The Economics of Artificial Intelligence

Spring 2024

ECON 390 (#22046) & ECON 691 (#22045)

Professors: Konrad Posch (kposch@usfca.edu) and Nihar Shah (nshah13@usfca.edu)
Dates: January 23 – March 7, 2024 (Course ends at Spring Break)
Time: Tuesdays and Thursdays, 7:00 – 9:00pm
Location: Main Campus, Room Lone Mountain Main 343
Number of Units: 2
Office Hours: Prof Posch (Tuesdays, 3-4pm, McLaren 107, Zoom by appointment)
Prof. Shah (Fridays, 11am-12pm, Zoom Room 875-1678-7441)

Course Description

Artificial intelligence is the most-discussed technology in the last year, owing to the breakout success of ChatGPT. But the societal discourse around it is poorly-understood and often prone to hyperbole.

This class is designed to introduce students to gain a wider understanding and appreciation for the economic and social implications of artificial intelligence. The course explores a wide range of topics concerning artificial intelligence, such as fairness considerations, employment dynamics, and regulation. Through student-led discussions, presentations, and writings oriented around questions and readings, students will develop nuanced perspectives on the upcoming challenges and opportunities.

This class prioritizes critical thinking, encourages diverse viewpoints, and enhances students' proficiency in analyzing and discussing core topics. By contrast, the class does not seek to establish "correct" answers, given the nascency and uncertainty of the field. Your goal in this course is to consider what authors, professors, and your fellow students think about the economic and social impact of AI and form that not into "the" correct answer but instead in "an" answer which makes sense to you and which you can support with evidence. All assignments in the course are designed and graded with this in mind.

For MSAE students, this class fulfills the 2-unit Professional Communication requirement.

Learning Outcomes

Students should learn the following by completing the course:

- How to think critically about the economic and social implications of artificial intelligence.
- How to evaluate risks and solutions associated with artificial intelligence.
- How to communicate insights and questions in written and oral form, both individually and in a large group setting.
- How to present and steer discussions, including how to encourage informed debate and how to find consensus among a variety of viewpoints.

Admission and Prerequisites

Undergraduate Economics students should have taken ECON 111.

Graduate Economic students should have taken ECON 501 or ECON 603.

Students outside of the Economics department are welcome. If you have not taken these prerequisites, please speak with the professors about your background. We welcome a diverse range of perspectives, but we would like to understand your interests and motivations to ensure you can get the most out of it.

Course Structure

Following the first week of introductory lectures (led by the professors), the course will spend each week covering a single topic related to artificial intelligence (e.g. regulation, employment, etc).

Each week, one group of students will lead both classes in that week. For the first class in that week, they will present an overview of the readings for that week. For the second class in that week, they will present an overview of the reactions from their fellow students.

All other students are expected to do the readings each week, and write a short reaction paper in advance of the first class for that week. (The students leading the classes that week are expected to do the readings, of course, but they are exempt from writing reaction papers for that week.)

Grading

The course will be graded based on the following framework:

50% – Reaction Papers (due at 12pm (noon) every Tuesday you are not presenting)

20% – Presentation of Readings (One Tuesday during the Term)

20% – Presentation of Reactions (One Thursday during the Term)

10% – Participation

Reaction Papers (due 12:00pm (Noon) every week you are not presenting)

Beginning Week 2, each week that a student is not presenting, they are expected to write a reaction paper of 300 - 500 words and submit it by 12pm PT (noon) prior to the Tuesday class through the relevant Canvas assignment.

The reaction paper should **be focused on a *single* argument** (e.g. a question that the readings did not address satisfactorily, a criticism of one of the reading's arguments, a point that the readings should have addressed, etc). Your response should be primarily and firmly grounded in the readings for the week including specific citations to them. While the best reaction papers may include external references or articles to develop this argument fully, this is encouraged but not required. Your primary assignment here is to read and engage deeply with the issues in the readings and then form your thoughts into an argument for what you think the class should discuss that week.

Reaction papers will be scored on three tiers: Excellent, Satisfactory, or Unsatisfactory.

