that the search function has a positive intercept term, which indicates that workers acquire a minimum number of tasks without searching. When the search function is specified as $\rho S + K$, the ratio of w^A/w becomes:

$$\frac{w^A}{w} = \frac{H_a(\bar{w}, a)(\rho S + K)}{H_a(\bar{w}, a)\rho(\rho S + K) + S} \tag{16}$$

In this case, S will not be easily cancelled out, and w^A/w will become a function of S, and therefore violate the exogeneity condition by depending on w. Similarly, if the search function takes any non-linear form to capture increasing/decreasing returns to the search effort, w^A/w will also become a function of S. The implications are clear: for a lower or higher level of effort, the rate of change in salary will also change.

These considerations delimit the application of our simplified theoretical model to a supply-only equilibrium, which, demand-side, takes into account only the final search shock experienced by each worker once the working week has ended. At the end of the working week, workers have already experienced the salary change w^A/w, and have already adjusted their labour supply for w^A, the hourly rate they settled with. What matters for each worker, *given* w^A/w and their utility function, is whether a different level of supply h^A and then S would have been more optimal, no matter if w^A/w would have been different under a different level of search.

This suggests that a non-linear search shock might be endogenous to the frictionless wage w, which is the main determinant of the search effort. Our empirical models explicitly address this problem by controlling for w. Further issues are related to the situation in which the shock is endogenous to X, i.e. when workers' skills correlate with the ability to predict the shock or to manipulate it and enter the search function.

These issues have some implications for our empirical approach, depending on two exogeneity conditions discussed in detail in Section 5.

C Additional statistics

Table C.2: Summary statistics by type of platform

	(1) On-location		(2) Online	
Control over pay:	No	Yes	No	Yes
Paid hours of work	13.360	12.398	12.178	15.159
	(13.367)	(11.524)	(11.501)	(14.138)
Actual hours of work	22.027	21.643	21.161	25.245
	(19.501)	(18.653)	(18.083)	(21.840)
Desired hours of work	17.521	18.760	20.338	21.192
	(13.766)	(14.164)	(15.355)	(15.702)
Frictionless hourly wage	41.530	39.680	20.290	44.656
	(121.350)	(185.425)	(88.932)	(275.838)
Actual hourly wage	20.518	17.181	10.231	16.196
	(68.078)	(80.985)	(36.682)	(77.789)
Gender: Female	0.338	0.404	0.447	0.383
	(0.474)	(0.492)	(0.498)	(0.487)
Married or living with a partner	0.592	0.667	0.630	0.649
	(0.493)	(0.473)	(0.483)	(0.478)
Age	33.983	33.690	36.945	33.129
	(12.344)	(12.192)	(12.730)	(11.489)
Foreign born	1.900	1.942	1.937	1.900
	(0.301)	(0.235)	(0.244)	(0.301)
Household size	3.087	3.152	2.960	3.168
	(1.260)	(1.222)	(1.211)	(1.186)
Number of children	0.754	0.724	0.621	0.715
	(0.999)	(0.857)	(0.856)	(0.915)
Education: Low (ISCED 1-2)	0.071	0.047	0.037	0.035
	(0.257)	(0.212)	(0.189)	(0.184)
Education: Medium (ISCED 3-4)	0.358	0.310	0.309	0.275
	(0.481)	(0.464)	(0.462)	(0.447)
Education: High (ISCED 5-8)	0.571	0.643	0.654	0.691
	(0.496)	(0.480)	(0.476)	(0.463)
Non-platform income	3780.139	2972.746	2803.045	2811.917
	(8489.877)	(6953.918)	(7282.389)	(6865.912)
Employed in trad. job	0.454	0.485	0.458	0.512
	(0.499)	(0.501)	(0.499)	(0.500)
Partner's income: higher	0.627	0.500	0.571	0.520
	(0.485)	(0.502)	(0.495)	(0.500)
Partner's income: similar	0.254	0.281	0.199	0.302
	(0.437)	(0.451)	(0.400)	(0.460)
Partner's income: lower	0.120	0.219	0.230	0.178
	(0.326)	(0.416)	(0.421)	(0.383)
Experience in platforms (yrs.)	4.604	4.287	4.684	4.795
	(4.401)	(3.665)	(4.342)	(4.206)
Regular platform worker	0.654	0.602	0.821	0.769
	(0.477)	(0.491)	(0.384)	(0.422)
Platform's control over work hours, None	0.487	0.427	0.655	0.503
	(0.501)	(0.496)	(0.476)	(0.501)
Observations	240	471	459	725

Notes: Mean coefficients; Standard Deviation in parentheses. Breakdown by platform type (On-location vs. Online) and remuneration schemes (workers having control over pay vs. workers having no control over pay). Partner's income statistics only estimated for the subsample of workers who are married/ living with a partner.

