WIDER Working Paper 2020/138

Updating great expectations

The effect of peer salary information on own-earnings forecasts

Sam Jones and Ricardo Santos*

October 2020

Abstract: How jobseekers set their earnings expectations is central to job search models. To study this process, we track the evolution of own-earnings forecasts over 18 months for a representative panel of university-leavers in Mozambique and estimate the impact of a wage information intervention. We sent participants differentiated messages about the average earnings of their peers, obtained from prior survey rounds. Demonstrating the stickiness of (initially optimistic) beliefs, we find an elasticity of own-wage expectations to this news of around 7 per cent in the short term and 16 per cent over the long term, which compares to a 22 per cent elasticity in response to unanticipated actual wage offers. We further find evidence of heterogeneous updating heuristics, where factors such as the initial level of optimism, cognitive skills, perceived reliability of the information, and valence of the news shape how wage expectations are updated. We recommend institutionalizing public information about earnings.

Key words: earnings expectations, information intervention, updating heuristics, Mozambique

JEL classification: D83, D84, E24

---

* UNU-WIDER; corresponding author: jones@wider.unu.edu

This study has been prepared within the project on Inclusive growth in Mozambique—scaling-up research and capacity, implemented in collaboration between UNU-WIDER, University of Copenhagen, University Eduardo Mondlane, and the Mozambican Ministry of Economics and Finance. The project is financed through specific programme contributions by the governments of Denmark, Finland, and Norway.

Information and requests: publications@wider.unu.edu

ISSN 1798-7237 ISBN 978-92-9256-895-5

https://doi.org/10.35188/UNU-WIDER/2020/895-5

Typescript prepared by Gary Smith.

United Nations University World Institute for Development Economics Research provides economic analysis and policy advice with the aim of promoting sustainable and equitable development. The Institute began operations in 1985 in Helsinki, Finland, as the first research and training centre of the United Nations University. Today it is a unique blend of think tank, research institute, and UN agency—providing a range of services from policy advice to governments as well as freely available original research.

The Institute is funded through income from an endowment fund with additional contributions to its work programme from Finland, Sweden, and the United Kingdom as well as earmarked contributions for specific projects from a variety of donors.

Katajanokanlaituri 6 B, 00160 Helsinki, Finland

1 Introduction

Theoretical approaches to job search have often assumed workers are well informed and cognitively unbiased. Early models (e.g., McCall 1970; Mortensen 1970) assumed individuals' wage expectations and, consequently, their reservation wages were fixed over time, having incorporated all labour market information. Later matching models, such as that of Mortensen and Pissarides (1994), have continued to rely on this formulation. Thus, a predominant theoretical assumption is that workers know the true distribution of wages.

Notwithstanding the theoretical merits of this simplification, empirical evidence suggests job search is rife with uncertainty. Kiefer and Neumann (1979) already found that jobseekers' reservation wages tend to decrease along the search period, suggesting they revise their wage expectations; and this observation has been repeated elsewhere (see Burdett and Vishwanath 1988; Cooper and Kuhn 2020). For instance, Krueger and Mueller (2016) find moderate declines in reservation wages (of 2–7 per cent) over the course of a year in New Jersey; Conlon et al. (2018) show that expected earnings increase by 47 cents for every US$1 shock in unanticipated salary offers; and Zafar (2011) find that salary expectations among undergraduates in the USA tend to become more accurate over time, at least in their pursued majors. The prevalence of labour market information frictions has also been underlined by studies showing how jobseeking strategies respond to new information. Altmann et al. (2018) show that targeted information about job search reduces the risk of unemployment; and Belot et al. (2018) show that providing seekers with information about a range of relevant alternative occupations widens the range of jobs they consider and increases their interview rates.

In light of the above, a relevant question is not just *whether* wage expectations might adjust over time in response to new information, but rather what form this process takes. Indeed, to improve the matching process it is of fundamental importance to better understand how jobseekers construct their wage expectations. Here, there is less consensus; but at least two generic findings merit comment. First, individuals tend to process the same information in different ways, which cannot only be attributed to initial differences in uncertainty. With respect to salaries, Wiswall and Zafar (2015) run a laboratory experiment providing US university students with information about the distribution of earnings for specific groups in the population. They find that own-beliefs about future earnings (at age 30) respond significantly to new information, but there is substantial heterogeneity in the students' updating heuristics, including a large share of so-called non-updaters. This aligns with other studies that find the process of updating beliefs is often asymmetric with respect to the direction of the information signal received, and reveals basic cognitive biases (Bénabou and Tirole 2016; Cooper and Kuhn 2020; De Paola et al. 2001; Mobius et al. 2011). Consequently, and as Bonilla-Mejía et al. (2019) demonstrate, there is no guarantee that poorly informed individuals will consistently update their priors in response to new information.

Second, what kind of information is provided matters. Information perceived as more personally relevant or less ambiguous tends to be attributed greater weight in the revision of beliefs. DellaVigna (2009) surveys literature covering how the format (simplicity) of information can affect its salience. Abbiati et al. (2018) propose a model of divergent responses to information about the returns to post-secondary vocational education associated with the perceived status of relevant career paths. In the domain of education, Herber (2018) shows that providing disadvantaged students customized information about access to scholarships increases applications, while more general information did not. And with respect to the labour market, Wiswall and Zafar (2015) found that providing major-specific information about earnings was associated with larger corrections of initial estimation errors compared to a non-specific alternative treatment. Thus, while information interventions can be relevant (and low cost), careful design is critical to enhance impact (also Bleemer and Zafar 2018).

The present study investigates how own-earnings expectations respond to information about peer wage outcomes during transitions from education into employment in a low-income country. To do so, we combine elements of the information experiment of Wiswall and Zafar (2015) with longitudinal data similar to that of Conlon et al. (2018). Concretely, we track a large sample of Mozambican university graduates as they enter the labour market, following the same individuals from the survey baseline in 2017, timed to coincide with their final year of studies, over six quarterly rounds to the end of 2019. Inspired by Kriechel and Pfann (2006), in rounds two through six of the follow-up surveys we sent the participants information by SMS regarding the labour market outcomes of their peers, namely rate of employment and wage, as observed in the immediately preceding survey round. We randomly allocated students to treatment arms that aggregated the prior survey information in three different ways: (1) across all surveyed students; (2) among peers from the same university; and (3) among peers from the same field of study (major).

Our design contributes to the literature by investigating the impact on own-wage expectations of information about peer earnings outcomes, identified from experimentally manipulated differences in participants' information sets received during early employment transitions. To our knowledge, in itself this represents the first investigation of this kind in a developing country. Our approach also has three further distinctions. First, we consider the evolution of beliefs over a relatively long period of time (18 months). Second, we apply different information treatments, which vary in specificity, quality, and potential relevance. This allows us to study the magnitude of adjustments to wage beliefs in response to different types of information, as well as any corresponding variation in updating heuristics, such as between jobseekers versus those already in work. Third, we focus on a representative sample of graduates in a low-income setting, where the formal economy is small, public earnings information is not generally available and social norms generally preclude open discussion of wages. As such, individuals may be inexperienced in interpreting statistical information, potentially weakening the value of information interventions of this sort.

Section 2 begins by setting out our general analytical framework. Here, we extend a basic 'update towards signal' model of beliefs updating (Chambers and Healy 2012), well-suited to a laboratory information experiment, to take into account both the presence of new unobserved private information, such as that gained from ongoing job search, as well as potential differences regarding either how the experimentally provided wage information is processed by participants or the weights applied to news shocks. Our most complete specification thus provides a basis for testing between these two alternative forms of response heterogeneity.

Section 3 describes the details of the experiment and survey design, and Section 4 briefly summarizes the data. In particular, we highlight the high *ex ante* wage expectations among participants, discussed further by Jones et al. (2020; see also Gradín et al. 2020), as well as preliminary evidence of somewhat greater reductions in these expectations over the course of the survey among members of the treatment groups compared to the control group. Section 5 presents our econometric analysis. Looking on both a static and dynamic basis, we find significant responses to the peer wage information, amounting to an elasticity of own-wage expectations to these news shocks of around 7 per cent in the short term and 16 per cent over the long run. While this response magnitude may appear moderate (conservative), it nonetheless compares to estimates of a 22 per cent elasticity in response to unanticipated actual wage offers. In other words, high own-wage expectations are persistent.

Further analysis of heterogeneity reveals a complex pattern, including differences in both information processing and weighting of the information treatments. We also find evidence of asymmetric responses to news. Responses to positive unanticipated actual wage offers appear to be incorporated into posterior expectations on almost a one-to-one basis. And while the short-run elasticity peer wage news does not seem to vary with its valence on average, there is evidence of very substantial asymmetric responses to positive general and field-specific messages. We conclude that provision of (more) information about wage outcomes is likely to be relevant for graduate jobseekers. Even so, the type of information needs to considered carefully (e.g. university-specific information alone may not be helpful) and, as found elsewhere (e.g. Sharot et al. 2011), exposure to more or better information alone is unlikely to fully eliminate 'unrealistic optimism' in expected wages.

2 Analytical framework

As noted in the Introduction, this paper focuses on how jobseekers revise their own-wage expectations in response to new information about the earnings of their peers. Before describing the experiment in more detail, we now set out a basic framework to analyse this process. We start with a canonical set-up in which the expected value of a future outcome is conditioned on the information set available to the individual. Here, we define $\tilde{w}_i$ as the nominal earnings of individual i at a specified point in the future (say, T). Her corresponding self-reported belief over this outcome, elicited at time $t < T$, is thus: $\mathrm{E}(\tilde{w}_i \mid \Omega_{it})$, which we denote hereafter as w_{it}, and where Ω_{it} represents the relevant information set available to the individual at time t. Following Zafar (2011), on the assumption that information available to the individual in one period remains accessible in subsequent periods (it does not depreciate), it follows that:

$$w_{it+1} = \mathrm{E}(\tilde{w}_i \mid \Omega_{it+1}) = \mathrm{E}(\tilde{w}_i \mid \Omega_{it}, \varphi_{it+1})$$

where φ_{it+1} represents new information received in period $t+1$. Since the expected value of the same belief in the subsequent period must be unchanged if no new information arrives:

$$\begin{aligned} \mathrm{E}(w_{it+1} \mid \Omega_{it}) &= \mathrm{E}(\ \mathrm{E}(\tilde{w}_i \mid \Omega_{it+1}) \mid \Omega_{it}) \\ &= \mathrm{E}(\tilde{w}_i \mid \Omega_{it}) \equiv w_{it} \end{aligned}$$

it follows the sum total of novel information (news) obtained in period $t+1$ must constitute the change in belief:

$$\begin{aligned} w_{it+1} - w_{it} &= \mathrm{E}(\tilde{w}_i \mid \Omega_{it}, \varphi_{it+1}) - \mathrm{E}(\tilde{w}_i \mid \Omega_{it}) \\ &\equiv \varepsilon_{it+1} \end{aligned} \tag{1}$$

Of course, Equation (1) is just a definition. To add analytical content, suppose the participant is exposed to some (noisy) information about peer earnings, denoted x. As such, for individual i, define the news content of the signal received in $t+1$ as: $x_{t+1} - w_{it}$. This allows us to split the total news term into two distinct components:

$$\varepsilon_{it+1} = \beta(x_{t+1} - w_{it}) + \nu_{it+1} \tag{2}$$

where β captures the perceived relevance or salience of the peer earnings signal and any remaining news is captured by ν. Inserting Equation (2) into Equation (1) implies:

$$\begin{aligned} w_{it+1} &= w_{it} + \beta(x_{t+1} - w_{it}) + \nu_{it+1} \\ &= (1-\beta)w_{it} + \beta x_{t+1} + \nu_{it+1} \end{aligned} \tag{3}$$

Under the reasonable assumption $\beta > 0$, the second expression conforms to a conventional 'update towards the signal' (UTS) model of how beliefs respond to new information (Chambers and Healy 2012; Kandel and Zilberfarb 1999)—that is, the posterior belief is a weighted average of the prior belief and the signal.[1]

[1] Classical Bayesian updating is a special case of this model, where the weights reflect the inverse of their variances (see, e.g. Viscusi and O'Connor 1984).

The above model constitutes our basic point of departure. However, it is limited in two directions. First, in many controlled (experimental) settings, only a short time tends to elapse between the elicitation of prior beliefs, provision of new information, and elicitation of posterior beliefs. Outside the laboratory, Haaland et al. (2020) note that eliciting updated beliefs through follow-up surveys is relatively recent, pioneered by studies such as those of Kuziemko et al. (2015) and Cavallo et al. (2017). But even in these cases, most follow-up surveys usually take place 1–8 weeks after the initial information provision.[2] With a short time elapsing between elicitation of prior beliefs, new information provision, and follow-up survey, it is reasonable to assume that ν can be treated as white noise, allowing the researcher to directly focus on changes in beliefs purely as a function of the experimentally manipulated news (e.g., Haaland et al. 2020; Wiswall and Zafar 2015). In our setting (see below) this assumption is problematic—a priori there is little reason to expect ν will be centred on zero; it is also reasonable to assume individuals obtain additional information concerning wages between rounds, such as through job search, some of which will be private (e.g. associated with their search effort). Denoting this privately sourced wage information as z_i, a more complete model for updating of beliefs is:

$$w_{it+1} = w_{it} + \beta(x_{it+1} - w_{it}) + \delta(z_{it+1} - w_{it}) + \varepsilon_{it+1} \tag{4a}$$

$$= (1-\delta)w_{it} + \beta(x_{it+1} - w_{it}) + \delta z_{it+1} + \varepsilon_{it+1} \tag{4b}$$

We now subscript the peer earnings information by i to account for differences in content across individuals. Also, prefiguring our information experiment, we impose $x_{it+1} = w_{it}$ for individuals who do not receive any information about peer earnings, implying they receive no news signal of this type. This obviates the need to include treatment dummy interaction terms.

