

Kevin Byran

AI for Economics Education

On Thursday, December 12, Kevin Bryan joined Markus' Academy for a conversation on "AI for Economics Education." Kevin Bryan is an Associate Professor at the University of Toronto.

A few highlights from the discussion.¹

- **A summary in three bullets**
 - Education-focused AI can address three key challenges: (1) personalizing education for students, (2) personalizing support for faculty, and (3) ensuring AI tools are used effectively by students
 - Raw models like ChatGPT can be counterproductive, as they may enable cheating and produce hallucinations. Crucially, they can provide answers misaligned with professors' teaching priorities, particularly in advanced courses. The goal should be to structure LLMs so that they are aligned with a professor's particular class
 - Restricted models like AlldayTA will solve the three challenges by automating tasks such as creating question banks, grading exams, and tracking student progress, enabling professors and course assistants to concentrate on planning courses and developing content
- **[0:00] Markus' introduction**
 - The first wave of education democratization came with Massive Open Online Courses (MOOCs), such as Khan Academy and Coursera. These platforms were massive in reach but kept students in a passive role
 - With AI we have the opportunity to individually tailor the educational experience and make it interactive. In a recent MA webinar, [Justin Wolfers](#) showed how LLMs enable personalization at scale
 - However, questions remain about whether AI-based education can foster social interaction among students and achieve the "gamification" needed to overcome students' commitment challenges
- **[4:23] Pros and cons of AI in education**
 - Education-focused AI can address three key challenges: (1) personalizing education for students, (2) personalizing support for faculty, and (3) ensuring AI tools are used effectively by students.
 - The traditional approach to higher level education is: (1) professors plan the course and lectures, (2) standardized readings, homeworks and exams are given to all students, (3) students pass courses with a potentially subpar performance (e.g., a "7/10")
 - The only step in this process that is not obsolete is the first: course planning

¹ Summary produced by Pablo Balsinde (PhD student, Stockholm School of Economics)

- The education literature has shown the value of mastery learning, where students should master a topic before moving on to the next, but implementing this approach is expensive
 - An education-focused AI model could make mastery learning feasible by tracking each student's progress and generating personalized assignments. The key problem is that generic "raw" chatbots (like ChatGPT) can be counterproductive
 - In a field experiment by Bastani et al. (2024) students with access to GPT-4 did 18% worse on assignments than the control group. In another experiment in Kenya, Otis (2024) found that entrepreneurs using AI saw 15% less revenue because AI usage crowded out social learning, such as engaging with local business mentors
 - There are several reasons for why general-purpose AI models can lead to counterproductive educational outcomes
 - They may create a multitasking problem, where using AI takes time away from other more valuable tasks like studying
 - Learning can also suffer if a professor's priorities differ from those in the model, which represents the "average" knowledge from the general corpus of written language. You can think of students using ChatGPT as them going to the local library and randomly picking a textbook to study for your class
 - The goal then is to structure LLMs in such a way that they match what professors want their students to learn, enabling professors to focus on planning the course and developing content
- **[15:01] A path to Bloom's Two Sigma**
 - Bloom (1984) established the "Two sigma problem." He argued that individualized tutoring for the same duration as a lecture yielded learning gains of 2 standard deviations. The challenge lies in developing cost-efficient tools that can replicate the effectiveness of one-on-one tutoring
 - We are heading towards education- and tutorial-focused AI models (rather than foundation models) that can achieve individualized tracking, expand course-specific materials like question banks, and grade exams
 - The key problem is that we need to pull students away from ChatGPT. Unconstrained AI tools can produce incorrect outputs; it is like having students using an arbitrary textbook where one in every 20 pages contains incorrect proofs
 - Additionally, raw AI tools often provide responses misaligned with a professor's priorities, and it is unrealistic to expect professors to monitor and adapt to the evolving accuracy of such general models in teaching
 - **[23:28] The free alternatives to raw models**
 - In his seminar [Justin Wolfers](#) explained how "GPTs" can be built using class materials, but highlighted that they may not promote long-term learning as they still provide students with direct solutions when asked to do so. They will also not be individualized enough for the particular class
 - NotebookLM is similar. It addresses the problem of aligning the model with what the professor is trying to teach, but it lacks the ability to restrict how the

content is used. It is better suited for individuals working on long-term projects who need to query and organize their materials.

- Many are already spending millions to build custom models to teach general classes like high school math, for example Khanmigo or Duoling Max. However the economies of scale will not be there to do this for advanced college courses
- As an example of what can be achieved with some enhancements to the API and the user experience, Bryan developed a small [model](#) to summarize economics papers. The output follows a predefined structure, highlighting key points such as the paper's findings, assumptions, and data

- **[34:44] AlldayTA**

- Together with Joshua Gans, Bryan has developed [AlldayTA](#): an AI teaching assistant fully personalized to a given class which is restricted to a "tutorial mode" to prevent cheating
- The tool enables students to go beyond the documents and lectures provided while still adhering to the professor's teaching priorities
- It can also assist professors by creating practice exams and highlighting topics where students struggle or frequently ask questions to the model.
- It can also perform other tasks for professors like creating practice exams, and giving them a sense of the topics that students are having difficulties with and asking more questions to the model
- It can also perform other tasks for professors, like creating practice exams and providing insights into the topics students are struggling with and frequently asking the model about
- To prevent hallucinations, the AI cites the class document most relevant to the content it generated. Yet once the assistant has been trained based on the class materials these are deleted and are unaccessible to AlldayTA
- AlldayTA balances affordability and specialization, avoiding having to create a highly specialized but expensive custom model with Khanmigo while also not just relying on raw ChatGPT. It costs 1\$ per student per term
- With these new education-focused models professors will be able to concentrate on course planning and delivering lectures. AI is unlikely to replace professors and textbooks, as they make key decisions on prioritizing content and determining the order of material presentation

- **[51:45] Q&A**

- We saw that MOOCs were more effective with quantitative classes. It is likely that AI assistants will be best in qualitative classes: AI excels at handling conceptual understanding, reframing concepts, and handling unstructured information
- It is unlikely that these educational models will take an audio form. Taking text to audio is still expensive, and today is only worth it in other settings outside of education like answering (non-urgent) 911 calls
- MOOCs were about scaling and teaching in the developing world. In a few years AI will be fully integrated into all classrooms, and will revolutionize teaching in a way that hasn't happened in 300 years. Today we still give college lectures the way David Hume taught

Timestamps:

[0:00] Markus' introduction

[4:23] Pros and cons of AI in education

[15:01] A path to Bloom's Two Sigma

[23:28] The free alternatives to raw models

[34:44] AlldayTA

[51:45] Q&A