萨金特教授讲座-2025 年 11 月 18 日

问： 您认为美联储独立性的降低会在多大程度上影响美国的长期增长？这其中的潜在含义是什么？如果他们失去制定货币政策而不受政府影响的独立性，会带来什么后果？

答： 我们可以回顾历史。美国在其建国后的 125 多年里都没有中央银行。当时由国会和总统直接管理货币政策，但他们不喜欢这样做，因为这常常需要为了稳定和低通胀而做出不受欢迎的决定。他们创建美联储，正是为了下放这一权力，建立一个可以因艰难决策而受到指责的机构。这种授权本身也是对国会的一种约束。例如，如果政客希望降低利率但美联储拒绝，政府可能就不得不提高税收或进行其他困难的财政调整。

如果一位总统直接控制了美联储，并迫使其追求政治上权宜但对经济不明智的政策（例如过度印钞），结果很可能就是更高的通货膨胀。关键在于，总统届时将无法再为这些负面后果责怪他人。这就破坏了该机构存在的根本目的。关键在于要思考其中的激励和约束体系。美联储的独立性是针对"时间不一致性问题"的一种制度性解决方案，削弱它则可能使政策环境回归到更不可预测、更不稳定的状态。

问： 作为一名学生，我想知道如何开始一个研究问题。我该如何找到一个要研究的问题？另外，AI 工具将如何影响我们的研究方法？

答： 这是我们所有人都经常面临的一个根本问题：如何分配我们有限的时间和智力精力。是专注于像 AI 这样的"热门话题"，还是专注于基础原理，这是一个真实的困境。我的方法是个人化的。我倾向于研究基础的、永恒的原理。我书架上经常反复翻阅的书通常都有几十年的历史，它们处理的数据集很小，但却精彩地阐释了普适的原理。这些原理——比如激励、约束和预期——适用于任何国家和任何时代。

关于 AI，你必须问自己一个关键的元问题：你的问题是你自己想出来的，还是你依赖 AI 提出的？这里没有道德评判，但这对于自我认知至关重要。危险在于，AI 可能成为你自身思考的替代品，而不是补充。我在教学中见过；完全依赖外部工具而不去钻研基础材料的学生，最终什么也学不到。AI 的核心是曲线拟合和统计学习——这些正是我们所教授的基础知识的一部分。要想受益，你必须让 AI 成为你技能的*补充*，而不是*替代品*。它应该增强你的思考能力，而不是控制思考过程本身。目标是利用这个工具来加深你对基本原理的理解，而不仅仅是生成输出。

问： 发展中经济体的政府应如何逐步建立公信力？在这些经济体中，家庭和公司正在学习过程中，可能不信任官方的模型和预测。中央银行和财政部在理解真实家庭和公司如何形成预期方面，通常最可能犯什么错误？

答： 这触及了预期、公信力和声誉的核心。让我们仔细定义这些术语。"声誉"不是你内在拥有的东西；它存在于他人的心中。是他人对于你在各种情况下会采取什么行动的信念或预期。你的声誉是由你的历史行为塑造的——人们处理这些"数据"来形成对你的看法。

考虑一个简单的例子：我向我妻子承诺晚上 6 点回家帮忙做晚饭。5 点半时，一个朋友邀请我去酒吧。我面临一个选择：遵守承诺，或者为了即时享乐而违约。如果我违约了，我妻子会更新她对我可信度的看法。关键在于，我*知道*她在进行这种数据处理。所以，我的决定是在预期其会如何影响我在她心中的声誉的前提下做出的。这就创造了一种动态互动。

这就是公信力背后的纯理论。它关乎一个具有均衡态的多期博弈。政府和中央银行与公众恰恰就处于这种博弈中。要建立公信力，他们不仅必须做出承诺，还要履行承诺，并且要明白公众一直在观察并更新他们的信念。违背承诺的后果是，公众将相应地调整其预期和行为，这使得未来的政策效果下降。官员们常常错在低估了公众从过去行动中学习的能力。这无关复杂的模型，而在于与既定目标一致的、始终如一的、可预测的行为。

