

Minds and Machines

Lecture 3

David Chalmers
NYU Philosophy

Previously

- Technophilosophy
- Ancient philosophy: questions about knowledge, reality, and value.

Today

- Skepticism and simulation
- *Reality+* chapters 2 and 3

Office Hours

- Office hours this week: Thursday 5-6pm

Philosophy

- Philosophy = (literally) love of wisdom.
- In practice: philosophy asks fundamental questions about reality, knowledge, and value.

Areas of Philosophy

- Metaphysics: theory of reality
 - What exists?
- Epistemology: theory of knowledge
 - What can we know?
- Value theory: theory of value
 - What is good?

Metaphysics

- What is the nature of reality?
- What exists? (ontology)
- What is the mind? (philosophy of mind)
- Is there a God? (philosophy of religion)

Value Theory

- What is good?
- What are right and wrong? (ethics)
- What is a just society? (political philosophy)
- What is beauty? (aesthetics)
- What is a life worth living?

Epistemology

- What is knowledge?
- Can we know anything at all? (skepticism)
- What can science tell us about the world?
(philosophy of science)

Knowledge

- We take ourselves to know many things.
- We know ...

Knowledge

- We take ourselves to know many things:
- We know
 - that we're in New York
 - that water is made of H₂O
 - that Biden is president

Questions

- Do we really know these things?
- Could we be wrong?
- Can you be sure that you know what you think you know?

Knowledge and Belief

- If you know something, you believe it.
- But some things that we believe, we don't know.
 - We've all made mistakes! Or believed things without sufficient reasons.
- So knowledge requires more than belief. What more does it require?

Justified True Belief

- Traditional answer (Plato): knowledge requires *justified true belief*.
- Your belief must be true (mistaken beliefs aren't knowledge)
- Your belief must be justified (beliefs held without good reasons aren't knowledge).
- Maybe even more is required (a long story).

Skepticism

- Skepticism: We don't know anything.

Varieties of Skepticism

- Global skepticism: We don't know anything at all.
- Local skepticism: We don't know anything about a specific domain
 - Future skepticism: We don't know anything about the future
 - Climate-change skepticism: We don't know anything about climate change.

External-World Skepticism

- We don't know anything about the external world.
- We don't even know that anything outside ourselves exists.
- This is consistent with knowing some things about ourselves, and perhaps about logic, mathematics, etc.

Skepticism vs Denialism

- The denialist says: there is no external world (cf. there is no climate change).
- The skeptic isn't a denialist: she just says, we don't *know* that there's an external world.


Skeptical Hypotheses

- Zhuangzi: You might be dreaming!
- Descartes: You might be being fooled by an evil demon!
- Putnam: You might be a brain in a vat!
- Contemporary: You might be in a computer simulation!

Simulation Hypothesis

- We'll work especially with the simulation hypothesis:
 - We are living in a computer simulation
 - I.e. we are and always have been receiving our sensory inputs from an artificially designed computer simulation of a world.


I'm sitting in a chair at NYU, attending a class on Minds and Machines



- <https://www.youtube.com/watch?v=nLRGI47Cudg>

Pure and Impure Sims

- Pure sim hypothesis: We are simulated creatures in a simulation.
- Impure sim hypothesis: We are unsimulated creatures connected to a simulation
- Biosim hypothesis: We are biological creatures connected to a simulation



**TOM the
DANCING
BUG**

BY
RUBEN
BOLLING

**THE
OUTER
REACHES
OF
PLOT
TWISTS**

*Please don't tell your
friends the shocking
beginning, middle and
ending!*





SUBMITTED FOR YOUR APPROVAL: WES SHELBY. AN ORDINARY MAN ABOUT TO UNDERGO A RECORD-SETTING NUMBER OF REALITY-BENDING AND UTTERLY SHOCKING PLOT TWISTS!

SOMETHING ABOUT MY LIFE FEELS VERY SPOOKY, SLOWLY PACED AND POORLY LIT.

BUT WHAT IS IT?







BUT... COULD THERE BE YET ANOTHER STARTLING REVELATION? AH, IT'S OBVIOUS NOW--

I'M DEAD!

SHOCKING? YES, BUT NOW THIS GHOST-BRAIN MUST WAKE UP TO THE REALIZATION THAT ITS FALSE-REALITY SHOCKS WERE ACTUALLY MERE ILLUSIONS...!

CONJURED FIGMENTS OF A HACK SCREEN-WRITER'S LIMITED IMAGINATION!

GOOD LORD! I'M FICTIONAL!

THE END?

Question

- Can you know whether you're not in a computer simulation?
 - If so: why and how?
 - If not: why not?

Evidence You're Not in a Simulation?

- Are your hands evidence that you're not in a simulation?
- Other people?
- The vast universe?
- Your own consciousness?

Evidence You Are in a Simulation? (Sim Signs)

- Glitches in the matrix?
- You're in a class about simulations?
- ...

SIGNS YOU MAY BE A SIM



Arguments

- There is a long tradition of arguing for skepticism in philosophy.
- Q: How do you argue for a philosophical conclusion?
- One way: via a philosophical argument.

A Formal Argument

1. Socrates is human
2. All humans are mortal.

3. Socrates is mortal.

Formal Arguments

- In philosophy a formal argument has *premises* and a *conclusion*.
- What connects the premises to the conclusion is “Therefore” or “so”.
- The conclusion is supposed to follow from the premises.

Validity and Soundness

- If the conclusion follows from the premises, the argument is *valid*.
- If the argument is valid and the premises are true, the argument is *sound*.
- For any sound argument, the conclusion is true!

An Argument about Simulations

1. In a simulation, everything would seem the same as it does to me.
 2. If everything would seem the same to me in a simulation, I can't know I'm not in a simulation
-
3. I can't know I'm not in a simulation.

Assessing the Argument

- The argument seems to be valid: the conclusion follows from the premises.
- So if you accept the premises, you must accept the conclusion.
- To reject the conclusion, you must reject one of the premises: which one?

External-World Skepticism?

- If I don't know that I'm not a brain in a vat, then it's hard to see how I can know anything about the external world.
- External-world skepticism!

Argument for Skepticism

1. I don't know whether I'm in a simulation.
 2. If I don't know whether I'm in a simulation, I don't know that I have hands.
-
3. I don't know that I have hands
[or that I'm at NYU, or that Biden is president, or
anything about external reality]

Assessing the Argument

- The argument seems to be valid: the conclusion follows from the premises.
- So if you accept the premises, you must accept the conclusion.
- To reject the conclusion, you must reject one of the premises: which one?

Other Skeptical Scenarios

- I'm being deceived by an evil genius (Descartes).
- I'm dreaming right now (Descartes).
- I'm in a Matrix-style computer simulation.
- I'm connected to a virtual reality device right now.

Next Class

- 9/12: Descartes on skepticism (*First and Second Meditations*).