

AI AND EDUCATION

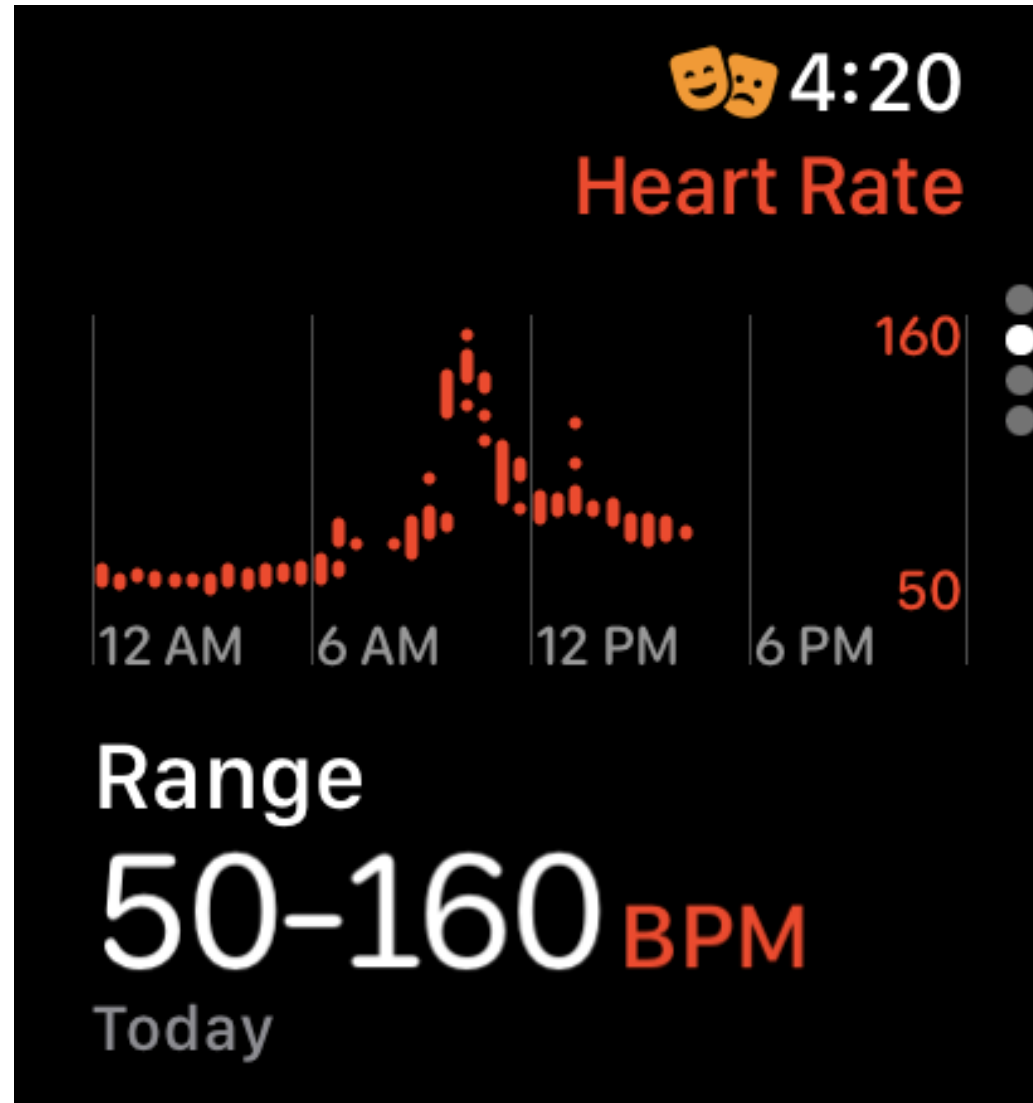
| Eric C So | MIT Sloan School of Management



A QUICK INTRODUCTION

- Eric So
 - Tenured Professor in the Global Economics and Management group
 - I teach a Finance elective titled Alphanomics
 - Friends/colleagues/students call me “E-So”
 - ...
 - Lead Faculty for AI Hub for Teaching and Learning

JULY 19, 2023



JULY 19, 2023: A FRAZZLED EMAIL



Dear Rodrigo, Ezra, and Dave,

I just mentioned to Rodrigo that I spent the morning feeling a bit frazzled.

I tried out a new AI tool and it honestly scared/troubled me. You may already know about this, but I created the attached slides to share because I think many people around Sloan don't know about it yet and need to prepare. Can you please take a quick look?

This is related to next week's town hall meeting. Here are my candid thoughts phrased as questions:

1. Trying it out this morning made me realize that education and the broader economy are about to change very fast. I think that trying to push back on the technology is infeasible, and that we should learn to somehow embrace the technology. But it's unclear to me how to do it. How should we think about designing assignments and fairly grading students now that the tech is widely available?
2. Should we be changing what we teach? Future employees won't be required to do nearly as much coding, and AI seems to already be able to qualitatively assess the results of data analysis. I am not sure it makes sense to spend a lot of time on teaching tasks that AI will be able to do in the near future. We can block AI access in the classroom but the reality is that it will be accessible when they enter the workforce.

JULY 19, 2023: A FRAZZLED EMAIL



Dear Rodrigo, Ezra, and Dave,

I just mentioned to Rodrigo that I spent the morning feeling a bit frazzled.

I tried out a new AI tool and it honestly scared/troubled me. You may already know about this, but I created the attached slides to share because I think many people around Sloan don't know about it yet and need to prepare. Can you please take a quick look?

This is related to next week's town hall meeting. Here are my candid thoughts phrased as questions:

1. Trying it out this morning made me realize that education and the broader economy are about to change very fast. I think that trying to push back on the technology is infeasible, and that we should learn to somehow embrace the technology. But it's unclear to me how to do it. How should we think about designing assignments and fairly grading students now that the tech is widely available?
2. Should we be changing what we teach? Future employees won't be required to do nearly as much coding, and AI seems to already be able to qualitatively assess the results of data analysis. I am not sure it makes sense to spend a lot of time on teaching tasks that AI will be able to do in the near future. We can block AI access in the classroom but the reality is that it will be accessible when they enter the workforce.