Presentation of Readings (Tuesdays)

Beginning in Week 2, each Tuesday's class will begin with a group of **up to five** students who will present an overview of the readings. This presentation includes two components:

1. First, **for each reading**, the group should prepare a short overview of the major argument and claims.
2. Second, the group should conclude their presentation with two or three general discussion questions which address the theme of the week and place multiple readings into conversation with each other.

Throughout the class, the students will then be expected to lead the discussion and encourage broader participation. (The professors will of course participate actively too.)

Presentations will be scored across three dimensions:

- quality of explanation,
- thoughtfulness of discussion questions,
- ability to lead discussions.

Each dimension will be scored on the same three tiers: Excellent, Satisfactory, or Unsatisfactory. All members of the group will receive the same grade.

Sign-ups for each topic will be handled in class on Thursday of Week 1.

Note: any group that has graduate students will be asked to also cover the optional readings for that week.

Presentation of Reactions (Thursday)

Beginning in Week 2, each Thursday's class will begin with the same group of students from Tuesday who will present an overview of the **student reaction papers** in the Thursday class. *(Note: Non-presenting students will submit their reaction papers through Canvas each Tuesday as explained above and the instructors will send the relevant papers to the presenting group after class each Tuesday to prepare for each Thursday.)*

In particular, the group should identify **three to five particular arguments worth discussing in class**, after consolidating all the student reactions from their papers and Tuesday's discussions. For each argument:

1. First, the group should present an overview of the argument, adding context and potentially supplementing with outside sources as needed.
2. Second, the students should steer a discussion around the argument. As before, the student group should both lead the discussion and encourage broader participation from the class.

Presentations will be scored across two dimensions: quality of explanation and ability to lead discussions. Each dimension will be scored on the same three tiers: Excellent, Satisfactory, or Unsatisfactory. All members of the group will receive the same grade.

Participation

Students are expected to participate in each and every discussion. The final component of the grade is thus participation. Students who are not active may be called on directly by the professors. Students who still do not participate meaningfully in the discussion after that point will have their grade penalized accordingly.

Extra Credit

We will not provide individual extra credit plans, as any such opportunities must be made available to all students and we do not expect to alter the course plan to include them. Your best strategy for success in this course is to engage with the material both in readings, in reaction papers, and in class discussions.

Grade Disputes

If you believe you have been graded unfairly, you must submit a formal grade dispute within one week of your grade being assigned. The grade dispute must be a written email, explaining what component of the grade you disagree with and why.

Course Reading Access

This course covers topics that are *extremely* new, and there is no comprehensive course textbook that addresses them all. The readings will thus be an assorted mix of papers, articles, blogs, and podcasts. All required readings will be made available through the course Canvas site. Many are open-access and linked in the syllabus below while others are available to you through the Gleeson Library's digital services or will be provided as fair use PDFs. **The authoritative list of and access to readings will be through the modules section of the Canvas website.**

https://usfca.instructure.com/courses/1619807/modules

If you have difficulty accessing a reading, please be sure to read the Gleeson library guide on how to set up access (linked at the top of the modules page). If you are still unable to access a reading, *please email the instructors ASAP so we can fix it for everyone!*

The course also includes optional readings for each week, some of which may not be open-access. (The optional readings do not have to be included in either the reaction papers or presentations, but they may be useful as additional source material.)

Readings may be altered at the discretion of the professors, given the speed at which the literature is evolving. Any changes will be communicated at least one week in advance of the associated class. We will make no significant change to the amount or difficulty of readings from those listed below, although we may change readings if current events or new publications cover the topics better.

Course Topics and Readings

Week 1: Introduction and Technical Foundations

January 23: Introduction

a. Course Themes and Course Expectations
b. Core Readings
   i. "Chapter 4. What Is to Be Done? Prescriptions and Recommendations" in Brynjolfsson, Erik, and Andrew McAfee. 2012. *Race Against the Machine: How the Digital Revolution Is Accelerating Innovation, Driving Productivity, and Irreversibly Transforming Employment and the Economy*. Lexington, Mass: Digital Frontier Press.
   ii. Andreesen, Marc (June 6, 2023). "Why AI Will Save the World." Link.
c. Optional Readings
   i. None listed.