问： 随着人们逐渐习惯使用 AI 进行决策，您的理性预期理论原本考虑的是人类的理性判断。这一理论应如何进行调整？最终可能形成什么样的稳定状态？

答： 这个问题与之前关于 AI 的问题紧密相关。当你问我这些问题而我回答时，我本质上是在执行一种人类智能的版本。我在倾听，将相似的问题归类，并构思回答——这也是 AI 所做的事情，只不过是在机器内部进行。AI 的基础是统计学：它关乎选择一类函数（例如，直线、多项式）以及一种将它们拟合到数据的方法（例如，最小二乘法）。

理性预期是一个关于人们如何有效利用*信息*的假说。信息论中（由图灵和香农提出）对"新息"的核心定义是一个*意外*——在给定可用信息的情况下无法预测的东西。如果 AI 仅仅是帮助人们更快、更准确地处理现有信息，它可能会使预期在传统意义上*更加理性*，因为人们（或他们使用的 AI）变得更擅长预测。

然而，如果 AI 开始产生新的策略，或者创造出一个复杂的、相互作用的 AI 智能体系统，那么"信息集"的性质以及学习过程可能会发生根本性的改变。稳定状态将是这个人机混合系统中的一种新均衡。理论不会被抛弃，但需要被扩展，以模拟这种更复杂的、多层次的学习和预期形成过程。关键在于对这个系统进行建模，包括人类和 AI 智能体，以及它们相互交织的信念和学习规则。

问： 如果我们想在几年后开始职业生涯或创业，在此期间我们应培养什么样的思维模式来面对未来的挑战？我们如何才能将所学的经济学和逻辑与真实的商业世界最好地结合起来？

答： 你所需要的思维模式正是我们一直在讨论的：经济学的思维方式。伟大的经济学家常说，其核心并不复杂。它就是关于理解激励和约束。在社会领域，这包括人们相互之间的*信念*和*预期*这股强大的力量，它们创造了声誉和公信力。

你不一定需要学习海量的事实，但你需要学会如何*思考*。学习那些基础的、普适的原理。帮助你理解我关于承诺回家那个简单例子的原理，同样能帮助你分析中央银行的公信力、国际贸易协定或社会安全网的稳定性。

当你面临一个商业问题或政策问题时，要将其简化。激励是什么？约束是什么？不同参与者的信念和预期是什么？他们的行动如何在均衡中相互作用？这个框架非常强大且普遍适用。你所学的经济学提供了系统分析这些情况的逻辑结构，使你能够超越轶事或肤浅的解释。这种分析的严谨性将使你在商业世界中获得优势。

问：你认为下一项重大的创新会是什么？
答：在回答之前，我反问你一句：你所说的“创新”指什么？如果你问统计学家，他们会提到两位极其聪明的人——艾伦·图灵和克劳德·香农，正是他们奠定了信息论的基础。

当时他们需要从数据中刻画何为“信息”。凭借当时最强的统计知识，他们提出：信息就是“惊奇”，即

相对于你原有的先验（Bayesian prior）而言无法预测的部分。

假设你给我一个庞大数据集，再额外提供一个观测值。如果这个新数据正落在我预期的位置，它就没有“信息量”。只有当它偏离预期、迫使我更新看法时，这种“惊奇”才是真正的信息，也就是统计学意义上的“创新”。

那么，经济学的下一项“大创新”是什么？这让我想起庞加莱的一个故事。1900 年前后，人们问他：20 世纪最重要的创新会是什么？他坦诚地说：“我不知道。”原因很简单：若你站在 1800 年，根本无法想象电、磁、电磁统一、元素周期表等概念。那时的人们甚至无法提出正确的问题，更不用说预测答案。最大的创新往往是那些我们连想象都无法想象的东西。

对你们学生而言，困难也一样，对我亦然。我不会坐下来想着“我要创造全新东西”。我做的是：阅读文献。例如读到王鹏飞的一篇论文，其中的某些技巧对他来说或许不新，但对我却是新的。我可能想到：“他这里也许可以换个方法”，或者“这个技巧可以用于另一个类似问题”，而那个问题正是我熟悉的。于是我便去探索或与作者交流。