D Monte Carlo results

D.1 Idiosyncratic variation in the frictionless wage

In this section, we evaluate the robustness of our empirical model through Monte Carlo simulations. We have mentioned in Section 4 that the model can tolerate some degree of idiosyncratic variation in the frictionless wage w. In this subsection, we quantify how much variation can be tolerated to derive a "rule of thumb" condition under which the results can be considered reliable.

Variations in the nominal wage in-between the desired and actual supply equation will be absorbed by α_2, so that the parameter will feature a random disturbance term. If the sample size is not large enough, our results may be affected by the heterogeneity in the expectations

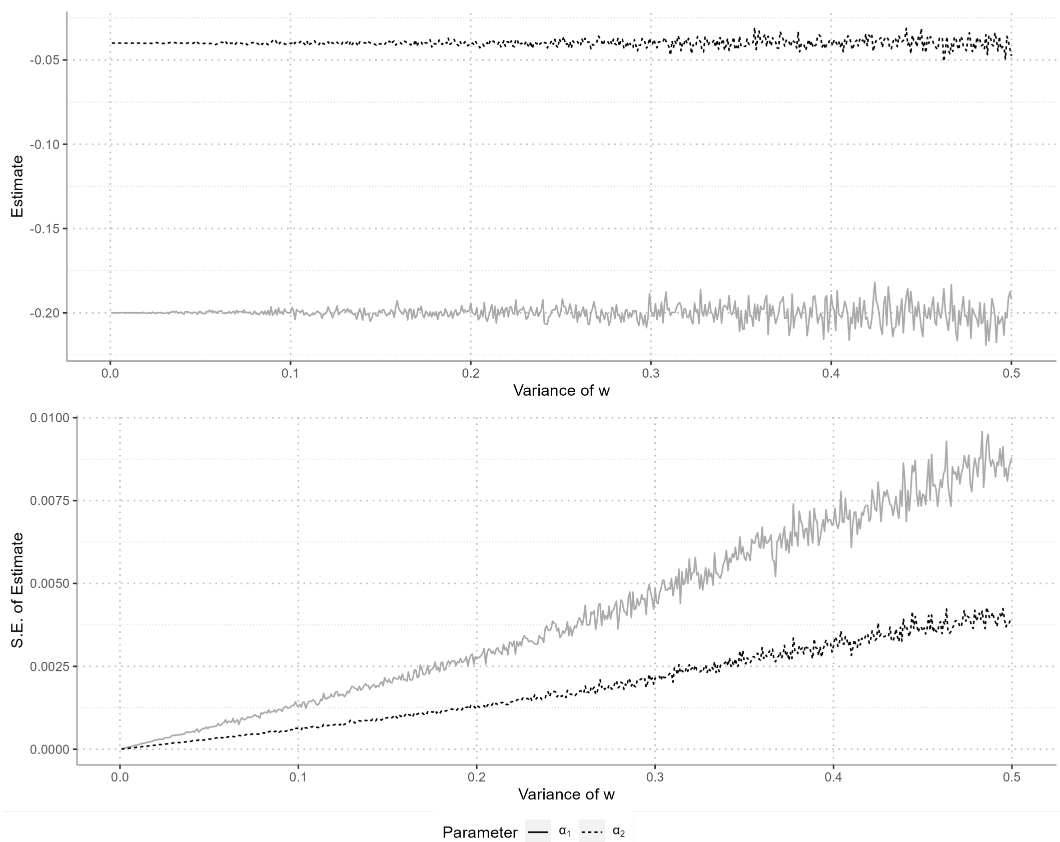

FIGURE D.1: MONTE CARLO SIMULATION - ESTIMATES, FIRST-DIFF. MODEL

Notes: Monte Carlo results for elasticity estimates and the standard error of the estimate for the differenced supply model.

