Received: 17 February 2024 | Accepted: 17 April 2024

DOI: 10.1111/1467-8489.12567

The Australian Journal of Agricultural and Resource Economics | AARES

NOTE

Using artificial intelligence for economic research: An agricultural odyssey

Andrew Leigh

Parliament of Australia, Canberra, Australia

Correspondence.
Email: andrew.leigh.mp@aph.gov.au

Abstract

Generative artificial intelligence tools have been shown to substantially increase productivity in a range of different contexts. I discuss the potential and limitations of the current models, and the evidence on how economic researchers can best make use of generative artificial intelligence in their work. To illustrate these points, I show how the data analysis tools of ChatGPT can be used to address a specific question: the accuracy of agricultural forecasts—and discuss the strengths and weaknesses of artificial intelligence in data cleaning, data analysis and producing graphs and illustrations.

KEYWORDS

agricultural forecasting, artificial intelligence, data analysis

JEL CLASSIFICATION

A11, C45, I23, O33

1 | INTRODUCTION

One of the pleasures of being an economist is analysing real-world problems. Yet, the tools and techniques for conducting applied economic research are changing fast. One of the most interesting developments has been in the area of generative artificial intelligence. Advances in these models over recent years have been remarkable, and they are being deployed in a wide range of situations.

In this article, I discuss some of the opportunities and challenges that artificial intelligence engines present for applied economic research, and how those show up when we use them to analyse the world.

This is vital for agricultural and resource economists because how we conduct economic research will drive our responses to critical issues confronting the Australasian and global

This is an edited version of a speech delivered to the 68th Annual Conference of the Australasian Agricultural and Resource Economics Society at the Australian National University on 7 February 2024. My thanks to Frances Kitt for valuable assistance in crafting the introductory sections, and to editor David Stern and an anonymous referee for valuable comments on an earlier draft. No words in this article were drafted by artificial intelligence, though it did help me tweak the title.

community, such as biosecurity, climate change, environmental degradation and energy system transitions.

The remainder of this article is structured as follows. In Section 2, I review some of the emerging literature about the potential productivity gains and use cases of generative artificial intelligence. In Section 3, I present a simple case study of how artificial intelligence might be used to analyse agricultural forecasting data. The final section concludes.

2 | USING ARTIFICIAL INTELLIGENCE TOOLS

Generative artificial intelligence engines, also known as large language models, include ChatGPT, Bard, Claude and LLaMA. They are powerful artificial intelligence systems which can generate and synthesise text at the tap of a button. They are trained on vast data sets. They are geared to respond to human instructions and deliver the outputs that humans need.

My favourite example of how artificial intelligence can boost productivity comes from a large randomised trial involving 758 BCG management consultants (7 per cent of the company's global workforce), carried out by Fabrizio Dell'Acqua and coauthors (Dell'Acqua et al., 2023). The researchers asked consultants to do various tasks for a fictional shoe company. There were creative tasks ('Propose at least 10 ideas for a new shoe targeting an underserved market or sport'), analytical tasks ('Segment the footwear industry market based on users'), writing and marketing tasks ('Draft a press release marketing copy for your product') and persuasiveness tasks ('Pen an inspirational memo to employees detailing why your product would outshine competitors').

Half the consultants were asked to do the tasks as normal, while the other half were asked to use ChatGPT. Those who were randomly selected to use artificial intelligence were not just better—they were massively better. Consultants using artificial intelligence completed tasks 25 per cent faster and produced results that were 40 per cent higher quality. That's like the kind of difference you might expect to see between a new hire and an experienced staff member. Figure 1 (from Mollick, 2023) illustrates the magnitude of the effect.

The researchers found that artificial intelligence was a skill leveller. Those who scored lowest when their skills were assessed at the start of the experiment experienced the greatest gains when using artificial intelligence. Top performers benefited too, but not by as much. The researchers also found instances in which they deliberately gave tasks to participants that were beyond the frontiers of artificial intelligence. In those instances, people who were randomly