如果我每天都想着“未来五年我必须做一件大事”，我根本动不了。

问：您为什么来到这里？
答：第一，我这些年参加过很多在中国举办的顶级学术会议，尤其是与 AI 和机器学习相关的。过去一年，我们学校也举办了几场高质量会议，来访学者都在做最前沿的研究。我们从他们那里学习机器学习基础及其在经济学中的应用。这里的学术活动非常活跃，这是我来的原因之一。

第二，更个人的原因。当我四十岁时，中国仍然很贫穷，深圳几乎不存在。但在我半生的时间里，这里经历了惊人的转型。我想理解这一过程。

“中国特色社会主义”到底意味着什么？并不是翻翻《论语》就能找到答案。上世纪七八十年代，中国有一套非常强的意识形态——私人产权不好、私营企业不好、利润不好、企业家不被信任。在我自己的国家，美国，也有人持类似观点。

但有趣的是，当时的中国领导人尽管深信这些理念，却仍然面对事实：他们看到数据——贫困、产业落后、科技不足——并与美国、德国和新加坡进行比较。新加坡尤其关键：80%的人口来自华南贫困地区，但新加坡迅速崛起。邓小平访问后，新加坡领导人告诉他：“如果你采用我们的某些经济原则，你们会好得多。”邓小平说：“那就试试。”

这就是市场经济在中国的起点，中国也从新加坡学习了治理、信任机制和反腐体系。

所有这些都依赖“数据的启发”。但必须强调：数据本身不告诉你任何东西，关键是你如何用模型、用理论去解释它。所有大模型、大数据系统背后都有某种统计或理论结构，只是有时没有被明说。你们会在统计课里学到这一点。

问：你如何看待中美关系，尤其在人工智能领域的差异？
答：我认为世界各地的人本质相同，差异主要来自政府与制度，而非人本身。

在 AI 方面，中国的进展非常惊人，在一些领域甚至领先。例如，中国一些车载 AI 系统极其先进。而这

些系统背后依赖的就是经济学也使用的数学工具：最优控制、滤波、实时求解 Hamilton–Jacobi–Bellman 方程等。

本质上，强化学习不过是针对特定函数族的递归最小二乘算法。汽车中的模型实质是对环境做预测，同时进行控制与决策——这和经济学预测行为、博弈互动如出一辙。

过去八九年，全球经济政策中出现一些令人困惑的现象。经济学家几乎一致认为贸易战是愚蠢的，关税有害。但有美国总统说“贸易战很容易赢”。要么他无知，要么他在开玩笑，但无论哪种情况，政策背后真正的推动力是国内的利益集团：关税损害国家整体利益，却保护落后企业，使其支持相关政治人物。

从经济学角度，这显然不是高效的制度安排。历史也表明：关税一升，首先出现的是走私，因为关税创造了走私利润。

问：Sutton 认为大语言模型无法获得真正的智能，你同意吗？
答：美国正投入巨额资金尝试通过扩大模型规模来获得 AGI。但目前的大模型主要还是模式识别、类比生成，并不能解决真正的智能问题。企业也无法解释“规模扩张为何必然带来智能跃迁”。

Yann LeCun 则认为现有 LLM 采用了错误的智能范式，因此无论投入多少都不可能达到 AGI。他指出，即便是猫狗的许多理解能力，LLM 也完全没有。若他的观点正确，目前的巨额投入更像一次泡沫——类似 1999 年的互联网泡沫，最后可能只剩一两家公司，甚至一无所获。

对于你们做研究的人来说，这一点非常重要：别人会告诉你什么有价值、什么没价值。但你必须自己判断。当我年轻时，有名的经济学家告诉我别做理性预期，他们认为已无研究空间。但我仍然坚持，因为我觉得有趣，且尚未完全理解。

研究就是风险。明斯基曾对我说：当律师很安全，做研究则不是。每个想法都是赌注，没人保证它被认可。但这就是学术的本质。

问：你对想读 PhD 的学生有什么建议？
答：举个例子，有个学生来 PHBS 读博时几乎不懂经济学，但五年后成了很优秀的经济学家。他不会承认，但现在有时还会教我东西。他敢读任何论文，不怕数学，甚至享受数学。