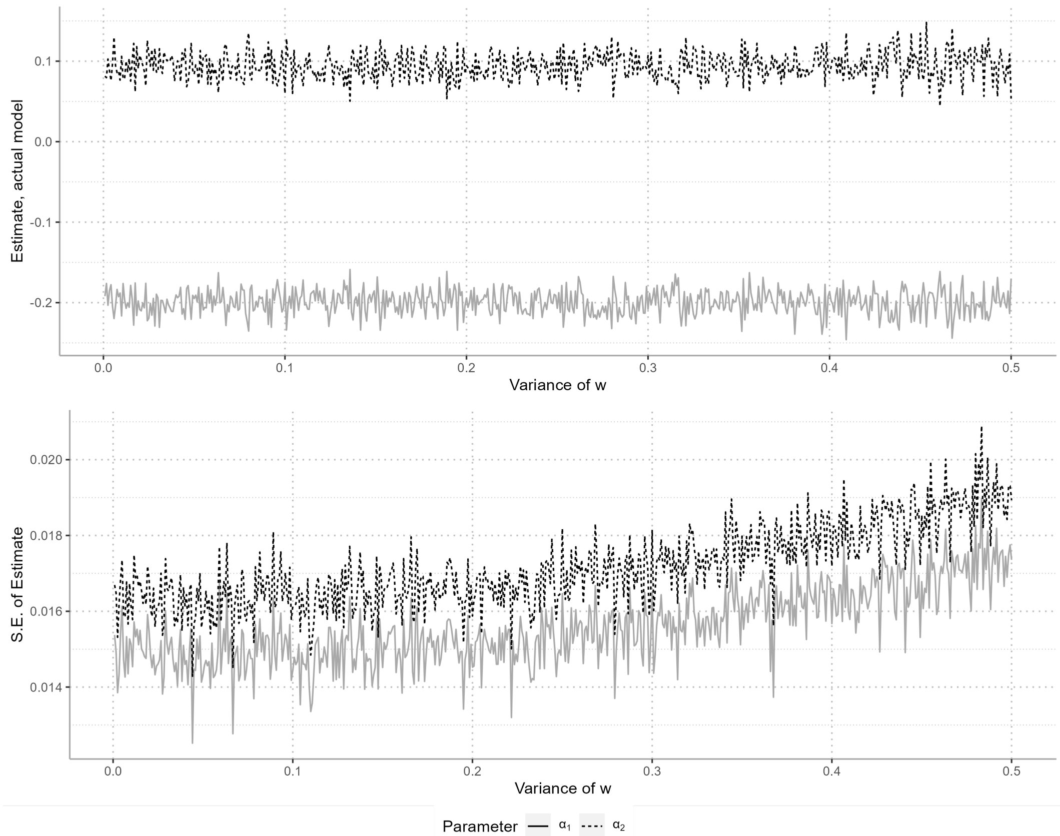

Figure D.2: Monte Carlo Simulation - Estimates, Actual model

Notes: Monte Carlo results for elasticity estimates and the standard error of the estimate for the actual supply model.

of the search shock, and estimates for both α_1 and α_2 will become increasingly unreliable with larger variance in expectations.

The problem is illustrated in Figures D.1 and D.2 in which we show, through Monte Carlo simulations, how the elasticity parameters react to variation in wage across the first-differenced and actual supply models. Assuming these disturbances are normally distributed, we model the disturbance as a multiplicative shock $w_0 = w * \mathcal{N}(1, \sigma)$ on the frictionless wage.[33] Estimates are produced for both the actual and differenced supply, with the priors $\alpha_1 = -0.20$ and $\alpha_2 = -0.04$ (mirroring our results), the search shock $\rho = \mathcal{N}(2, 0.4)$ and the standard assumption that both supply and the frictionless wage are also affected by an

[33]The same disturbances can equivalently affect w_1, or both terms. Our considerations are unaffected because the disturbance will always be absorbed by α_2. We have also experimented with different functional forms of the disturbance (logarithmic and additive) with similar results.

