Day 5 - Summary

[JULY 31, 2026]

[Saieh Hall, Room 146]
[5757 S. University Avenue, Chicago, IL 60637]

Lecture Summaries

Week 1, Day 5 Sessions

Session 1 | 9:00 AM – 10:30 AM

Lecture Summary – Artificial Intelligence, Innovation, and Adjacencies

In this lecture, Prof. James Evans examined how artificial intelligence might reshape the pathways along which discovery and production actually move. He began from the observation that knowing something is a collective production: researchers rarely gather their own data, build their own theories, generate their own predictions, and act on them, but instead trade with others at every step. The relevant unit is therefore often the national innovation system, and countries do not jump from capabilities they have to capabilities they lack. They move through spaces of adjacency. Much of his work builds data-driven generative models, or digital doubles, of individuals, laboratories, fields, and whole innovation systems, on the reasoning that such models are needed wherever first-principles models are insufficient, which is always the case in the social sciences and true even in chemistry once molecules become complex.

These models can be built for different purposes. A digital twin reproduces a system as it is and supports conditional prediction, which is what leadership requires: knowing what follows if a particular intervention is made. But simulation can also explore systems that do not exist, which Prof. Evans divided playfully into cultural aliens, which are unfamiliar but intelligible and buildable, and cognitive aliens, which are black boxes resistant to explanation. Which is acceptable depends on the domain: for the origins of language or the universe an inaccessible model defeats the purpose, whereas in biomedical research a model may be worth having if it demonstrably improves outcomes, as AlphaFold predicts protein structure without anyone knowing quite how. He noted

that large models are now increasingly built by other large models, so that much of the work at leading AI conferences has become mechanistic interpretability, the attempt to understand systems we have produced but not designed.

Examining millions of reasoning traces, his group found that models trained to reason do not simply talk themselves through a problem but spin up internal conversations: non-reasoning models average a single persona, reasoning models about three, and amplifying an internal feature that marks surprise increases the number of personas and doubles performance. Large models, he argued, behave less like individual agents than like complex institutions, and are best understood as cultural and social technologies. This has consequences. Dialogue with AI has been shown to reduce conspiracy beliefs durably, and moderator agents can lower polarization, yet models left to converse on average polarize more than people do. Models tuned to be helpful, friendly, and open acquired a systematic political skew that no one designed. And when people consult a single balanced assistant on genuinely novel problems, they surrender cognitive agency, perform worse, and feel better, whereas deliberately opposed models make them engage more, perform better, and feel worse.

The final section turned to adjacencies directly. Embedding the keywords, methods, and problems of science in a high-dimensional space and letting it evolve by purely local random motion predicts the following year's discoveries with roughly ninety-five percent discrimination, because search is local: researchers bump into what is near them, constrained by the equipment in their laboratory and the components in their supply chain. Adding people to the space predicts who will make a given materials discovery with a precision near one half, and the predictions sharpen at the level of universities and countries, so that whether a discovery occurs in the United States, Germany, or China follows from the adjacencies available there. Building a system that makes adjacencies possible therefore matters more than recruiting outstanding individuals. Rare improbable jumps produce the most consequential work, and distant discoveries typically start below the nearby state

of the art but have far higher ceilings. AI can help rewire these pathways: simulating routes that cross institutionalized boundaries between fields, and between the academy and industry, improved the discovery of energetic materials and the repurposing of drugs, with the effect confirmed experimentally in work with Argonne National Laboratory.