A second limitation is that we have presumed a direct and homogeneous updating process. However, as Wiswall and Zafar (2015) elaborate, a wide range of heuristics actually might be employed, reflecting different ways in which information signals are processed. Here, individuals may map information about peer earnings to their own-wage expectations in different ways. To account for this parsimoniously, we can generalize the experimentally manipulated news term as: $f_i(x_{it+1}) - w_{it}$, where $f(\cdot)$ allows for differences in the mapping or salience of the message (see DellaVigna 2009), with our earlier assumption of an identity function now being a special case. By way of a simple example, consider: $f_i(x_{it+1}) = (1+\theta_i)x_{it+1} = (s_i/\bar{s})x_{it+1}$, where s_i is the ability of the individual recipient and $\bar{s}$ is the average ability in the same population, both of which we take as known constants.[3] So, for some $s_i < \bar{s}$ the processed news may be in the negative domain, even while the observed news is positive—that is, $x_{it+1} > w_{it} \not\Longrightarrow f_i(x_{it+1}) > w_{it}$. As such, heterogeneous and even apparently counter-intuitive responses to new information may reflect different information-processing rules.

Relating this back to our updating equation, and retaining the simple assumption of an unknown scalar processing heuristic, note:

$$(1+\theta_i)x_{it+1} - w_{it} = (x_{it+1} - w_{it}) + \theta_i x_{it+1} \tag{5}$$

Employing this expression in place of the observed news term within Equation (4a) thus yields the more general framework:

$$\Delta w_{it+1} = \beta\theta_i x_{it+1} + \beta(x_{it+1} - w_{it}) + \delta(z_{it+1} - w_{it}) + \varepsilon_{it+1} \tag{6}$$

where the presence of information-processing heterogeneity is indicated directly by the coefficient: $\beta\theta_i \neq 0$. This form of heterogeneity is distinctive to that in which the perceived relevance or precision

[2] Fehr et al. (2019) is an outlier, implementing their follow-up survey one year after initial information provision.

[3] This construct reflects the notion that wages are likely to reflect both the quantity and price of skills; it has the attractive property that $\mathrm{E}(\theta_i) = 0$. A richer specification would allow $\bar{s}$ to vary with the sample used to construct x_{it+1}, which would align more closely with the relative skill updating heuristic suggested by Wiswall and Zafar (2015). However, more complex formulations cannot be estimated within a linear econometric framework.

of the processed news shock varies, which would be reflected by differences in the weighting factor, β. To distinguish between these forms of heterogeneity we thus employ the following specification:

$$\Delta w_{it+1} = \underbrace{(\beta+\beta_i)\theta_i}_{\varphi_i} x_{it+1} + (\beta+\beta_i)(x_{it+1} - w_{it}) + \delta(z_{it+1} - w_{it}) + \varepsilon_{iit+1} \tag{7}$$

noting that the presence of information-processing heterogeneity on its own would be apparent if all β_i were zero but at least one φ_i was non-zero: $(\exists i : \varphi_i \neq 0) \wedge (\forall i : \beta_i = 0)$; while the unique presence of weighting-heterogeneity weights would be indicated by the opposite case: $(\forall i : \varphi_i = 0) \wedge (\exists i : \beta \neq 0)$. Thus, separate Wald tests of whether all β_i are jointly equal to zero, or whether all φ_i are jointly equal to zero, should be a sufficient basis to determine which form of heterogeneity is present, if any.

Lastly, econometric estimation of our proposed framework raises some challenges. Equation (4a) highlights that the peer and private news terms contain a common component (w_{it}) and therefore should be correlated. Thus, for $\delta \neq 0, |z_{it+1} - w_{it}| > 0$, excluding the private news term from the model would bias the regression estimates for β, even if the peer information term is strictly exogenous. Furthermore, as long as $\mathrm{E}(z_{it+1} \cdot w_{it}) \neq 0$, which generally cannot be ruled out, this concern remains pertinent in the transformed version given by Equation (4b). For instance, z_{it} may be a function of fixed characteristics, the market value of which are gradually discovered over time (or are prone to change), and also will be reflected in w_{it} as per the proposed law of motion of beliefs (Equation 4a).

These considerations point to the inherent difficulties of estimating learning models outside of strictly controlled conditions. Nonetheless, while we cannot ignore the private information term, we can proceed on the assumption that it can be approximated as: $z_{it} = f(u_i, v_{it}, t)$, where u_i are a set of observed time-invariant individual characteristics and v_{it} are time-varying individual outcomes (e.g. work experience). Linearizing the unknown functional form for z_{it} and inserting in Equation (4b) implies for the homogeneous updating case:

$$w_{it+1} = (1-\delta)w_{it} + \beta(x_{it+1} - w_{it}) + \underbrace{\gamma_{0t+1} + u_i'\gamma_1 + v_{it+1}'\gamma_2}_{=\mathrm{E}(z_{t+1}|u_i, v_{it+1}, t+1)} + \varepsilon_{it+1} \tag{8}$$

An alternative, somewhat more parsimonious approach is to generate an estimate z_{it} in a preliminary stage and employ these predicted values in a regression on the form of Equation (4a):

$$\Delta w_{it+1} = \beta(x_{it+1} - w_{it}) + \delta(\hat{z}_{it+1} - w_{it}) + \mu + \lambda_{t+1} + \xi_{it+1} \tag{9}$$

where we have expanded the error term, $\varepsilon_{it+1} = \mu + \lambda_{t+1} + \xi_{it+1}$, to account for both drift and common round-specific shocks or other systematic public news (e.g. wage inflation). In addition, it would be convenient to allow the drift term to vary across individuals. However, since the information treatment is plausibly exogenous (by randomization) and the lagged dependent variable counts among the regressors, which is expected to induce Nickell bias when combined with individual fixed effects, we do not pursue such an extension. Even so, in a preliminary analysis (Section 5.1), we investigate the conditional average effect associated with exposure to the information treatments (coded as dummy variables). This static specification permits inclusion of individual fixed effects.

3 Experimental design

The context for our information experiment is a longitudinal school-to-work survey of Mozambican university students, comprising students in their final year of undergraduate studies in 2017, who were then tracked on a quarterly basis by telephone over 18 months (i.e. over six rounds) to the end of 2019.[4] The

[4] For further details on the design and results of the survey, see Jones et al. (2020).

survey was conducted at six of the largest universities, which, jointly, teach approximately three-quarters of the entire population of Mozambican university students. It encompassed a total of 27 faculties, surveying students from 106 different courses.[5] In total, 2,174 finalists were surveyed at the baseline (1,024 women and 1,150 men), which we later classified into seven distinct study fields: Education; Languages and Humanities; Social Sciences, Management and Law; Natural Sciences; Engineering, Industry, and Construction; Agriculture; and Health and Welfare.[6] Of those surveyed, 2,100 agreed to remain in the follow-up sample and take part in the survey experiment. These constitute our baseline sample.[7]

Following Section 2, we investigate how beliefs about future wages respond to current information about realized wages, varying the granularity of information provided. To do so, in each round of the survey we elicited beliefs about future wages. Concretely, in the baseline, as well as the first two follow-up telephone rounds, we asked about the wage they expected to receive at the end of their first year of work; and, given the impending termination of the survey, in the remaining rounds (3–5) we elicited the wage they expected to receive in June 2019, in all cases assuming they were working.[8] In each telephone survey round we also elicited detailed information about their work status, including their specific position, type of contract, and current wage.

Given the scarcity of reliable data on graduate wages, both in general as well as specifically for new labour market entrants, we used information on realized wages from prior rounds of the telephone survey to design three distinct information treatments. Namely:

1. *General message:* summarizes wage information from the entire sample—for example, 'Survey results at Dec.1st: of all graduates in Mozambique (class of 2017), 59% are working and their average wage = 14,000 Mts / mes.'

2. *University-specific message:* summarizes wage information from the sub-sample of participants who attended the same university as the recipient—for example, 'Survey results at Dec.1st: of all graduates from your university (class of 2017), 52% are working and their average wage = 24,000 Mts / month.'

3. *Field-specific message:* summarizes wage information from the sub-sample of participants in the same study field as the participant—for example, 'Survey results at Dec.1st: of all graduates from your area of studies (class of 2017), 50% are working and their average wage = 13,500 Mts / month.'

The messages were sent by SMS at the beginning of each telephone survey round, excluding the first. Note that the specific information contained in each SMS varied by survey round; in the second and

---

[5] The universities were: Universidade Eduardo Mondlane (UEM), Universidade Pedagógica (UP), Universidade Politécnica (AP), and Universidade São Tomás de Moçambique (USTM) in their Maputo campuses; Universidade Católica de Moçambique (UCM) and Universidade Zambeze (UZ) in their Beira campuses.

[6] The survey baseline sample size, for a theoretical proportion p of 50 per cent respondents finding a job, with a confidence interval $(1-\alpha)$ of 95 per cent and as per Cochran (1977), allows for error margins of 1.7 per cent in the total sample, of up to 2.4 per cent in each gender sub-sample, and up to 6.3 per cent in the least represented study field.

[7] All participants were aged 18 or above and gave written informed consent to participate in the baseline survey and follow-up telephone rounds; they were also free to desist at any time. The data collected was anonymized. In the absence of a full institutional review board at both UNU-WIDER and local institutions at the time of the surveys, approval for undertaking the research was received from each participating university prior to fieldwork and upon discussion of the research and survey design and procedures.

[8] The question in the first set of rounds was: *After working for one year, how much do you expect to be earning per month (after tax)?*'. In the following rounds, it was: '*Looking ahead to June 2019, how much do you expect to be earning per month (after tax) at that time?*' (own translations from Portuguese). Use of month/round fixed effects in our later econometric analysis controls for these differences in time horizons.

third types of message, the information varied by individual according to the specific university they had attended or their field of study. Mirroring variation in actual wages (see below), this design provides for substantial variation in the underlying wage information received. Figure 1 indicates the empirical distributions of the peer salary information sent by SMS, indicating variation both within and between the different message types. Notably, the field-specific message shows the largest variation (twice that of the general message), consistent with segmentation in earnings across occupations.

Figure 1: Distribution of peer salary information, by type of treatment message

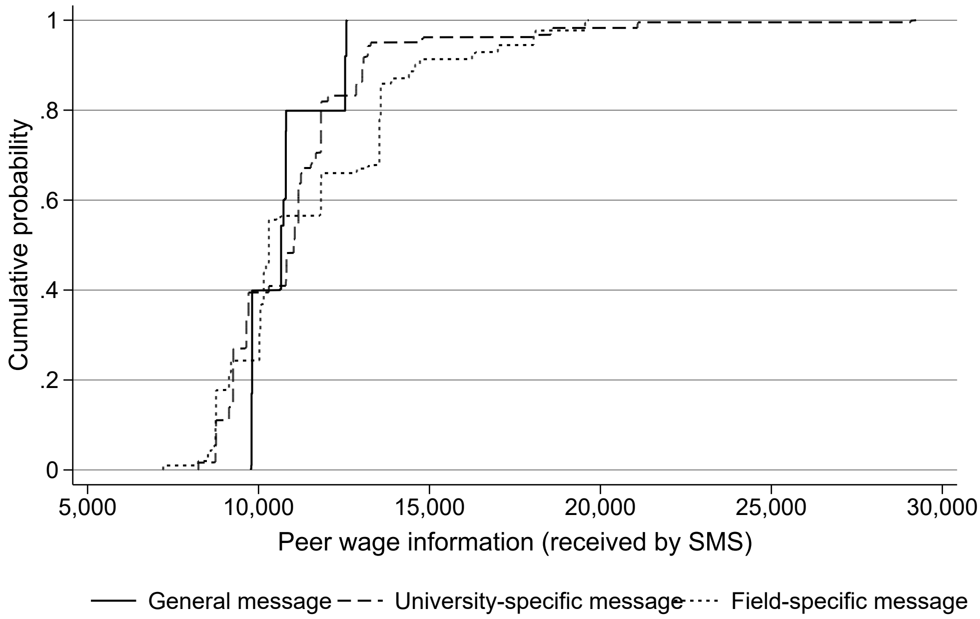

Note: the figure plots the (cumulative) empirical distributions of the different peer earnings information (mean salaries in the prior round) sent by SMS.

Source: authors' construction.

In terms of exposure to the treatments, directly after the baseline survey individuals were randomly allocated to one of five treatment arms, distinguished by the type of messages they would receive. Specifically:

- The *control* group: received no information messages.
- Group 1 (*general*): received the general (all-student) message in all relevant rounds.
- Group 2 (*university*): received the university-specific message in all relevant rounds.
- Group 3 (*field*): received the field-specific message in all relevant rounds.
- Group 4 (*mixed*): received the general message in round two, the university-specific message in rounds three and four, and the field-specific message in rounds five and six.

Randomization across the five experimental arms was stratified by study field and gender, with a target of around 500 participants in the control group and 400 in each treatment arm. The first row of Table 1 presents the initial allocation (undertaken after the baseline was concluded), while the remaining rows give a sense of attrition over the follow-up rounds. As it shows, just a small share of participants could not be contacted in the follow-up rounds, amounting to a loss of 12 per cent of the baseline sample after 18 months (six rounds). We note some small differences in attrition across the arms, a closer examination of which is suggestive of very moderately (yet statistically significant) lower attrition among participants receiving the information treatment (see Appendix A1). To address any bias this might create, we adjusted the sample weights to ensure treatment and control arms represent a fixed share of observations within each strata in each round (e.g. Chen et al. 2015), based on the proportion of participants observed in at least one of the follow-up rounds.[9] Overall the realized experiment is well-powered—using the sub-samples obtained in the last round, a *t*-test of differences in means between treatment and control groups (jointly) would be expected to detect a proportional difference (e.g. in log wages) of around 5.25 per cent, with a power of 80 per cent at the 5 per cent significance level.