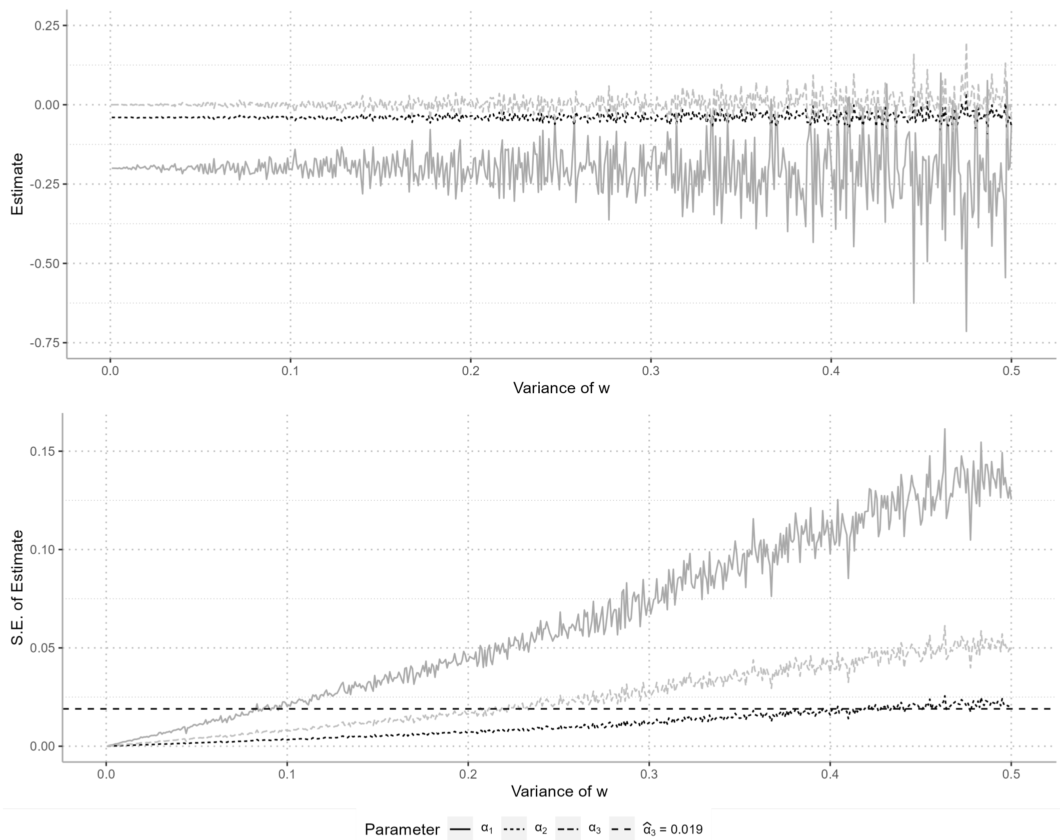

FIGURE D.3: MONTE CARLO SIMULATION - ESTIMATES, FIRST-DIFF. MODEL WITH WAGE INTERACTION

Notes: Monte Carlo results for elasticity estimates and the standard error of the estimate for the differenced supply model, including interactions between the wage terms.

unobserved skill term.[34] We produce 600 replicates with increasing variance in wage σ, holding the sample size fixed at 1580, replicating our empirical setting. In the simulation, the variance of σ is allowed to take values between 0 and 0.5 (which is equal to 50% of a given frictionless wage).

Figures D.1 and D.2 also show that estimates for both parameters become increasingly more imprecise as the variance of expectations increases. Notably, the mean of the estimated parameters always matches our priors in the differenced model (Figure D.1), which is not the case for the actual supply model (Figure D.2), for which, as discussed in Section 4, only

[34] Distributed as $X = \mathcal{N}(4, 1)$. A one-point increase in X affects the supply by $\delta = -0.045$ and the frictionless wage by $\delta_w = 1$. The frictionless wage is distributed as $w = \mathcal{N}(10, 2)$. These priors are arbitrary and have been tuned to reproduce the results from our paper, but they are instrumental in showing how endogeneity in X does not affect our main conclusions.

