Epistemic Logic notes

PHIL 735 Week 2
Matt Kotzen and Jim Pryor

September 6, 2023

1 Basics of Modal Logic

- Modal logics are formal systems that allow us to capture the logic of a variety of different "necessitation" (or, perhaps, necessitation-like) relations.
- We're going to focus on "normal" modal logics, the weakest of which is System K (for Kripke).
  - System K is just propositional logic plus $\Box$, which this system understands to obey particular constraints. There are multiple equivalent ways to axiomatize System K:
    - The standard way to axiomatize System K:
      - (Rule of Necessitation): If p is a theorem of K, then so is $\Box p$
      - (Axiom K): $\Box(p \supset q) \supset (\Box p \supset \Box q)$
    - Another (perhaps more perspicuous) way to axiomatize System K:
      - (Rule of Equivalents) If $p \subset\supset q$ is a theorem of K, then $\Box p \subset\supset \Box q$ is a theorem too
      - (Axiom N) $\Box\top$
      - (Axiom M) $\Box(p \wedge q) \supset (\Box p \wedge \Box q)$
      - (Axiom Converse-M) $(\Box p \wedge \Box q) \supset \Box(p \wedge q)$
  - We understand $\Diamond p$ to be the "dual" of $\Box$ — i.e., to mean $\neg\Box\neg p$
- Should we also assume (Axiom T): $\Box p \supset p$?
  - If yes, we get System T
  - This looks to depend on what kind of necessitation relation we're dealing with
    - Metaphysical necessity: Yes.
    - Deontic necessity: Almost certainly no.

        * Other kinds of necessity?

- If we reject (Axiom T), we need to decide whether we want (Axiom D): $\Box p \supset \Diamond p$
    - If yes, we get System D
    - Again, this may depend on the necessitation relation under consideration
        * Metaphysical necessity? Yes.
        * Deontic necessity? Almost certainly yes.
        * Other kinds of necessity?
    - Some deontic logicians also insist on $\Box(\Box p \supset p)$
- Suppose we accept (Axiom T)
    - Should we also accept (Axiom 4): $\Box p \supset \Box\Box p$?
        * If yes, we get System $S4$
        * In $S4$, iterated $\Box$'s collapse to one $\Box$, and iterated $\Diamond$'s collapse to one $\Diamond$
    - How about (Axiom B): $p \supset \Box\Diamond p$ (i.e., $p \supset \Box\neg\Box\neg p$, or $\neg\Box\neg\Box q \supset q$, or $\Diamond\Box q \supset q$)?
        * If yes, we get System B
    - How about (Axiom 5): $\Diamond p \supset \Box\Diamond p$ (i.e., $\neg\Box q \supset \Box\neg\Box q$)?
        * If yes, we get System $S5$
        * In $S5$, all that matters is the last operator in a sequence of $\Box$'s and $\Diamond$'s
- It turns out that, once you've accepted (Axiom T), accepting (Axiom 5) is equivalent to accepting both (Axiom B) and (Axiom 4)
- You can also add (Axiom 4), or (Axiom 5), or both, to System K or to System D. You can also add (Axiom B) to System K or to System D, though that won't be very relevant for our purposes.

2 Model Theory

- In order to give a semantics for $\Box$ (and hence $\Diamond$, on our understanding of $\Diamond$), we need:
    - A set of possible worlds, Ω
    - R, a binary relation over Ω (i.e., R's graph is a subset of Ω x Ω), known as the accessibility relation
    - I, which tells you which atomic propositions are true at each world
- The first two things are together called a "frame," and the three things together are called a "model"

- A model plus a choice of a world "makes true" or "satisfies" a formula ϕ
  - We write this $[[\phi]]_{M,w} = \text{true}$
  - Or: $M, w \vDash \phi$
- And we assume that:
  - $M, w \vDash \phi \wedge \psi$ iff $M, w \vDash \phi$ and $M, w \vDash \psi$
  - $M, w \vDash \phi \vee \psi$ iff $M, w \vDash \phi$ or $M, w \vDash \psi$
  - $M, w \vDash \neg\phi$ iff $M, w \nvDash \phi$
  - etc.
- And here's the crucial part:
  - $M, w \vDash \Box\phi$ iff, for each $w, u \in \Omega$ such that $w\mathrm{R}u$: $M, u \vDash \phi$
  - i.e., $\Box\phi$ is true at w iff ϕ is true at every world accessible from w
- A consequence of the above is that $\Diamond\phi$ will be true at w iff ϕ is true at *any* world accessible from w
- OK, now, let's go back to the modal logics we discussed above
  - Assuming (Axiom D) is equivalent to assuming that access is serial
  - Assuming (Axiom T) is equivalent to assuming that access is reflexive
  - Assuming (Axiom 4) is equivalent to assuming that access is transitive
  - Assuming (Axiom B) is equivalent to assuming that access is symmetric
  - Assuming (Axiom 5) is equivalent to assuming that access is R-Euclidean
- Entailments among these properties of relations:
  - R-Euclideanness alone doesn't entail symmetry, transitivity, or reflexivity
  - But reflexivity and R-Euclideanness together entail both symmetry and transitivity
    - Proof?
  - Reflexivity alone doesn't entail R-Euclideanness, transitivity, or symmetry; but reflexivity alone does entail seriality
  - Neither symmetry nor transitivity entails each other, even when we assume reflexivity.
  - So System $S5$ is (strictly) stronger than System $S4$, System B, and System T; System $S4$ and System B are each (strictly) stronger than System T; System $S4$ is neither stronger than nor weaker than System B; and System T is (strictly) strongly than System D. Each of these systems is (strictly) stronger than System K.