January 25: Technical Foundation

a. How does a Large Language Model work? What are the core components and steps needed for it to function?
b. Core Readings
   i. Wolfram, Stephen (February 14, 2023). "What is ChatGPT Doing… and Why Does It Work?" Stephen Wolfram Writings. Link.
   ii. "New Scaling Laws for Large Language Models" (April 1, 2022). Less Wrong blog. Link.
c. Optional Readings
   i. Athey, Susan (2018). "The Impact of Machine Learning on Economics." *The Economics of Artificial Intelligence: An Agenda*. Link.
   ii. Hoffman, Jordan, et al (March 29, 2022). "Training Compute-Optimal Large Language Models." arXiv. Link.
   iii. Mullainathan, Sendhil, and Jann Spiess (2017). "Machine Learning: An Applied Econometric Approach." *Journal of Economic Perspectives*, 31 (2): 87-106. Link.

Week 2: (January 30 and February 1) Experience of the Makers

a. What should the people and companies that make AI consider in their development? Specifically, how should they design the business and product strategy, how should they make choices between open versus closed source development, and how should they set up the corporate governance of the firm?
b. Core Readings
   i. Varian, Hal (2019). "Artificial Intelligence, Economics, and Industrial Organization." *The Economics of Artificial Intelligence: An Agenda*. Link.
   ii. Goldman, Sharon (November 3, 2023). "Forget ChatGPT, why Llama and Open Source AI win 2023." *VentureBeat*. Link.
   iii. Harris, David Evan (January 12, 2024). "Open Source AI is Uniquely Dangerous." IEEE Spectrum. Link.

iv. Tallarita, Roberto (December 5, 2023). “AI Is Testing the Limits of Corporate Governance. *Harvard Business Review*. Link.
v. Posner, Eric (January 11, 2024). “AI revolution likely to cement Big Tech monopoly.” *The Asset*. Link.

c. Optional Readings
i. Liesenfeld, Andreas, Alianda Lopez, Mark Dingemanse (July 8, 2023). "Opening up ChatGPT: Tracking openness, transparency, and accountability in instruction-tuned text generators." arXiv. Link.
ii. Narechania, Tejas and Ganesh Sitaraman (January 17, 2024). "An Antimonopoly Approach to Governing Artificial Intelligence." SSRN. Link.

Week 3: (February 6 and 8) Experience of the Takers

a. What can/should/do companies who bundle AI “off the shelf” into their products/services think about when integrating AI into their current and future product offerings? What can/should/do customers/consumers who are subjected to AI either as inputs (data to be processed) or outputs (customers for AI-produced products/services) think about when making their purchasing decisions?
b. Core Readings
i. Anderson, Chris. 2008. “The End of Theory: The Data Deluge Makes the Scientific Method Obsolete.” *Wired*, June 23, 2008. link.
ii. “Introduction” and “Chapter 1- Bomb Parts: What is a Model?” in O’Neil, Cathy. 2016. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Crown.
iii. Farrell, Henry, and Marion Fourcade. 2023. “The Moral Economy of High-Tech Modernism.” *Daedalus* 152 (1): 225–35. Link.
c. Optional Readings
i. Mullainathan, Sendhil, and Ziad Obermeyer. 2017. “Does Machine Learning Automate Moral Hazard and Error?” *American Economic Review* 107 (5): 476–80. link.
ii. Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Ashesh Rambachan. 2018. “Algorithmic Fairness.” *AEA Papers and Proceedings* 108 (May): 22–27. link.