Session 2 | 11:00 AM – 12:30 PM

Lecture Summary – AI for Development: The Case of Digital Agriculture

In this lecture, Nobel Laureate Michael Kremer argued that rather than dwelling on the risks of artificial intelligence, there is an immediate opportunity to invest in realizing its potential for development. Frontier laboratories are producing remarkable capabilities, but capturing their benefits in low and middle income countries requires complementary investments that firms have little commercial incentive to make, which leaves a clear role for public policy. He illustrated the argument through digital agriculture, and specifically AI weather forecasts for farmers, drawing on the work of the Development Innovation Lab, an innovation commission whose recommendations led to the Aim for Scale initiative, and the Human-Centered Weather Forecast Initiative at the University of Chicago. The underlying decision is when to plant. In rain-fed tropical agriculture, planting on a false start destroys both seed and fertilizer, so farmers wait to be certain and sow later than agronomists would advise, and climate change has eroded the traditional rules of thumb they once relied on.

The evidence that farmers act on forecasts is strong. Half a dozen randomized trials and several careful observational studies find behavioral responses: farmers in Pakistan adjusted pesticide application and irrigation ahead of heavy rain, forecasts of monsoon onset in India changed not only timing but crop choice and total agricultural investment, and flood forecasts in Bihar improved household preparedness while anticipatory transfers in Bangladesh and Nepal allowed families to protect assets before disaster rather than receive relief afterward. Estimated benefits in Benin ran between one hundred and three hundred fifty dollars per farmer per year, and gains were larger in

unusual years than normal ones, which is what one should expect. Since dissemination through government systems costs pennies per farmer, the benefit-cost ratios are very large.

What has changed is the underlying science. Open-access global AI weather models from Google, NVIDIA, and the European Centre for Medium-Range Weather Forecasts are accurate for many phenomena, have extended useful lead times, and run roughly eighty thousand times faster than the supercomputer simulations they replace, which makes local tailoring feasible. They are not optimized for smallholders, however, because the commercial return is small. Testing these models against an agriculturally relevant definition of monsoon onset, the team found that blending two AI models with a long history of Indian rain gauge data outperformed any single model and forecast onset up to four weeks ahead. Katie Kowal, a weather scientist leading the parallel work with the Ethiopian Weather Service, described what makes such forecasts usable: relationships that give access to the observational data a government trusts, calibration to local conditions, and conversion from model grid cells into the administrative units in which farmers actually live.

Delivery matters as much as accuracy. Farmers initially read a seventy percent likelihood of onset as a prediction of seventy millimeters of rain, and simply rephrasing it as seventy out of a hundred sharply raised comprehension, while simpler message design lifted it from twenty-seven to forty percent. In 2026 India's agriculture ministry disseminated these forecasts across sixteen states and more than three thousand sub-districts, reaching some fifty million farmers, and multilateral development banks have followed, with the Asian Development Bank committing three hundred million dollars and the Inter-American Development Bank integrating weather into its digital agriculture portfolio. Prof. Kremer cautioned that competitions scored on prediction alone invite overfitting when outcomes are rare, and that procurement rules weighted toward price serve software badly. The general lesson, which he extended to education technology and to accountability in public services, is that the technology is necessary but not sufficient: realizing its value requires testing, dissemination, and integration into government systems.

Session 3 | 1:30 PM – 3:00 PM

Session Summary – Participant Country Presentations

The final session was given over to participant presentations. Pervin Dadashova, Director of the Financial Stability Department at the National Bank of Ukraine, described economic management under full-scale war, the role of inflation targeting in anchoring expectations and support instruments ranging from subsidized loans to a special tax and legal regime for the defense industry.

Fabio Terra, Chief of Staff at Brazil's Ministry of Finance, talked about the assets Brazil can build from, including an almost entirely renewable and nationally integrated electricity system, large reserves of critical minerals, and a world-leading agricultural sector.

The program concluded on a lighter note as Robin Dolezal and Enrico Stivella of the Growth Academy shared fascinating stories about Chicago and the University of Chicago, spanning history, scientific breakthroughs, and iconic moments in film. A fitting finale was the reminder that *When Harry Met Sally* begins with two students meeting at the University of Chicago before embarking on a lifelong journey together—a warm and symbolic ending to a week defined not only by new ideas, but also by the friendships and professional connections formed along the way.