Table 1: Observations across survey rounds, by experimental arm

	Experimental arm						
Round	Control	General	University	Field	Mixed	Total	% base
Baseline	504	397	400	398	401	2,100	1.00
1	486	393	390	390	392	2,051	0.98
2	473	389	379	377	386	2,004	0.95
3	456	381	376	375	383	1,971	0.94
4	439	377	367	366	372	1,921	0.91
5	428	372	361	359	360	1,880	0.90
6	423	366	353	350	357	1,849	0.88

Note: cells report the raw number of observations by experimental arm and round number; the final column gives the overall follow-up rate relative to the baseline sample.

Source: authors' construction based on survey results.

4 Descriptive statistics

Table 2 presents the sample's baseline characteristics, showing means across the five experimental arms and, in the last column, the results of a joint F-test to check whether the means in the treatment arms differ from the control arm. From the first panel, we note that gender parity is relatively high, participants were around 25 years old, over 42 per cent considered themselves to have an 'excellent' academic level compared to their peers, around 10 per cent considered themselves to be proficient in English, and around 60 per cent had some prior work experience. From the remaining panels, which refer to university and study fields respectively, we note that the sample is dominated by graduates in the fields of social sciences and education. Nonetheless, there is no evidence of significant differences in baseline characteristics across the experimental arms.

It is useful to verify whether participants' expected monthly salaries (one year after starting work) were statistically equivalent over the experimental arms at baseline.[10] As can be seen from the first row of Table 3, which follows the structure of the previous table, as well as the density distributions in Figure 2(a), expected wages were not significantly different across the arms in 2017, confirmed using a Kolmogorov–Smirnov test for the equality of the distribution functions for the treated (considered jointly) and control groups (prob. = 0.15). The same figure also indicates the wide range of baseline expected salaries—for example, the interquartile range runs from 21,000 to 38,750 MZN (equal to around US\$300 and US\$550, respectively). Those familiar with the Mozambican labour market would recognize these represent very 'sanguine' expectations, even for university graduates. For instance, estimates from the most recent household survey (2014/15) indicate that the average reported salary of men with tertiary education was around 30,000 MZN, while the average salary for women with similar qualifications was 26,500 MZN.

[9] These adjusted weights are used henceforth (unless indicated otherwise). Quantitatively, the adjustments are small (always less than ±10 per cent or ±3 per cent on average) and none of our results are significantly affected by this procedure (full details available on request). In the robustness analysis (Section 5.3), we run our preferred model for a balanced sample of individuals observed in all rounds.

[10] Hereafter, all salary values, both expected and realized, are stated and analysed in November 2019 prices, based on the national consumer price index.

These wages incorporate returns to average experience and, thereby, are expected to be substantially larger than wages at entry.

Table 2: Baseline average characteristics across experimental arms

	Experimental arm					
	Control	General	University	Field	Mixed	F-test
Female	0.44	0.44	0.45	0.45	0.44	0.983
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	.
Age	26.11	25.63	25.84	25.83	26.18	0.724
	(0.12)	(0.13)	(0.13)	(0.12)	(0.13)	.
Proficient in English	0.09	0.09	0.09	0.12	0.12	0.680
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	.
Excellent academics	0.41	0.47	0.48	0.41	0.44	0.219
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	.
Prior work experience	0.62	0.61	0.61	0.58	0.60	0.858
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	.
UEM	0.36	0.36	0.35	0.34	0.34	0.991
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	.
UCM	0.08	0.10	0.09	0.09	0.09	0.950
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	.
UZ	0.11	0.11	0.10	0.09	0.11	0.965
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	.
USTM	0.05	0.05	0.07	0.06	0.03	0.066
	(0.00)	(0.00)	(0.01)	(0.01)	(0.00)	.
UP	0.34	0.31	0.36	0.38	0.38	0.491
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	.
AP	0.06	0.07	0.03	0.04	0.04	0.381
	(0.00)	(0.01)	(0.00)	(0.00)	(0.00)	.
Education	0.31	0.30	0.32	0.32	0.30	0.993
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	.
Humanities	0.02	0.02	0.02	0.01	0.02	0.999
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	.
Social sciences	0.44	0.45	0.45	0.44	0.45	0.996
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	.
Natural sciences	0.04	0.04	0.04	0.04	0.04	0.994
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	.
Engineering	0.08	0.08	0.08	0.08	0.08	0.999
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	.
Agriculture	0.05	0.05	0.05	0.05	0.06	0.967
	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)	.
Health	0.06	0.06	0.06	0.06	0.06	0.997
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	.

Note: cells show means and mean standard errors (in parentheses) across different experimental arms of baseline variables, based on one observation for each individual observed at least once in the follow-up rounds, $N = 2,069$; 'F-test' reports the probability that the means in the treatment arms jointly differ from the control arm. UEM–AP are universities (dummy variables); study fields are given in the final panel (dummy variables).

Source: authors' construction based on survey results.

To get a further sense of the (*ex post*) elevated nature of the participants' baseline expectations, we consider the highest wage they reported to have received during the follow-up period. Figure 2(c) illustrates the distribution of these realized wages ($N = 1,528$), and Figure 2(d) shows its value relative to the baseline expected wage (in their first job). Notably, the 75th percentile of the highest realized wage (20,250 MZN) is slightly below the 25th percentile of the baseline expectations (adjusting for inflation); and for the median participant reporting a wage, their highest realized wage is about half of what they had expected in real terms.

Table 3: Evolution of wage expectations across experimental arms ('000s of MZN)

	Experimental arm						
	Control	General	University	Field	Mixed	F-test	Obs.
Baseline	34.30	32.93	32.71	31.13	32.35	0.407	2,065
	(0.86)	(0.90)	(0.81)	(0.78)	(0.83)	.	.
Round 1	31.25	29.41	30.74	27.92	31.85	0.038	1,967
	(0.83)	(0.76)	(0.80)	(0.74)	(0.92)	.	.
Round 2	30.06	29.35	29.30	26.83	30.59	0.165	1,925
	(0.77)	(0.77)	(0.81)	(0.70)	(0.87)	.	.
Round 3	29.19	25.75	27.28	24.56	27.31	0.001	1,884
	(0.72)	(0.70)	(0.74)	(0.66)	(0.78)	.	.
Round 4	27.53	24.81	25.49	23.04	25.39	0.002	1,807
	(0.70)	(0.66)	(0.70)	(0.68)	(0.72)	.	.
Round 5	25.85	23.12	23.25	21.31	22.72	0.001	1,695
	(0.69)	(0.67)	(0.68)	(0.68)	(0.66)	.	.

Note: following the structure of Table 2, cells show mean own-earnings expectations by survey round, in thousands of MZN (at constant prices); mean standard errors in parentheses; F-test reports the probability that the means in the treatment arms jointly differ from the control arm; information treatment provided from start of round 2 onwards; number of observations differ from Table 1 due to missing wage expectations.

Source: authors' construction based on survey results.

In light of the above, one might reasonably expect participants to revise their salary expectations downward as they come into contact with the labour market. This is supported by Table 3, which reveals the average expected wage follows a declining trend across all experimental arms during the telephone surveys, which in the control group is equal to about 5 per cent per round (quarter). However, comparing across the treatment and control groups, the same table indicates moderately larger absolute declines in the former relative to the latter, at least in the later rounds (3–5), which are statistically significant at the 1 per cent level.[11] Comparing endline expectations, defined as the last observation per participant, Figure 2(b) confirms treatment and control group distributions have diverged, with the mass of the former now compressed and shifted to the left of the latter; and a Kolmogorov–Smirnov test rejects that the treatment and control endline distributions are equal (prob. = 0.00). This provides initial evidence of a material effect from the information treatments on expectations formation. Nonetheless, the presence of some pre-treatment differences in mean expectations across the groups, such as observed in round 1 (see Table 3), implies it will be useful to control for any fixed (unobserved) group-wise differences in subsequent analysis.

Last, the suggestion of a relationship between exposure to the information treatments and own-wage beliefs is supported by a preliminary non-parametric analysis, depicted in Figure 3. On the y-axis we plot the residuals from a regression of the change in log wage beliefs on a detailed set of controls (as described in Section 5); and on the x-axis we plot the residuals from a regression of the log peer news shock ($x_{it} - w_{it-1}$) on the same controls. The solid line is the local polynomial fit between the two sets of residuals for the treated participants only. The plot indicates that information on peer wages is associated with belief-updating in the direction of the signal. However, it also suggests the response may be asymmetric, with a stronger correction towards positive signals. We follow up this insight in Section 5.4.

[11] Missing (lagged) salary expectations, primarily due to temporary attrition from one or other follow-up round, are imputed by combining each individual's observed mean updating rate with the mean updating rate of her peers. Controls for imputation are included in subsequent regressions to account for imputation error; we show results are robust to exclusion of these observations in Section 5.3.

Figure 2: Expected and realized wage distributions

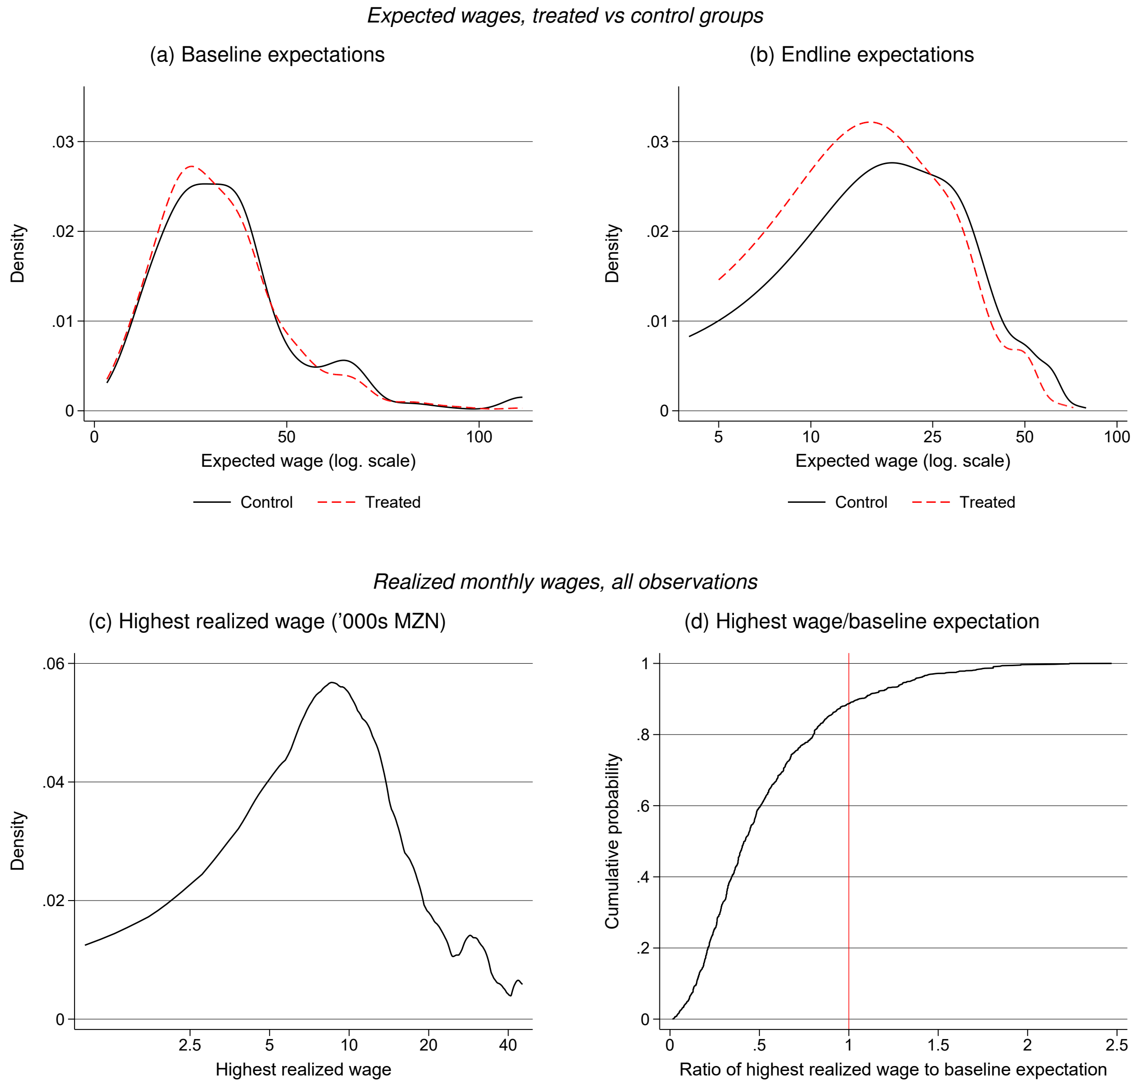

Note: panels (a) and (b) plot empirical distributions of earnings expectations between treatment and control groups at baseline and endline (round 5), respectively; panels (c) and (d) plot empirical distributions of highest realized wages (per individual) and the ratio of the same wage to the baseline expected wage, respectively.

Source: authors' construction, based on survey results.

Figure 3: Non-parametric response to peer wage information

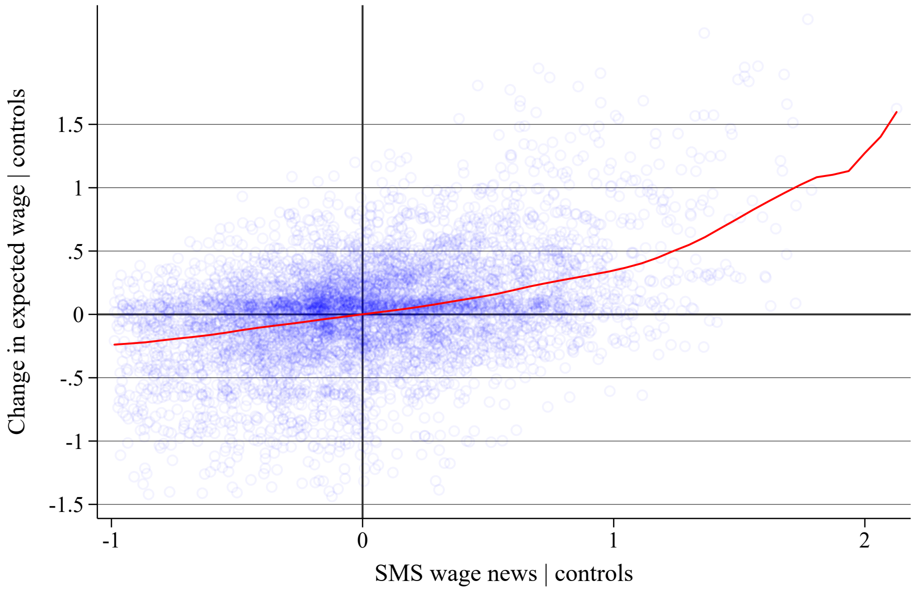

Note: the scatter plot is of the SMS wage news variable vs the change in expected earnings, both residualized against a set of controls. The solid line is the non-parametric regression fit.

