Week 1: Epistemic justification; Foundationalism vs. coherentism

1. Preliminaries: Epistemic justification
    a. 'What am I (epistemically) justified in believing': what *ought* I to believe/what is it *reasonable* or *rational* for me to believe?
    b. Epistemic vs pragmatic justification: Pascal's Wager provides a pragmatic, not an epistemic, justification for believing in a particular kind of god.
    c. Two putative 'adequacy conditions' on accounts of justification
        i. The metajustification criterion: We must have some reason for thinking that beliefs that meets the account's criteria for counting as 'justified' are (as a result) likely to be *true*.
        ii. The knowledge-directedness criterion: Knowledge must be analysable as JTB + X, for some suitable property X.
    d. A policy warning!
2. Introducing the foundationalism-coherentism debate
    a. Some justified beliefs are justified *by other beliefs*.
    b. If we follow the 'justificatory chains' backwards, one of four things must happen:
        i. The chains terminate at *unjustified* beliefs. (Total scepticism)
        ii. The chains continue forever. (Also scepticism? Cf. Sosa, "The raft and the pyramid...")
        iii. The chains terminate at beliefs that are justified, but that are not justified *by other beliefs*. (Foundationalism)
        iv. The chains eventually 'loop back on themselves'. (Coherentism)
    c. We will assume that scepticism is false.
3. Foundationalism
    a. Classical foundationalism: the foundational beliefs are beliefs *about one's own present experiences*, and the non-foundational justified beliefs are related to the foundational ones *by deductive inference*.
        i. This leads to a radical form of scepticism.
    b. Wide foundations: allow foundational status to e.g.
        i. Perceptually justified propositions *about the external world*;
        ii. The contents of memory;
        iii. The contents of testimony (?)
    c. Ampliative inference: allow some forms of inference that are not deductively valid to (nevertheless) transmit justification.
    d. Some more permissive ('moderate') versions of foundationalism:
        i. Type 2: wide foundations, but only deductive inference
        ii. Type 3: Narrow foundations, but ampliative inference
        iii. Type 4: Wide foundations and ampliative inference
            1. Types 3 and 4 (plausibly) don't lead to scepticism.
    e. Objections to moderate foundationalism

        i. Objection 1: No form of foundationalism other than the (sceptical) classical one can meet the meta-justification requirement.
            1. Possible reply: externalism [on which more in week 2!]
        ii. Objection 2: It's just not plausible that any of the alleged wide-foundational beliefs is justified *regardless of which other beliefs we hold* – e.g., regardless of our beliefs about the reliability or otherwise of our perceptual faculties/memory/informant.
            1. Possible reply: externalism, again

4. Coherentism (See e.g. Bonjour, 'The elements of coherentism')
    a. The basic idea: Not all circular arguments are *viciously* circular – some confer justification on all members of the circle.
        i. A pressing question: what's the difference, then?
    b. Intuitive objection: No circular argument can confer justification, because in order for one belief to justify another, the first must be justified *before* the second.
        i. Reply: reject the temporal metaphor.
    c. The coherentist's proposal: for a belief to be justified *just is* for it to be part of a 'coherent' overall body of beliefs.
    d. Clarifying 'coherence'
        i. Not just logical consistency
        ii. Not as much as: 'each belief entails and is entailed by every other'
        iii. Something like: each belief is *probabilified* by the remainder taken together.
            1. (So the coherentist needs ampliative inference too...)
    e. Objections to coherentism
        i. Objection 1: There are alternative and mutually incompatible coherent systems. The coherentist has nothing to say about what justifies believing one whole system rather than another.
        ii. Objection 2: coherentism cannot meet the metajustification requirement.
        iii. Objection 3: coherentism contains no place for input from experience.

5. Some questions
    a. Are you a foundationalist or a coherentist?
    b. What's the difference, if there is one, between a viciously circular argument and a virtuously circular one?
    c. Is the meta-justification criterion an adequacy condition on accounts of justification?
    d. Is there any [adequate] account of justification according to which scientific belief counts as (epistemically) justified, but religious belief does not?
    e. Is conservatism a legitimate principle of rational enquiry?