Week 4: (February 13 and 15) Employment

a. How will AI shape the future of work? Specifically, what are the skills, tasks, and positions that AI will augment and enhance; and which are the skills, tasks, and positions that AI will substitute and replace?
b. Core Readings
i. Acemoglu, Daron, and Simon Johnson (February 6, 2023). “What's Wrong with ChatGPT?” Project Syndicate. Link.
ii. Baily, Martin Neil, Erik Brynjolfsson, and Anton Korinek (May 10, 2023). “Machines of mind: The case for an AI-powered productivity boom.” Brookings Commentary. Link.
iii. Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond (April 25, 2023). “Generative AI at Work.” Link.

iv. Evans, Benedict (July 2, 2023). "AI and the automation of work." Link.

c. Optional Readings
   i. Acemoglu, Daron (August 2021). "Harms of AI." Working Paper. Link. *Note: it is acceptable to skip over the math.*
   ii. Kreitmeir, David, and Paul Raschky (April 20, 2023). "The Unintended Consequences of Censoring Digital Technology – Evidence from Italy's ChatGPT Ban." Link.

Week 5: (February 20 and 22) Regulation

a. Regulations and governance make markets work. How can/should/has/will AI technology be integrated into the economic regulatory structures that make markets possible? What are the implications of regulation of AI and regulation by AI?

b. Core Readings
   i. "Chapter 1 - Introduction: More than Mere Deadweight" in Posch, Konrad Edward Ian. 2023. "More than Mere Deadweight: The Variety of Regulatory Imaginaries That Shape How Regulators, Innovators, and Entrepreneurs Coproduce Disruptive Technological Innovation." Ph.D. Dissertation, United States -- California: University of California, Berkeley. Link.
   ii. Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Cass R Sunstein. 2018. "Discrimination in the Age of Algorithms." *Journal of Legal Analysis* 10 (December): 113–74. Link.
   iii. Burrell, Jenna, and Marion Fourcade. 2021. "The Society of Algorithms." *Annual Review of Sociology* 47 (1): 213–37. Link.
   iv. Rambachan, Ashesh, Jon Kleinberg, Jens Ludwig, and Sendhil Mullainathan. 2020. "An Economic Perspective on Algorithmic Fairness." *AEA Papers and Proceedings* 110 (May): 91–95. https://doi.org/10.1257/pandp.20201036.

c. Optional Readings
   i. "Chapter 2 - A Genealogy of Disruptive Innovation and Regulatory Imaginaries: How Disruptive Technological Innovation Cast Regulation as a Villain" in Posch, Konrad Edward Ian. 2023. "More than Mere Deadweight: The Variety of Regulatory Imaginaries That Shape How Regulators, Innovators, and Entrepreneurs Coproduce Disruptive Technological Innovation." Ph.D. Dissertation, United States -- California: University of California, Berkeley. Link.
   ii. Öjehag-Pettersson, Andreas, Vanja Carlssson, and Malin Rönnblom. 2023. "Political Studies of Automated Governing: A Bird's Eye (Re)View." *Regulation & Governance* n/a (n/a). https://doi.org/10.1111/rego.12569.
   iii. Browse the Special Issue of *Regulation & Governance* on "Algorithmic Regulation" link to special issue.
      - If you are a graduate student assigned to present on this week (and thus need to cover the optional readings), you can focus on just the introductory reading: Ulbricht, Lena, and Karen Yeung. 2022. "Algorithmic Regulation: A Maturing Concept for Investigating Regulation of and through Algorithms." *Regulation & Governance* 16 (1): 3–22. https://doi.org/10.1111/rego.12437.