这些就是好的特质。你不需要数学天才才能学经济。我自己数学也不好，只是“够用”。你不需要在每门数学课都拿 A。很多数学可以在研究生阶段按需补学。

经济学不是数学。你可以数学很好却经济学很差，因为经济学关注的是声誉、激励、行为等问题，这些往往先于形式化定义而存在。一些数学家会急着要“严格定义”，但经济学的问题常在定义之前。

Seminar with Thomas J. Sargent — 18 November 2025

Q: How much do you think a lessening degree of independence of the Federal Reserve might impact the long-run growth of the US? What's the implication of that, potentially them losing their independence in making monetary policy without influence from the government?

A: We can look at history. The United States did not have a central bank for its first 125+ years. Congress and the President directly managed monetary policy, but they disliked it because it often involved making unpopular decisions necessary for stability and low inflation. They created the Federal Reserve precisely to delegate this authority, creating an institution that could be blamed for tough choices. This delegation acts as a constraint on Congress itself. For example, if a politician wants lower interest rates but the Fed refuses, the government might then have to raise taxes or make other difficult fiscal adjustments.

If a President gains direct control over the Fed and forces it to pursue politically expedient but economically unsound policies (like excessive money printing), the result would likely be higher inflation. Crucially, the President would then have no one else to blame for the negative consequences. This undermines the very purpose of the institution. The key is to think about the system of incentives and constraints. The Fed's independence was an institutional solution to a time-inconsistency problem, and weakening it risks returning to a less predictable and less stable policy environment.

Q: As a student, I just want to know how to start with a research problem. How do I find a problem to study? Also, how will AI tools influence our methods of research?

A: This is a fundamental question we all face constantly: how to allocate our scarce time and intellectual energy. The dilemma of whether to focus on a "hot topic" like AI or on foundational principles is a real one. My approach is personal. I am drawn to fundamental, timeless principles. The books I keep returning to on my shelf are often decades old, dealing with small datasets but illustrating universal principles brilliantly. These principles—like incentives, constraints, and expectations—are applicable to any country and any time period.

Regarding AI, you must ask yourself a critical meta-question: Did you think of your question yourself, or did you rely on AI to suggest it? There's no moral judgment here, but it's crucial for self-awareness. The danger is that AI can become a substitute for your own thinking rather than a complement. I've seen this in teaching; students who rely entirely on external tools without engaging with the foundational material end up learning nothing. AI, at its core, is about curve-fitting and statistical learning—topics that are part of the very foundations we teach. To benefit, you must make AI a *compliment* to your skills, not a *replacement*. It should augment your ability to think, not take control of the thinking process itself. The goal is to use the tool to deepen your understanding of fundamental principles, not just to generate outputs.

Q: How should governments in developing economies think about building credibility over time, where households and firms are learning and may not trust official models and forecasts? What do central banks and finance ministries often get wrong about how real households and firms form expectations?

A: This gets to the heart of expectations, credibility, and reputation. Let's define our terms carefully. A "reputation" is not something you possess internally; it resides in the minds of others. It is their belief or expectation about what you will do in various situations. Your reputation is shaped by your historical behavior—people process this “data” to form a view of you.

Consider a simple example: I promise my wife I will be home at 6 PM to help with dinner. At 5:30 PM, a friend invites me to a bar. I have a choice: keep my promise or break it for immediate enjoyment. If I break it, my wife will update her beliefs about my trustworthiness. Now, the key is that I *know* she is doing this data processing. So, my decision is made in anticipation of how it will affect my reputation with her. This creates a dynamic interaction.

This is the pure theory behind credibility. It’s about a multi-period game with equilibrium. Governments and central banks are in precisely this game with the public. To build credibility, they must not only make promises but also follow through, knowing that the public is constantly watching and updating its beliefs. The consequence of breaking promises is that the public will adjust its expectations and behavior accordingly, making future policy less effective. What officials often get wrong is underestimating how intelligently the public learns from past actions. It’s not about complex models; it’s about consistent, predictable behavior that aligns with stated goals.