Week 2: Internalism and externalism about epistemic justification

1. The internalist intuition
    a. A test case: believing that there will be a street party this weekend
    b. The fundamental internalist intuition: There cannot be two individuals, identical 'from the skin in', such that the first individual's beliefs are justified, while those of the second are not. (Cf. the introductory paragraphs of Wedgwood's *Internalism explained*)
    c. Making internalism precise:
        i. Accessibilism: Only things to which an agent has *direct introspective access* can be justifiers for that agent's beliefs.
        ii. Mentalism: Only an agent's *own mental states* can be justifiers for that agent's beliefs.
2. The externalist move
    a. The externalist *denies* the fundamental internalist intuition.
    b. E.g. of an externalist account of justification: Process reliabilism
        i. Process reliabilism (roughly): A belief is justified iff it is *(in fact)* caused by a reliable process.
        ii. Clarification: "reliable" does not mean *100%* reliable.
        iii. A (serious) objection to process reliabilism: What is "the process" leading to a given belief? There seem to be too many equally good answers.
            1. This is the "generality problem" for reliabilism. (Cf Feldman and Conee, 'The generality problem for reliabilism', *Philosophical Studies* 89, pp.1-29, 1998.)
3. Some more test cases: Alice's beliefs that the president is in New York today (Bonjour, *Externalist theories of empirical knowledge*, sec. III)
    a. Suppose Alice believes that the president is in New York, merely on the basis of a hunch that he is. Q: Is her belief justified?
        i. ... if she believes, with no reason, that she has a clairvoyant power?
        ii. ...if she believes, with good reason, that she has a clairvoyant power, but in fact she has no such power?
        iii. ...if she neither believes nor has any reason to believe that she has a clairvoyant power, but in fact she does?
        iv. ...if she believes she has a clairvoyant power and indeed she does?
    b. The internalist is likely to answer 'yes' to (ii) and 'no' to (iii); vice versa for the externalist.
4. The allure of an externalist account
    a. It is easier for an externalist than for an internalist to defend common sense against the skeptic [about justification].
    b. It is easier for an externalist than for an internalist account to meet the "meta-justification requirement" (cf. week 1).

   c. It is easier for an externalist than for an internalist to defend foundationalism (again, cf. week 1).
5. Internalist objections to externalist accounts of justification
   a. Objection 1: An externalist notion of justification is irrelevant to epistemic *rationality/blame/duty/responsibility*.
   b. Objection 2: An externalist notion of justification provides no useful *guidance* or *advice* to an agent trying to decide what to believe.
6. Objections to internalism
   a. Objection 1: Internalism has no coherent motivation. (Goldman, *Internalism exposed*)
      i. Either the would-be internalist requires justifiers to be directly accessible to introspection, or she merely requires that they be accessible. But in the first case, her requirement leads to scepticism, and in the second case, her requirement fails to justify internalism.
   b. Objection 2: Nothing meets the internalist's "luminosity" condition (Williamson, Knowledge and its limits, ch. 4)
      i. Definition: Say that a condition (e.g., the condition of feeling cold) is "luminous" iff: any time an agent is in that condition, she's also in a position to *know* that she's in that condition.
      ii. Safety requirement on knowledge: If S knows that p, then S's belief that p could not easily have been false.
      iii. Williamson uses Sorites-like series to argue that the claim that a given condition (e.g. feeling cold) is luminous, when supplemented by the safety requirement on knowledge, leads to contradiction.
      iv. Williamson's conclusion: *no* (non-trivial) condition can be luminous.
      v. The intended moral: The thing the internalist wants is impossible. So she should revise her (theoretical) desires.
7. On pluralism
   a. A natural idea: We don't need to choose between internalism and externalism. We can just acknowledge that there are both internalist and externalist notions of justification, and use both as appropriate.
   b. This may be right, but
      i. We have yet to see what the externalist notion of justification is useful for;
      ii. The anti-internalism arguments argued not merely that internalist justification failed to coincide with some intuitive notion, but (respectively) that internalist justification is *unmotivated* or *impossible*. Objections of this form impugn pluralism just as much as they impugn a monist internalism. So pluralists still need replies to those objections.