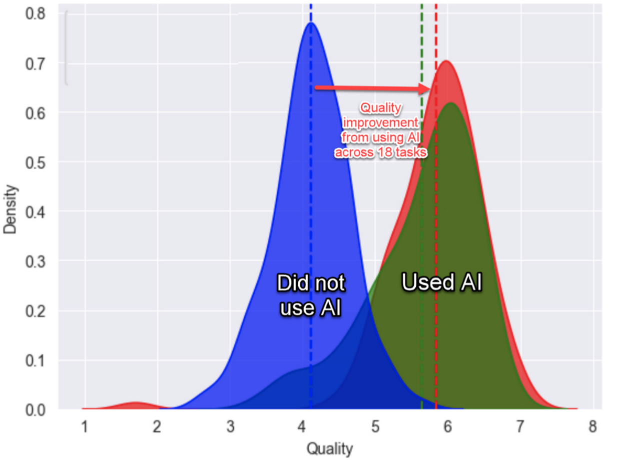

FIGURE 1 Impact of artificial intelligence on the productivity of management consultants.

selected to use artificial intelligence performed worse—a phenomenon that the researchers dubbed 'falling asleep at the wheel'.

Economics researchers can also stand to benefit from the productivity dividends of artificial intelligence models. In an insightful survey paper on using artificial intelligence for economic research, Anton Korinek notes that artificial intelligence can be especially useful for economics researchers at tasks such as synthesising, editing and evaluating text, generating catchy titles and headlines and promotional tweets. Artificial intelligence models are also useful in coding, particularly translating code, and in key steps in data analysis, such as extracting data from text and reformatting data (Korinek, 2023).

In other tasks, artificial intelligence models can also be useful, although they do require more active oversight. Such tasks include providing feedback on text, offering counterarguments, setting up mathematical models, explaining concepts and debugging code (Korinek, 2023).

In a further set of economic research tasks, using artificial intelligence models requires significant oversight. These include deriving equations and literature research (Korinek, 2023). If you've used artificial intelligence to generate academic citations, you'll know the problem here. The current models do not know that an academic citation is a single thing. So they hallucinate, generating made-up citations with believable authors, titles and journal sources, and then adding randomly generated volume and page numbers. Current artificial intelligence models do not just do this for some references—they do it for all references. This will invariably change in coming years as the models evolve. But for now, it's a case of *caveat oeconomista*.

As Korinek suggests: 'A useful model for interacting with large language models is to treat them like an intern who is highly motivated, eager to help, and smart in specific domains, but who has just walked into the job, lacks the context of what you are doing, and is prone to certain types of errors'.

However, if you've just starting using them, the current artificial intelligence models are the worst you'll ever encounter. Coming models will use more computing power and larger training datasets. One indicator of the scale of resources flowing into the sector is that in 2023, venture-capital investors put over US$36 billion into generative artificial intelligence, more than twice as much as the previous year (Scriven, 2023).

The trouble with discussing artificial intelligence and productivity is that the discussion can quickly turn vague. What do we mean by a large language model? How precisely do these tools improve productivity? What exactly does artificial intelligence help us do better? If you haven't been using an artificial intelligence model on your computer or smartphone, then it may not be obvious how they can help you do your job better.

So rather than surveying high-level studies, a better way of making my point might be to demonstrate what the tools can do. Budding writers are often given the advice 'don't tell it, show it'. So that's what I will do in the next section. The specifics of the example don't matter—what I'm aiming to do is to show how artificial intelligence can be deployed to produce better work in less time.

3 | ARTIFICIAL INTELLIGENCE DATA ANALYSIS IN ACTION

So, let's see what an artificial intelligence model can do.

To start, I logged on to the website of the Australian Bureau of Agricultural and Resource Economics and Sciences (ABARES), and downloaded the Historical Agricultural Forecast Database. This is an impressive database, containing tens of thousands of forecasts for different agricultural commodities. I uploaded the Excel file into ChatGPT4 (for which I paid a A$33 monthly subscription fee) and asked it to describe the dataset. It read the data directory and data pages together, quickly told me in plain English what each variable was, and pointed out a few notable features of the dataset, such as the fact that forecast years range from 2000 to 2022, and wheat is the most common commodity forecast.

Providing my commands in plain English, I asked it to drop forecasts for sub-regions, and create variables for the percentage error in the forecast and the forecast horizon. Then, I asked it to plot forecast errors against the forecast horizon. Figure 2, based on 14,622 observations, shows that forecasts get better as the year approaches.

Now let's see how the forecasts compare across commodities. To keep things simple, I'll focus from this point onwards only on forecasts that are made for the next year.