Source: authors' construction based on survey results.

5 Results

5.1 Average treatment effects

To begin, we elaborate on the descriptive analysis of Table 3 and verify whether evidence for a conditional average effect associated with the information treatments holds. To do so, we abstract from dynamics and focus on observations from the first and fifth telephone follow-up rounds only (i.e. before and after the treatment, amounting to a conventional difference-in-difference analysis). We regress own-wage beliefs, hereafter expressed in natural logarithms, on exposure-to-treatment dummy variables, the coefficients on which thereby capture the average effects of the (different) information treatments on treated participants (ATTs). Table 4 reports the results. Sub-columns (a) capture the ATTs associated with exposure to *any* of the treatments, and sub-columns (b) differentiate between the three types of treatment messages received. Column (I) omits any controls other than month-by-round fixed effects. Columns (II)–(V) progressively add controls, including (in order): time-varying controls; time-invariant controls, taken from the baseline survey; experimental-arm fixed effects; and individual fixed effects.

The time-varying controls include (*inter alia*) information on the current working status of the individual, the elapsed time (in months) between being contacted in each round, and a dummy variable capturing whether the elicited wage expectation pertains to full-time work. We also include the 'SMS employment rate', which refers to the information received in the treatment messages regarding the share of peers in each specified group who are currently working (see Section 3). This is used to control for possible differences in the perceived credibility of the primary salary message.[12] Also, in the telephone surveys we asked whether individuals had heard any information from friends or colleagues about the earlier rounds of the survey; we assign affirmative answers to the dummy variable 'spillover'. The baseline control variables are extensive and include: the participants' age, gender, performance on simplified psychometric tests, baseline expected wage, and their field of study and university attended. Due to space limitations, we do not discuss these controls further.

Table 4: Analysis of average information treatment effects

	(Ia)	(Ib)	(IIa)	(IIb)	(IIIa)	(IIIb)	(IVa)	(IVb)	(Va)	(Vb)
Treated	-0.14^{***}		-0.11^{***}		-0.11^{***}		-0.13^{***}		-0.15^{***}	
	(0.04)		(0.04)		(0.04)		(0.04)		(0.04)	
Gen. treatment		-0.11^{**}		-0.08^{*}		-0.07		-0.09		-0.10^{*}
		(0.05)		(0.05)		(0.05)		(0.05)		(0.05)
Uni. treatment		-0.10^{**}		-0.10^{**}		-0.10^{**}		-0.15^{***}		-0.17^{***}
		(0.05)		(0.05)		(0.05)		(0.05)		(0.05)
Field treatment		-0.17^{***}		-0.14^{***}		-0.12^{***}		-0.14^{***}		-0.16^{***}
		(0.04)		(0.04)		(0.04)		(0.05)		(0.05)
Working			-0.02	-0.02	-0.03	-0.03	-0.03	-0.03	-0.13^{*}	-0.14^{*}
			(0.07)	(0.07)	(0.07)	(0.07)	(0.07)	(0.07)	(0.08)	(0.08)
Experience			0.05	0.05	0.02	0.02	0.02	0.02	0.09^{**}	0.09^{**}
			(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.04)	(0.04)
Full-time expect.			0.14^{***}	0.14^{***}	0.12^{***}	0.12^{***}	0.12^{***}	0.12^{***}	0.14^{**}	0.14^{**}
			(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.06)	(0.06)
Spillover			0.03	0.03	0.04	0.04	0.04	0.04	0.05	0.06
			(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.04)	(0.04)
SMS employ. rate			-0.01^{***}	-0.01^{***}	-0.00^{*}	-0.00^{*}	-0.00^{*}	-0.00^{*}	-0.00	-0.00
			(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Elapsed time			0.10^{**}	0.10^{**}	0.02	0.02	0.01	0.01	0.02	0.02
			(0.04)	(0.04)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
Obs.	3,591	3,591	3,577	3,577	3,577	3,577	3,577	3,577	3,324	3,324
R^2 (adj.)	0.06	0.06	0.16	0.16	0.25	0.25	0.26	0.26	0.46	0.46
RMSE	0.56	0.56	0.53	0.53	0.50	0.50	0.50	0.50	0.42	0.42
Time-vary. controls	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Baseline controls	No	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Exper.-arm FEs	No	No	No	No	No	No	Yes	Yes	Yes	Yes
Individual FEs	No	No	No	No	No	No	No	No	Yes	Yes

Note: significance: * 10 per cent, ** 5 per cent, *** 1 per cent. The dependent variable is log own-earnings expectations; observations from first and fifth rounds only; all specifications include month $\times$ round fixed effects; column (I) is a naive specification; column (II) adds time-varying controls; column (III) adds observed baseline controls; column (IV) adds experimental-arm fixed effects; column (V) uses individual fixed effects (absorbing all other fixed controls); sub-columns (a) use a single treatment dummy; sub-columns (b) include treatment dummies for the specific message type received; all treatment dummies are zero in the first telephone survey round; standard errors clustered at the individual level shown in parentheses; selected coefficients shown.

Source: authors' construction based on survey results.

Three main findings emerge. First, in keeping with the descriptive statistics, we confirm the information treatments had a material impact on own-earnings beliefs. Regardless of the specification, sub-columns (a) suggest that treated individuals on average revised their earnings expectations downwards by around 13 per cent more than the control group (over the course of the follow-up survey), consistent with the observation that initial expectations were almost always optimistic. Second, while the different treatments all seem to have influenced expectations in a similar direction, the magnitude of their effects vary somewhat. In particular, once baseline controls and treatment-arm fixed effects are added (e.g. column IV), the general message appears to have been least effective (on average) and occasionally cannot be distinguished from a zero effect. In contrast, the two more granular information treatment dummies are larger in magnitude and always different from zero. Third, as expected from the randomized design, the inclusion of (more) controls does not dramatically alter the ATT estimates. Indeed, all point estimates remain within the 90 per cent confidence interval of the unconditional estimates given in column (I). Thus, we can be confident the information treatment did influence own-earnings beliefs.

[12] We have investigated the presence of interaction effects between the SMS employment rate and the SMS salary information, but we found no systematic relationships. We return to this below.

5.2 Dynamic effects

Moving to the law of motion for these beliefs, Table 5 reports our main results. Sub-columns (a) continue to summarize the effect of exposure to any treatment via a single dummy variable.[13] Sub-columns (b) replace the treatment dummy with the overall peer earnings news shock, defined at the individual level as: $x_{it} - w_{it-1}$, and which is zero when there is no treatment exposure.[14] Sub-columns (c) then differentiate the news shock across the three different SMS message types. The specifications in super-column (I) include the same full set of time-invariant and time-varying individual controls, group-wise fixed effects, and period fixed effects applied in column (IV) of Table 4. Thus, columns (Ib)–(Ic) correspond to estimates of Equation (8), where these controls stand in for private beliefs.

Columns (II) and (III) follow Equation (9) and use a pre-generated estimate of the private news term in place of the main set of controls. This private shock is derived in a preliminary step, based on a regression of reported earnings beliefs against the main individual-level controls used in prior specifications (i.e. columns I) and time effects. To avoid any possible contamination from exposure to the treatment, this regression is only run for individuals not exposed to the treatment, but predicted values are then estimated for all. Due to the presence of this generated regressor in these specifications, the main regression estimates are based on a bootstrap procedure, combining both the preliminary first step and the final regression. In these specifications we also directly include a few remaining time-varying controls that were not used to generate the private information term (not shown in the table). These account for other sources of systematic noise in expectations, which conceptually enter the error term, and include the employment rate SMS information, the spillover dummy variable, and whether the elicited expectation pertains to full- or part-time work, among others.

The main findings broadly resonate with those of the previous section. Consistent with our framework, exposure to the information treatment leads to significant yet moderate revisions in earnings beliefs in the direction of the news content (signal) given by the SMS wage information. Column (Ia) reports that, on average, recipients of peer earnings information reduced their own expectations by about –0.07 in the short-run. Based on a simplified version of Equation (8) in which we replace the news signal with the treatment dummy, the implied long-run effect is given by the ratio: $\hat{\beta}/\hat{\delta} = -0.16$, which is highly consistent with the ATT estimates in Table 4.[15] Leveraging the continuous nature of the shock term, column (Ib) suggests the short-run elasticity of own-earnings beliefs to the peer salary news shock is around 7 per cent on average, or 15 per cent over the long run, both of which are statistically different from zero.[16] Looking at the coefficients associated with the different types of shocks in column (Ic), differences between the point estimates continue to be small and the footer of the table indicates we cannot reject the hypothesis that these coefficients are all equal. Nonetheless, we confirm they are jointly different from zero and that the field-specific message is consistently associated with the largest effect (8 per cent in the short run; 17 per cent in the long run).

---

[13] This is useful not only as a comparison to the previous estimates, but also because omission or misspecification of dynamics will generally bias estimates in longitudinal panels (Pirotte 1999; Van den Doel and Kiviet 1994), implying (static) average treatment effect estimates may not precisely capture long-run impacts.

[14] For the first round, we use the baseline expected wage to stand in for the prior belief (w_{it-1}). However, recognizing this specific quantity was elicited under different conditions to that in the follow-up surveys—namely, on a face-to-face basis, while individuals were still in education, and in some cases up to a year previously—we allow the coefficient on the prior belief to vary according to whether it derives from the baseline versus subsequent follow-up surveys. In later robustness analysis we show that results are essentially unchanged if we just exclude the first round from the analysis. The advantage of the more complete specification is that we can identify treatment-group fixed effects separately from the exposure to treatment variables.

[15] As per Equation (8), we take $\hat{\delta}$ as one minus the estimated coefficient on the lagged dependent variable.

[16] This long-run estimate is now derived from Equation (4b), which suggests $w_{it+1} \approx (1-\delta-\beta)w_{it} + \beta x_{t+1} + \delta z_{it+1}$, implying the long-run coefficient on x_{t+1} can be derived as: $\beta/(\delta+\beta)$.

Table 5: Dynamic analysis of information treatments

	(Ia)	(Ib)	(Ic)	(IIa)	(IIb)	(IIc)	(IIIa)	(IIIb)	(IIIc)
Prior belief	0.56***	0.61***	0.61***	0.97***	1.01***	1.01***			
	(0.01)	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)			
Treated	–0.07**			–0.04*			–0.04*		
	(0.03)			(0.02)			(0.02)		
SMS wage news		0.07***			0.05**			0.05***	
		(0.02)			(0.03)			(0.02)	
Gen. SMS wage news			0.06**			0.05**			0.04**
			(0.02)			(0.02)			(0.02)
Uni. SMS wage news			0.06**			0.04**			0.04*
			(0.02)			(0.02)			(0.02)
Field SMS wage news			0.08***			0.08***			0.08***
			(0.03)			(0.03)			(0.03)
Private news (estd.)				0.40***	0.39***	0.39***	0.43***	0.39***	0.39***
				(0.03)	(0.04)	(0.04)	(0.01)	(0.02)	(0.02)
Constant	4.07***	3.56***	3.56***	0.29	–0.07	–0.06	0.01	0.01	0.01
	(0.17)	(0.23)	(0.23)	(0.24)	(0.29)	(0.35)	(0.03)	(0.03)	(0.03)
Obs.	9,053	9,053	9,053	9,053	9,053	9,053	9,053	9,053	9,053
AIC	9,760	9,751	9,754	9,746	9,737	9,739	9,756	9,733	9,739
R^2 (adj.)	0.435	0.436	0.436	0.435	0.435	0.435	0.269	0.270	0.270
RMSE	0.414	0.414	0.414	0.414	0.414	0.414	0.414	0.414	0.414
Gen. = Field = Uni.			0.759			0.262			0.275
Jointly zero			0.018			0.013			0.045

Note: significance: * 10 per cent, ** 5 per cent, *** 1 per cent. In columns (I) and (II), the dependent variable is own-earnings beliefs (log); in column (III) it is the first difference in the log of own-earnings beliefs; sub-columns (a) use a general treatment dummy; sub-columns (b) are the SMS wage news shock (all treatments combined); and sub-columns (c) distinguish between the message types; in place of extensive time-varying and baseline controls included in columns (I), columns (II) and (III) use a pre-generated proxy for the private news shock term; all specifications include experimental-arm fixed effects and month × round fixed effects; treatments are all zero in the first telephone survey round; standard errors (in parentheses) are clustered at the individual level; estimates in columns (II) and (III) derived from a bootstrap procedure (due to presence of a generated regressor).

Source: authors' construction based on survey results.

Comparing the results in column (II) against those of (I), the former specification performs well—that is, the main estimates of interest are qualitatively unchanged; goodness-of-fit statistics are similar; and, perhaps more critically, the estimated coefficient on the prior (lagged) belief is not different from 1. This result supports Equation (9), which suggests that once all material (systematic) news shocks are incorporated, the coefficient on the prior belief should be unity. Given the suitability of this approach to dealing with unobserved (private) news, column (III) shifts to the difference specification, in which the prior belief coefficient is constrained to unity by assumption. As expected, these results are almost exactly equivalent to the specifications in column (II) and the magnitude of field-specific news is double that of the other messages (but only borderline significantly different from them).