Week 3: Scepticism

1. Sceptical hypotheses
    a. A sceptical hypothesis (H) has the following features:
        i. It's logically possible;
        ii. If it were true, then (1) you would be having all the same experiences as those you're actually having, but (2) many of your ordinary beliefs (O – e.g. that you're in a lecture, that you have hands) would be false.
    b. Examples of sceptical hypotheses: you're dreaming, you're a brain in a vat
    c. Sceptics about knowledge (respectively, about justification) think that the existence of such hypotheses threaten our ordinary beliefs' claims to be cases of knowledge (resp. to be justified).
2. Scepticism about knowledge: the "argument from ignorance"
    (P1) You don't know that not-H.
    (P2) If you don't know that not-H, then you don't know that O.
    Therefore,
    (C) You don't know that O.
3. Three reasons not to accept the sceptic's conclusion (C)
    a. Methodological concern: the sceptic is applying the methodology of reflective equilibrium (http://plato.stanford.edu/entries/reflective-equilibrium/) incorrectly.
    b. We don't want (?) to be sceptics about *justification*. But whatever answer we give to justification-scepticism may apply to knowledge-scepticism too.
    c. Accepting (C) just amounts to throwing away a useful distinction.
4. Denying (P2): The 'Nozickean' response
    a. Nozickean accounts of knowledge (cf Nozick, *Knowledge and Scepticism*) hold that *sensitivity* is necessary for knowledge.
    b. Definition: S's belief that P is sensitive iff: if P had been false, S would not have believed that P. (S does not believe that P in the *closest* not-P world.)
    c. Applying this to scepticism: Your belief that you're not a BIV is not sensitive. Hence (P1) is true. But your belief that you have hands *is* sensitive. There's no reason to accept (P2).
5. Denying (P1): 'Moorean' and 'neo-Moorean' accounts
    a. Moore's notorious 'short way' with the sceptic (Moore, *Proof of an external world*): Here is a hand. Therefore there is an external world.
    b. 'Neo-Moorean' accounts
        i. Safety-based accounts
            1. Safety theorists of knowledge hold that *safety* (but *not* sensitivity) is necessary for knowledge.
            2. Roughly: S's belief that P is safe iff that belief *could not easily* have been false – i.e., there is no *close* possible world in which S continues to believe P but P is false.

3. Both your belief that O and your belief that not-H are safe.

ii. The semantic externalist response to the sceptic: more below!

6. Scepticism about justification

(P1') You're not justified in believing that not-H.
(P2') If you're not justified in believing that not-H, then you're not justified in believing that O.
Therefore,
(C') You're not justified in believing that O.

7. Accepting (C') and denying (P2') both seem(?) to be bad options.
8. The argument for (P1')
   a. Any justification that one is not (e.g.) a BIV must be either a priori or a posteriori.
   b. There is no a priori justification for believing that one is not a BIV.
   c. There is no a posteriori justification for believing that one is not a BIV.
   Therefore,
   d. There is no justification for believing that one is not a BIV.
9. Denying (P1')
   a. 'A priorist' responses
      i. Basic a priorism: the sceptical hypothesis is a priori less likely than the ordinary-world alternative.
      ii. Semantic externalism, again
   b. 'A posteriorist' responses
      i. Externalism about justification
      ii. 'Dogmatism' (Pryor, *The sceptic and the dogmatist*): Our perceptual experiences have external-world propositions as their content, and we are justified in believing that content in the absence of any positive reason not to.
10. More on semantic externalism (see also: http://plato.stanford.edu/entries/scepticism-content-externalism/)
    a. Intuitively, one might think that *what someone believes* depends only on goings-on *inside her head*.
    b. Semantic externalists deny this. They hold that the content of one's thoughts depends in part on the external causal history of those thoughts.
    a. Arguing for semantic externalism: Putnam's "Twin Earth" thought experiment (Putnam, *The meaning of 'meaning',* in his *Mind, Language and Reality: Philosophical Papers Volume 2*)
    c. Applying semantic externalism to scepticism: If we were brains in vats, our sentence "I am not a brain in a vat" would nevertheless express a truth – because our term "brain in a vat" would then mean something like: *simulation of a brain in a vat*.
    d. But even if this line of thought succeeds in proving that we are not brains in vats, it cannot refute more radical sceptical hypotheses, e.g. that we are *recently envatted* brains.

Week 4: The analysis of knowledge

1. The value of knowledge
    a. There seems to be *something* of value that is present in cases of knowledge, but missing in cases of 'lucky guessing'.
    b. One question: what is it that's missing? (This is 'the value problem' - http://plato.stanford.edu/entries/knowledge-value/)
    c. A prior question: what *is* knowledge, anyway?
2. The 'traditional account': knowledge is justified true belief
    a. ('Gettier') counterexamples: Smith and Jones, Stopped Clock... (Gettier, *Is justified true belief knowledge?*)
3. 'JTB + X' accounts
    a. Gettier examples show only that JTB isn't *sufficient* for knowledge. Perhaps we can find an adequate analysis by adding some fourth condition (X).
    b. The causal theory (Goldman, *A causal theory of knowing*): S knows that p iff S has a (justified?) true belief that p, *and* S's belief that p is caused by the fact that p.
        i. Counterexample: Fake Barn County
    c. The 'No false lemmas' theory: S knows that p iff S has a justified true belief that p, and S's justification for p does not proceed via any false propositions ('lemmas').
        i. Zagzebski's point (*The inescapability of Gettier problems*): Any "TB + Y" account will be subject to Gettier counterexamples, provided that the conditions Y do not *logically entail* that the belief is true: simply construct a case in which conditions Y obtain but the belief fails to be true, and then modify that case so that *by accident* the belief is true after all.
        ii. Applying this strategy to create a counterexample to No False Lemmas:
            1. James and the dog?
            2. Dr Walker and the virus
4. Nozick's "truth-tracking" account (Nozick, *Knowledge and skepticism*)
    a. The basic account: S knows that p iff
        i. P is true;
        ii. S believes that p;
        iii. If P had been false, then S would not have believed that P ('sensitivity');
        iv. If things had been slightly different but P had still been true, then S would still have believed that p (roughly, 'safety' – cf. also below).
    b. Counterexample: Nozick's grandmother
    c. Modification of the sensitivity condition: If P had been false, then S would not have believed *via the method he actually uses* that P.

