

## *PHIL S183: Philosophy and Technology – DRAFT syllabus*

Instructor: Harry R. Lloyd

Contact: [harry.lloyd@yale.edu](mailto:harry.lloyd@yale.edu)

Office hours: tbd

Dates: Session B, July 1 – August 2, 2024

Location: tbd

Meeting times: MW 9.00 – 12.15

### 1: Course description

The last few years have witnessed significant improvements in artificial intelligence and virtual reality technologies. Within our lifetimes, it may become possible to spend time in extremely realistic computer simulated worlds, populated by artificially intelligent agents. This prospect forces us to grapple with a host of important questions, including: whether artificially intelligent agents will ever achieve consciousness and moral status; whether we could survive ‘uploading’ our minds into a virtual reality; whether or not this would be good for us; and whether we can know if we are already living inside a simulation. In this class we will discuss all of those questions, as well as several ethical dilemmas posed by present-day technologies like dating apps and self-driving cars. (Is it immoral to ‘swipe left’ on potential romantic partners based on physical appearance? How should self-driving cars handle ethical trade-offs?) The class will not presume any background in philosophy or technology studies, and will use technological developments as a springboard into philosophy.

### 2: Course structure

Assessments will be as follows:

- 25% reading responses
  - 24 hours before the start of each class, students will submit a response to the readings via Canvas. These responses should not summarise the readings. Rather, the student should aim to engage substantively and philosophically with one or more of the assigned authors. For instance, you might do one or more of the following things: (1) ask critical questions about the readings; (2) raise an objection to the author’s arguments; (3) suggest a new argument in favour of the author’s conclusion, or a way in which the author’s arguments could be strengthened; (4) suggest that the author’s arguments have interesting implications for other topics or areas that the author does not discuss. The typical length might be 250-300 words – one good, solid paragraph.
- 25% participation & quizzes
  - Student must attend and participate in every meeting of the class. Each meeting will include time for student discussion. Students should come to class having studied the assigned readings, and ready to discuss and debate with each other. Participation grading will partly reflect class citizenship (are you listening and responding to other people’s ideas? Are you contributing actively to the discussion? Are you giving others the space to talk?). I will also give credit for the

philosophical quality of your contributions. But don't be nervous about this! I will *award* marks for good quality contributions; but I won't ever *subtract* marks for mistakes or blunders.

- I will also make time for a few unannounced quizzes throughout the semester. For instance, at the start of a couple of meetings, I might hand around a flash quiz on the last meeting's readings or lecture.
- 25% midterm take-home exam
- 25% final exam

### 3: Schedule of readings

I have listed the readings for each week in the order that I recommend you read them. There will be two textbooks for the course:

- Michael Huemer. 2021. *Knowledge, Reality, and Value: A Mostly Common Sense Guide to Philosophy*
- David J. Chalmers. 2022. *Reality+: Virtual Worlds and The Problems of Philosophy*

You should purchase copies of these two texts before the start of the summer session (both should be inexpensive online). All other readings will be uploaded to the class Canvas page.

#### **Monday 1<sup>st</sup> July, Class 1: Introduction to analytic philosophy and technology**

- Huemer, *Knowledge, Reality, and Value*
  - Chapters 1 and 2 [~25 pages, readable]
- Chalmers, *Reality+*
  - Chapters 1-3 [~60 pages, readable]

#### **Wednesday 3<sup>rd</sup> July, Class 2: Are we living in a simulation?**

- Chalmers, *Reality+*
  - Chapters 4 and 5 [~40 pages, readable]
- Huemer, *Knowledge, Reality, and Value*
  - Chapter 6 [~20 pages, readable]

#### **Monday 8<sup>th</sup> July, Class 3: Is the external world real?**

- Preston Greene. 2020. The termination risks of simulation science. *Erkenntnis*, 85, 489-509 [~20 pages, academic paper]
- Chalmers, *Reality+*
  - Chapters 6 and 10 [~35 pages, readable]

#### **Wednesday 10<sup>th</sup> July, Class 4: Can AIs have consciousness?**

- Chalmers, *Reality+*
  - Chapter 15 [~20 pages, readable]
- Eric Schwitzgebel. 2015. If materialism is true, the United States is probably conscious. *Philosophical Studies*, 17, 1697-721 [~20 pages, academic paper]

- Ned Block. 1978. The China brain. From ‘Troubles with Functionalism’ [3 pages, paper excerpt]

### **Monday 15<sup>th</sup> July, Class 5: Could we survive mind uploading?**

- Huemer, *Knowledge, Reality, and Value*
  - Chapter 12 [~15 pages, readable]
- David J. Chalmers. 2014. Mind uploading: a philosophical analysis. From *Intelligence Unbound: The Future of Uploaded and Machine Minds* [~15 pages, readable]
- Shelly Kagan. 2012. *Death*
  - Chapters 6-7 [~70 pages, readable]

### **Wednesday 17<sup>th</sup> July, Class 6: Would living in VR be good for us?**

- Chalmers, *Reality+*
  - Chapters 17 and 20, and online appendix ‘Free will in the experience machine and in virtual reality?’ [~35 pages, readable]
- Derek Parfit. 1984. *Reasons and Persons*
  - Appendix I [~10 pages, academic book]
- Ben Bradley. 2015. *Well-Being*
  - Excerpts on hedonism: §§2.1, 2.2, 8 [~15 pages, readable]

### **Monday 22<sup>nd</sup> July, Class 7: Can AIs have moral standing?**

- Chalmers, *Reality+*
  - Chapter 18 [~20 pages, readable]
- Simon Goldstein and Cameron Domenico Kirk-Giannini. Unpublished paper. AI wellbeing. [~25 pages, academic paper]

### **Wednesday 24<sup>th</sup> July, Class 8: Self-driving cars and trolley problems**

- Huemer, *Knowledge, Reality, and Value*
  - Chapters 14 and 15 [~30 pages, readable]
- Judith Jarvis Thomson. 1976. Killing, letting die, and the trolley problem. *Monist*, 59, 204-17 [~15 pages, academic paper]
- William Ratoff. 2022. Self-driving cars and the right to drive. *Philosophy and Technology*, 35, 57. [~15 pages, academic paper]

### **Monday 29<sup>th</sup> July, Class 9: Is it immoral to ‘swipe left’ based on looks?**

- William D’Alessandro. 2023. Is it bad to prefer attractive partners? *Journal of the American Philosophical Association*, 9.2, 335-54. [~20 pages, academic paper]
- John Danaher. 2020. A defence of sexual inclusion. *Social Theory and Practice*, 46, 467-96 [~30 pages, academic paper]
- Tena Thau. 2020. Expanding the romantic circle. *Ethical Theory and Moral Practice*, 23, 915-29 [~15 pages, academic paper]
- Ross Douhat. 2018. The redistribution of sex. *New York Times* [short opinion piece]

**Wednesday 31<sup>st</sup> July, Class 10: Final exam, no readings**

- ---