Overall, the results in column (IIIb) suggest that the weights attached to peer earnings news shocks (based on the SMS information), private news shocks, and prior beliefs equal around 0.05, 0.39, and 0.56 respectively. In this light, prior beliefs appear to be highly persistent (sticky) and the public information we provided about earnings played a minor yet non-trivial role in the revision of future earnings beliefs (on average). From a Bayesian perspective, this would imply the information treatment was not considered to be particularly precise (reliable) relative to prior beliefs. Also, the magnitude of the private news shocks suggest an important role for learning-by-experience—that is individuals learn about feasible salaries through job search and exposure to different jobs, which represent the dominant time-varying components of this term.

5.3 Robustness

Before proceeding, we validate the robustness of these dynamic results by applying a set of sample restrictions. Panels (a) and (b) of Appendix Table A3 adopt the parsimonious difference specifications from Table 5 columns (IIIb) and (IIIc), respectively. With these, we consider the following restrictions: in column (I) we exclude observations from the first telephone round, which relies on the baseline expected wage to construct the current news shock;[17] column (II) excludes all imputed wage expectations; column (III) excludes outliers, defined as individuals reporting a first difference in log expected wages that is either below the 1st or above the 99th percentile; column (IV) excludes outliers, now defined as individuals with a standard deviation of log expected wage that is either below the 5th or above the 95th percentile; column (V) excludes all outliers according to either of the previous definitions; column (VI) restricts attention to the balanced panel of individuals appearing in all rounds; column (VII) only includes treated individuals interviewed (by phone) within 25 days of receiving the information SMS; column (VIII) includes treated individuals interviewed (by phone) 25 days after receiving the information SMS;[18] column (IX) excludes individuals from the mixed information treatment arm (i.e. those who received different types of messages across the rounds); and column (X) includes all controls and *only* those individuals treated from the mixed treatment arm.

Looking across these results, panel (a) indicates no substantive deviations from the full sample results. This supports the insight that our information messages played a moderate role, on average, in the process of beliefs updating. In panel (b) we observe some variation in the main coefficient estimates, with a number of estimates for the general and university-specific messages falling below the minimum detectable effect threshold (approx. 0.05). However, the field-specific message continues to attract the largest and most robust estimated elasticity, and tests of the hypothesis that the coefficients are jointly equal to zero are generally rejected. Nonetheless, due to the relative imprecision of these estimates, the hypothesis that the three types of peer news shocks have an equal effect generally cannot be rejected.

5.4 Heterogeneous updating

Following Section 2, we now use Equation (7) to identify whether either of two sources of heterogeneity are present. These are: differences in information processing, captured by parameter θ_i; and differences in how peer news shocks are weighted (reflecting their salience), captured by parameter β_i. In terms of implementation, the limited number of times each participant was observed generally precludes accurate individual-level estimation of these heterogeneous parameters. Rather, we follow Armantier et al. (2016) and consider differences across sub-groups, obtained from interacting a series of control variables (denoted Y, including both dummy and continuous variables centred on their means) with both the peer news shock and the SMS information variable. These interaction terms are set to zero for all non-treated observations.

Table 6 sets out the main results, handling the different treatment arms homogeneously as in column (IIIb) of Table 5, while Appendix Table A5 extends the same specification to account for the disaggregated treatment news shocks. Column (1) is the reference or average model;[19] the remaining columns add different interaction terms, namely: (2) a dummy for wage earners; (3) a dummy for jobseekers; (4)

---

[17] Appendix Table A4 also replicates the estimates in Table 5 excluding all observations from round 1, but dropping the treatment-arm fixed effects since they are collinear with the information treatments.

[18] This split point reflects the median interview time after receiving the SMS among treated participants.

[19] In principle, we might include x_{it} as an additional regressor to obtain the reference estimate for φ. However, such estimates are not different from zero (available on request); excluding this term imposes a precise zero, in line with the assumption under Equation (7).

the round number, as a proxy to time since finishing university; (5) a measure of the reliability of the wage information sent by SMS; given by the share of survey respondents within each group who are in work, standardized to mean zero and standard deviation of 1; (6) women; (7) self-reported to be proficient in English; (8) Raven's test score; (9) baseline expected wage; and (10) the prior period's expected wage.[20] However, recognizing that wage earners are likely to respond to unanticipated real wage offers, given by the difference between their actual wage and their prior expected wage, column (2) onwards adds this additional news component. Also, so as to allow for substitution between the private and actual news shocks, we also add an interaction term between a dummy variable taking a value of 1 if the individual reports receiving a wage and the (generated) private news shock.

Three main points emerge. First, as shown in column (2), unexpected actual wage shocks are highly material (coefficient = 0.22) and partially substitute for some portion of the private news term. While this estimate does not imply anything close to one-for-one updating—that is, unanticipated wage offers are not fully incorporated into posterior expectations—it is consistent with previous research (e.g. Conlon et al. 2018) and supports the sticky nature of prior wage expectations. At the same time, it helps interpret the relative magnitude of responses to peer wage news shocks (our treatments). Here, the latter represents about about one-third of the realized wage shock, confirming our thesis that peer wage information is hardly trivial in revisions of own-wage expectations.

Second, looking across the columns, there is evidence for both forms of heterogeneity. Five of the 10 specifications including interaction terms indicate some form of heterogeneity is present, as per the joint probability test reported in the footer. In columns (4) and (5), we only observe heterogeneity in the weights applied to the peer earnings news ($\varphi_i = 0$), while differences in information processing are found in columns (7), (8), and (10). For instance, news received in later rounds is treated as much less salient than in the earlier treatment rounds (e.g. $\hat{\beta} \approx 0.10$ in round 2 vs $\hat{\beta} \approx 0.04$ in round 5). However, a one standard deviation increase in the reliability of the information message, proxied by the share of peers employed, maps to an increase of around one-third in the weight attributed to the peer earnings news. Above all, then, this points to a quite complex and varied pattern of responses to the information treatment.

Third, as speculated in Section 2, differences in ability appear to correspond to differences in information processing. However, the direction in which this operates runs counter to what one might expect. Lower-skill participants, as proxied by self-reported English-language skills and short-form Raven's test scores, tend to scale the peer earnings information upward. So, for optimistic prior wage expectations, this either compresses the magnitude of the news variable or even switches its sign, in turn reducing the magnitude of adjustment towards the raw signal for a given weight. For instance, consider the estimates in column (8) with values: $x_{it+1} = 10, w_{it} = 10.5$. For individuals with mean ability, the expected adjustment in wage expectations driven by this news term equals –0.04. For individuals with a Raven's score one standard deviation above the mean, the expected adjustment, given by $\beta([1+\varphi]x_{it+1} - w_{it})$, is –0.14. The expected adjustment for individuals with a Raven's score one standard deviation below the mean is +0.06 (i.e. in the positive domain).

Fourth, as shown in column (10), there is also some heterogeneity according to the individuals' prior wage expectations. The dominant pattern here is that individuals with above-average prior expectations tend to scale upward the peer wage earning information, similar to or perhaps even reflecting the response of lower-ability participants. Not dissimilar to Zafar (2011), this points to a distinction between 'highly optimistic' personality types, who tend to show low responses to (negative) information, versus individuals holding initially more realistic beliefs and concomitantly tending to be more open

[20] In the case of dummy variables (columns 2, 3, and 6), the column names refer to the group for which the variable takes a value of 1. Raven's test scores are IRT (item response theory) scores, based on a series of questions administered during the baseline survey.

to revising expectations over time. Moreover, the fact that responsiveness to news shocks appears to be a function of the individuals' prior beliefs contradicts the assumption of invariance under classical Bayesian updating (Mobius et al. 2011). Put differently, cognitive biases appear important, but absent experimental conditions, it is not possible to precisely determine which specific biases account for the observed patterns.

An alternative source of heterogeneity, not considered so far, would be non-linear responses to news, such as asymmetric updating (e.g. Barron 2020). To allow for this, Table 7 extends the previous specification, adding two interaction terms that capture differences in the response to both peer earnings and actual wage news when these shocks are in the positive domain. An immediate insight is that the previous findings regarding heterogeneity in both information processing and weighting broadly persist, implying they cannot be explained by misspecification of the updating process as linear. Second, although the responses to positive peer earnings shocks appear larger than when the news signal is negative (0.08 vs 0.04 in column 1), this difference is not significant. However, column (2) reveals a much stronger response to positive unanticipated wage offers than to negative—that is, responses to negative unanticipated wage shocks just one-quarter of the magnitude of responses to positive shocks, which enter posterior expectations on an almost one-for-one basis.

Disaggregated versions of the same estimates are reported in Tables A5 and A6. While these broadly confirm the above results, they also point to even more complex patterns of responses according to the type of information treatment received. For instance, column (1) of Table A5 suggests earners receiving the general message tended to significantly down-scale the peer earnings information, while variation in responses according to Raven's score tests are dominated by individuals receiving the field-specific message (column 8). However, there is not always sufficient power to determine these differences decisively, and tests for the equality of the coefficients across the same interaction terms involving different messages are often not different from zero (not shown; more information available on request).

A more robust finding is distinctive asymmetric responses to the different types of information. While Table 7 shows no significant overall asymmetry, Table A6 (column 1) reveals a large positive and significant coefficient on positive news shocks for recipients of the general and field-specific messages (0.42 and 0.29); this asymmetry remains present once interaction terms to account for heterogeneity are included (other columns), albeit most robustly for the field-specific messages. Interestingly, there is not even a hint of asymmetric responses to the university-specific message, suggesting that for this message the linear assumption is reasonable. To validate these insights and allow for a more general non-linear response to peer wage news, we re-run the specification in column (2) of Table 7, now introducing cubic polynomials in each of the three peer news shocks. Figures 4(a)–(d) plot the results, showing the predicted fit of changes in expected wages for different news signals (x-axis). Panel (a) is the least squares estimate, while panels (b)–(d) apply a quantile regression at the 50th, 25th, and 75th percentiles, respectively. Plots (a) and (b) confirm the stronger response to positive news for both the general and field-specific messages. Notably, at the 25th percentile of the change in own-wage expectations (c), we see very strong responses to negative news across the two granular messages, and a flat response in general to the university-specific message. Only at the 75th percentile does it seem that responses to the different messages are more similar.

Table 6: Heterogeneity analysis of information treatments (aggregated treatments)

Interaction term (Y) $\rightarrow$	(1) –	(2) Earning	(3) Seeking	(4) Round	(5) Reliability	(6) Female	(7) English	(8) Raven's	(9) Base w	(10) Prior w
SMS wage news	0.05***	0.08***	0.08***	0.08***	0.09***	0.09***	0.09***	0.08***	0.08***	0.10***
	(0.02)	(0.01)	(0.01)	(0.01)	(0.03)	(0.01)	(0.02)	(0.02)	(0.01)	(0.02)
SMS wage news $\times Y$		–0.01	0.03*	–0.02**	0.03**	0.00	–0.02	0.01	–0.03	–0.04
		(0.03)	(0.02)	(0.01)	(0.01)	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)
Salary info. $\times Y$		–0.04	0.03	0.03	–0.00	0.00	–0.21***	–0.13***	–0.03	0.10*
		(0.07)	(0.07)	(0.03)	(0.03)	(0.06)	(0.08)	(0.03)	(0.08)	(0.06)
Private news	0.39***	0.35***	0.36***	0.36***	0.36***	0.36***	0.36***	0.36***	0.36***	0.35***
	(0.02)	(0.02)	(0.03)	(0.02)	(0.04)	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)
Private news $\times$ earns		–0.11***	–0.12***	–0.12***	–0.15***	–0.12***	–0.12***	–0.13***	–0.12***	–0.11**
		(0.04)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.02)	(0.03)	(0.04)
Actual wage news		0.22***	0.22***	0.21***	0.22***	0.21***	0.21***	0.22***	0.21***	0.21***
		(0.01)	(0.01)	(0.02)	(0.02)	(0.01)	(0.01)	(0.02)	(0.01)	(0.02)
Some hetero.		0.69	0.20	0.06	0.06	0.97	0.01	0.00	0.31	0.05
Processing hetero.		0.52	0.68	0.36	1.00	0.99	0.01	0.00	0.76	0.09
Weighting hetero.		0.61	0.09	0.04	0.02	0.80	0.40	0.11	0.16	0.10

Note: significance: * 10 per cent, ** 5 per cent, *** 1 per cent. The first column presents the reference model, based on column (IIIb) in Table 5; in each subsequent column, we add interaction terms (Y) with both the peer news shock and the SMS information variable, which are set to zero for all non-treated observations; interaction variables are, by column: (2) wage earners; (3) individuals seeking a (new) job; (4) round number, equivalent to time since finishing university; (5) the standardized share of peers' working, as received via the SMS treatment; (6) women; (7) English proficiency; (8th) Raven's test score; (9th) baseline expected wage; (10th) prior period expected wage; 'some heterogeneity' is a joint test of the significance of the interaction terms; 'processing heterogeneity' is the probability that the SMS information interaction term is zero; 'weighting heterogeneity' is the probability that the wage news interaction term is zero; estimates calculated via a bootstrap procedure with standard errors clustered at the individual level (in parentheses).

Source: authors' construction based on survey results.