   d. 'Counterexample': Toy Siren
      i. This shows only that sensitivity is not necessary for knowledge.
      ii. The issue is (though) whether, in that case, the sensitivity condition is actually redundant, and possibly damaging. (This is what 'safety theorists' think.)
5. 'Safety' theory (see also: http://www.iep.utm.edu/safety-c/)
   a. A first pass at the safety condition: S's belief that P is safe iff that belief could not easily have been false.
   b. More precisely: It could not easily have been the case that: S formed a belief as to whether or not P, via the method that she actually uses, and the resulting belief was false.
      i. 'Could not easily have been the case': is not the case in possible world that is '*close*' to the actual world [whichever that in fact is!].
   c. Some safety theorists think that "true belief + safety" is an adequate *analysis* of knowledge; others (e.g. Williamson) merely think that safety is a *necessary* condition for knowledge.
6. Some 'meta-level' reflections on the project
   a. This is a project of 'conceptual analysis'. (http://plato.stanford.edu/entries/concepts/#ConConAna) Opinions differ over whether such projects are interesting/important/tractable/etc.
      i. The critique from 'experimental philosophy': see e.g. Weinberg, Nichols & Stich, *Normativity and epistemic intuitions*, in Philosophical Topics 29, 2001.
   b. Three reactions to the apparent failure to find an adequate analysis of knowledge
      i. Optimism: We just need to keep trying.
      ii. Dismissivism: The prospects for a successful analysis look dim, but the notion of knowledge is unimportant anyway. The important questions are all in 'knowledge-free epistemology'.
      iii. Primitivism: The prospects for a successful analysis look dim, but that doesn't mean there's anything wrong with the notion of knowledge, or that we can't usefully theorise about it. E.g.
         1. Is the 'KK principle' (roughly: if you know, then you know that you know) true?
         2. Is some form of safety necessary for knowledge?
         3. Is knowledge the norm of assertion?
         4. 'Knowledge-first' epistemology: e.g. attempts to explain justification in terms of knowledge

Appendix: A zoo of counterexamples to analyses of knowledge

Smith and Jones (~ Gettier, *Is justified true belief knowledge?*)*:* Smith drives to work in a Ford. Jones sees Smith arriving in this way every day, and infers, quite reasonably, that Smith owns a Ford. Jones then further infers that either Smith owns a Ford or Brown is in Barcelona, but with no belief or reason for believing that Brown is out of the country. As things turn out, Smith doesn't *own* a Ford (he drives a rental car), but by sheer coincidence Brown has just taken a daytrip to Barcelona. Intuition: Jones does not know that: either Smith owns a Ford or Brown is in Barcelona.

Stopped Clock: Jerry looks at the clock, which reads 'noon', and accordingly forms the belief that it's noon. However, unbeknownst to Jerry, the clock in question stopped exactly twelve hours ago. *Intuition:* Jerry does not know that it's noon.

Fake Barn County (~ Goldman, *Discrimination and perceptual knowledge):* Henry is driving through the countryside. He sees a barn from the road, and accordingly forms the belief that there is a barn. What Henry has seen really is a barn, so there is nothing directly amiss with this episode of belief-formation. However, unbeknownst to Henry, the county through which he's driving is liberally populated with fake barn facades – structures that look exactly like barns from the road, but are not barns. *Intuition*: Henry does not know that there is a barn.

James and the dog: Suppose that James sees something he takes to be a dog in a nearby field, and accordingly forms the belief that there is a dog in the field. The thing James sees, however, is actually a wolf. Nevertheless, there is a dog in the field in question – just not one that James can see. *Intuition*: James does not know that there's a dog in the field.

Dr Walker and the virus (~ Zagzebski, *The inescapability of Gettier problems*): Bennett presents to Dr Walker with advanced flu-like symptoms. Dr Walker runs a blood test that reveals abnormally high levels of a certain antibody, and concludes on that basis that Bennett has a particular virus. In fact Bennett does have the virus in question, but has only very recently contracted it; the high observed levels of antibody are due to some entirely unrelated cause. *Intuition*: Walker does not know that Bennett has the virus.