Figure 3 shows the livestock forecasts, which range from an 18 per cent error in mutton forecasts to a 6 per cent error in poultry forecasts. Note that this chart was produced using regular language, not by coding.

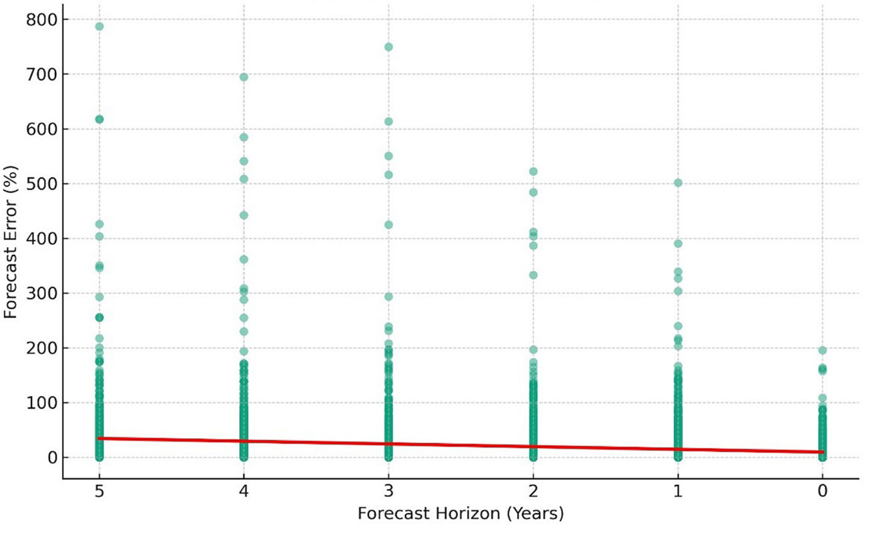

FIGURE 2 Forecast error versus forecast horizon.

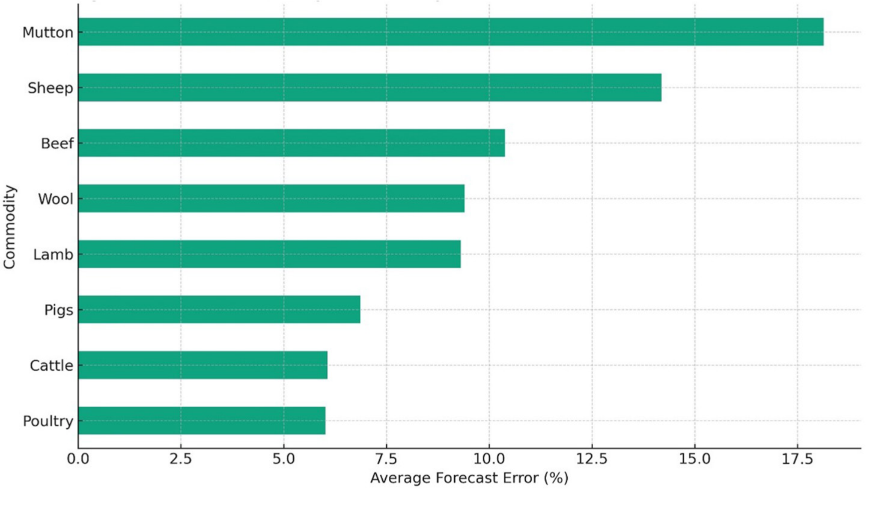

FIGURE 3 Average 1-year forecast error for livestock commodities.

It appears that poultry forecasters are more accurate than mutton forecasters. Perhaps you might forget that fact. So let's ask ChatGPT to make a digital art image that helps it stick in your mind. The result is Figure 4.

FIGURE 4 Poultry forecasts are more accurate than mutton forecasts.