THE ASSIGNMENT



MIT Sloan School of Management

Alphanomics

Overview: Analyzing the Analyst

You are tasked with analyzing the behavior of sell-side analysts in terms of their buy/sell recommendations. The assignment objective is to use data to rigorously characterize analysts' use of buy/sell recommendations and assess their ability to help inform your portfolio allocation decisions. To conduct this task, you have been given a dataset of stock-month level observations summarizing analysts' buy-sell recommendations and one-month ahead stock returns, which is downloaded from the Institutional Brokers Estimate System (IBES). The dataset is titled "AnalystData_20XX.csv"

Using the provided data, answer the questions below. For each portion of this assignment, please provide the results of any empirical tests that you conduct, as well as any relevant programming code (if applicable).

Variable Descriptions:

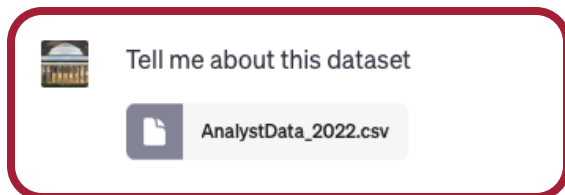
- Year
- Month
- YearMonth = Year-month indicator
- OFTIC = Ticker symbol of stock
- MeanRec = the average analyst recommendation numerically, ranging in value from 1 to 5. Lower values of MeanRec (i.e., close to 1) indicate a 'Strong Buy' whereas higher values of MeanRec (i.e., close to 5) indicate a 'Strong Sell'.
- Q_MeanRec = The variable Q_MeanRec equals one for firms with the lowest values of MeanRec (i.e., firms with a relatively high concentration of buy recommendations) in a given calendar month, and Q_MeanRec equals five for firms with the highest values of MeanRec (i.e., firms with a relatively high concentration of sell recommendations) in a given calendar month.
- BuyPCT = Fraction of analyst covering the firm that issue a recommendation to buy the stock
- HoldPCT = Fraction of analyst covering the firm that issue a recommendation to hold the stock (consider this a neutral recommendation)
- SellPCT = Fraction of analyst covering the firm that issue a recommendation to sell the stock
- retlag12 = trailing 12-month market-adjusted stock returns, expressed as a % (cumulative over 12 months before the analyst-recommendation is observed)
- MR1 = one-month ahead market-adjusted stock returns (one month after the analyst-recommendation is observed), expressed as a % divided by 100

- Historically took MBA students ~6 hours in groups
- I tried to solve it using ChatGPT's new data analysis tool
 - You can upload datasets
 - It codes in Python based on plain-English instructions that I gave it
- The AI did my assignment in under 2 minutes (!!)





USING CHATGPT TO SOLVE THE ASSIGNMENT



Finished working Show work

The dataset consists of the following columns:

1. `Year`: The year of the data point.
2. `Month`: The month of the data point.
3. `YearMonth`: A combined year and month identifier.
4. `OFTIC`: An identifier, possibly a ticker symbol or company code.
5. `MeanRec`: The mean recommendation score for the company.
6. `Q_MeanRec`: A quantile or categorized version of the MeanRec.
7. `BuyPCT`: The percentage of analysts recommending to buy the stock.
8. `HoldPCT`: The percentage of analysts recommending to hold the stock.
9. `SellPCT`: The percentage of analysts recommending to sell the stock.
10. `retlag12`: A lagged return variable, possibly the 12-month lagged return.
11. `MR1`: An unknown variable that could be related to returns or market reactions.

This dataset appears to contain information on stock market analyst recommendations and some return data. However, without further context or a data dictionary, it is challenging to accurately interpret all the column names.

- With no other information, it listed the variables and guessed their content
- You can also ask it to briefly summarize the dataset and look for abnormalities

USING CHATGPT TO SOLVE THE ASSIGNMENT



1. Summarize and report the averages of the variables BUYPCT, HOLDPCT, and SELLPCT for the full sample. What do you notice about the frequency of sell recommendations relative to buy recommendations?



Finished working

Show work ▾



The average percentage of "Buy" recommendations for the full sample is approximately 55.37%.

The average percentage of "Hold" recommendations for the full sample is approximately 39.86%.

The average percentage of "Sell" recommendations for the full sample is approximately 4.77%.

What is noticeable from these results is that "Sell" recommendations are much less common than "Buy" recommendations. The frequency of "Sell" recommendations is also lower than "Hold" recommendations. This might suggest that analysts are generally more inclined to recommend buying or holding stocks rather than selling them. However, these are just averages and the actual distributions of these recommendations could vary across time and companies.

- I copied and pasted my HW prompt verbatim into ChatGPT
- It calculated everything correctly
- It also “interpreted the data” like a student would/should and answered perfectly