Nozick's grandmother: Tom, recently returned from a hazardous expedition, goes along on a family visit to his elderly grandmother. Having seen him, the grandmother forms the belief that Tom is alive and well. However, if Tom had not survived, the family would have reported that he was fine and concocted some excuse for his absence, to protect his grandmother from the bad news. *Intuition*: Tom's grandmother *does* know that Tom is alive and well.

Toy siren: Gemma believes that there is a police car outside on the basis of hearing a siren sound. In fact there is a police car outside, but the car is silent; the sound is made by a kid with a toy siren. *Variant 1*: The kid is oblivious to the presence of the police car – he's just playing a make-believe game. *Variant 2*: The kid is sounding his toy siren *in response to*

seeing the police car. *Intuition*: in neither variant of the case does Gemma know that there's a police car outside.

Murder trial: John is accused of murder. John's father, Adam, is convinced that John is innocent. Adam's belief is overdetermined: Adam has seen forensic evidence that establishes conclusively that John is innocent, but even if he hadn't, his faith in his son is such that he would never have believed John guilty regardless of the evidence. Meanwhile, John could easily have been the perpetrator. In this case, Adam could easily have falsely believed that his son is innocent. *Intuition*: Adam does know that John is innocent.

Week 5: Contextualism about 'knows'

1. The motivation for contextualism
    a. We seemed to have good arguments against scepticism. On the other hand: If scepticism is false, why is it so plausible?
    b. The contextualist claims that even though someone asserting "I know I have hands" in an *ordinary* context speaks truly, when the skeptic says "we don't know we have hands", he speaks truly too.
    c. A bad way of fleshing this out: 'loose use' vs 'strictly speaking'
2. Towards a better way: Context-sensitivity in language
    a. A *word* is context-sensitive iff: which *thing* (i.e. which object, relation, etc) it *refers to* depends on the 'context' in which it is uttered.
    b. 'Context' can just mean: the place/time/speaker of the utterance
        i. E.g. 'I', 'here', 'now'
    c. 'Context' can also mean: conversational context
        i. 'Tall' picks out a different property in the context of (a) a discussion about how to get a child over a stile vs (b) a basketball election.
    d. Context-sensitivity leads to 'contradictions that aren't contradictions'
        i. E.g. 'I am a student'/'I am not a student'
3. Contextualism about 'knows': the proposal
    a. "Knows" is context-sensitive in the way that "tall" and "flat" are
    b. In the context of a discussion of scepticism (resp. in an ordinary context), "knows" expresses a very sparse (resp. a relatively abundant) relation, so the sceptic's (resp. the ordinary subject's) 'knows'-sentences express truths.
    c. An error theory: the overenthusiastic skeptic assumes that ordinary knowledge-attributions are false *because* she doesn't realise that the semantics for "knows" is contextualist.
    d. Aside: This proposal doesn't remove the need for explaining *how* the ordinary-context referent of 'knowledge' avoids the sceptical argument.
4. Contextualism and particular analyses of knowledge
    a. Any analysis of knowledge can be 'contextualised'.
    b. Lewis (*Elusive knowledge*) sketches a way of doing this for the "relevant alternatives" analysis of knowledge.
    c. But we could also contextualise the JTB account, safety theory, etc.
5. Two clarifications
    a. Subject-sensitivity vs attributor-sensitivity: the contextualist's claim is that the referent of 'knows' is sensitive to the *attributor's* context, not the subject's.
    b. The need to state contextualism metalinguistically
        i. The contextualist does *not* hold that <u>whether or not S *knows* that P</u> depends on some context. What she holds is that <u>whether the *sentence* "S knows that P" expresses a true or a false *proposition*</u> depends on the context in which that sentence is uttered.

   ii. <u>These are different!!</u> This is important in avoiding misguided objections to contextualism.

