

README

Code and Data for “Not Faster, but Better: Generative AI and Teacher Skill in Peru”
by Carla Z. Glave, Carolina Lopez, Ezequiel Molina, and Román Andrés Zárate

Overview

This document provides the information needed to replicate all tables in “Not Faster, but Better: Generative AI and Teacher Skill in Peru”. The replication package is organized into the following folders:

- code/
- data/
- output/
- paper/

Files and Contents

All analysis is conducted in Stata.

- code/00_master.do: Master script. Running this file from top to bottom reproduces all tables in the paper. Before running, set the global repkg at the top of the file to the location of the replication package folder on your computer (see “How to Run” below). Estimated runtime is 5–30 minutes depending on hardware.
- code/01_clean.do: Builds all analysis variables from data/ai4t_merged.dta and saves data/ai4t_clean.dta. Called automatically by the master script.
- code/tables/: Contains one do file per table in the paper (12 scripts). Each script loads ai4t_clean.dta, estimates the corresponding specifications, and writes a LaTeX table to output/tables/.
- data/: Contains the two datasets used by the code. ai4t_merged.dta is the main teacher-level dataset, combining the school and teacher rosters, the baseline survey, weekly Microsoft Copilot usage records, the endline survey, and school-level diagnostic scores. All analysis starts from this file. ai4t_clean.dta is derived from it by 01_clean.do. The individual source datasets are confidential and are not distributed with this package.
- output/tables/: LaTeX table files (.tex) produced by the scripts in code/tables/. Contains 12 files corresponding to Tables 1–5 and A1–A7 in the paper.
- paper/replication.tex: A standalone LaTeX file that compiles all tables from the paper in order. Compile with pdflatex (run twice).

Software Requirements

- Stata 14 or later
- The following user-written Stata packages, which the master script installs automatically via SSC: lassopack, ivreg2, ranktest, pdslasso, texdoc

The tables shipped with the package were generated with Stata 19.0 MP, pdslasso 1.0.03 (package 1.3), rlasso 1.0.11 (lassopack 1.4.3), ivreg2 4.1.12, ranktest 2.0.04, and texdoc 2.4.0. Because pdslasso selects

control variables through the LASSO, borderline sub-treatment contrasts can shift slightly under other package versions; to reproduce the shipped tables exactly, use the versions listed here.

How to Run

1. Open `code/00_master.do` in Stata.
2. Set the path at the top of the file:

```
global repkg "path/to/your/Replication_Package"
```

3. Run `00_master.do`. Tables are saved to `output/tables/`. To produce a PDF with all tables, compile `paper/replication.tex` with `pdflatex` (run twice).

Relevant Sample

The intervention was implemented during the 2025 school year in Metropolitan Lima, Peru. The experimental sample comprises the 390 public schools in Lima that offer sixth grade of primary education (196 assigned to treatment and 194 to control) and their 1,269 sixth-grade teachers (660 treatment, 609 control). Teachers in treatment schools received a one-day in-person training on the use of generative AI for instructional tasks, a Microsoft Copilot license, and weekly WhatsApp follow-up messages over eight months. Two components were cross-randomized within the treatment group: the sender of the follow-up messages (an education specialist from DRELM in the authority arm, a fellow teacher in the peer arm) and their content (general instructional support or teaching-at-the-right-level practices). Randomization was stratified by UGEL (local school district) and by whether the school's baseline math performance was above or below the median; the strata are recorded in `strata_schools` (1–8).

Baseline information comes from administrative rosters and a post-training survey. Program take-up is measured with training attendance records and weekly Copilot license usage reports. Outcomes come from an endline survey collected in December 2025, completed by 607 teachers either in person or online; survey modality (`endline_online`) is controlled for in all survey-based regressions.

Variable Construction and Estimation

All variables are described in the main text of the paper. The main treatment indicator is `treat`; the cross-randomized arms are `treat_auth` and `treat_tar1`, recoded 0/1/2 in the table scripts (0 = control, 1 = first arm, 2 = second arm). Baseline covariates include teacher characteristics (age, gender, university degree, years of experience, additional employment) and school characteristics (enrollment, number of classes, network size, mean fourth-grade math and language scores, share of technology equipment). Missing covariate values are replaced with `-1` and flagged with a missing-value indicator that is added to the covariate vector. The AI tasks index is the sum of nine binary task indicators, standardized to have mean 0 and standard deviation 1 in the control group.

