NA310G Lab 3: the agent's analysis, annotated: the key

28 October 2026

As on your sheet: 1–3 and 6 planted, 4, 5 and 7 the agent's own. **red** a problem, **green** done right. Fixes in green inside each box. Line numbers as on your sheet.

A. Its results section (results_section.pdf)

0RESULTS

Data and descriptive statistics

The analysis uses the synthetic NA310G firm panel for 2018–2024. The dependent variable is the natural logarithm of salaries per employee (thousand SEK), defined as salaries divided by the number of employees. Salaries are recorded as negative costs in the accounts, so the sign is reversed. Firm-years with zero or missing employees, 6or with zero or missing salaries, are dropped because salaries per employee cannot be computed for them. Firm-years with missing controls are also dropped. This leaves 9,983 observations on 1,534 firms (9,983 of 10,492, or 95 per cent). The key explanatory variable, ai_firm, equals one if at least one of the firm's job advertisements in the year mentions AI skills. The controls are log employees, log total assets per employee and the equity/asset ratio. Year dummies are included in all models except column (1) of Table 1.

6Mean salary per employee is 557 thousand SEK (median 481, SD 317). AI-advertising firms are observed in 11.3 per cent of firm-years, and this share rises from 8.1 per cent in 2018 to 12.8 per cent in 2024. Firm-years with AI advertisements pay clearly more (mean 661 against 544 thousand SEK). The difference in mean log salaries is 0.213 (t = 13.7). These firms are also much larger, 7averaging 1,575 employees against 671. The raw gap is therefore not evidence of an effect of AI advertising, which motivates the regressions below. Only 376 firms ever advertise for AI skills, and 49 of them do so in every year they are observed. The fixed-effects estimate rests on the 327 firms that switch status .

Pooled OLS, fixed effects and random effects

Table 1 reports the main estimates. 3Standard errors are clustered by firm in all columns.

Table 1. Salary per employee and AI advertising (dependent variable: ln salary per employee)

	(1) Pooled OLS	(2) Pooled OLS + year	(3) Fixed effects	(4) Random effects
ai_firm	0.175***	0.171***	0.031***	0.043***
	(0.025)	(0.025)	(0.007)	(0.008)
ln employees	0.010	0.007	5−0.238***	−0.165***
	(0.010)	(0.010)	(0.005)	(0.007)
ln total assets per employee	0.152***	0.151***	0.040***	0.109***
	(0.006)	(0.006)	(0.005)	(0.005)
Equity/asset ratio	−0.043	−0.043	0.022	0.016
	(0.036)	(0.036)	(0.013)	(0.016)
Year effects	No	Yes	Yes	Yes
Firm effects	No	No	Yes	No
Observations	9,983	9,983	9,983	9,983
Firms	1,534	1,534	1,534	1,534
R-squared	0.344	0.353	0.408 (within)	0.206 (overall)

Notes: Coefficients with 3firm-clustered standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01. Synthetic data.

In pooled OLS, AI-advertising firms pay about 17–18 log points more per employee (roughly 19 per cent), and year effects barely change this. Fixed effects reduce the coefficient to 0.031 (1about 31 per cent), which is still significant at the 1 per cent level. Most of the pooled gap therefore reflects stable differences between firms and not 4changes in pay when a firm starts or stops advertising for AI skills. The random-effects estimate (0.043) lies between the two. Firm-specific unobserved heterogeneity is large: the firm effects account for 92 per cent of the residual variance in the fixed-effects model.

Model choice and diagnostic tests

The F test that all firm effects are zero rejects (F(1533, 8439) = 46.0, p < 0.001). The Breusch–Pagan Lagrange multiplier test 2does not reject the null of no random effect (χ^2 = 19,457, p < 0.001), so pooled OLS remains an acceptable alternative to random effects. The Hausman test rejects random effects in favour of fixed effects ($\chi^2(10)$ = 719.6, p < 0.001). The test uses the non-robust variance matrices because the classical test is not valid with clustered errors. A Mundlak test that adds the firm means of the regressors , which allows clustered errors, gives the same conclusion ($\chi^2(4)$ = 560.2, p < 0.001). The random-effects assumption that firm effects are uncorrelated with the regressors is clearly violated. The fixed-effects estimate of 0.031 is my preferred estimate.

The other diagnostics point to problems that clustered standard errors address. The Breusch–Pagan/Cook–Weisberg test rejects homoskedasticity ($\chi^2(1)$ = 27.9, p < 0.001). The Wooldridge test strongly rejects the absence of first-order serial correlation (F(1, 1479) = 700.2, p < 0.001). The Ramsey RESET test rejects the pooled specification (F(3, 9969) = 16.4, p < 0.001), and the fixed-effects residuals are not normal because of heavy tails (skewness/kurtosis test p < 0.001). With around 1,500 firms, non-normality matters little for inference. Multicollinearity is not a concern (mean VIF 1.44, maximum 1.71).