Week 6: (February 27 and 29) Hardware and Physical Dimensions

a. The "magic" of AI may be the software, but it must run on hardware and generally quite a lot of expensive and energy-intensive hardware with serious implications for climate change, entrepreneurial access, and ownership. How does this affect the uses, implementation, and impacts of AI technology? How might we develop AI technology and industries to create better or worse social, economic, and ecological impacts?
b. Core Readings
   i. Alloway, Tracy, and Joe Weisenthal (May 8, 2023). "Inside the Battle for Chips That Will Power Artificial Intelligence." Odd Lots. Paywalled Official Link PocketCast Free Link Youtube Free Link
   ii. Cowls, Josh, Andreas Tsamados, Mariarosaria Taddeo, and Luciano Floridi. 2023. "The AI Gambit: Leveraging Artificial Intelligence to Combat Climate Change—Opportunities, Challenges, and Recommendations." *AI & SOCIETY* 38 (1): 283–307. link.
   iii. Dave, Paresh (August 23, 2024). "Nvidia Chip Shortages Leave AI Startups Scrambling for Computing Power." Wired. Link.
   iv. Patel, Dylan and Daniel Nishball (August 27, 2023). "Google Gemini Eats The World – Gemini Smashes GPT-4 By 5X, The GPU-Poors." SemiAnalysis Blog. Link. *Note: only read sections prior to the paywall.*
c. Optional Readings
   i. Wu, Carole-Jean et al (January 9, 2022). "Sustainable AI: Environment Implications, Challenges and Opportunities." Link.
   ii. Sundberg, Niklas. 2024. "Tackling AI's Climate Change Problem." *MITSloan Management Review* 65 (2): 38–41.

Week 7: (March 5 and 7) Existential Risk

a. Certain tech entrepreneurs are concerned that AI represents an existential threat to humanity in the vein of Skynet from the movie *Terminator 2: Judgement Day* or at least in the vein of *Westworld*. After all we have discussed and read this semester on the underlying technology (capabilities and limitations), the social and economic impacts (good and ill), and the history of technological change over the last decades, centuries, and millennia what do we think? Is AI truly an existential threat? If so, how so? Why is this different from everything that's ever happened before? Or is it not? And what should we do about all this?
b. Core Readings
   i. "Chapter 11: Engines of Destruction" in Drexler, K. Eric. 1986. *Engines of Creation*. Anchor Press/Doubleday.
   ii. Johnson, Deborah G., and Mario Verdicchio. 2017. "AI Anxiety." *Journal of the Association for Information Science and Technology* 68 (9): 2267–70. link.
   iii. Müller, Vincent C., and Michael Cannon. 2022. "Existential Risk from AI and Orthogonality: Can We Have It Both Ways?" *Ratio* 35 (1): 25–36. https://doi.org/10.1111/rati.12320. ***(Note: this was written by two Philosophy Professors and then published in a peer reviewed academic journal)***

iv. Carlsmith, Joseph. 2022. “Is Power-Seeking AI an Existential Risk?” arXiv.Org. June 16, 2022. https://arxiv.org/abs/2206.13353v1. ***(Note: this was written by a PhD Candidate in Philosophy who works for an NGO and then self-published on a website archive)***
v. Cunningham, Tom. “An AI Which Imitates Humans Can Beat Humans.” October 6, 2023. Link.

c. Optional Readings
i. Bowman, Samuel (April 2, 2023). "Eight Things to Know about Large Language Models." arXiv. Link.

Additional Course Policies

Class Attendance

The core of the course is discussion, and so attendance is critical. If you must miss a class, please email both professors in advance and get one of their approvals. In lieu of the missed class, you will be asked to write an additional 600 - 1000 word essay answering a question set by the professors, based on the in-class discussion missed.

In general, one class can be missed without risking a failing grade, as long as the student gets the excuse preapproved and the student subsequently completes the associated writing assignment. A failure to complete the assignment will result in a failing grade for the course.

Any class that is missed with an unapproved excuse will also result in a failing grade. Finally, multiple classes missed without deeply extenuating circumstances will result in a failing grade. If you know in advance that you will need to miss one or more classes, **speak with the professors as soon as you know, preferably in person after the first class.**

For the week that you are presenting, you must be present for both classes or it will not be possible for you to complete this major assignment which makes up a total of 40% of your grade.

Usage of ChatGPT and other AI Assistants

The usage of ChatGPT and other AI assistants in writing reaction papers and slides is ***permitted***. However, you remain fully responsible for being able to explain any written argument orally in class. In addition, you should ensure the accuracy of your submitted materials, as ChatGPT is known to make up citations and facts.

Most critically, all your written assignments must be grounded firmly in course material. While you should absolutely experiment with how ChatGPT can help you to form deeper insights, do not allow it to replace your own engagement with the material and **certainly** do not submit a paper written “blindly” by ChatGPT which fails to engage with the assigned readings and hallucinates similar (but non-existent or lower quality) citations.