Q: As people are gradually getting used to using AI for decision-making, your Rational Expectations theory originally considers human rational judgment. How should the theory be adjusted, and what kind of stable state could eventually form?

A: This question is deeply connected to the previous ones about AI. When you ask me these questions and I respond, I am essentially performing a version of human intelligence. I’m listening, grouping similar questions, and formulating a response—this is what AI also does, just inside a machine. The foundation of AI is statistics: it’s about choosing a class of functions (e.g., straight lines, polynomials) and a method to fit them to data (e.g., least squares).

Rational Expectations is a hypothesis about how people use *information* efficiently. The core definition of an “innovation” in information theory (by Turing and Shannon) is a *surprise*—something that could not have been predicted given the available information. If AI simply helps people process existing information faster and more accurately, it might make expectations *more rational* in the traditional sense, as people (or the AIs they use) become better at forecasting.

However, if AI starts to generate novel strategies or creates a complex, interacting system of AI agents, the nature of the “information set” and the process of learning could change fundamentally. The stable state would then be a new equilibrium in this human-AI hybrid system. The theory wouldn’t be discarded but would need to be expanded to model this more complex, multi-layered learning and expectation formation process. The key is to model the system, including both the human and AI agents, and their intertwined beliefs and learning rules.

Q: If we want to start our career or business in a few years, what kind of mindset should we build up during this time to face future challenges? How can we best combine the economics and logic we learn with the real business world?

A: The mindset you need is precisely the one we've been discussing: the economic way of thinking. Great economists often say that at the core, it's not that complicated. It's about understanding incentives and constraints. In the social world, this includes the powerful force of people's *beliefs* and *expectations* about each other, which create reputations and credibility.

You don't necessarily need to learn an enormous amount of facts, but you need to learn how to *think*. Learn the fundamental, universal principles. The principles that help you understand my simple example about promising my wife I'd be home are the same principles that can help you analyze central bank credibility, international trade agreements, or the stability of a social safety net.

When you face a business problem or a policy question, boil it down. What are the incentives? What are the constraints? What are the beliefs and expectations of the different actors? How do their actions interact in an equilibrium? This framework is incredibly powerful and universally applicable. The economics you learn provides the logical structure to analyze these situations systematically, moving beyond anecdotal or superficial explanations. This analytical rigor is what will give you an edge in the business world.

Q: What do you think is the next big innovation?

A: Well, before answering that, let me ask you: *what do you mean by innovation?* If you ask a statistician what "innovation" means, they'll point you to two extremely smart people—Alan Turing and Claude Shannon—who basically invented information theory.

They needed a way to define "information" from data. They knew a tremendous amount of statistics—more than almost anyone at the time—and they relied heavily on statistical ideas. What they came up with was this: information is a surprise, something you couldn't have predicted based on your prior beliefs, your Bayesian prior.

Suppose you have a big data set and then you give me one additional observation. If that new observation falls exactly where I expected it to fall, then it doesn't teach me much. It doesn't count as information.

But if the new observation lands somewhere unexpected—if it surprises me—then that forces me to revise what I thought before. That surprise is what Shannon and Turing would call information, and it's also what a statistician might call an "innovation."

So what's going to be the next big Innovation in economics?There's a nice example of what I mean. It comes from Henri Poincaré, one of the greatest mathematicians and physicists at the turn of the 20th century. He was a French scientist, and around 1900 people asked him: *"Tell us what you think the major innovations of the 20th century will be."* And Poincaré said, very honestly: *"I have no idea."* Then he went on to explain why. He said: imagine you were sitting in the year 1800. Think about the things that were discovered later—electricity, the unification of electricity and magnetism, the periodic table. In 1800, people didn't even know enough to formulate the questions that would lead to those discoveries. They only had a handful of known chemical elements. They had no conceptual framework for imagining what was coming. So they were completely surprised. That's the point: the biggest innovations are usually the ones nobody even knows how to predict.

