

Some Thoughts on AI and Research

Isaiah Andrews

April 3, 2026

AI model capabilities have been advancing very quickly (e.g., ChatGPT was released less than 4 years ago). While there are many questions about the social and economic impacts of these tools, it seems clear they will change the production function for research in economics, and for the broader set of jobs available for PhD economists inside and outside of academia. While I can't claim any deeper insight than anyone else into where this is going, I think it's worth thinking about and discussing. I originally wrote these notes for my PhD student advisees at MIT Economics, but am sharing them more widely in case others find them useful.

Because I'm very unsure where model capabilities will end up, I've found it helpful to think in terms of cases.

Case 1: Models capable of doing all intellectual tasks better than humans

- In this case, human capital investments seem likely to be irrelevant.
- Consequently, it doesn't seem like this case should play a big role in your current research and skill acquisition decisions, conditional on still focusing on research and skill acquisition for the time being, since all investments you make now will have a low return in this state.
- One response I've seen advocated in response to concern about this case is to switch from focusing on human capital acquisition to focusing on earning as much as possible now.

Case 2: Models get much better at what the current generations of models are OK at, but are still below "good" human performance in other areas

- What precisely it'd makes sense to invest in now seems to depend a lot on what the models end up being "bad" at.
- In particular, I'd expect returns to be high for whatever set of skills (a) are scarce and (b) the models are bad at.
- To the extent that anything complementary with these skills (e.g. next-generation model access) is also scarce, it would also be quite valuable.

To make this case more concrete, suppose the models end up being "bad" at taste, judgement, and problem selection.

- This is largely true of the current generation of models, in my view.
- But suppose they get "great" in the areas where they are currently "OK," e.g. coding, writing, proofs.
- In this case, there would be high returns to developing good taste and skill at problem selection.
- Also in this case, empirical research would change quite a lot, and the private and social returns for methodological research will come from addressing the issues people

currently have, not the ones they used to have, the ones that are discussed in the literature, etc.

Case 3: Model capabilities level off modestly above their current level

- Capabilities available right now dramatically change the production possibility frontier for both theoretical and applied research in economics.
- As these get incorporated into people's workflow, I expect we will see changes in both the types and volume of research produced.
- For instance, I've found that GPT 5.4 Pro is substantially better at convex analysis proofs than I am.
- Even assuming little or no further progress in model capabilities, I think failing to figure out how to take full advantage of these tools would leave you at a major disadvantage...
- ... However, given how fast model capabilities have been advancing and current investment levels, I also find this case implausible.

Recommendations

I'll condition on your wanting to keep focused on finishing your PhD, purely because I'm even less qualified to give advice in the other case. These recommendations are, I think, valid across Cases 2 and 3 above (and irrelevant in Case 1).

You should be actively thinking about, and planning around, this.

- You're currently in the middle of making an expensive (in opportunity cost), long-term human capital investment.
- A change in the production possibilities for research (whether from AI or anything else) changes the expected return for specific investments you can make now, which should probably also change what you invest in.
- Obviously you'll be making decisions without knowing what's going to happen, but you're also PhD students in economics, and decision under uncertainty is not a new problem.
- In this note I've tried to lay out a bit of how I'm thinking about this, but I have no special claim to knowledge here and you should be thinking about this too.

You should be learning how to use these tools effectively.

Hiring, publication, and tenure standards are set in equilibrium. If the production possibility frontier expands and you do not keep up, you will lose out in an absolute sense (i.e. having worse outcomes than had the new technology not been available). "No one uses these models for research" is not an equilibrium, and moreover would be bad if you think research has social value.

There are several possible dimensions to this:

- *Experimentation.* Returns to experimentation are high: these are new tools and are evolving quickly, so I doubt anyone has "solved" the problem of how to best use them (for research or for any other tasks). My current sense is that I'm using these tools more aggressively than many PhD students are. Given that I'm (a) a tenured professor and (b) kinda old, this suggests to me that many students are currently under-investing in exploring what these tools can do for you. On a related note, the better models (often accessible only through institutional or individual subscriptions, not the open web) are often much, much better than the free versions. It's more than worth the trouble of getting access to at least see if you get value from them.
- *Verification.* The models produce output that tends to look/sound plausible but (at least for now) can be wrong in ways that require expertise to detect. Learning when and how to audit model output is a core skill for working with these tools. Note that if Case 2 happens, this could mean that the returns to have good taste, good instincts on problem framing and selection, etc., could go up rather than down.
- *Division of labor.* AI changes the production function not just for solo work but for how you structure collaboration. It can substitute for some tasks you might otherwise need an RA for, change the optimal allocation of effort across coauthors, and make it feasible to take on projects you otherwise wouldn't. Thinking carefully about how these tools fit into the way you organize your research is as important as learning the tools themselves.

You should closely follow what people are currently doing in applied research

- Even more concretely, from year 3 on you should be attending at least 1 applied lunch + 1 applied seminar every week.
- This is a good idea for many reasons, and something I would have recommended regardless.
- However, I think the expected return from doing this will be even higher going forward, as we will likely see substantial changes in both how empirical research is produced and what sorts of research gets done.
- In both Cases 2 and 3 above, there will continue to be a high return for the ability to communicate clearly and effectively with applied researchers, and to understand what problems empirical research in economics is actually running into.
- Regular engagement with applied research and applied talks is, in my experience, the main way econometrics students get good at this.
- Note that this is in addition to weekly attendance at the econometrics lunch + seminar.