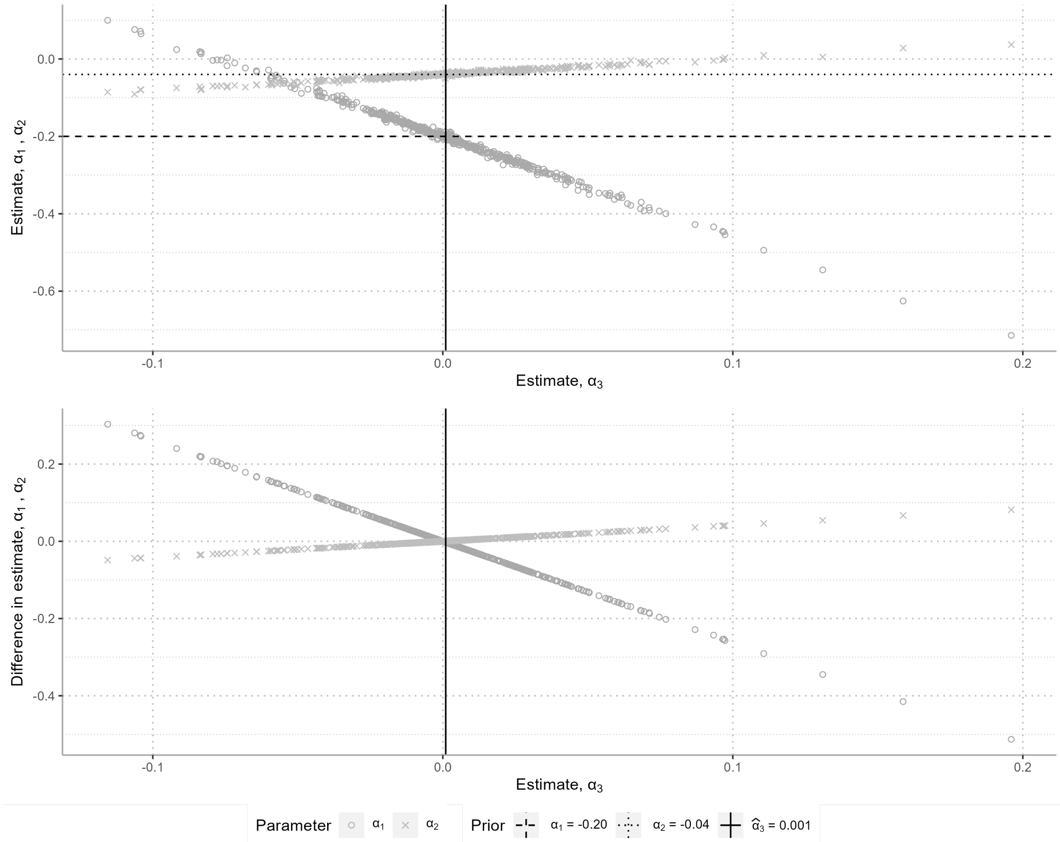

Figure D.4: Monte Carlo Simulation - Estimates, First-diff. model with wage interaction

Notes: Monte Carlo results for elasticity estimates and the standard error of the estimate for the differenced supply model, including interactions between the wage terms.

the actual supply elasticity is correctly identified. Results from the differenced model are effectively robust to variation in the unobserved ability X,[35] so that we can study deviations in frictionless wage net of all these other factors.

One important takeaway from these results is that estimates for the two elasticity parameters are fairly robust to idiosyncratic deviations of w, and are basically unaffected for deviations lower than 50% of the wage. The variation in the estimate remains, in most cases, less than proportional to the idiosyncratic variation in wage: with the chosen priors, the wage elasticity estimate remains negative and inelastic even for large values of σ.

Another notable difference between the two models is that the variation in the estimated α_2, which absorbs the variation in wage, is smaller in the differenced model, while it is roughly

[35] We can use different priors for unobserved ability, only to reach the same conclusion.