Table 7: Heterogeneity analysis of information treatments, with asymmetric responses (aggregated treatments)

Interaction term (Y) →	(1) –	(2) Unpaid	(3) Not sk.	(4) Round	(5) Emp. rate	(6) Female	(7) Academic	(8) Raven's	(9) Base w	(10) Prior w
SMS wage news	0.04**	0.08***	0.08***	0.07***	0.08**	0.08***	0.08***	0.08***	0.08***	0.09***
	(0.02)	(0.01)	(0.02)	(0.01)	(0.03)	(0.02)	(0.01)	(0.01)	(0.01)	(0.02)
SMS wage news [+]	0.04	0.03	0.03	0.03	0.05*	0.03	0.03	0.03*	0.02	0.02
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)	(0.03)
SMS wage news $\times Y$		–0.02	0.03**	–0.02**	0.03**	0.01	–0.02	0.01	–0.02	–0.03
		(0.02)	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)
Salary info. $\times Y$		–0.03	0.03	0.04	–0.00	–0.01	–0.22**	–0.13***	–0.02	0.11*
		(0.06)	(0.07)	(0.04)	(0.03)	(0.08)	(0.10)	(0.04)	(0.08)	(0.06)
Private news	0.38***	0.35***	0.36***	0.36***	0.36***	0.36***	0.36***	0.36***	0.36***	0.36***
	(0.02)	(0.03)	(0.02)	(0.02)	(0.04)	(0.03)	(0.02)	(0.02)	(0.02)	(0.03)
Private news × earns		–0.12***	–0.13***	–0.13***	–0.15***	–0.13***	–0.13***	–0.14***	–0.13***	–0.13***
		(0.03)	(0.02)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.02)	(0.03)
Actual wage news		0.20***	0.20***	0.19***	0.21***	0.19***	0.19***	0.20***	0.19***	0.19***
		(0.02)	(0.01)	(0.01)	(0.02)	(0.01)	(0.01)	(0.01)	(0.02)	(0.01)
Actual wage news [+]		0.60***	0.60***	0.57***	0.58***	0.60***	0.60***	0.60***	0.61***	0.60***
		(0.14)	(0.14)	(0.13)	(0.14)	(0.14)	(0.12)	(0.14)	(0.13)	(0.17)
Some hetero.		0.36	0.03	0.04	0.06	0.93	0.02	0.00	0.48	0.05
Processing hetero.		0.61	0.71	0.34	0.99	0.90	0.03	0.00	0.79	0.08
Weighting hetero.		0.31	0.01	0.02	0.02	0.70	0.44	0.22	0.23	0.27

Note: significance: * 10 per cent, ** 5 per cent, *** 1 per cent. This table replicates Table 6, adding interaction terms, denoted by '[+]', defined as the SMS wage news and actual wage news variables multiplied by a dummy variable taking a value of 1 if these variables are in the positive domain.

Source: authors' construction based on survey results.

Figure 4: Non-linear responses to SMS wage news

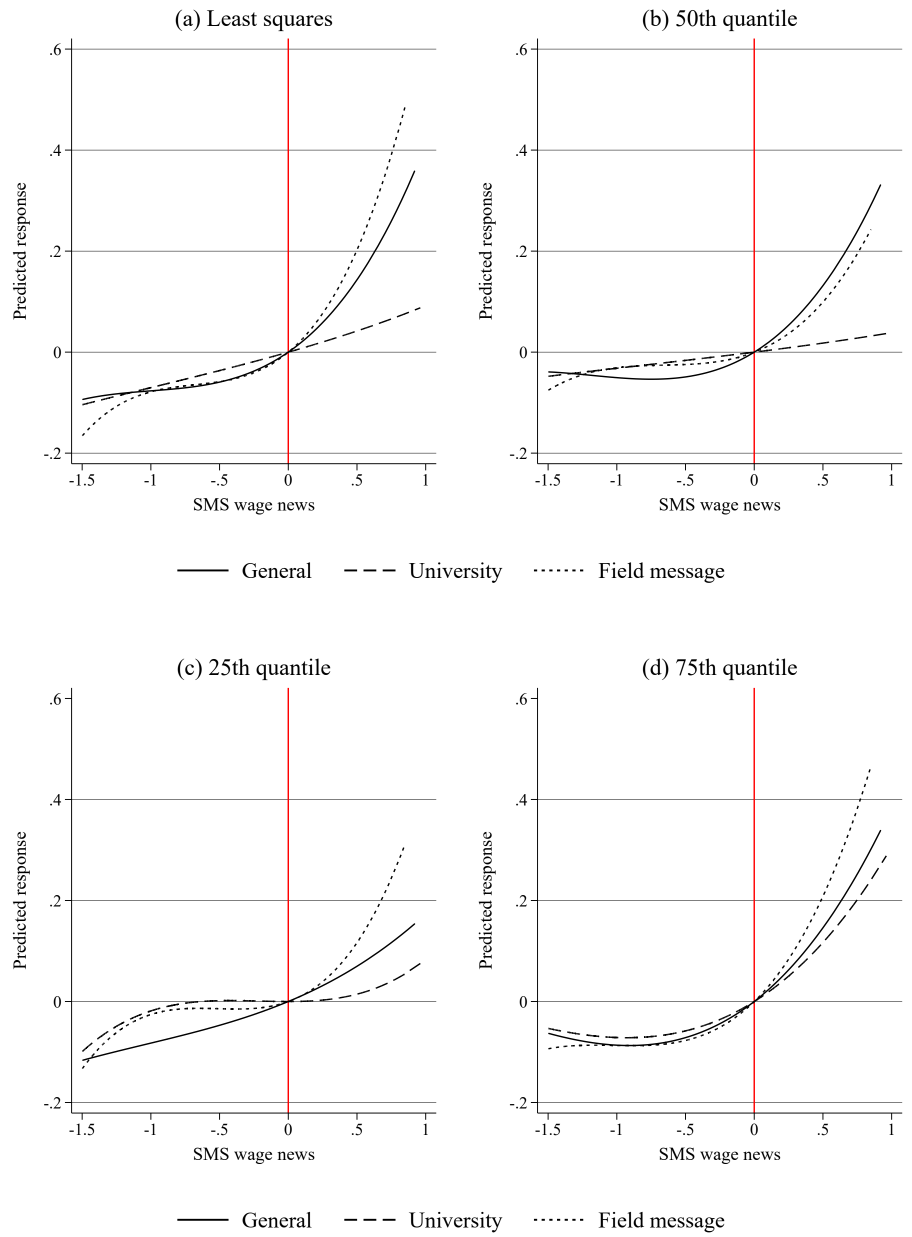

Note: both figures plot the predicted change in the expected log wage (y-axis) in response to news regarding peer wages, where the news is defined as the peer wage information minus the individuals' expected wages in the previous period; estimates are based on the model in column (I) of Table A3, extended to include a third degree polynomial in the SMS wage news treatments; panels (b)–(d) are quantile regressions of the same specification set at 50th, 25th, and 75th percentiles, respectively.

Source: authors' construction based on survey results.

There is no one unique plausible interpretation of these findings. At a general level, however, they confirm the hypothesis that the type of message matters—looking across the results, the field-specific message appears the most robust and salient of the three types; and the response to positive news is an order of magnitude (at least three times) larger than the response to negative news, but not for the university-specific message. One possibility here is that, in the absence of information regarding which fields of study were more represented among peers who found employment alongside an expectation of large differences in wages across occupations, individuals receiving the university-specific messages were less able to arrive at a consistent interpretation of peer earnings information. The same disadvantage might apply to the general message, explaining its weaker salience relative to the field message, but could be somewhat offset by its much broader coverage, thereby providing an informative overall reference point or anchor for graduate wages. The point is that how individuals update their own-earnings expectations based on partial information is not homogeneous, and different types of information appear to be used in different ways by different people.

6 Conclusion

We began this paper noting that much empirical research, even in high-income markets with extensive labour market data, finds that jobseekers are often poorly informed about their future employment prospects. As a consequence, they tend to report systematically inaccurate, typically optimistic, wage expectations. In turn, the authors of one recent paper conclude: 'Since expectations play a critical role in decision making under uncertainty ... the large errors in population beliefs in our sample of high-ability students and the low cost of information interventions suggest a role for information campaigns focused on providing accurate information on returns to schooling/occupations' (Wiswall and Zafar 2015: 157). This paper sought to respond to this challenge by examining how college graduates in Mozambique update their wage expectations in response to information regarding the average salaries their peers recently obtained. In doing so, we have contributed one of the first longitudinal real-world (field) experiments into the formation of wage expectations in a low-income setting; we also examined the extent to which different types of information (e.g. according to the field of study) may have differential effects.

Our headline finding was that receipt of the information treatment yielded an approximate 13 per cent overall decline in the expected wage relative to non-recipients, which is consistent with the observation that (initial) expectations were much higher than realized wages on average. In terms of what we called the peer earnings news shock, defined as the difference between the information about peer wages and the individuals' prior wage expectations, we similarly found a short-term elasticity of wage expectations to these shocks in the order of 7 per cent, or a long-term elasticity of around 16 per cent. Looking across the treatments, while it was not possible to precisely identify differences in effect magnitudes, the field-specific message consistently elicited the largest and most robust response of the three.

Following the large existing literature documenting deviations from classical Bayesian updating, we went on to investigate the extent and nature of heterogeneity in responses to new information. To do so, we extended a basic empirical model of wage updating to account for heterogeneity in information processing, as well as heterogeneity in the weight or salience attributed to peer news shocks. The main insight was that both forms of heterogeneity appeared present, but in complex ways—for example, responsiveness to the news tended to fall over time, while higher-ability participants tended to downscale the earnings information. We further found that responses to unanticipated wage offers were generally around three times larger than the peer wage shock, but still significantly less than unity.

Additional tests for the presence of asymmetric responses to news indicated that positive unanticipated wage offers were almost fully incorporated into posterior beliefs. On average, however, we found that positive peer wage news shocks were associated with much stronger responses than equivalent shocks on the negative domain, but only for the general and field-specific information types. We interpreted this variation in responses to the different message as indicative of the difficulty of consistently interpreting university-specific information, at least where there are large occupational differences in earnings that map to differences in study fields (e.g. engineering and medicine).

Two main general implications follow. First, while providing information about actual earnings of university graduates represents both a low-cost and relevant contribution to the formation of wage expectations (at least in our context), optimistic beliefs nonetheless remain remarkably persistent. This is particularly the case among individuals who held comparatively high initial expectations, as well as when the news content of this information is in the negative domain. As a result, it is not evident that information provision *per se* will lead to a substantial shift in wage expectations towards the true distribution, or different labour market choices. Having said that, we highlight that our study ran in a context of scarce public information about wages and employability. Jobseekers therefore are not likely to be well trained in interpreting third-party job market information or able to fully determine its trustworthiness. Notably, those more open to adjusting expectations and those with higher Raven's test scores tended to show higher response elasticities to the peer earnings news. Thus, institutionalizing public information on wages and job opportunities may increase the effectiveness of this information over time.

Second, given the varied response to different types of messages, it would seem valuable to provide detailed information on realized salaries, including both aggregate and occupation- or field-specific outcomes. This may be more useful than information at the institution level alone, but in the absence of regular public surveys this would require coordination across universities to share (pool) information on alumni salaries. Further research on responses to distributional information, rather than just salary averages, would also be helpful to improve the design of information interventions.

Finally, we would emphasize that our measure of elicited own-wage expectations is not equivalent to the reservation wage (for discussion, see Jones et al. 2020), which arguably may be more critical for jobseekers. Thus, a further test of the efficacy of information interventions concerns whether they prompted a swifter transition into work or changes in jobseeking strategies. This fell beyond the scope of the present study but merits future research.

References

Abbiati, G., G. Argentin, C. Barone, and A. Schizzerotto (2018). 'Information Barriers and Social Stratification in Higher Education: Evidence from a Field Experiment'. *British Journal of Sociology*, 69(4): 1248–70. https://doi.org/10.1111/1468-4446.12332

Altmann, S., A. Falk, S. Jäger, and F. Zimmermann (2018). 'Learning About Job Search: A Field Experiment with Job Seekers in Germany'. *Journal of Public Economics*, 164: 33–49. https://doi.org/10.1016/j.jpubeco.2018.05.003

Armantier, O., S. Nelson, G. Topa, W. van der Klaauw, and B. Zafar (2016). 'The Price Is Right: Updating Inflation Expectations in a Randomized Price Information Experiment'. *Review of Economics and Statistics*, 98(3): 503–23. https://doi.org/10.1162/REST_a_00499

Barron, K. (2020). 'Belief Updating: Does the "Good-News, Bad-News" Asymmetry Extend to Purely Financial Domains?' *Experimental Economics*. https://doi.org/10.1007/s10683-020-09653-z

Belot, M., P. Kircher, and P. Muller (2018). 'Providing Advice to Jobseekers at Low Cost: An Experimental Study on Online Advice'. *Review of Economic Studies*, 86(4): 1411–47. https://doi.org/10.1093/restud/rdy059

Bénabou, R., and J. Tirole (2016). 'Mindful Economics: The Production, Consumption, and Value of Beliefs'. *Journal of Economic Perspectives*, 30(3): 141–64. https://doi.org/10.1257/jep.30.3.141

Bleemer, Z., and B. Zafar (2018). 'Intended College Attendance: Evidence from an Experiment on College Returns and Costs'. *Journal of Public Economics*, 157: 184–211. https://doi.org/10.1016/j.jpubeco.2017.11.002

Bonilla-Mejía, L., N.L. Bottan, and A. Ham (2019). 'Information Policies and Higher Education Choices Experimental Evidence from Colombia'. *Journal of Behavioral and Experimental Economics*, 83: 101468. https://doi.org/10.1016/j.socec.2019.101468

Burdett, K., and T. Vishwanath (1988). 'Declining Reservation Wages and Learning'. *Review of Economic Studies*, 55(4): 655–65. https://doi.org/10.2307/2297410

Cavallo, A., G. Cruces, and R. Perez-Truglia (2017). 'Inflation Expectations, Learning, and Supermarket Prices: Evidence from Survey Experiments'. *American Economic Journal: Macroeconomics*, 9(3): 1–35. https://doi.org/10.1257/mac.20150147

Chambers, C.P., and P.J. Healy (2012). 'Updating Toward the Signal'. *Economic Theory*, 50(3): 765–86. https://doi.org/10.1007/s00199-010-0588-0

Chen, Q., A. Gelman, M. Tracy, F.H. Norris, and S. Galea (2015). 'Incorporating the Sampling Design in Weighting Adjustments for Panel Attrition'. *Statistics in Medicine*, 34(28): 3637–47. https://doi.org/10.1002/sim.6618

Cochran, W.G. (1977). *Sampling Techniques*, 3rd edition. New York: Wiley.