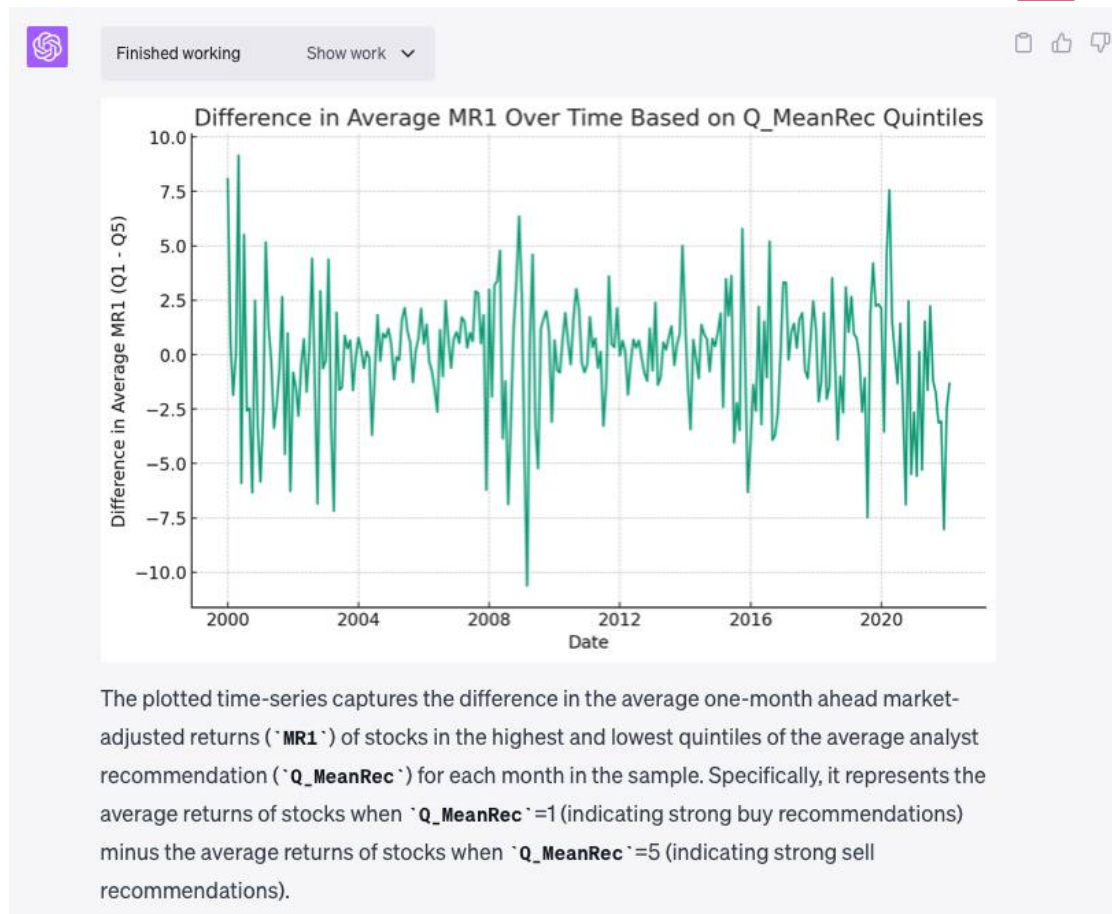


USING CHATGPT TO SOLVE THE ASSIGNMENT



For each calendar month in the sample, I have assigned firms into quintiles (i.e., five buckets) based on their average analyst-recommendation. Using the following steps, evaluate the efficacy of forming a monthly rebalanced portfolio that bets on firms with higher analyst recommendations and bets against firms with lower analyst recommendations.

For each year-months in the sample window, calculate the mean of one-month ahead market-adjusted return (variable MR1) for each quintile of Q_MeanRec. This should give you a time-series of 250+ months of data that tell you the average market-adjusted return of stocks in each quintile of Q_MeanRec. Next, for each month, calculate the difference in the average return of the lowest and highest quintile of Q_MeanRec. Because lower values of Q_MeanRec indicate that analysts are recommending buying the stock, you should calculate the difference in returns as: the average returns of stocks when Q_MeanRec=1 minus the average returns of stocks when Q_MeanRec=5. Plot the time-series of this difference as a graph. In words, what does this time-series capture and how does it relate to the assignment objective listed above?



➔ CI can create nice looking plots with data captions

USING CHATGPT TO SOLVE THE ASSIGNMENT



For each calendar month in the sample, I have assigned firms into quintiles (i.e., five buckets) based on their average analyst-recommendation. Using the following steps, evaluate the efficacy of forming a monthly rebalanced portfolio that bets on firms with higher analyst recommendations and bets against firms with lower analyst recommendations.

For each year-months in the sample window, calculate the mean of one-month ahead market-adjusted return (variable MR1) for each quintile of Q_MeanRec. This should give you a time-series of 250+ months of data that tell you the average market-adjusted return of stocks in each quintile of Q_MeanRec. Next, for each month, calculate the difference in the average return of the lowest and highest quintile of Q_MeanRec. Because lower values of Q_MeanRec indicate that analysts are recommending buying the stock, you should calculate the difference in returns as: the average returns of stocks when Q_MeanRec=1 minus the average returns of stocks when Q_MeanRec=5. Plot the time-series of this difference as a graph. In words, what does this time-series capture and how does it relate to the assignment objective listed above?

```
python Copy code

# Calculate the mean of MR1 for each quintile of Q_MeanRec for each year-month
avg_MR1_quintiles = data.groupby(['Year', 'Month', 'Q_MeanRec'])['MR1'].mean()

# Pivot the data to get each quintile as a separate column
pivot_avg_MR1_quintiles = avg_MR1_quintiles.pivot_table(index=['Year', 'Month'], columns='Q_MeanRec')

# Calculate the difference in averages between the 1st (buy) and 5th (sell) quintiles
pivot_avg_MR1_quintiles['Q1_Q5_diff'] = pivot_avg_MR1_quintiles[1] - pivot_avg_MR1_quintiles[5]

# Convert the 'Year' and 'Month' columns to datetime for plotting
pivot_avg_MR1_quintiles['Date'] = pd.to_datetime(pivot_avg_MR1_quintiles[['Year', 'Month']])

# Plot the difference in averages over time
plt.figure(figsize=(10, 6))
plt.plot(pivot_avg_MR1_quintiles['Date'], pivot_avg_MR1_quintiles['Q1_Q5_diff'])
plt.xlabel('Date')
plt.ylabel('Difference in Average MR1 (Q1 - Q5)')
plt.title('Difference in Average MR1 Over Time Based on Q_MeanRec Quintiles')
plt.show()
```

- CI will even give you the corresponding Python code (!!!)
- This means that asking students for their code will not solve the problem



JULY 19, 2023: A FRAZZLED EMAIL

Dear Rodrigo, Ezra, and Dave,

I just mentioned to Rodrigo that I spent the morning feeling a bit frazzled.