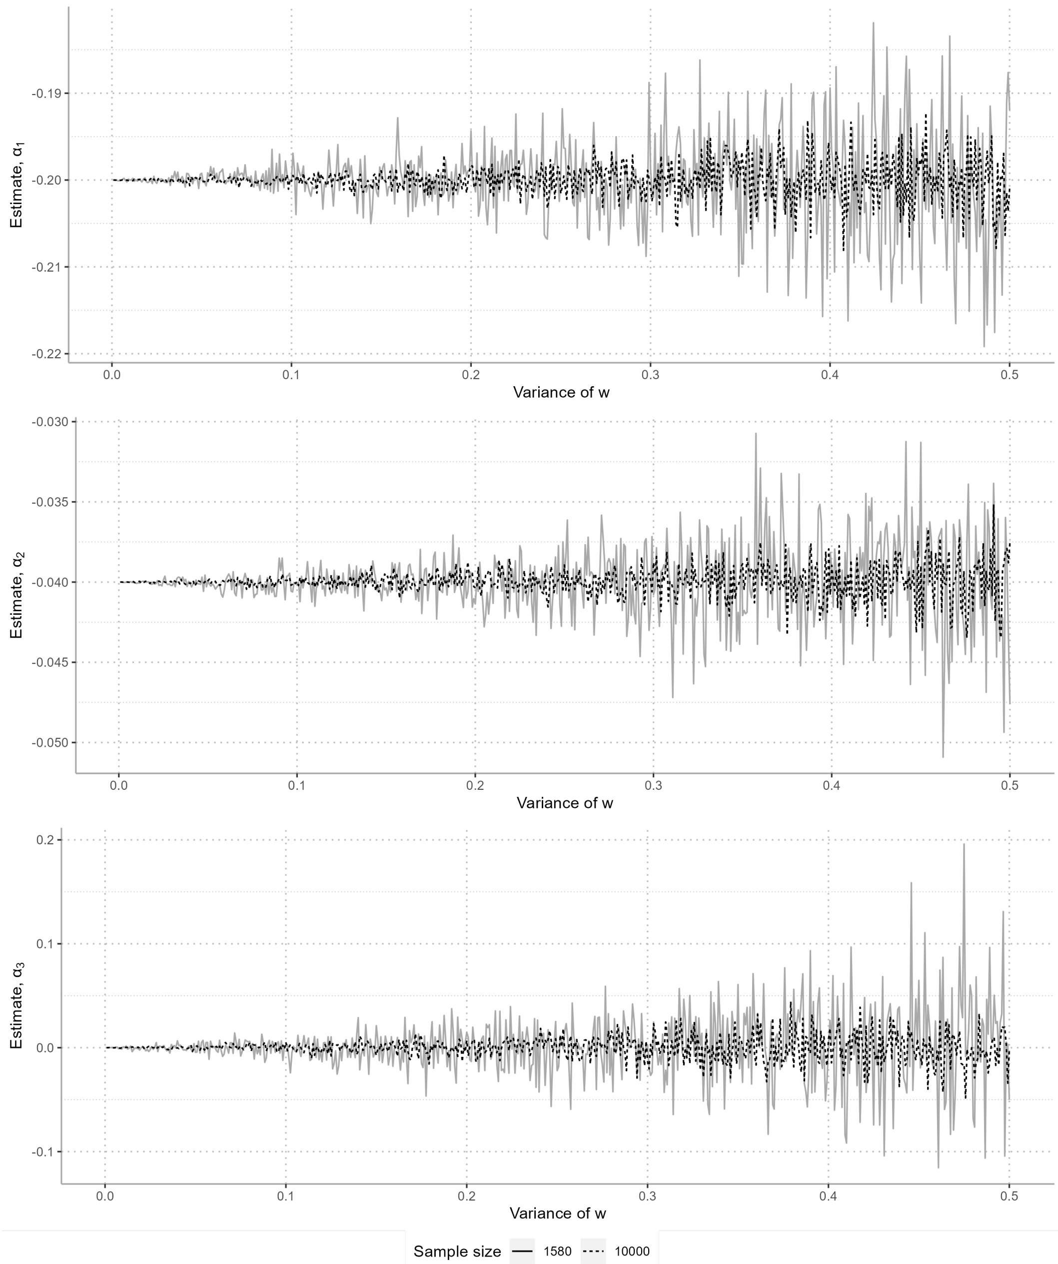

Figure D.5: Monte Carlo Simulation - Estimates, different sample sizes

Notes: Monte Carlo results for the α_1 (top, baseline first-diff. model), α_2 (middle, baseline first-diff. model), and α_3 (bottom, first-diff. model with wage interactions) estimates under different sample sizes.

identical to the variation in α_1 in the actual supply model. This is made even more evident when plotting the variation in the standard error of estimation of the two parameters between the models. The figures show that, in the differenced model, the standard error of the estimate is generally smaller than in the actual model and features much lower standard errors for α_2.

Most importantly, the variation in the standard error of the estimate is extremely limited and predictable.

Increases in the variability of w will also affect the elasticity α_3 of the interaction term between the actual and frictionless wage elasticity, as discussed in Section 5. The relationship between the interaction elasticity and the main elasticity estimates is key to evaluating the reliability of the results produced by our method.

This is shown in Figure D.3, which uses the same Monte Carlo simulation set-up from Figure D.1 but includes an interaction between the observed wage terms w^A/w and w, marked by α_3. The top figure shows how the estimate of α_3, which should be zero on average, becomes increasingly unstable as σ increases. The figure suggests that α_3 is less likely to be zero when the variance of w is high, but it can still be zero. The bottom figure suggests that the S.E. of the estimated interaction also increases linearly in σ. The dashed line sits at the standard error of the estimate (0.019) we retrieved from Table 2, Column (10). Given the sample size, we estimate an underlying variance of 0.20, a point to which we will return later.

Regardless of the variance of w, Figure D.4 provides key evidence of the relationship between the estimated elasticity of α_3 and the estimated elasticity of the other elasticity parameters α_1 and α_2. The top figure shows that whenever the estimated α_3 is close enough to zero, both α_1 and α_2 are close enough to their true elasticities, regardless of the frictionless wage variance. The vertical line sits at the estimated elasticity of α_3 from Table 2, Column (10). As the estimated elasticity is basically zero (0.001), we can conclude that the elasticity estimates we produced are unaffected by the wage variance, regardless of its level. This is exemplified by the bottom figure, which also shows that the variation in first-differenced estimates before and after the introduction of the wage term is basically zero if the interaction is also zero.

This suggests little uncertainty exists on the estimated elasticity parameters under a zero interaction term. The estimated standard error from Figure D.3 can then be considered reliable,[36] and an upper bound for σ can then be found at 0.20.

[36]The relationship between the standard error of the α_3 estimate and the variance of w would feature a different inclination under different elasticity parameters.