Conlon, J.J., L. Pilossoph, M. Wiswall, and B. Zafar (2018). 'Labor Market Search with Imperfect Information and Learning'. Working Paper w24988. Cambridge, MA: NBER. https://doi.org/10.3386/w24988

Cooper, M., and P. Kuhn (2020). 'Behavioral Job Search'. In K.F. Zimmermann (ed.) *Handbook of Labor, Human Resources and Population Economics*. Amsterdam: Springer. https://doi.org/10.1007/978-3-319-57365-6_116-1

De Paola, M., C. Lupi, and P. Ordine (2001). 'A Tale on Information and Wage Expectations'. ISAE Working Paper 22. Rome: Italian National Institute of Statistics (ISTAT).

DellaVigna, S. (2009). 'Psychology and Economics: Evidence from the Field'. *Journal of Economic Literature*, 47(2): 315–72. https://doi.org/10.1257/jel.47.2.315

Fehr, D., J. Mollerstrom, and R. Perez-Truglia (2019). 'Your Place in the World: The Demand for National and Global Redistribution'. Working Paper w26555. Cambridge, MA: NBER. https://doi.org/10.3386/w26555

Gradín, C., F. Mambo, Y. Paris, and R. Santos (2020). 'Unequal Expectations: Gender Inequality in Salary Expectations of University Students'. Working Paper 2020/78. Helsinki: UNU-WIDER.https://doi.org/10.35188/UNU-WIDER/2020/835-1

Haaland, I., C. Roth, and J. Wohlfart (2020). 'Designing Information Provision Experiments'. CESifo Working Paper 8406. Munich: CESifo.

Herber, S.P. (2018). 'The Role of Information in the Application for Highly Selective Scholarships: Evidence from a Randomized Field Experiment'. *Economics of Education Review*, 62: 287–301. https://doi.org/10.1016/j.econedurev.2017.12.001

Jones, S., R. Santos, and G. Xirinda (2020). 'Misinformed, Mismatched, or Misled? Explaining the Gap Between Expected and Realised Graduate Earnings in Mozambique'. Working Paper 2020/47. Helsinki: UNU-WIDER. https://doi.org/10.35188/UNU-WIDER/2020/804-7

Kandel, E., and B.-Z. Zilberfarb (1999). 'Differential Interpretation of Information in Inflation Forecasts'. *Review of Economics and Statistics*, 81(2): 217–26. https://doi.org/10.1162/003465399558193

Kiefer, N.M., and G.R. Neumann (1979). 'An Empirical Job-Search Model, with a Test of the Constant Reservation-Wage Hypothesis'. *Journal of Political Economy*, 87(1): 89–107. https://doi.org/10.1086/260741

Kriechel, B., and G.A. Pfann (2006). 'Learning to Update Your Reservation Wage While Looking for a New Job'. *Portuguese Economic Journal*, 5(2): 135–48. https://doi.org/10.1007/s10258-006-0008-3

Krueger, A.B., and A.I. Mueller (2016). 'A Contribution to the Empirics of Reservation Wages'. *American Economic Journal: Economic Policy*, 8(1): 142–79. https://doi.org/10.1257/pol.20140211

Kuziemko, I., M.I. Norton, E. Saez, and S. Stantcheva (2015). 'How Elastic Are Preferences for Redistribution? Evidence from Randomized Survey Experiments'. *American Economic Review*, 105(4): 1478–508. https://doi.org/10.1257/aer.20130360

McCall, J.J. (1970). 'Economics of Information and Job Search'. *Quarterly Journal of Economics*, 84(1): 113–26. https://doi.org/10.2307/1879403

Mobius, M.M., M. Niederle, P. Niehaus, and T.S. Rosenblat (2011). 'Managing Self-Confidence: Theory and Experimental Evidence'. Working Paper w17014. Cambridge, MA: NBER. https://doi.org/10.3386/w17014

Mortensen, D.T. (1970). ‘Job Search, the Duration of Unemployment, and the Phillips Curve’. *American Economic Review*, 60(5): 847–62.

Mortensen, D.T., and C.A. Pissarides (1994). ‘Job Creation and Job Destruction in the Theory of Unemployment’. *Review of Economic Studies*, 61(3): 397–415. https://doi.org/10.2307/2297896

Pirotte, A. (1999). ‘Convergence of the Static Estimation Toward the Long Run Effects of Dynamic Panel Data Models’. *Economics Letters*, 63(2): 151–58. https://doi.org/10.1016/S0165-1765(99)00023-3

Sharot, T., C.W. Korn, and R.J. Dolan (2011). ‘How Unrealistic Optimism Is Maintained in the Face of Reality’. *Nature Neuroscience*, 14(11): 1475–79. https://doi.org/10.1038/nn.2949

Van den Doel, I., and J.F. Kiviet (1994). ‘Asymptotic Consequences of Neglected Dynamics in Individual Effects Models’. *Statistica Neerlandica*, 48(1): 71–85. https://doi.org/10.1111/j.1467-9574.1994.tb01432.x

Viscusi, W.K., and C.J. O’Connor (1984). ‘Adaptive Responses to Chemical Labeling: Are Workers Bayesian Decision Makers?’. *American Economic Review*, 74(5): 942–56.

Wiswall, M., and B. Zafar (2015). ‘How Do College Students Respond to Public Information About Earnings?’. *Journal of Human Capital*, 9(2): 117–69. https://doi.org/10.1086/681542

Zafar, B. (2011). ‘How Do College Students Form Expectations?’. *Journal of Labor Economics*, 29(2): 301–48. https://doi.org/10.1086/658091

Appendix

A1 Attrition analysis

The level of attrition sustained by the experiment, previously depicted in Table 1, is highly satisfactory. Even so, here we investigate whether differential attrition might bias the experimental results. To look at this, we define 'early leavers' as those who dropped out of the survey before the fifth round, which is the final round in which we elicited wage expectations. As can be seen in the first three panels of Table A1, early leavers appear statistically similar to non-leavers in most observed characteristics at baseline. However, there is a clear indication of gender-based differences as well as some differences between universities and study areas, which likely also reflect gender differences.

The fourth panel of Table A1 gives further assurance that (later) attrition does not correlate with baseline characteristics—that is, both the baseline expected wage and the expected wage observed in the first follow-up survey round are not different between early leavers and others (see also Figure A1(a)). Even so, the final two rows of the table do suggest differential attrition according to exposure to the treatment. Even after controlling for period-specific effects, the last observed expected wage is generally smaller among non-survey leavers (see also Figure A1(b)). A candidate explanation is greater attrition within the control group, which is confirmed in the final row.

To examine this further, we regressed the probability of being observed in the final (fifth) round on a full set of baseline characteristics (as described in the text) plus dummy variables capturing the treatment arm to which the individual was allocated. The results, shown in column (Ia) of Table A2 based on a probit model (odds ratios reported), confirm that being part of the treatment group significantly increased the odds of remaining in the survey; and column (Ib) indicates this was primarily driven by recipients of the general message. Columns (IIa) and (IIb) repeat the same analysis using an ordered probit estimator, where the dependent variable is now the number of the final round in which the individual was observed. The same key results emerge, confirming lower attrition among individuals exposed to the treatment. *Prima facie*, this would suggest the information contained in the treatments were perceived to be of some value.

Table A1: Baseline characteristics (means), comparison of survey leavers vs non-leavers

	Non-leavers	Leavers	F-test	Obs.
Female	0.43	0.70	0.000	2,069
	(0.00)	(0.02)	.	.
Age	25.91	26.41	0.376	2,069
	(0.06)	(0.30)	.	.
Prior work experience	0.61	0.53	0.048	2,069
	(0.00)	(0.02)	.	.
UEM	0.35	0.39	0.222	2,069
	(0.00)	(0.02)	.	.
UCM	0.09	0.07	0.299	2,069
	(0.00)	(0.01)	.	.
UZ	0.10	0.08	0.132	2,069
	(0.00)	(0.01)	.	.
USTM	0.05	0.05	0.369	2,069
	(0.00)	(0.01)	.	.
UP	0.36	0.27	0.038	2,069
	(0.00)	(0.02)	.	.
AP	0.05	0.13	0.000	2,069
	(0.00)	(0.01)	.	.
Education	0.31	0.24	0.177	2,069
	(0.00)	(0.02)	.	.
Humanities	0.02	0.02	0.479	2,069
	(0.00)	(0.01)	.	.
Social sciences	0.44	0.55	0.012	2,069
	(0.00)	(0.02)	.	.
Natural sciences	0.04	0.02	0.004	2,069
	(0.00)	(0.01)	.	.
Engineering	0.08	0.07	0.616	2,069
	(0.00)	(0.01)	.	.
Agriculture	0.06	0.04	0.114	2,069
	(0.00)	(0.01)	.	.
Health	0.06	0.05	0.704	2,069
	(0.00)	(0.01)	.	.
Expected wage (baseline)	32.96	32.15	0.592	2,065
	(0.40)	(1.23)	.	.
Expected wage (first obs.)	30.14	32.22	0.821	1,981
	(0.38)	(1.33)	.	.
Expected wage (last obs.)	23.31	32.16	0.100	1,860
	(0.30)	(1.38)	.	.
Treatment group	0.77	0.68	0.001	2,069
	(0.00)	(0.02)	.	.

Notes: first two columns report the means of each row variable (baseline characteristic), split between those participants who did and did not remain in the sample until the fifth round; standard errors of means in parentheses; 'F-test' reports the probability that the means are different, based on a regression controlling for first and last rounds in which the individual was observed. UEM–AP refer to universities (dummy variables); expected wages are in '000s of MZN.

Source: authors' construction based on survey results.

Figure A1: Wage expectations (in ’000s MZN) by first and last round observed

(a) *First observed round*

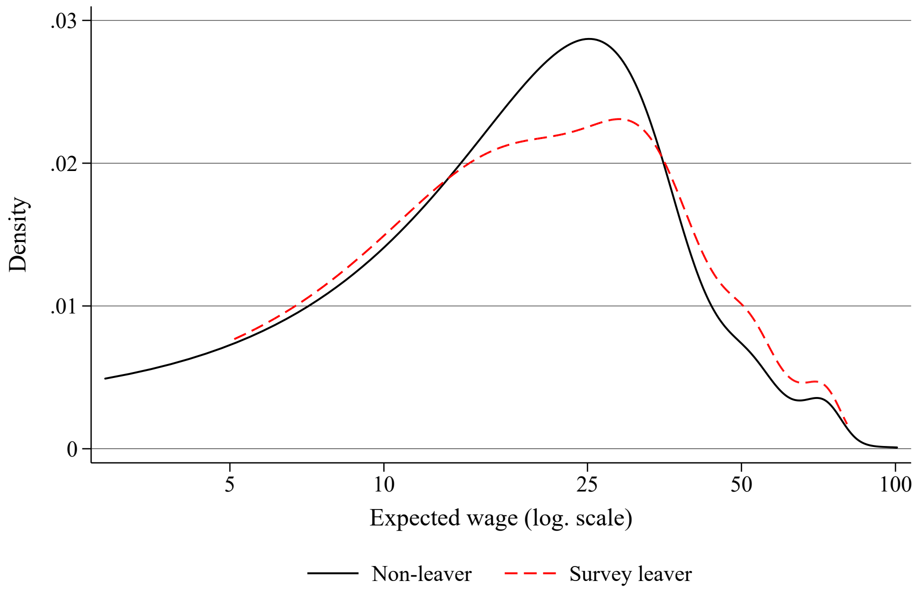

(b) *Last observed round*

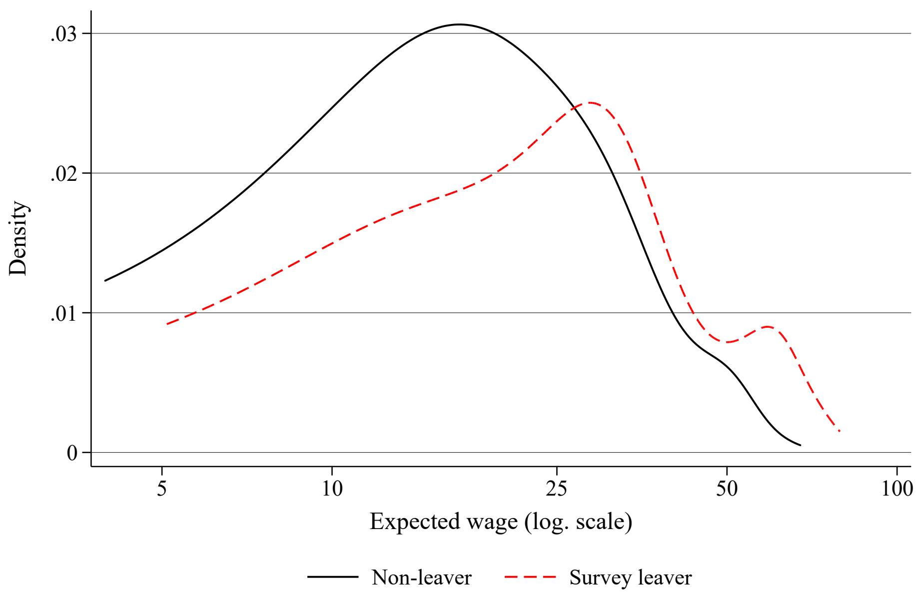

Note: the figures plot kernel density functions of expected wages based on the first and last observations per individual, split between those that did and did not drop out before the fifth round.

Source: authors’ construction based on survey results.