Treatment effects are estimated by double post-LASSO (`pdslasso`), which selects controls from the covariate set; strata fixed effects and the endline-modality indicator are partialled out rather than subjected to selection. Standard errors are clustered at the school level, the unit of treatment assignment. Balance tables use OLS with strata fixed effects and school-level clustering.

De-identification

Anonymous identifiers were assigned during data preparation, after the confidential source files were merged on their original keys (national ID, school code, license e-mail, district code); those keys were then dropped. The anonymous identifiers are `anon_school_id`, `anon_uget`, and `anon_teacher_id`, assigned by random shuffling with fixed seeds (12345, 99999, and 67890, respectively) so that the assignment is reproducible from the source data.

The following were removed from the distributed dataset: teacher and school names, national ID numbers, e-mail addresses and phone numbers, IP addresses and geolocation, open-ended text responses (including the raw prompts, which are retained only as graded scores), exact timestamps, district and locality names, and direct school identifiers. In addition, school-level variables that are not used in the analysis but could be matched against Peru's public school census (exact equipment counts, detailed diagnostic-score summaries, and the school internet indicator) were removed.

Two residual considerations are documented for transparency. First, exact school enrollment and school mean test scores are retained because they are analysis covariates. Second, schools contribute between 1 and 9 study teachers, so within-school cells are small. For these reasons the dataset is distributed under licensed access rather than posted publicly (see the Data Availability Statement).

Manual corrections to individual teacher records were applied before anonymization and are not visible in the replication data. The control group has no Copilot usage data by design; missing Copilot variables for control teachers are set to zero or "No adoptador" in `01_clean.do`.

Data Availability Statement

The data used in this paper are primary data collected by the authors through a randomized controlled trial conducted between March and December 2025 in 390 public schools in Metropolitan Lima, Peru. The datasets were not obtained from third-party sources and cannot be reproduced from publicly available materials.

The replication package includes a de-identified version of the dataset required to reproduce the results. The dataset is publicly available through the World Bank Microdata Library (Licensed Access) at: <https://microdata.worldbank.org/catalog/8514>. The entry is currently under embargo and will become accessible once the paper is published in a peer-reviewed journal.

Detailed information for each dataset:

- File name:** `ai4t_merged.dta`
Description: Teacher-level analysis dataset. Combines the school and teacher rosters, baseline survey, weekly Microsoft Copilot usage records, endline survey, and school-level diagnostic scores for 1,269 sixth-grade teachers across 390 public schools in Metropolitan Lima. All analysis starts from this file; `ai4t_clean.dta` is derived from it by `01_clean.do`.
Citation: Glave, Carla Z., Carolina Lopez, Ezequiel Molina, and Román Andrés Zárate (2026). "Not Faster, but Better: Generative AI and Teacher Skill in Peru." AEA RCT Registry: AEARCTR-0016336.
Access: Public access via World Bank Microdata Library (Licensed Access) at

<https://microdatalib.worldbank.org/index.php/catalog/20061>. Data collection took place March 2025 – December 2025. The entry is currently under embargo and will become accessible once the paper is published in a peer-reviewed journal.

Access restrictions: Licensed access. The dataset has been de-identified as described above; school, district, and teacher identifiers have been replaced with anonymous sequential codes.

List of Exhibits

The following table maps each output file generated by the replication code to the corresponding table in the paper.

Output file	Paper exhibit	Description
output/tables/table1.tex	Table 1	Descriptive statistics and balance
output/tables/table2.tex	Table 2	First stage: take-up of the AI training program
output/tables/table3.tex	Table 3	Time allocation
output/tables/table4.tex	Table 4	Prompt quality
output/tables/table5.tex	Table 5	Prompt scenarios
output/tables/tableA1.tex	Table A1	Balance: authority cross-randomization
output/tables/tableA2.tex	Table A2	Balance: TaRL cross-randomization
output/tables/tableA3.tex	Table A3	Time allocation by cross-randomized arms
output/tables/tableA4.tex	Table A4	Prompt quality by cross-randomized arms
output/tables/tableA5.tex	Table A5	Prompt scenarios by cross-randomized arms
output/tables/tableA6.tex	Table A6	Equity beliefs and willingness to pay
output/tables/tableA7.tex	Table A7	Prompt quality: heterogeneity by age

References

Glave, Carla Z., Carolina Lopez, Ezequiel Molina, and Román Andrés Zárate. “Not Faster, but Better: Generative AI and Teacher Skill in Peru.” AEA RCT Registry: AEARCTR-0016336. IRB Protocol: University of Toronto, #48010.