Finally, if you use ChatGPT to construct your reaction papers or slides, please note that explicitly at the end of the paper or presentation, with a line stating that “ChatGPT was used in the construction of these materials.”

Program Learning Outcomes

This course is a component of the MS Applied Economics. It will contribute to the following program learning outcomes:

PLO-3 Applied Economic Theory: Students will be able to understand and apply economic theory to understand how businesses and other organizations interact with each other and with users/customers/clients and use this understanding to guide data analysis.

PLO-4 Economic Problem Solving: Students will be able to solve real-world data-driven business and policy problems working with economists, policy makers, data scientists and business practitioners.

PLO-5 Economic Communication: Students will be able to communicate their research approach and findings at an excellent level, both in writing and verbally.

University-wide Policies and Legal Declarations

Communication

All course communications, like all other USF communications, will be sent to your USF official email address. You are therefore strongly encouraged to monitor that email account.

Studies with Disabilities

The University of San Francisco is committed to providing equal access to students with disabilities. If you are a student with a disability, or if you think you may have a disability, please contact Student Disability Services (SDS) at sds@usfca.edu or 415 422-2613, to speak with a disability specialist. (All communication with SDS is private and confidential.) If you are eligible for accommodations, please request that your accommodation letter be sent to me as soon as possible; students are encouraged to contact SDS at the beginning of the semester, as accommodations are not retroactive. Once I have been notified by SDS of your accommodations we can discuss your accommodations and ensure your access to this class or clinical setting. For more information please visit the SDS website: https://www.usfca.edu/student-disability-services.

Behavioral Expectations

All students are expected to behave in accordance with the Student Conduct Code and other University policies (see http://www.usfca.edu/fogcutter/). Students whose behavior is disruptive or who fail to comply with the instructor may be dismissed from the class for the remainder of the class period and may need to meet with the instructor or Dean prior to returning to the next class period. If necessary, referrals may also be made to the Student Conduct process for violations of the Student Conduct Code.

Academic Integrity

As a Jesuit institution committed to cura personalis -- the care and education of the whole person -- USF has an obligation to embody and foster the values of honesty and integrity. USF upholds the standards of honesty and integrity from all members of the academic community. All students are

expected to know and adhere to the University's Honor Code. You can find the full text of the code online at http://myusf.usfca.edu/academic-integrity/. The policy covers:

- Plagiarism — intentionally or unintentionally representing the words or ideas of another person as your own; failure to properly cite references; manufacturing references.
- Working with another person when independent work is required.
- Submission of the same paper in more than one course without the specific permission of each instructor.
- Submitting a paper written by another person or obtained from the Internet.

Counseling and Psychological Services (CAPS)

CAPS' diverse staff offers brief individual, couple, and group counseling to student members of our community. CAPS services are confidential and free of charge. Call (415) 422-6352 for an initial consultation appointment. Telephone consultation through CAPS After Hours is available Monday - Friday from 5:00 p.m. to 8:30 a.m., 24 hours during weekends and holidays; call the above number and press 2. Further information can be found at https://myusf.usfca.edu/student-health-safety/caps.

Confidentiality, Mandatory Reporting, and Sexual Assault

As instructors, one of our responsibilities is to help create a safe learning environment on our campus. We also have a mandatory reporting responsibility related to our role as faculty. We are required to share information regarding sexual misconduct or information about a crime that may have occurred on USF's campus with the University. Here are some useful resources related to sexual misconduct:

- To report any sexual misconduct, students may visit the Title IX coordinator (UC 5th floor) or see many other options by visiting usfca.edu/student_life/safer.
- Students may speak to someone confidentially or report a sexual assault confidentially by contacting Counseling and Psychological Services at (415) 422-6352.
- To find out more about reporting a sexual assault at USF, visit USFs Callisto website at: usfca.callistocampus.org.
- For an off-campus resource, contact San Francisco Women Against Rape (SFWAR) (415) 647-7273 (sfwar.org).