I tried out a new AI tool and it honestly scared/troubled me. You may already know about this, but I created the attached slides to share because I think many people around Sloan don't know about it yet and need to prepare. Can you please take a quick look?

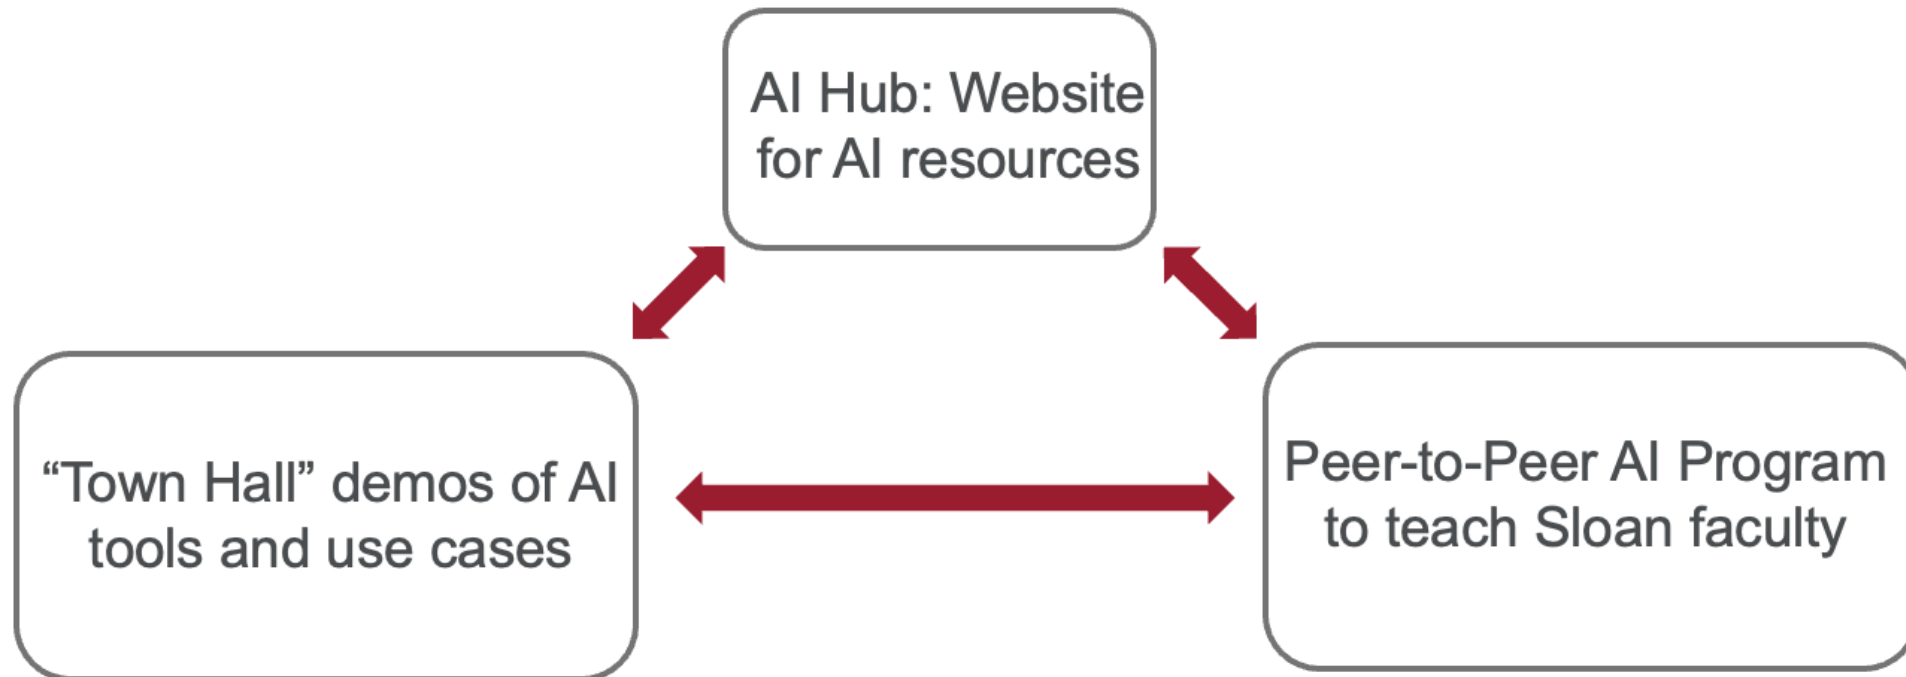
This is related to next week's town hall meeting. Here are my candid thoughts phrased as questions:

1. Trying it out this morning made me realize that education and the broader economy are about to change very fast. I think that trying to push back on the technology is infeasible, and that we should learn to somehow embrace the technology. But it's unclear to me how to do it. How should we think about designing assignments and fairly grading students now that the tech is widely available?
2. Should we be changing what we teach? Future employees won't be required to do nearly as much coding, and AI seems to already be able to qualitatively assess the results of data analysis. I am not sure it makes sense to spend a lot of time on teaching tasks that AI will be able to do in the near future. We can block AI access in the classroom but the reality is that it will be accessible when they enter the workforce.



NO GOOD DEED...

I was “volun-told” to lead a 3-front effort to prepare Sloan for teaching + learning with gen AI:



STEP 1: AI RESOURCE HUB AT MIT



Home [AI Hub](#) Trainings Tools Teaching Spaces How-to Guides Support

Generative AI for Teaching & Learning

Explore our resources to unlock AI's potential for enhancing teaching and learning at MIT Sloan.

LEARN MORE



Explore the Resources

Welcome to your hub for all things AI. Here, you can explore ways generative AI can augment your teaching and enrich learning at MIT Sloan.

MEET OUR TEAM →



AI Basics

Foundational AI concepts, key terms, and practical guides tailored for higher education.



AI Tools

An evolving list of tools and plugins that you can leverage in your classroom and beyond.



Teach & Learn

Subject-specific use cases, actionable guides, and ways to teach and learn in the era of AI.