0 Seven problems. 1–3 and 6 we planted (27 Oct; numbers, tests, code, data), kinds taught in Labs 1–2. 4, 5 and 7 the agent's own. Before planting, all 137 numbers were in its log. And the course rule: the AI ran the analysis and wrote the text, so it could never be handed in.

6 Not done. C shows salary per employee with Min 0 and N 10,017 (text: 9,983): 34 firm-years with employees and no salaries are kept (D, line 56).

6 Includes the 34 zeros, as do the group means in this line. The regressions do not (the log of 0 is missing): two samples, one text.
Fix: Drop them in D, rerun: mean 559 (median 483, SD 316), 664 against 546.

7 A skewed variable read by its mean. Employees: mean 773, SD 1,167, max 31,026 (C); a few giants drive "1,575 against 671".
Fix: Report medians (`summarize, detail`: 435 overall), or compare log employees.

3 Not true for column (3). The FE line in its code (B, line 26) has no `vce(cluster firmID)`.
Fix: Cluster column (3) too (B, line 26); then the sentence is true.

5 Employees on both sides. Outcome: $\ln(\text{salaries/employees}) = \ln \text{salaries} - \ln \text{employees}$, and log employees is a control (assets *per employee* too). A headcount recorded too high lowers pay per head and raises employees in the same row: a negative coefficient with no economics in it. Never explained.
Fix: Explain it in one sentence, or leave the control out (the brief's controls: log employees and log capital).

3 The same claim in the note: false for column (3).

1 A log outcome, read wrong. A dummy's coefficient times 100 is about the per cent: 0.031 is 3.1 per cent ($100(e^{0.031} - 1) = 3.1$). It cannot be 31: that would be above the pooled 19. Line 61 says 3 itself.
Fix: …to 0.031 (about 3 per cent), which …

4 Composition, not a raise. The outcome is an average. A firm that hires a few well-paid AI specialists raises it with nobody paid more; firm data cannot tell the two apart.

2 A test read backwards. Null: no random effect (no firm effects). $p < 0.001$: reject. Firm effects exist, so pooled OLS is out, as the F test in the same line says.
Fix: The Breusch–Pagan LM test rejects the null of no random effect ($\chi^2 = 19{,}457$, $p < 0.001$). Pooled OLS is therefore inappropriate.

Instrumental-variables attempt

ai_firm may be endogenous: better-paying or faster-growing firms may be more likely to recruit for AI skills, and the proxy measures advertising, not adoption. As an instrument I use the AI exposure (DAIOE percentile rank) of the firm's most advertised occupation, merged from the occupation file. This is available only for firm-years with advertisements, which cuts the sample to 7,141 observations on 1,347 firms. The two DAIOE measures in the file are highly correlated (0.95), so only one instrument is used and the model is just identified. Overidentification cannot be tested.

The instrument is relevant in the pooled data. The cluster-robust first-stage F is 63.2 and falls to 34.6 once O*NET controls are added. In the within-firm specification it is irrelevant: the coefficient is 0.0000 with p = 0.87 and F = 0.03. Occupation exposure rarely changes enough within a firm to identify anything.

Table 2. IV estimates, firms with advertisements (dependent variable: ln salary per employee)

	OLS	FE	2SLS	2SLS + O*NET	FE 2SLS
ai_firm	0.158***	0.027***	−0.331*	1.242***	7.282
	(0.026)	(0.009)	(0.178)	(0.277)	(46.384)
First-stage F			63.2	34.6	0.03

Notes: Same controls and clustering as Table 1. The O*NET controls are the cognitive-abilities and social-skills scores of the occupation. * p<0.10, ** p<0.05, *** p<0.01.

The IV results are not credible and I do not use them to draw conclusions. The 2SLS estimate is −0.33, the sign reverses to +1.24 when two occupation characteristics are added, and the within-firm estimate is imprecise (7.3, SE 46). An estimate that changes this much when observable controls are added suggests that the exclusion restriction fails. Occupation exposure should affect pay only through ai_firm, but exposure is itself linked to occupational skill content and pay. Among firm-years with no AI advertisements, higher exposure is associated with lower salaries per employee (coefficient −0.0010, p = 0.001), which is a direct effect that the instrument is not allowed to have. The sample is also selected on advertising. I conclude that the occupation file does not provide a valid instrument here.