Finally, it can be worth examining this relationship between σ and the other parameters under different sample sizes. We do so in figure D.5 in which we superimpose estimates with a sample size of 10,000. The figures unambiguously show that the larger the sample size, the smaller the effect of heterogeneity in expectations on the estimates and their standard errors. The main takeaway is that larger deviations in expectations can be accommodated with larger samples. With a smaller sample size, the results become less reliable, and this can explain the larger standard errors and variation in the estimates for heterogeneous elasticities from table 3, correlating with the size of each platform-specific cell.

In conclusion, our Monte Carlo results have three implications. The first one is that independently of the variance of frictionless wage, the parameters will always be consistently estimated as long as the interaction term α_3 is approximately zero. The second one is that once we are comfortable enough with the estimated elasticity parameters, we can use the standard error of the α_3 estimate to derive an approximation of the variance of the wage. The rule of thumb is that if the standard error is low enough, we can confidently argue that the results are unaffected by heterogeneity in expectations. The third one is that the proposed model increases in accuracy with larger sample sizes, even if the removal of the individual characteristics vector has freed up several degrees of freedom. This means the model can tolerate a fair degree of idiosyncratic wage variation if the sample size is large enough. Further extensions of the model which wish to estimate wage elasticities on much smaller samples, such as semi-parametric approaches, could attempt to observe and factor in the unobserved idiosyncratic variation in wages during the reference period.

D.2 Replication of the validity conditions

In this section, we perform a simulation exercise to reproduce the conditions under which model validity fails. This exercise only aims to show how our results may differ if any of the conditions elaborated in Section 5 is not verified.

The results are reproduced in Table D.3, and show estimates and standard error of α_1 and α_2 for several replicates of the model under different sample sizes. Each replicate is generated for each change in sample size, so each replicate is shared between all models of the same sample size. By default, we assume that w is endogenous and let the frictionless wage have a stochastic idiosyncratic component so that $\sigma = \mathcal{N}(1, 0.20)$, which approximates our wage variance estimate obtained in Appendix D.1. The other priors are calibrated to roughly reproduce our results, as discussed in the subsection above. Each trio of columns reports results for the desired ((1)-(3)), actual ((4)-(6)), first-difference ((7)-(9)), and interacted first-difference models ((10)-(12)), respectively.

The first set of rows reports results for the baseline model, under which w is endogenous. The results are no different from the ones discussed in Appendix D.1, and the variance of the estimated elasticity parameters is well within the boundaries dictated by the variance of the idiosyncratic term of w. Note that the interaction elasticity α_3 can still be non-zero, and the further it is from this value, the more inconsistent the estimates of α_1 and α_2 will grow. It should be noted that these inconsistency issues are entirely caused by the variation in the frictionless term to which we tuned the model.

In the second set of rows, we let the search shock be endogenous to individual experience, leading to "weak" endogeneity. As discussed in the main text, this strongly affects all elasticities estimated in the actual hours model, while the elasticities of the first-differenced model are completely unaffected. The results are unchanged after introducing interactions between wage terms in the last set of columns. We study the "strong" endogeneity implications in the third set of rows. In this case, the shock correlates only with a subset of predictors, affecting only one of the two supply equations. Note how the endogeneity cannot be detected by comparing the actual and first-differenced supply equation. However, since this problem is connected with measurement error, the issue can be detected by checking for changes in the estimated coefficients after introducing the introduction term.

In the last three sets of rows, we investigate the role of measurement error beyond its effects on the idiosyncratic component. Measurement error in X is shown in the fourth set of rows and affects all estimates significantly. However, it is easily detectable by looking at the interaction coefficients. Measurement error in w_0 and w_0 (fifth and sixth set of rows) generates bias in α_2 only. Estimates for α_1 in the differenced model are basically unaffected because, as discussed, variation in w is absorbed by the α_2 term only. This variation will be undetectable. The main implication is simply that the interpretation of α_2 will need to be updated to reflect the component of variation in wage w that is shared between all individuals.