Policy & Ethics

Considerations, principles, and guidelines for responsible AI use in academic settings.



Training & Support

Training, expert insights, and collaborative discussions at MIT Sloan.



Research & Trends

Keep pace with the latest research, insights, and shifts in AI-driven education.

<https://mitsloanedtech.mit.edu/ai/>



STEP 2: GEN AI CLASS FOR FACULTY

- **Goal 1: Use AI to help run a class more efficiently & potentially better**
 - Custom AI chatbots and tutors
 - AI grading assistance to give more feedback
 - Redesign assignments with AI in mind
- **Goal 2: Equip faculty with skills to teach students how to use AI**
 - Effective prompting with AI
 - Using AI to code and visualize data
 - Building more complex AI workflows



STEP 2: GEN AI CLASS FOR FACULTY



Format: In-person sessions:

- Group discussions on current and emerging AI tools
- Hands-on projects focusing on AI tools and aspects of teaching

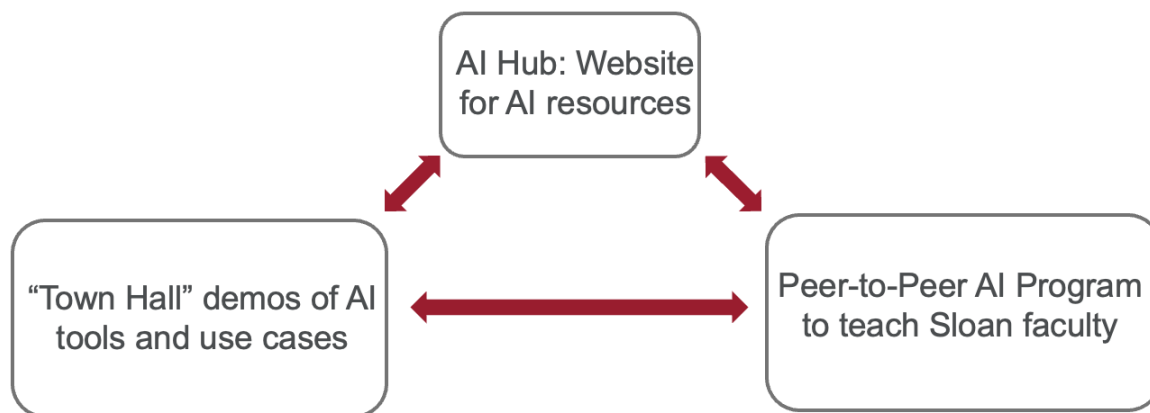
Participants: 20+ MIT Sloan faculty

- Mix of tenure-track and teaching faculty
- Representation of faculty from all of the Sloan “core courses ”





WE ARE PROUD BUT NOT COMPLACENT



Our next steps:

1. Re-running and expanding our P2P program
2. An "AI Bootcamp" for incoming MBA students
3. Expanded class offerings focused on AI
4. AI Labs: partnering with firms



BIG PICTURE PERSPECTIVE

- There is a risk of overhyping the near- to immediate-term implications of AI
 - Many credible folks compare the development of AI to the discovery of fire, which is polarizing
- With that said, I am confident AI is here to stay
 - Thinking practically, AI appears poised to transform society over the next two decades
 - We need to prepare. Our students need to prepare



OUTLINE

- AI at MIT Sloan
- Power and Pitfalls of Current AI Models
- Looking Back and Looking Ahead

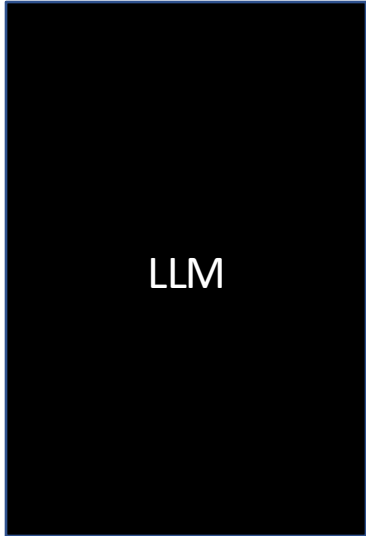
REDUX PRIMER ON LLMS



- An important part of what makes humans intelligent is that we can use the past to predict what might happen in the future
- AI systems, like LLMs, attempt to mimic human intelligence by creating a prediction machine
- They are trained on trillions of words from books, internet, etc.



REDUX PRIMER ON LLMS



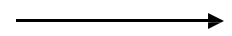


REDUX PRIMER ON LLMS

Input

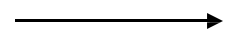
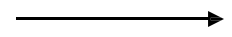
Output

The cat sat on the



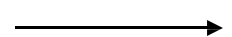
mat

Elementary, my dear



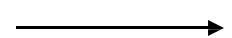
Watson

Four score and seven



years

...



...



REDUX PRIMER ON LLMS

The cat sat on the



LLM



Next Word	Probability ("how often")
aardvark	0.0
...	
fridge	0.05
...	
mat	0.2
...	
...	
zebra	0.0



REDUX PRIMER ON LLMS

The LLM then appends that word to the original prompt and the process continues

LLMs often create outputs to match these probabilities, which explains why outputs can vary with the same prompt

Next Word	Probability ("how often")
aardvark	0.0
...	
fridge	0.05
...	
mat	0.2
...	
...	
zebra	0.0



REDUX PRIMER ON LLMS

RA

explain the difference between polysemy and synonymy



Polysemy and synonymy are two concepts that are often used in linguistics to describe the relationships between words. Polysemy refers to the phenomenon where a single word has multiple related meanings, while synonymy refers to the phenomenon where two or more words have the same or nearly the same meaning.



LLM

I USE AI EVERYDAY: PRESENTING



Are you communicating the right way?