Summary

4 AI-advertising firms pay substantially more per employee. Most of the pooled gap (about 17 log points) reflects differences between firms, and the within-firm estimate is a modest 3 per cent (0.031, SE 0.007). Because the Hausman test favours fixed effects, this is my preferred estimate. It still assumes that nothing time-varying and firm-specific, such as a change in business strategy, drives both AI advertising and pay. The instrumental-variables analysis does not settle this. Because the data are synthetic, the numbers describe the data-generating model and say nothing about actual Swedish firms.

4 4 again: “pay more” reads as a raise.
Fix: Within a firm, years with AI ads have about 3% higher average salary per employee: higher pay, or a shift to better-paid staff. These data cannot separate the two.

B. Its regression code (dofiles/4_regressions.do)

```
14 xtset firmID year
15
16 * Standard errors are clustered by firm everywhere: errors are serially correlated within firm.
17 local X ai_firm ln_emp ln_kpe ny_solid
18
19 * (1) Pooled OLS, no year effects
20 regress ln_sal `X', vce(cluster firmID)
21 estimates store m1
22 * (2) Pooled OLS with year effects
23 regress ln_sal `X' i.year, vce(cluster firmID)
24 estimates store m2
25 * (3) Fixed effects (within), year effects
26 xtreg ln_sal `X' i.year, fe
27 estimates store m3
28 * (4) Random effects, year effects
29 xtreg ln_sal `X' i.year, re vce(cluster firmID)
30 estimates store m4
```

3 The comment says "everywhere". Read the code, not the comment.

5 Log employees, and log assets *per employee*: employees on both sides.
Fix: `local X ai_firm ln_emp ln_cap`, and one sentence on what `ln_emp` does.

OK Clustered here, and in columns (2) and (4).

3 **Planted: not clustered**, against the text (A, lines 11 and 31). SE 0.007 instead of 0.008; the gap is small here, but the method described is not the method run.
Fix: `xtreg ln_sal `X' i.year, fe vce(cluster firmID)`

C. Its descriptive statistics (output/tables/descriptives.txt)

Variable	N	Mean	SD	Min	Max
Salary per employee (KSEK) (salary_pe)	10017	557.274	316.674	0.000	3408.022
ln salary per employee (KSEK) (ln_sal)	9983	6.200	0.495	4.108	8.134
AI adoption proxy: 1 if at least one ad mentions AI skills (ai_firm)	10017	0.113	0.316	0.000	1.000
Employees (bslov_antanst)	10017	772.840	1166.637	7.000	31026.000
ln employees (ln_emp)	10017	6.131	0.974	1.946	10.343
ln total assets per employee (ln_kpe)	10017	7.122	1.869	0.334	13.091
Equity/asset ratio (ny_solid)	10017	0.369	0.235	0.000	0.980

6 **Min 0: firms with staff and no pay.** N 10,017 here, 9,983 in the regressions (ln_sal, next line). Look at min, max and N before anything else.

7 **Very skewed:** SD above the mean, max 40 times the mean. Means mislead; use medians or logs.

D. Its sample and variables (dofiles/2_sample_vars.do)

```
* Dependent variable: salary per employee, KSEK. Salaries are recorded negative -> flip the sign.
* Needs employees > 0 and non-missing salaries.
* (Salaries are costs, so after the sign flip every value is zero or positive.)
gen salary_pe = -rr04a_loner / bslov_antanst if bslov_antanst > 0
count if missing(salary_pe)
summarize salary_pe
gen ln_sal = ln(salary_pe) if salary_pe > 0
label variable ln_sal "ln salary per employee (KSEK)"
label variable salary_pe "Salary per employee (KSEK)"

* Controls: size, capital intensity, financial strength. In logs because levels are very skewed.
gen ln_emp  = ln(bslov_antanst) if bslov_antanst > 0
gen ln_kpe  = ln(br09_tillgsu / bslov_antanst) if br09_tillgsu > 0 & bslov_antanst > 0
label variable ln_emp "ln employees"
label variable ln_kpe "ln total assets per employee"
label variable ny_solid "Equity/asset ratio"
* Value added per employee is NOT used as a control: it contains wage costs by construction.

* Winsorising is not used; instead the log specification handles the right skew.
* Sample: non-missing outcome and all controls
gen byte insample = !missing(salary_pe, ln_emp, ln_kpe, ny_solid, ai_firm)
tab insample
keep if insample
display "N analysis sample: " _N
```

6 The missing check. A plain `summarize` shows Min 0 and moves on.
Fix: `count if salary_pe == 0 // 34: staff, no pay`
`drop if salary_pe == 0 // not credible`

6 Keeps rows with salary 0, because `salary_pe` is not missing there.
Fix: `gen byte insample = !missing(ln_sal, ln_emp, ln_kpe, ny_solid, ai_firm)`