Review of: "AI Adoption and Firm Demand for Workers and Skills: Insights from Online Job Postings"

Michael Stops[1]

1 Institute for Employment Research, Nürnberg, Germany

Potential competing interests: No potential competing interests to declare.

General remarks

This work sheds (descriptive) light on AI adoption in establishments in Australia by utilising Australian job ads for the years 1996 to 2023.

The paper finds a relatively low share of AI vacancies compared with other countries and, also due to a more restrictive measurement, also compared to other studies for Australia. The paper further finds that AI-adopting firms reveal, in the observation period, a stronger employment growth than other firms, even for occupations that are not strongly exposed to AI. Furthermore, the authors report that the number of required skills increased in occupations that are highly exposed to AI, particularly when they are asked for by AI-adopting firms. The authors conclude that AI tools rather complement human tasks instead of replacing them.

The paper complements the view of AI adoption in the world for the case of Australia. The analyses are based on very granular and highly innovative data for which analysis approaches are currently widely discussed and deliver insights that can hardly or even not be generated with other data.

The following thoughts and questions may be a motivation to further develop the paper:

Definition of AI tools and examples

What are “AI tools [that] were only capable of performing routine tasks and rule-based tasks”? It would be helpful to make a few examples. From my understanding, this description refers more to classic software without any AI features. Particularly, rule-based tasks can normally be implemented with “classic” algorithms.

The state of AI adoption

Currently, the paper seems to assume that the demand for AI tools, particularly in the early years, was related to well-functioning and established AI tools. But I think this was mainly not the case; many firms started to think about the development or implementation of AI tools in an experimental setup with open results. And the firms, like in other

countries, may have reconsidered their needs in order to adapt to AI tools since the publication of ChatGPT 3.5 in November 2022, with the consequence that firms, at least partly, may have stopped their own development efforts in favor of adopting or collecting experiences with the provided Large Language Models by a couple of firms since then.

Australia

For a comprehensive and clear understanding, I miss some information on the Australian economic and labour market development, as well as other relevant time periods like the Covid-19 lockdowns; e.g., in the introduction, the data section, or in the analysis section. E.g., the authors refer to lockdown phases in Australia, but they don't mention the exact period. Considering the lockdown periods in other countries, e.g., in the EU, contradicts the authors' findings (based on Fig.2.).

What I also miss are the authors' arguments for why the European ESCO system mirrors the Australian skill structure better than, e.g., O*NET skill structures. There may be good arguments, and it would be interesting to learn more about those reasons.

Hypotheses

I don't see a clear derivation of the hypotheses from the literature or the theory. This can be solved by a clear assignment of the arguments in theory or from the literature that end with the derived hypotheses. This would be better than listing hypotheses at the end of the text and forcing the reader to find the arguments for each of the hypotheses in the text before.

Methodological aspects

- I don't understand the empirical criteria for a firm to be considered AI-adopting; part of the shift of ambiguous terms from “specific AI terms” to “generic AI terms” does not solve any problems with ambiguity; and also, the requirement that two (ambiguous) terms must be mentioned does not solve the “false positive” problem for job ads that mention two “false positive” terms; maybe this decision is the main reason for different results, e.g., in the OECD study.
- The authors don't mention whether they search for AI terms in the whole job ad texts or just in a part, e.g., where the skills requirements are described.
- I miss information on how the industry is identified in the job ads (see section 2.2.2, where information for the other firm properties is given, i.e., geography and firm size).
- I miss some information on the scatter plot in section 3 (Fig. 4) – currently, I would assume that each point denotes the number of job ads posted by a firm in T1 and T2? If so, this would mean that there are a few firms with 10 thousands of job postings? Some more details on how the plot was generated would be helpful. Related to that, I would miss confidence intervals around the regression lines to get a notion of how strong the differences are. Besides this, I would

consider a robustness check without the extreme values.

- Fig. 5 and Fig. 6 should illustrate a “moderation” effect of AI adopters

Discussion

- Acemoglu et al.[38] don't refer to the concept of AI-adopting firms as it is presented in the current paper. That's why I have concerns that their results are restricted compared to the results in this paper.
- Particularly, the practical implications should refer to the selected observation period in which many firms made their first experience with AI and AI tools based on Large Language Models appeared; it deserves further research if and how AI tools will be established in firms and what the effects are when AI is broadly utilised.

Formal aspects

- I could not follow Figure 1, and I think the reason is that the figure presents the study design AND the findings: how is it possible to derive empirical findings from a figure with bubbles that each represent “a job posting from one of two groups of firms”? It's hard for the reader to see what the authors derive from that figure in the current section 1.3..
- The current structure of the paper raises a lot of questions that are answered quite late in the paper. E.g., section 1 mentions “AI-adopting” firms, “AI-exposed” occupations, and “AI demand” without any definition of what is meant by those terms. This also makes it hard to understand the presentation of the results in 1.3. Maybe examples would help, e.g., for “some formerly non-AI occupations transitioning to become AI-skilled occupations”.
- The references are not always clear; it seems that no publication years are expected. This can be misleading for references of authors that are highly productive, and at least the reviewer has to check whether the authors referred to further studies of the same author and just forgot to mention those studies in the reference list, or they referred to a reference mentioned before (in the worst case, before a lot of previous paragraphs).

Some further minor remarks:

- The first sentence in the current section 1.3 can/should be deleted; maybe the subsections 1.1, 1.2., 1.3 can be dissolved since an introduction does not necessarily rely on subsections.