THE FIRST MINUTE



HOW TO START CONVERSATIONS
THAT GET RESULTS

CHRIS FENNING

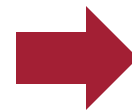


Google AI Studio

Chat

The_first_minute_Intro.pdf 25613 Tokens

create lecture notes that I can copy and paste into powerpoint slides. Include examples whenever possible. Target the material to PHD students at MIT Sloan School of Management. Make the examples relevant to PhD research and teaching business school students. After each major topic, create a task for my PhD students to turn to the person next to them and reflect on what we just learned and how they would apply to their lives as researchers and educators



(aistudio.google.com)

Chat

Slide 1: Mastering Communication as Researchers & Educators

A Guide to Effective Communication for PhD Students at MIT Sloan

This presentation will equip you with essential communication skills for:

- **Research Dissemination:** Clearly present your work at conferences, in publications, and during grant applications.
- **Effective Teaching:** Explain complex concepts and engage with students in a meaningful way.
- **Collaboration:** Foster productive discussions with colleagues and advisors.

Key Focus: Leveraging the power of the "First Minute" through Framing and Structured Summaries

Slide 2: The Power of the First Minute

- **Cognitive Load:** When we communicate, our audience's brains work hard to process information and understand our intent.
- **Framing:** The first 15 seconds of a conversation are crucial for setting context, clarifying intent, and delivering the key message.
- **Impact:** Effective framing reduces cognitive load, sets clear expectations, and paves the way for a productive exchange.



I USE AI EVERYDAY: AS A CO-WRITER

ChatGPT 4o 

Create a funny monologue in the spirit of Stephen Colbert's late night show. Make the focus of the monologue to welcome MIT Sloan alumni back for a reunion weekend



Welcome back, MIT Sloan alumni! It's reunion weekend, the perfect time for those who went from students living on instant ramen to professionals indulging in artisan avocado toast. We've missed you—well, mostly your tuition fees, but you get the point.

I USE AI EVERYDAY: MULTI-MODAL ASSISTANT



ChatGPT 4o ▾



Explain this photo



This photo shows a person holding a piece of paper with the handwritten message "HELLO MIT." The person is smiling and appears to be in an indoor setting, possibly an office or classroom. The individual is wearing a gray zip-up jacket and is positioned next to a round table and a chair. The setting looks clean and well-lit, suggesting a professional or academic environment. The message on the paper indicates a greeting directed at MIT, possibly signifying a connection or communication with the Massachusetts Institute of Technology.



AI IS HELPFUL AS A CREATIVE TOOL

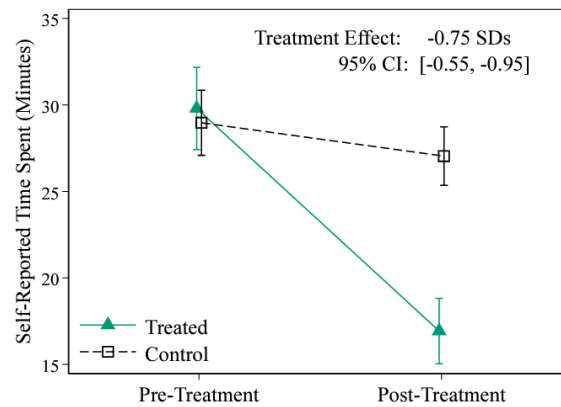
- Some people are surprised to hear AI is quite powerful for creativity
- But, breakthroughs commonly occur by connecting seemingly unrelated but pre-existing ideas
- AI is so powerful as a creative tool because it is trained specifically to make connections!
 - Remember that LLMs are trained by studying connections between “tokens” that potentially seem unrelated to humans but represent a deeper meaning

GEN-AI RAISES HUMAN PRODUCTIVITY

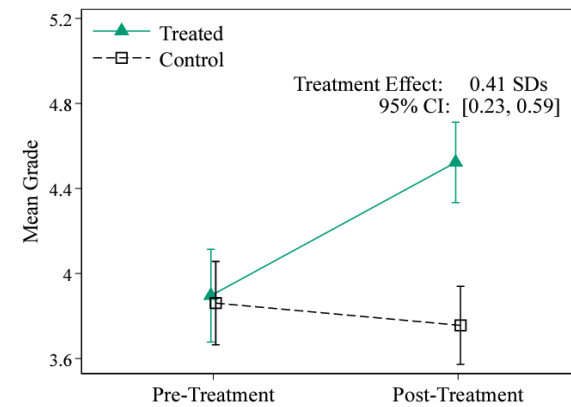
Writing tasks: Noy and Zhang (2023) show that lower-productivity workers, given access to ChatGPT, improve performance in writing tasks.



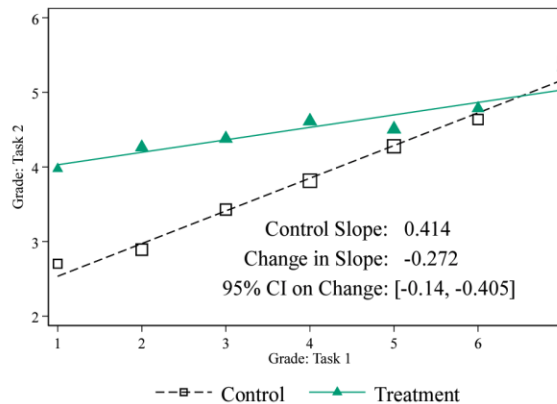
Time taken decreases



...improve grades...



It can also help level the playing field.



...and lessen grade inequality.

GEN-AI RAISES HUMAN PRODUCTIVITY



Programming: Peng et al. (2023) show software engineers “treated” with GitHub Copilot code twice as fast.