You know what the problem is for you as a student? It's actually the same problem I have. I deal with it every day, because I’m in the same boat as you.Here’s what I do: I don’t try to sit down and think, *“I’m going to create something totally new,”* or *“I’m going to invent a brand-new tool.”* That’s not how it works. What I do instead is read a paper—say, a paper by Pengfei Wang. The ideas in that paper aren’t new to him, but they’re new to me. So I read the paper. And a couple of things might happen: I might think, *“This is really good.”* But I might also think, *“He could have done something a little different here. He didn’t realize that the tricks he used could also be applied to another, very similar problem.”* And I happen to know that other problem. So maybe I start working on it. Or maybe I go talk to him about it. That’s what I do. And if I spent all my time worrying about what big important thing I should do over the next five years, I’d completely freeze up. It would shut me down.

Q: How did you come here?

A: I came here for a couple of reasons.

First, over the years, I’ve attended several scientific conferences in China—really world-class ones—especially in fields related to AI and machine learning. And in the past year, we’ve also had two or three excellent conferences here at this school. Many of the scholars who came are doing frontier research in AI and machine learning.

They are selected by different Chinese academic groups, and when they come, we learn from them—about the foundations of machine learning and how these tools can be applied in economics. We had a wonderful conference just a couple of months ago, and this December we’re having another truly world-class one. So there’s a lot of intellectual activity happening here. That’s one reason.

The second reason is more personal. When I was 40, China was still a very poor country. Shenzhen basically didn’t exist. And yet, in just half of my lifetime, this extraordinary transformation has happened. Many “miracles,” really. And I want to understand how that happened.

China has something it calls “socialism with Chinese characteristics.” What exactly does that mean? Those “Chinese characteristics” aren’t something you can just look up in Confucius. They had something very specific in mind.

Back when I was 30 or 40, China had a strong ideological view about how to run the economy: private property is bad, private firms are bad, profits are bad, entrepreneurs are antisocial. That framework was firmly in place—and for many years. And in my own country, there are still people who believe that. The new mayor of New York agrees with some of that. So the idea hasn’t disappeared.

But here’s what I find fascinating: Despite being lifelong believers in that old ideology, Chinese leaders at the time looked around and said, “This isn’t working.” So they turned to data. They looked at China’s level of wealth and poverty, its scientific output, its industrial capacity—cars, machinery, all sorts of things. And then they compared it to countries like Germany, the United States, and also Singapore.

Singapore was especially influential. Around 80% of its population came from this region of China—poor, undernourished people who left here with almost nothing. And yet Singapore was beginning to grow rapidly. Deng Xiaoping went there, spoke to its leaders, and they told him: “If you adopt some of the economic principles we use,

you'll do much better. You already have the people." And Deng Xiaoping said: "Let's try it."
That's how the market economy started being introduced into China.

They also learned from Singapore about institutions—things like trust, governance, how to reduce corruption. Singapore became a kind of model.

All of that was *learning from data*. Big data, in a sense. And here's the main point:Data alone doesn't tell you anything. You need a model. You need a theory to interpret the data.

Every large language model, every big-data system—you name it—is always learning within the structure of some underlying statistical or theoretical model. Sometimes people don't tell you what that model is, but it's in there. And that's something you'll learn about in your statistics classes.

Q: What are your thoughts on U.S.–China relations, or the differences between the two countries, especially regarding AI?

My starting point is that people everywhere are fundamentally the same. I don't really think in categories like "Chinese" or "American." People care about similar things.So if differences exist, they mostly come from different governments and institutions, not different kinds of people.

When it comes to AI, China is doing extraordinary work. In some areas, China is ahead. For example, the AI systems in some Chinese cars are extremely advanced. What impresses me is that these systems rely on the same mathematics economists use—optimal control, filtering, and even solving Hamilton–Jacobi–Bellman equations in real time.

That's the key point: behind all the impressive AI applications is a very concrete set of mathematical tools. Machine learning terms—like "reinforcement learning"—may sound new, but at their core they're just recursive least-squares algorithms designed for specific function classes. The models inside those cars are basically applying filtering and control theory while making predictions about the environment.