6. Linguistic criticism of contextualism (Stanley, *On the linguistic basis...*)
   a. Stanley's argument:
      i. 'Knows' does not behave, linguistically, in the same way as 'tall' and 'flat'.
         1. E.g. the sentence "John very knows that P" is ungrammatical.
      ii. Therefore, the fact that 'tall' and 'flat' are context-sensitive provides no support for contextualism.
   b. Reply: the contextualist's point, in appealing to the context-sensitivity of 'tall' and 'flat', is to establish the *possibility* of a contextualist semantics for 'knows', not a *high a priori likelihood* for the contextualist claim.
7. Hawthorne's criticisms of contextualism (*Knowledge and lotteries,* ch. 2)
   a. Hawthorne appeals to a variety of principles that (1) seem to be both intuitive and explanatory and (2) that the contextualist must deny.
   b. 'The assertion constraint': If S asserts that P while not knowing P, S is thereby a proper subject of criticism.
      i. This seems to explain why I shouldn't assert "your lottery ticket did not win" (if I have no inside information about the lottery draw).
      ii. The contextualist cannot uphold the assertion constraint.
      iii. Hawthorne takes this to constitute an objection to contextualism.
   c. Reply:
      i. Note that the contextualist can uphold a 'nearby' principle: If S asserts that P while "I know that P" would be false if uttered by S at that time, S is thereby a proper subject of criticism.
      ii. This principle equally well explains the data.
      iii. And it is not clear that this principle is any less 'intuitive' than the original.
   d. A similar story plays out in the case of 'the practical reasoning constraint': the principle that one should not rely, in one's practical reasoning, on premises that one does not know to be true.
8. Epilogue: Subject-sensitive invariantism
   a. An example: Hannah at the bank
   b. 'Subject-sensitive invariantism' (SSI) holds that whether or not a subject knows a given proposition depends in part on the practical importance of the proposition to the subject.
   c. This differs from contextualism because according to SSI,
      i. it is the *subject*, not the attributor, whose interests count;
      ii. there is just a single knowledge relation (this is an *invariantist* proposal).

Week 6: The *a priori*

1. A priori justification
    a. Rough definition: A *justification* is a priori iff it is 'independent of experience'.
    b. A *proposition* is a priori iff there is some a priori justification for it.
    c. 'Experience': needs to include sensory experience and also proprioception, introspection etc., but needs to *exclude* the 'experience' of grasping a proof.
    d. 'Independently of experience': must not rule out cases in which some experience is needed for *understanding* the proposition in question.
    e. More precise definition: A justification is a priori iff a person can have that justification without having had any experiences *beyond* those needed for grasping the proposition in question.
2. The puzzle of a priori justification
    a. Rationalism: we have a 'faculty of rational intuition' which somehow 'grasps' necessary truths. (How else *could* it work?)
    b. Many of us find this hard to swallow:
        i. It's mysterious.
        ii. There's no causal story to explain how it could work.
        iii. The counterfactuals that would go some way towards trusting it (e.g. 'if 2+2 hadn't equalled 4, I wouldn't have believed that it did) are merely *trivially* true, since the antecedents are necessary.
        iv. Accepting 'rational intuition' as a source of justification would licence dogmatic adherence to prejudices.
    c. On the other hand, we *need* a priori justification, on pain of scepticism:
        i. In some domains, there is a priori justification or no justification at all. (Maths, logic; Ethics, metaphysics, epistemology... all of philosophy.)
        ii. Even in the everyday-empirical and scientific domains, if there is no a priori justification, then justification is restricted to the contents of immediate experience (à la classical foundationalism, but worse).
3. A spectrum of possible positions
    a. Radical rationalism: Rational intuition is an infallible source of justification.
    b. Moderate rationalism: Rational intuition is a source of justification, but it's fallible.
    c. Moderate empiricism: We have some a priori justification, but only of non-substantive truths, for which no mysterious 'rational intuition' is required.
    d. Radical empiricism: There is no a priori justification.
4. Against radical rationalism
    a. Why *would* anyone think that rational intuition would be infallible?
    b. Anyway, it's clear that it isn't: consider e.g. a priori paradoxes.
5. Radical empiricism (Quine)

   a. Confirmational holism: when a theory's prediction turns out to be true, what is confirmed is the *whole collection* of statements that went into deriving that prediction.
      i. Quine's addendum: *this includes maths and logic*.
   b. Thus even the allegedly 'a priori' parts of our belief-system are ultimately justified by their connections to experience. But this just is to say that no justification is really 'a priori'.

6. Interlude: 'a priori' vs 'necessary' vs 'analytic'
   a. An *epistemological* distinction: a priori/a posteriori. (See above!)
   b. A *metaphysical* distinction: necessary/contingent. (A proposition is *necessary* iff it's *true in all possible worlds*.)
   c. A *semantic* distinction: analytic/synthetic (Harder to define! See http://plato.stanford.edu/entries/analytic-synthetic/)
   d. These distinctions(!) are distinct, and (probably) not coextensional.
      i. Against 'a priori = necessary' and 'analytic=necessary': see http://plato.stanford.edu/ entries/rigid-designators/; and Kripke
      ii. Against the thesis that a priori=analytic: see below
7. Against moderate empiricism
   a. The moderate empiricist holds
      i. That the only a priori propositions are *analytic* propositions;
      ii. That a priori justification for a priori propositions is *unproblematic*, while for other propositions it would be unacceptably mysterious.
   b. The rough idea: a proposition is analytic iff it is 'true in virtue of meaning'; 'one needs only to understand such a proposition in order to see that it is true.'
      i. Ayer's example: "Either some ants are parasitic or none are"
      ii. The meanings of "either/or", "some" and "none" (in some sense...) guarantee that this is true.
   c. Criticism of this idea
      i. Distinguish between *knowing which proposition a given sentence expresses*, and *knowing, of the proposition thus expressed, that that proposition is true*.
      ii. Talk of "meanings" helps with the first task, but is irrelevant to the second.
         1. Ayer still needs to tell us what justification we have for believing the logical truths, and appeal to analyticity *cannot help here*.
      iii. The answer might be: the logical truths are just obvious. But that is (apparently!) to invoke rational intuition.
   d. Anyway, not all of the propositions for which we need a priori justification *are* analytic (cf. 2.c.ii, above).
8. Conclusion: Barring radical empiricism, moderate rationalism survives as the (somewhat unsatisfactory) default, on pain of skepticism.