Table D.3: Monte Carlo results

	$log(h^D)$			$log(h^A)$			$log(h^A/h^D)$			$log(h^A/h^D)$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Endogenous w (baseline)											
α_1	-0.075	-0.079	-0.077	-0.214	-0.213	-0.2	-0.202	-0.201	-0.201	-0.147	-0.186	-0.189
	(0.005)	(0.003)	(0.002)	(0.011)	(0.007)	(0.005)	(0.002)	(0.002)	(0.001)	(0.033)	(0.024)	(0.018)
α_2				0.099	0.094	0.084	-0.04	-0.04	-0.04	-0.049	-0.042	-0.041
				(0.012)	(0.008)	(0.006)	(0.001)	(0.001)	(0.001)	(0.005)	(0.004)	(0.003)
α_3										-0.021	-0.006	-0.005
										(0.013)	(0.009)	(0.007)
	Weakly endogenous ρ ($\rho^* = \rho + 0.05X$)											
α_1	-0.075	-0.079	-0.077	-0.144	-0.148	-0.128	-0.199	-0.2	-0.198	-0.202	-0.205	-0.207
	(0.005)	(0.003)	(0.002)	(0.013)	(0.009)	(0.006)	(0.003)	(0.002)	(0.001)	(0.04)	(0.029)	(0.021)
α_2				0.026	0.027	0.01	-0.041	-0.04	-0.04	-0.041	-0.04	-0.039
				(0.014)	(0.01)	(0.007)	(0.001)	(0.001)	(0.001)	(0.006)	(0.004)	(0.003)
α_3										0.001	0.002	0.003
										(0.015)	(0.011)	(0.008)
	Strongly endogenous ρ ($\rho^* = \rho + 0.05(X^A - X)'$)											
α_1	-0.075	-0.079	-0.077	-0.152	-0.135	-0.122	-0.136	-0.119	-0.124	-0.097	-0.03	-0.006
	(0.005)	(0.003)	(0.002)	(0.019)	(0.013)	(0.01)	(0.014)	(0.01)	(0.007)	(0.208)	(0.153)	(0.114)
α_2				0.042	0.013	0.01	-0.035	-0.043	-0.035	-0.041	-0.056	-0.052
				(0.02)	(0.014)	(0.01)	(0.005)	(0.004)	(0.003)	(0.031)	(0.023)	(0.017)
α_3										-0.015	-0.034	-0.045
										(0.08)	(0.058)	(0.043)

	$log(h^D)$			$log(h^A)$			$log(h^A/h^D)$			$log(h^A/h^D)$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	D. M. error ($\delta^D = 0.2$)											
α_1	-0.075	-0.079	-0.077	-0.259	-0.253	-0.195	-0.247	-0.241	-0.197	-0.251	0.115	0.107
	(0.005)	(0.003)	(0.002)	(0.048)	(0.033)	(0.024)	(0.037)	(0.025)	(0.018)	(0.541)	(0.391)	(0.293)
α_2				0.577	0.552	0.504	0.393	0.378	0.385	0.394	0.322	0.337
				(0.053)	(0.036)	(0.026)	(0.017)	(0.012)	(0.008)	(0.088)	(0.063)	(0.047)
α_3										0.002	-0.136	-0.116
										(0.207)	(0.149)	(0.111)
	D. M. error ($w_1 = w - 5$, $\delta^D = 0.2$)											
α_1	-0.075	-0.079	-0.077	-0.207	-0.217	-0.199	-0.194	-0.206	-0.201	-0.145	-0.213	-0.147
	(0.005)	(0.003)	(0.002)	(0.011)	(0.008)	(0.006)	(0.004)	(0.003)	(0.002)	(0.06)	(0.043)	(0.033)
α_2				0.065	0.071	0.052	-0.069	-0.068	-0.071	-0.077	-0.067	-0.08
				(0.013)	(0.009)	(0.006)	(0.002)	(0.001)	(0.001)	(0.01)	(0.007)	(0.005)
α_3										-0.019	0.003	-0.021
										(0.023)	(0.017)	(0.013)
	D. M. error ($w_0 = w - 5$, $\delta^D = 0.2$)											
α_1	-0.203	-0.208	-0.207	-0.216	-0.21	-0.199	-0.204	-0.2	-0.202	-0.28	-0.221	-0.288
	(0.005)	(0.004)	(0.003)	(0.011)	(0.008)	(0.005)	(0.003)	(0.002)	(0.001)	(0.041)	(0.029)	(0.022)
α_2				0.1	0.092	0.083	0.089	0.089	0.09	0.101	0.093	0.103
				(0.012)	(0.008)	(0.006)	(0.001)	(0.001)	(0.001)	(0.007)	(0.005)	(0.004)
α_3										0.029	0.008	0.033
										(0.016)	(0.011)	(0.008)
Observations	2500	5000	10000	2500	5000	10000	2500	5000	10000	2500	5000	10000

Notes: Mean estimates. Priors: $\alpha_1 = -0.20$, $\alpha_2 = -0.04$, $w = \mathcal{N}(10, 2)$, $\delta = -0.045$, $\delta_w = 1$.