Driving is a good example: the system has to anticipate where pedestrians and other cars are going, forming expectations—exactly the kind of problem economists think about. Sometimes it even becomes a game-theoretic interaction among multiple agents.

Over the last eight or nine years, I've been puzzled by what's happening in global economic policy. Economists around the world overwhelmingly agree that trade wars are foolish and that tariffs are harmful. Yet a U.S. president said, "It's easy to win a trade war."

So we have to ask: *Is he joking?* If he's serious, then he's simply ignorant of economics. If he's joking, then *why is he doing it*?

When you look deeper, you see what's going on: the country as a whole loses from a trade war, but certain groups inside the country gain. Tariffs lower overall efficiency, but they protect inefficient producers from competition.

Those protected firms—and their workers—value that protection and are willing to support the politician who gives it to them.

But from the standpoint of most economists, this is not a good way to organize an economy.

Historically, we know trade barriers are harmful because whenever tariffs go up, one of the first things that emerges is smuggling. Tariffs create profits for smugglers. So, in that sense, smugglers are the ones who “vote” for trade wars.

You guys have any follow up questions or questions that we have not, of course, Sergeant has asked.

Q: Richard S. Sutton(The father of Reinforcement learning) mentioned that large language models can’t reach real intelligence because they lack a model of the world and can’t learn from experience. Do you agree?

The United States is betting enormous amounts of money—trillions—on scaling up LLMs. Much of that investment is financed by borrowing, and it’s being used to buy chips, build data centers, and train ever-larger models. The hope is that by making models bigger, they will somehow “jump” into general intelligence.

But these models still do mainly pattern recognition and analogy making. They don’t solve the subtle cognitive problems we need for real intelligence. And when you ask companies why they believe scaling alone will work, they can’t really explain it.

There is also Yann LeCun, a leading machine-learning scientist from France, now at NYU and formerly at Meta. He argues strongly that current LLMs are based on a wrong model of intelligence, and therefore no amount of money will make them reach AGI. He points out that even a cat or dog has forms of understanding LLMs completely lack. LeCun believes a fundamentally different architecture is needed.

This matters because the major U.S. tech companies—Amazon, Google, OpenAI, Anthropic—are all in a race doing essentially the same thing. It resembles a bubble: massive spending, same strategy, same expectation of a breakthrough. Financially, it looks a lot like the dot-com bubble of 1999–2000. Historically, such situations end in consolidation—maybe one company survives, maybe none, depending on whether the underlying model is correct.

If LeCun is right, most of them will fail.

For your own research, this is a reminder: people will always tell you what’s worth working on and what isn’t. When I was young, well-known economists told me not to work on rational expectations—they thought everything important was already known. But I pursued it because *I* found it interesting and didn’t fully understand it yet.

That’s what research is. It’s risky and uncertain—Hyman Minsky once told me this when I was your age. He said: being a lawyer is safe. Being a researcher isn’t. Every idea you pursue is a gamble—you don’t know if it’s good, and even if it’s good, you don’t know if anyone will appreciate it. Some people love that uncertainty; others hate it.

But that uncertainty *is* the nature of academic work.

Q: Give some advice for students wish to pursue PHD

A: Let me give you an example. There was a student here who got a PhD from PHBS. When he first arrived, he knew very little economics. But five years later, he somehow became a very good economist. He won't admit it, but he even teaches me things now. He shows me papers, walks me through them, and he has read a lot in his field.

What I noticed is: he isn't afraid of math or hard work, and he actually enjoys it. He has built up solid mathematical tools, and he's not scared of reading any paper. Some papers even I hesitate to read, but he just dives in and discusses them.

Those are useful characteristics.

You do **not** need to be very good at math to do economics. I'm not good at math. You just need to be good enough. You don't need A's in every math class. You can learn the math you need later and still use it in economics.

Economics is **not** math. You can be a great mathematician and a terrible economist. Economics is about the kinds of questions we were talking about—reputation, incentives, behavior—things that are not always written in precise mathematical form. Some mathematicians become impatient and say, "Give me a definition." But economics often starts before the definitions are formalized.