Week 7: Perception and testimony

1. The "classical view" of perception: What we are immediately given in experience is a *sense-datum*. We *infer* the existence and properties of the external world.
2. Veridical perception vs illusion and hallucination
    a. Illusion: experience presents a (genuine) external object as having some *property* that the object does not in fact have.
    b. Hallucination: experience is as if there is an external object having certain properties, but one is not in perceptual contact with *any object at all*.
    c. [Veridical] perception: experience correctly represents both the existence and the properties of the object.
3. Direct vs indirect realism
    a. Indirect (or representational) realism: When one perceives an ordinary external object, one does so in virtue of more directly perceiving an intermediate internal object – a 'sense-datum'.
    b. Direct realism: One perceives external objects 'directly', i.e., not in virtue of perceiving anything else.
    c. The argument from illusion (for indirect realism)
        i. When one is subject to an illusion, there is some property F such that: it seems to one that the external object is F, while in fact it is not F. (Premise)
        ii. When it seems to one that something is F, one directly perceives something that really is F. (Premise)
        iii. In cases of illusion, one directly perceives something other than the ordinary external object. (From i, ii)
        iv. If what one directly perceives is something other than the ordinary external object in cases of illusion, then the same is true in cases of veridical perception. (Premise)
        v. In cases of veridical perception, what one directly perceives is a second object, not the ordinary external object. (From iii, iv)
    d. Against indirect realism
        i. It is true that we perceive external objects in virtue of having visual experiences (or sense-data, or Ideas...). But it is just an unnatural use of 'perceives' to say that we *perceive* these experiences.
4. The characterisation of an experience
    a. Two ways to describe visual experience
        i. As a pixel-by-pixel colour array
        ii. In terms of ordinary objects: "a visual experience as of a tree"
    b. Strawson's claim: one cannot *accurately* describe experience in way (i).
5. Disjunctivism
    a. 'Conjunctivists' analyse a case of veridical perception as a conjunction: subjective experience AND certain external factors.

   b. 'Disjunctivists' analyse the subjective character of visual experience as a disjunction: veridical perception OR hallucination.
   c. Query: what sort of 'priority' or 'fundamentality' is the dispute supposed to be about? (Cf Hawthorne and Kovakovich, *Disjunctivism*)
6. The justification of testimony
   a. The vast majority of what we believe is justified, if at all, by the testimony of others. How (and when) is reliance on that testimony justified?
   b. *Reductionism* about the justificatory power of testimony: when trusting testimony is justified, that is so in consequence of more general principles of justification (concerning perception, memory and inference), applied to the special case of testimony.
   c. *Antireductionism:* One neither *can* nor *needs to* reduce the justifying power of testimony to a testimony-free base. Testimony is an independent source of justification.
   d. Coady against reductionism[1]
      i. We haven't directly observed anything like enough correlations between testimony and observed fact to justify trusting testimony via a *non-circular* inductive argument.
      ii. Anyway, the attempted presupposes that we could have observed no conformity at all between testimony and truth; but that's false.
      iii. The only reason for thinking reductionism true is that experiences could undermine beliefs acquired by testimony. But testimony could undermine beliefs acquired by observation too, so the inference to reductionism must be invalid.
7. The justification of perception
   a. Any attempt to justify trusting *perception* on the basis of (naive, enumerative) induction would be circular [too].
   b. Alternative ways to justify trusting perception: IBE, a priorism, externalism.
      i. These could be (equally well or badly?) applied to testimony too.
8. The metaphysics of testimony
   a. A rough analogue of the Strawson discussion?
      i. The reductionist emphasises a sense in which when one is told that P, what one is fundamentally given is a certain *sound or shape* (rather than a proposition).
   b. A 'disjunctivist' about testimony might hold that being told that P is a fundamentally disjunctive state of affairs (*either* veridical testimony or false testimony).
   c. As in the case of perception, none of these distinctions seem to help with the sceptical worries.