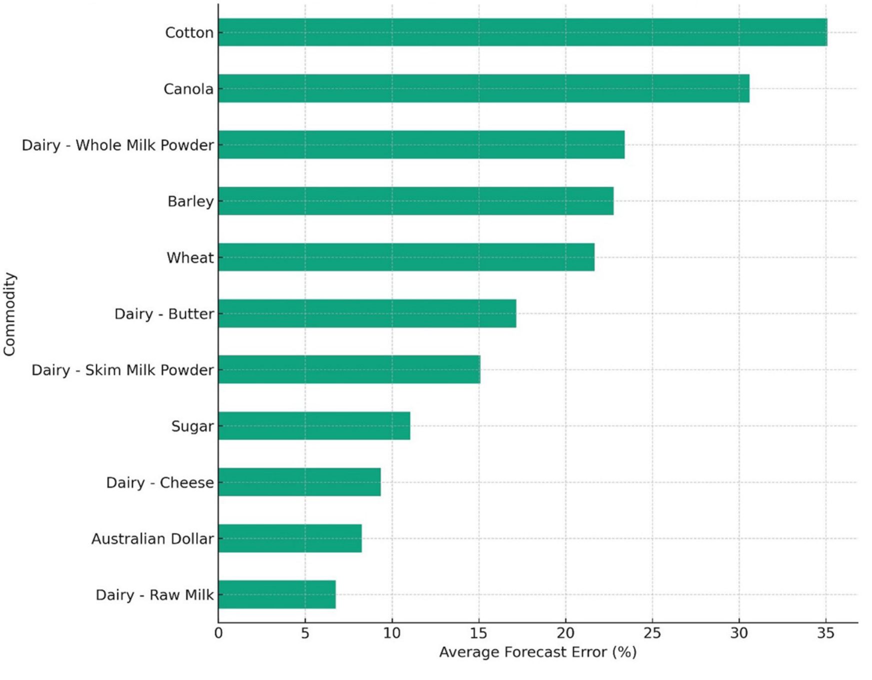

FIGURE 5 Average 1-year forecast error for crop commodities.

Next, I turn to look at crop forecasts. In doing so, it is worth noting that when I asked ChatGPT to produce graphs for crops and livestock, I didn't tell it which was which. I left it to the artificial intelligence model to figure out whether poultry was crop or livestock. Basically, it got it right.

Figure 5 shows crop forecasts.

As the chart shows, the average errors for crops tend to be larger than for livestock. For example, the average error in 1-year ahead cotton forecasts is around 35 per cent. It is also worth noting that ChatGPT has mistakenly decided that dairy is a crop, as is the Australian dollar. Accepting that classification, Figure 5 shows that forecasts for raw milk have an average forecast error of just 6 per cent.

Might readers like a way to remember that dairy forecasts are more accurate than cotton forecasts? Figure 6 depicts an apt artificial intelligence image, generated in the style of anime.

How are forecasts changing over time? One of the things that artificial intelligence does especially well is to run the same analysis repeatedly for different subsets of the data. So I asked it to look at each commodity in turn, and see which set of forecasts showed the greatest improvement over time, and which showed the largest deterioration. I then asked it to draw two graphs—one showing the commodity where forecasts are getting better at the most rapid rate, and the other showing the commodity where forecasts are deteriorating most rapidly.

As you can see from Figure 7, forecast errors for cotton are getting bigger, while forecast errors for wool are getting smaller. Since bigger errors are bad, that suggests that cotton forecasts are going downhill, while wool forecasts are improving.

A busy economist might easily forget that result, so Figure 8 shows another artificial intelligence graphic that summarises the research result.

Poor cotton.

FIGURE 6 Dairy forecasts are more accurate than cotton forecasts.

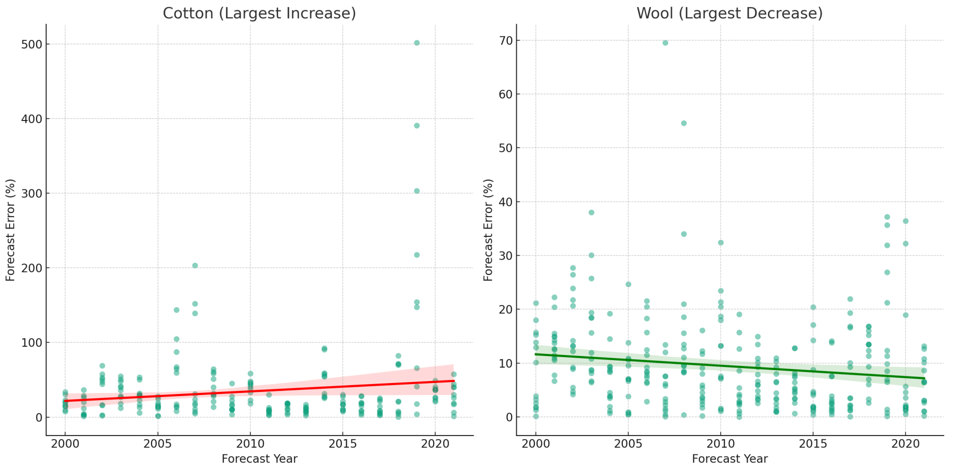

FIGURE 7 Changes in forecast errors over time for cotton and wool.