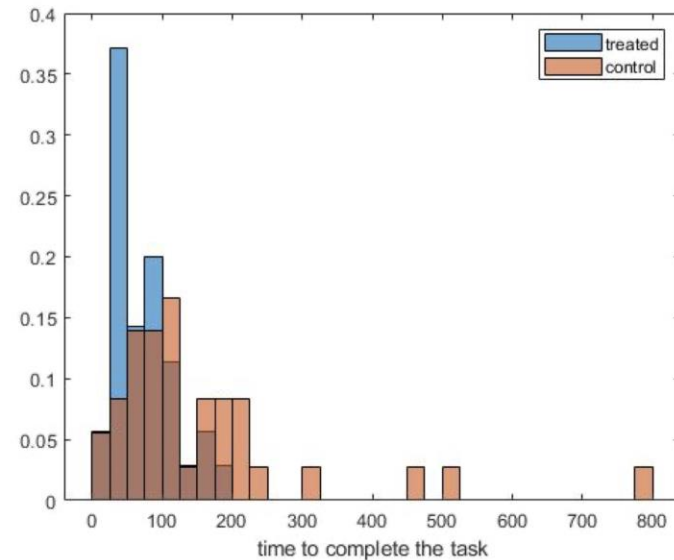


Figure 6: Time to task completion

Note: Distribution of time to task completion between treated (blue) and control (orange) groups

Coders “treated” with an AI pair programmer were more likely to complete tasks and did so 55.8% faster

Coders increasingly outsource mundane work and shift focus to creative and more complicated problems



AI CHALLENGES

- **AI is already quite useful, but the news is not all good.**
- **For example, the mainstream arrival of AI poses a significant challenge to educators.**
 - Ethan Mollick describes “The Homework Apocalypse”



AI CHALLENGES

- **Can you detect AI writing?**
 - No.
- **Aren't there services designed to detect AI writing? Yes, but...**
 - To the extent that they work at all, they can be defeated by making slight changes to text.
 - Detectors also have high false positive rates, and may falsely flag people of using AI, especially non-native speakers
- **The implication for educators is that students are already using, and will continue to use, AI**
 - Shift toward more in-person evaluation and paper exams
 - Discerning good-AI output from bad-AI output and AI skills



AI CHALLENGES

- Research shows that LLM outputs can be problematic
- Because LLMs are trained on the internet are reinforced by human feedback, then tend to:
 - Display biases against women and minorities
 - Reinforce stereotypes
 - Hallucinate
- There is also evidence that use of AI can, in some circumstances, make people lazy or careless



AI CHALLENGES

- Age of Manipulation: AI can massively expand how information is presented in misleading ways.



- Philosopher Hannah Arendt: “If everybody lies to you, the consequence is not that you believe the lies, but rather that nobody believes anything any longer.”



MIT RESEARCH ON AI: IDE.MIT.EDU

The screenshot shows the homepage of the MIT Initiative on the Digital Economy (IDE). At the top left is the MIT Digital logo with the text 'MIT INITIATIVE ON THE DIGITAL ECONOMY'. To the right is a 'MENU' icon. The main heading is 'The MIT Initiative on the Digital Economy (IDE)' followed by the tagline 'Shaping a brighter digital future.' Below this is a red-bordered button labeled 'About the IDE'. A video player features a portrait of Andy McAfee, identified as 'IDE Co-Director & Tech for Good Research Lead', with a 'Play Video' button. The video background shows a man in a suit in a hallway with a colorful floor.

**Join us select Wednesdays this Spring
2024**

Explore All Research Groups

AI, Marketplaces, And Labor Economics

Building A Distributed Economy

Human-First AI

Misinformation & Fake News

New Data Analytics

Tech For Good

Technologies That Create Prosperity

Generative AI And Decentralization

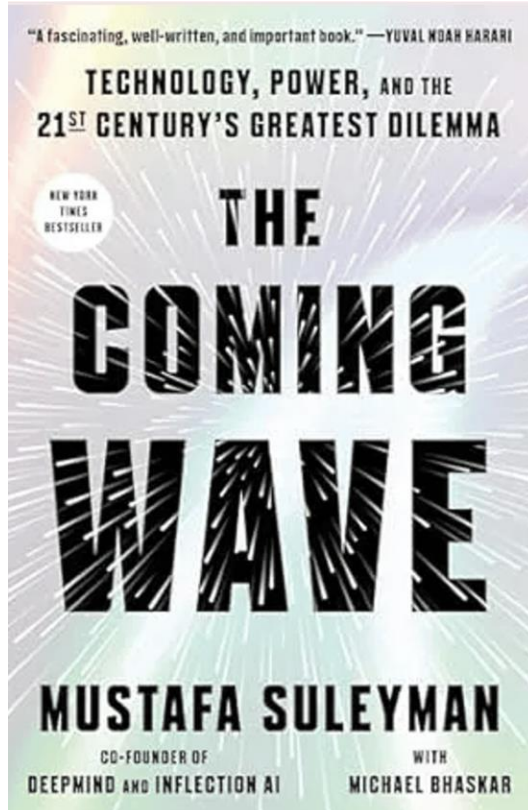


OUTLINE

- AI at MIT Sloan
- Power and Pitfalls of Current AI Models
- Looking Back and Looking Ahead



LESSONS FROM HISTORY



- Throughout human history, technological waves have followed a similar pattern
 - E.g., factory systems, the printing process, electricity, and the internet
 - As new technology emerges, people use it to expand the production of goods, make better products, and/or produce more efficiently at lower prices
- Competition drives more demand for tech that produce these goods
 - As a result, the tech itself becomes cheaper, more powerful, and more widely available



LESSONS FROM HISTORY



- The J-Curve: Many technologies start out as niche but quickly proliferate
- Computing began as the focus of obscure academic papers and niche laboratories
 - Computing, like other technological waves, quickly transformed the nature of work
- Spectacular advances in computers enriched many over the last few decades, but also likely contributed to increased inequality