3 Logic of Knowledge

- Actual vs. ideal agents
- If $\Box$ is interpreted as "knows," then $\Diamond$ gets interpreted as "doesn't know to be false"
- Do we want (Axiom T) for knowledge? Almost certainly yes.
- Do we want (Axiom 4) for knowledge? Plausible.
- Do we want (Axiom B) for knowledge? Implausible.
- Assuming we want (Axiom T) and (Axiom 4) and don't want (Axiom B), that means we can't have (Axiom 5) (which we may not have wanted anyway)
- So we probably want something like $S4$, perhaps with some additional axioms that don't get us all the way up to $S5$. Some possible additional axioms:
  - (Axiom .2): $\Diamond\Box p \supset \Box\Diamond p$. Corresponds to assuming that access is "one-step confluent"—i.e., that if $x\mathrm{R}y$ and $x\mathrm{R}z$, then $\exists u(y\mathrm{R}u \wedge z\mathrm{R}u)$.
  - (Axiom .3): Lots of different ways to formulate, but one is: $\Box(\Box p \supset q) \vee \Box(\Box q \supset p)$. Corresponds to assuming that access is non-R-branching—i.e., that if $x\mathrm{R}y$ and $x\mathrm{R}z$ and $y \neq z$, then $y\mathrm{R}z$ or $z\mathrm{R}y$.
    - Another way to formulate (Axiom .3): $(\Diamond p \wedge \Diamond q) \supset [\Diamond(p \wedge \Diamond q) \vee \Diamond(p \wedge q) \vee \Diamond(\Diamond p \wedge q)]$
  - Note that, if we are assuming (Axiom T), then (Axiom .3) is strictly stronger than (Axiom .2)
    - But if we are not assuming (Axiom T), then (Axiom .3) is independent of (Axiom .2)
    - Regardless, accepting (Axiom .2) and (Axiom .3) is equivalent to accepting: (Axiom .3.2) $(\Diamond p \wedge \Diamond\Box q) \supset \Box(\Diamond p \vee q)$. This corresponds to assuming that access is semi-Euclidean: if $x\mathrm{R}y$ and $x\mathrm{R}z$, then $z\mathrm{R}x$ or $y\mathrm{R}z$.
  - (Axiom .4): $(p \wedge \Diamond\Box p) \supset \Box p$. As far as we're aware, this doesn't correspond to an easily articulable constraint on R.
  - System $S4$ supplemented with (Axiom .2) is called System $S4.2$, and similarly for (Axiom .3) and (Axiom .4)

4 Logic of Belief

- If $\Box$ is interpreted as "believes," then $\Diamond$ gets interpreted as "doesn't believe to be false"

- Knowledge is generally thought to be factive, so (Axiom T) was plausible for knowledge. But belief isn't factive, so (Axiom T) is implausible for belief. So we're "below" System T in the hierarchy of modal logics.
- Do we want (Axiom D) for belief? Plausible.
- Do we want (Axiom 4) for belief? Plausible.
- Do we want (Axiom B) for belief? Implausible.
- Do we want (Axiom 5) for belief? Maybe. Note that, since we're not assuming (Axiom T), we can have (Axiom 5) without committing to (Axiom B).
- So, System $D4$ (sometimes also called $KD4$) is one plausible logic for belief. If we accept (Axiom 5) too, we get System $D45$ (sometimes also called System $KD45$).
  - Do we also want (Axiom .2), (Axiom .3), and/or (Axiom .4) for belief? Recall that if we're rejecting (Axiom T), then (Axiom .2) and (Axiom .3) are independent.

5 Relation Between Knowledge and Belief

- (Axiom KB1): $Kp \supset Bp$. Very plausible.
- (Axiom KB2): $Bp \supset KBp$. Plausible.
- (Axiom KB3): $Bp \supset BKp$. AKA "Moore Principle." More controversial.
- Note that if knoweldge obeys (Axiom 5) and belief obeys (Axiom D), then (Axiom KB1) allows us to prove $BKp \supset Kp$, which is very implausible.
  - Assume BKp. Since B obeys (Axiom D), $BKp \supset \neg B\neg Kp$. By modus ponens, $\neg B\neg Kp$. By (Axiom KB1), $K\neg Kp \supset B\neg Kp$. By modus tollens, $\neg K\neg Kp$. Since K obeys (Axiom 5), $\neg Kp \supset K\neg Kp$. By modus tollens, Kp.
- Perhaps even worse: once we have $BKp \supset Kp$, (Axiom KB3) let us derive $Bp \supset Kp$.
  - Assume Bp. By (Axiom KB3), BKp. From above, Kp.
- Then, assuming only that knowledge obeys (Axiom T), we get $Bp \supset p$, which is *really* bad.
- If we want (Axiom 5) for knowledge, this puts pressure on us to give up (Axiom KB1). But the better lesson is probably that we should just reject (Axiom 5) for knowledge.
- Do we want what Meyer calls "cross-over negative introspection"—i.e., (KB4): $\neg Bp \supset K\neg Bp$ and/or (KB5): $\neg Kp \supset B\neg Kp$?