[1] Coady, 'Testimony and observation', *American Philosophical Quarterly* 10 (1973), 149-55. For criticism, see Fricker, 'Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony', *Mind* 104 (1995): 393-411

Week 8: Degrees of belief

1. Full belief and degrees of belief
    a. The notion of full belief recognises 3 options: S believes that P, S believes that not-P, S suspends judgment as to whether or not P.
    b. Alternative: S's *degree* of belief that P is x, for any number x between 0 and 1.
    c. For many propositions, the 'degree' notion is more appropriate: e.g. weather forecasting, playing Russian roulette, the 'preface paradox'.
2. Measuring degrees of belief: betting quotients
    a. A bet on proposition P, with stake S and quotient q:

	P	Not-P
Bettor pays...	qS	qS
Bettor wins...	S	0
Bettor's net gain	1-qS	qS

    b. Definition: an agent's betting quotient for P is the largest number q such that the agent is willing to enter into a bet on P (as bettor) with quotient q.
    c. A rational agent's betting quotient for P = his degree of belief that P.
3. (Bayesian) principles governing degrees of belief
    a. Probabilism: Degrees of belief satisfy the axioms of probability.
    b. Conditionalisation: On learning new evidence E, update by *conditionalisation on E*.
        i. Example: The letter from Barbados

	E&B	E&W	J&B	J&W
Prior DOB	0.5	0.2	0.1	0.2

            1. After learning B, my degree of belief in E should be 5/6.
        ii. Formal statement of the rule:
            1. Define '|' by: $p(A|B) = p(A\&B)/P(B)$.
            2. Suppose that prior degrees of belief are given by p_1.
            3. Then, after learning E, ('posterior') degrees of belief should be given by p_2, where, for any proposition C, $p_2(C) = p_1(C|E)$.
    c. Both Probabilism and Conditionalisation can be supported by 'Dutch Book' arguments.
4. What degrees of belief are good for (I): inductive logic
    a. 'First-year logic' gives precise rules for *deductive* reasoning.
    b. Bayesian theory is the analogue for *inductive/probabilistic* reasoning.
5. What degrees of belief are good for (II): rational decision under uncertainty
    a. Q: What should I do if I *don't know* which action will lead to the best outcome?
    b. Bad Answer 1: 'Do whatever you *think* will lead to the best outcome.'
        i. This is ridiculously rash.
    c. Bad Answer 2: 'Do whatever makes the worst possible outcome least bad.'
        i. This is ridiculously cautious.

   d. Correct answer: maximise *expected* utility (or expected value):
      i. Assign a number (a 'utility') to each possible outcome, measuring its goodness or badness;
      ii. Calculate the 'expected utility' of each action: (probability of outcome 1) x (utility of outcome 1) = (probability of outcome 2) x (utility of outcome 2) + ...
      iii. Choose the action with the highest expected utility.
      iv. But the 'probabilities' here are *the agent's degrees of belief*.
   e. The analogue in ethics: 'maximise expected moral value' is the right way to formulate 'subjective consequentialism'.
6. Some criticisms of the Bayesian approach
   a. The 'problem of logical omniscience': Bayesianism assumes that the agent has credence 1 in every logical truth. That's too demanding for real agents.
   b. Even *ideal* agents needn't have *precise* degrees of belief.
   c. The Conditionalisation rule assumes that 'evidence' always takes the form of some new proposition becoming *certain*. Not all evidence (or: no evidence?) is like this.
   d. The 'problem of the priors': Bayesianism just shifts all the hard questions into the 'choice of prior', and has nothing to say to constrain the latter.
7. What becomes of full belief?
   a. A 'worry': once we have the partial-belief notion, there doesn't seem(?) to be any work left for a concept of full belief to do.
   b. Revisionism: Just stop talking about full belief, then!
   c. Threshold views
      i. A (too) extreme view: S believes that P iff S's degree of belief that P is 1.
      ii. More moderate threshold views: S believes that P iff S's degree of belief that P is >x, for some specified x<1.
         1. First worry: any value for the threshold will be arbitrary.
         2. Second worry: the lottery paradox
            a. The moderate threshold account seems to yield the result that a perfectly rational agent can have logically inconsistent beliefs.
   d. Alternatives
      i. Link full-belief to *assertion*. (Kaplan, *Decision theory as philosophy*, pp.107-111.)
      ii. Link full-belief to what the subject is *disposed to "treat as true" in reasoning*. (Ross and Schroeder, http://www-bcf.usc.edu/~jacobmro/ppr/Belief_Credence_and_Pragmatic_Encroachment.pdf)