Table A2: Correlates of attrition

	Non-leaver		Final round no.	
	(Ia)	(Ib)	(IIa)	(IIb)
Treated	1.57***		1.78***	
	(0.18)		(0.20)	
General SMS		1.35*		1.37**
		(0.22)		(0.22)
University SMS		1.10		1.09
		(0.17)		(0.16)
Field SMS		1.14		1.10
		(0.18)		(0.17)
Mixed SMS		1.03		1.04
		(0.16)		(0.16)
Female	0.51***	0.52***	0.52***	0.53***
	(0.05)	(0.05)	(0.05)	(0.05)
Age	0.97	0.97	0.98	0.98
	(0.05)	(0.05)	(0.06)	(0.06)
Speaks English	0.67***	0.68***	0.68***	0.71**
	(0.10)	(0.10)	(0.09)	(0.10)
Education	1.30*	1.29*	1.30**	1.29**
	(0.17)	(0.17)	(0.17)	(0.16)
Humanities	1.14	1.16	1.17	1.20
	(0.23)	(0.23)	(0.23)	(0.23)
Natural sciences	1.57***	1.56***	1.55***	1.53***
	(0.26)	(0.25)	(0.25)	(0.24)
Engineering	1.25	1.22	1.27	1.21
	(0.25)	(0.24)	(0.24)	(0.23)
Agriculture	1.32	1.32	1.34*	1.34*
	(0.24)	(0.24)	(0.23)	(0.23)
Health	1.10	1.11	1.10	1.12
	(0.26)	(0.26)	(0.27)	(0.26)
Obs.	2,069	2,069	2,069	2,069

Note: significance: * 10 per cent, ** 5 per cent, *** 1 per cent. Odds ratios reported; a probit model is applied in column (I), where the dependent variable is binomial, taking a value of 1 if the participant is observed in the fifth round of the experiment; an ordered probit is used in column (II), where the dependent variable is the last round in which the participant is observed; sub-columns (a) use an aggregate treatment dummy (taking a value of 1 if the participant was allocated to any treatment arm); sub-columns (b) distinguish between the treatment arms; selected coefficients shown; robust standard errors in parentheses.

Source: authors' construction based on survey results.

A2 Additional tables

Table A3: Robustness analysis based on sample restrictions

	(I)	(II)	(III)	(IV)	(V)	(VI)	(VII)	(VIII)	(IX)	(X)
(a) Average info. treatment shock										
SMS wage news	0.07***	0.04*	0.06***	0.05***	0.07***	0.05**	0.04**	0.07***	0.05***	0.05**
	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Private news (estd.)	0.36***	0.41***	0.31***	0.35***	0.32***	0.39***	0.38***	0.39***	0.39***	0.39***
	(0.02)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.03)
Obs.	7,154	7,875	8,216	8,118	7,622	8,355	6,475	6,102	7,301	4,907
(b) Separate info. treatment shock										
Gen. SMS wage news	0.07**	0.02	0.06***	0.05**	0.06***	0.04	0.04*	0.03	0.04**	0.03
	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)	(0.02)	(0.03)	(0.02)	(0.03)
Uni. SMS wage news	0.04*	0.05***	0.06***	0.04*	0.06***	0.03	0.02	0.07**	0.04*	0.04
	(0.02)	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)	(0.02)	(0.03)
Field SMS wage news	0.10***	0.07***	0.07***	0.07**	0.07***	0.08**	0.05***	0.10***	0.08***	0.10**
	(0.03)	(0.02)	(0.02)	(0.03)	(0.02)	(0.03)	(0.02)	(0.04)	(0.02)	(0.04)
Private news (estd.)	0.36***	0.41***	0.31***	0.35***	0.32***	0.39***	0.38***	0.39***	0.39***	0.39***
	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)	(0.02)	(0.02)
Gen. = Field = Uni.	0.173	0.069	0.996	0.512	0.877	0.318	0.280	0.125	0.119	0.457
Jointly zero	0.016	0.000	0.000	0.032	0.001	0.086	0.015	0.045	0.001	0.046

Note: significance: * 10 per cent, ** 5 per cent, *** 1 per cent. Estimates follow the difference specification as per Table 5, column (III); column (I) drops observations from the first follow-up round; column (II) excludes imputed own-earnings beliefs; column (III) excludes individuals reporting a change in expected wages below the 1st or above the 99th percentile; column (IV) excludes individuals whose standard deviation of log expected wages is below the 5th or above the 95th percentile; column (V) excludes all individuals dropped from either column (III) or (IV); column (VI) includes only individuals appearing in all rounds; column (VII) includes treated individuals who were interviewed (by phone) within 25 days of receiving the information SMS; column (VIII) includes treated individuals who were interviewed (by phone) 25 days after receiving the information SMS; column (IX) excludes individuals from the mixed information treatment arm; column (X) includes all controls and only individuals from the mixed treatment arm; estimates calculated via a bootstrap procedure.

Source: authors' construction based on survey results.

Table A4: Dynamic analysis of information treatments, excluding round 1

	(Ia)	(Ib)	(Ic)	(IIa)	(IIb)	(IIc)	(IIIa)	(IIIb)	(IIIc)
Prior belief	0.58***	0.62***	0.62***	0.94***	0.98***	0.98***			
	(0.01)	(0.02)	(0.02)	(0.04)	(0.03)	(0.02)			
Treated	–0.04***			–0.05***			–0.05***		
	(0.02)			(0.01)			(0.01)		
SMS wage news		0.06***			0.06***			0.06***	
		(0.01)			(0.01)			(0.02)	
Gen. SMS wage news			0.04**			0.05**			0.05***
			(0.02)			(0.02)			(0.02)
Uni. SMS wage news			0.05***			0.05**			0.05***
			(0.02)			(0.02)			(0.02)
Field SMS wage news			0.09***			0.09***			0.10***
			(0.02)			(0.02)			(0.02)
Private news (estd.)				0.35***	0.35***	0.35***	0.41***	0.37***	0.37***
				(0.05)	(0.04)	(0.03)	(0.02)	(0.02)	(0.02)
Constant	3.88***	3.45***	3.47***	0.52	0.13	0.13	–0.09***	–0.08**	–0.09***
	(0.18)	(0.20)	(0.20)	(0.37)	(0.27)	(0.23)	(0.03)	(0.04)	(0.03)
Obs.	7,153	7,153	7,153	7,154	7,154	7,154	7,154	7,154	7,154
AIC	6,919	6,907	6,901	7,007	6,997	6,991	7,016	6,995	6,991
R^2 (adj.)	0.488	0.489	0.489	0.480	0.481	0.482	0.216	0.218	0.219
RMSE	0.391	0.391	0.391	0.394	0.394	0.394	0.394	0.394	0.394
Gen. = Field = Uni.			0.040			0.019			0.093
Jointly zero			0.000			0.000			0.000

Note: significance: * 10 per cent, ** 5 per cent, *** 1 per cent. This table replicates Table 5, dropping all observations from round 1.

Source: authors' construction based on survey results.

Table A5: Heterogeneity analysis of information treatments

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Interaction term (Y) $\rightarrow$	–	Unpaid	Not sk.	Round	Emp. rate	Female	Academic	Raven's	Base w	Prior w
Gen. SMS wage news	0.04**	0.08***	0.08***	0.08***	0.10**	0.08***	0.08***	0.08***	0.08***	0.10***
	(0.02)	(0.02)	(0.01)	(0.02)	(0.04)	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)
Uni. SMS wage news	0.04**	0.08***	0.08***	0.07***	0.08**	0.08***	0.08***	0.08***	0.07***	0.08***
	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)
Field SMS wage news	0.08***	0.11***	0.11***	0.10***	0.13***	0.12***	0.12***	0.11***	0.12***	0.13***
	(0.03)	(0.03)	(0.02)	(0.03)	(0.04)	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)
Gen. SMS wage news $\times Y$		–0.01	0.01	–0.01	–0.09	–0.02	–0.00	0.00	–0.03	–0.06*
		(0.03)	(0.02)	(0.01)	(0.32)	(0.02)	(0.04)	(0.01)	(0.02)	(0.03)
Uni. SMS wage news $\times Y$		–0.02	0.07***	–0.02	0.02*	0.00	–0.02	0.02	0.01	–0.02
		(0.02)	(0.03)	(0.02)	(0.01)	(0.02)	(0.03)	(0.01)	(0.02)	(0.03)
Field SMS wage news $\times Y$		–0.00	0.01	–0.02*	0.04**	0.04	–0.03	0.02	–0.07	–0.04
		(0.03)	(0.03)	(0.01)	(0.02)	(0.03)	(0.03)	(0.02)	(0.04)	(0.04)
Gen. SMS info. $\times Y$		–0.61***	0.37	0.02	–3.77	0.26	–0.65*	–0.08	–0.48**	0.37
		(0.20)	(0.27)	(0.11)	(13.57)	(0.25)	(0.35)	(0.13)	(0.21)	(0.23)
Uni. SMS info. $\times Y$		–0.04	0.01	–0.00	0.04	–0.00	–0.04	–0.12**	0.06	0.02
		(0.08)	(0.10)	(0.05)	(0.03)	(0.09)	(0.18)	(0.05)	(0.10)	(0.08)
Field SMS info. $\times Y$		0.05	–0.01	0.06	–0.11	–0.06	–0.28*	–0.13***	–0.01	0.15
		(0.10)	(0.09)	(0.04)	(0.08)	(0.08)	(0.15)	(0.04)	(0.09)	(0.12)
Some hetero.		0.02	0.05	0.08	0.01	0.07	0.00	0.00	0.00	0.42
Processing hetero.		0.01	0.58	0.42	0.45	0.61	0.07	0.00	0.10	0.25
Weighting hetero.		0.89	0.06	0.17	0.04	0.03	0.67	0.33	0.02	0.40

Note: significance: * 10 per cent, ** 5 per cent, *** 1 per cent. This table replicates Table 6, now distinguishing between the three types of information treatments.

Source: authors' construction based on survey results.

Table A6: Heterogeneity analysis of information treatments

Interaction term (Y) →	(1) –	(2) Unpaid	(3) Not sk.	(4) Round	(5) Emp. rate	(6) Female	(7) Academic	(8) Raven's	(9) Base w	(10) Prior w
Gen. SMS wage news	0.02	0.07***	0.07***	0.06***	0.07*	0.07***	0.07***	0.07***	0.07***	0.09***
	(0.02)	(0.01)	(0.02)	(0.01)	(0.04)	(0.02)	(0.01)	(0.02)	(0.02)	(0.03)
Uni. SMS wage news	0.03	0.07***	0.07***	0.06***	0.06*	0.07***	0.07***	0.07***	0.07***	0.08***
	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)	(0.01)	(0.02)	(0.02)	(0.02)	(0.03)
Field SMS wage news	0.06**	0.10***	0.10***	0.09***	0.10**	0.11***	0.10***	0.10***	0.11***	0.11***
	(0.03)	(0.02)	(0.02)	(0.02)	(0.04)	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)
Gen. SMS wage news [+]	0.42***	0.28	0.26	0.30*	0.31*	0.26	0.25*	0.25	0.22	0.36*
	(0.16)	(0.19)	(0.20)	(0.16)	(0.17)	(0.16)	(0.15)	(0.18)	(0.20)	(0.19)
Uni. SMS wage news [+]	0.04	–0.01	–0.02	0.02	0.10	0.00	0.00	0.01	–0.00	0.09
	(0.15)	(0.16)	(0.14)	(0.13)	(0.17)	(0.17)	(0.16)	(0.22)	(0.13)	(0.17)
Field SMS wage news [+]	0.29**	0.27*	0.28	0.26**	0.32**	0.27**	0.27*	0.27***	0.22	0.40*
	(0.15)	(0.15)	(0.19)	(0.13)	(0.15)	(0.13)	(0.14)	(0.10)	(0.15)	(0.21)
Gen. SMS wage news $\times Y$		–0.01	0.01	–0.02	–0.07	–0.02	0.00	0.00	–0.02	0.00
		(0.03)	(0.02)	(0.01)	(0.29)	(0.02)	(0.04)	(0.01)	(0.02)	(0.04)
Uni. SMS wage news $\times Y$		–0.02	0.07***	–0.03	0.03**	0.00	–0.02	0.02	0.01	0.02
		(0.03)	(0.02)	(0.02)	(0.01)	(0.02)	(0.03)	(0.01)	(0.02)	(0.04)
Field SMS wage news $\times Y$		–0.01	0.01	–0.03*	0.04**	0.04	–0.03	0.02	–0.05	0.04
		(0.04)	(0.03)	(0.02)	(0.02)	(0.03)	(0.05)	(0.02)	(0.03)	(0.05)
Gen. SMS info. $\times Y$		–0.63***	0.39	–0.04	–1.13	0.24	–0.65*	–0.07	–0.45*	0.45*
		(0.23)	(0.25)	(0.10)	(9.77)	(0.23)	(0.34)	(0.15)	(0.25)	(0.27)
Uni. SMS info. $\times Y$		–0.04	0.02	0.01	0.03	0.00	–0.04	–0.11**	0.06	0.05
		(0.12)	(0.08)	(0.05)	(0.03)	(0.10)	(0.16)	(0.05)	(0.12)	(0.11)
Field SMS info. $\times Y$		0.08	–0.03	0.06	–0.09	–0.07	–0.29*	–0.13**	0.01	0.21**
		(0.11)	(0.12)	(0.05)	(0.08)	(0.11)	(0.15)	(0.06)	(0.11)	(0.10)
Some hetero.		0.00	0.00	0.22	0.17	0.46	0.16	0.07	0.09	0.39
Processing hetero.		0.00	0.38	0.64	0.48	0.73	0.12	0.03	0.27	0.10
Weighting hetero.		0.84	0.00	0.16	0.08	0.20	0.88	0.35	0.14	0.92

Note: significance: * 10 per cent, ** 5 per cent, *** 1 per cent. This table replicates Table 7, now distinguishing between the three types of information treatments.

Source: authors' construction based on survey results.