LESSONS FROM HISTORY

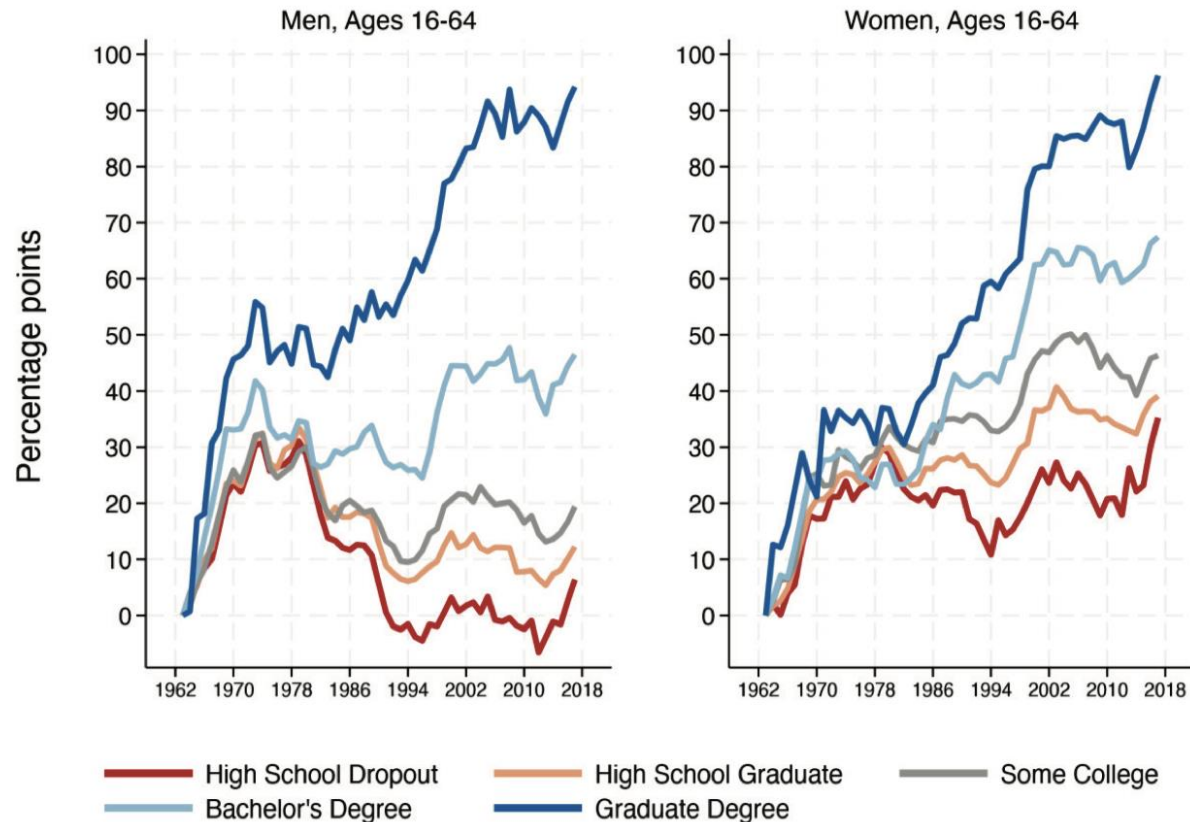


Figure 2: Cumulative percentage point changes in real mean weekly earnings of full-time, full-year workers ages 18–64, United States, 1963–2017

History suggests AI will change the nature of work

- Skill-based technologies drove disparities in wage growth across high- vs low skill workers
- AI could continue these polarizing trends for people who know how to use AI effectively



A FRAMEWORK FOR ADAPTING TO AI

- I underscore to my students the importance of staying up to date in a quickly evolving world of AI
- When asked, I have given students the following advice/perspective:
 - There is seemingly undue focus on being replaced by AI. The more immediate concern is being replaced by a person that better knows how to use AI
 - Central Message: Staying up to date with AI is a growing imperative



ADAPTING TO AI: A FRAMEWORK

- To help people and organizations adapt to AI, I introduced the ADAPT Framework for AI Integration
- A – **Assess**: Build awareness of AI tools and their abilities
- D – **Define**: Define potential use-cases and goals for incorporating AI
- A – **Allocate**: Allocate resources (time, attention, and funds) for accessing AI tools
- P – **Pilot**: Use AI for all tasks for which it is legally and ethically permissible
- I – **Transform**: Foster an environment of experimentation and sharing, and scale up your pilot projects

AI AT WORK

- Consider forming an ADAPT Task Force that helps organize efforts around AI
- The goal is to discuss and attempt applications of AI that can be showcased to others:
 - Creates a discussion forum for technological/fiscal/ethical issues about using AI
 - Reduces redundant efforts
 - Helps map the “jagged frontier” and disseminate AI knowledge
 - Sifts through the noise from people that know what the firm is seeking to accomplish



- AI offers significant promise to reshape work by giving us more time for tasks that are more interesting, meaningful, and higher value.
- The optimistic vision of us being more productive & fulfilled is not inevitable.
- Leaders must decide whether and how to commit resources to build AI into organizations in a way that is mutual beneficial, safe, and ethical.



PLANNING FOR THE FUTURE

- It is important to plan for the future with the perspective that AI is rapidly evolving
- Many AI skeptics are focused on the flaws that LLMs have today: hallucinations, short context windows, slow answers, and so on.
 - But AI is advancing rapidly, and some of these concerns may soon vanish
- So don't just build for what is possible today, but what is possible in [the near future].

WRAPPING UP



AI AND EDUCATION AT MIT



Delivery

- AI tutors and chatbots
- More in-depth feedback
- More time for content creation and mentoring

Evaluation

- In-person exams and tasks
- Asking more of students
- Discerning good vs bad AI

Skills

- Effective prompting
- LangChain and custom processes
- Fine-tuning and building models

(ERIC) SO WHAT? SO WHAT'S NEXT?



- **So What?**
 - The AI boom represents a pivotal shift in tech likely more impactful than the internet
 - The internet reduced the cost of broadcasting information, whereas AI reduces the cost of acting on it. So far, there is no obvious upper limit on what's possible.
- **So What's Next?**
 - There is a risk of overhyping the technology, but I think the bigger risk is not appreciating the magnitude of the coming changes.
 - We need to prepare our students for the future of work. We want your help!



THANK YOU

- AI Hub: <https://mitsloanedtech.mit.edu/ai/>
- IDE Research Group: <https://ide.mit.edu>
- Contact Eric So: eso@mit.edu

THANK YOU

Feel free to contact me: eso@mit.edu