FIGURE 8 Wool forecasts are improving and cotton forecasts are worsening.

4 | CONCLUSION

What can we learn from this cute exercise? First, artificial intelligence can be helpful at all stages of the research process. The models are quick to understand the features of a dataset and easy to program.

Second, artificial intelligence doesn't take away the need for a bit of training. When I first looked at the forecast dataset, I asked the engine to simply tell me the interesting patterns that it saw. It knew that the forecast variables and the actual variables were related, but it didn't immediately figure out that it needed to look at the absolute percentage difference to have a meaningful benchmark. It was also happy to combine forecasts from different time horizons, making comparisons across commodities meaningless. Because my analysis took only a few hours, there are probably other coding errors that I've made, which an experienced user of the database would avoid. Artificial intelligence is no substitute for knowing the data, and understanding the kind of exercise you're trying to perform.

In addition, artificial intelligence models currently lack the precision needed for rigorous, replicable research. As a referee noted when reviewing this paper, 'I agree that AI can do a better job in data description, but I am not sure whether it can do very sophisticated statistical analysis'. It is still too early for researchers to uninstall their copies of Stata, R, MATLAB or SPSS. The safest approach for a major research project would be to use artificial intelligence data analysis tools to explore the interrelationships in the data, but then carry out the actual analysis using a traditional software package.

Third, artificial intelligence is great at graphs. When I'm coding in Stata, I need to specify every detail of the labels and style of the graphs. By contrast, ChatGPT often made judgement calls that accorded with what seemed most reasonable. Its graphs popped out quickly and were almost ready to insert into a presentation.

Fourth, artificial intelligence is useful for creating images. If you are looking to liven up a dry presentation, but your artistic abilities are as elusive as a unicorn holding a paintbrush, then artificial intelligence can help. In the past, it has been difficult to find stock images that fit the specifics of an economics presentation, but the ability of artificial intelligence engines to create bespoke imagery opens up new possibilities, particularly when presenting to generalist audiences.

Fifth, artificial intelligence makes it easy to look at different aspects of the data. One of the skills of a good applied economist is to know your data and the patterns within it. By making analysis quicker, the researcher can spend more time understanding quirks and relationships that might have been missed.

For researchers, artificial intelligence has the potential to democratise data analysis. Just as translation apps make it easier to have a rudimentary exchange with someone who does not speak your language, so too artificial intelligence data analysis tools allow someone without coding skills to explore the patterns in their data. From data analysts to digital artists, artificial intelligence is reshaping our world.

FUNDING INFORMATION

No specific funding was received for conducting this research.

DATA AVAILABILITY STATEMENT

All data analysis in this study was based on the ABARES Historical Agricultural Forecast Database (downloaded on 28 January 2024), which is available at https://www.agriculture.gov.au/abares/research-topics/agricultural-outlook/historical-forecasts. My log from ChatGPT4 is available here.

ORCID

Andrew Leigh https://orcid.org/0000-0002-5639-0509

REFERENCES

Dell'Acqua, F., McFowland, E., Mollick, E.R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S. et al. (2023) *Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality*. Harvard Business School Technology and Operations Management Unit Working Paper 24-013, Harvard Business School, Cambridge, MA.

Korinek, A. (2023) Generative AI for economic research: use cases and implications for economists. *Journal of Economic Literature*, 61(4), 1281–1317.
Mollick, E. (2023) *Centaurs and cyborgs on the jagged frontier*, Blog post, 16 September. Available from: https://www.oneusefulthing.org/p/centaurs-and-cyborgs-on-the-jagged
Scriven, G. (2023) Generative AI will go mainstream in 2024, *The Economist*, 13 November.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Leigh, A. (2024) Using artificial intelligence for economic research: An agricultural odyssey. *Australian Journal of Agricultural and Resource Economics*, 00, 1–9. Available from: https://doi.org/10.1111/1467-8489.12567