

Learning and Doing: Toward a Unified Account of Rationality in Belief, Desire, and Action

John Locke Lectures 2018

Dedicated to Derek Parfit (1942-2017)

Lecture 4:

“Skill in Guidance by Norms and Values”

Peter Railton





(University of Michigan)

Oxford, May 2018

No seminar tomorrow


- I'll be travelling.
- We'll resume the following week.

What do these have in common?




Thinking and doing

- William James
- “My thinking is first and last and always for my doing.”
(*Principles of Psychology*, I: 333).




Thinking and doing

- William James
- “My thinking is first and last and always for my doing.”
(*Principles of Psychology*, I: 333).
- He could have added: “And all my doing stretches forward, not backward, in time.”



Thinking and doing

- William James
- “My thinking is first and last and always for my doing.”
(*Principles of Psychology*, I: 333).
- He could have added: “And all my doing stretches forward, not backward, in time.”
- So thinking that is for doing is *prospective*.



Today's roadmap

- (1) Review of the prospective, regulative, learning character of desire and belief, and how this relates to acting intentionally.
- (2) Development of an idea of implicit intelligence and skill in deliberation and action-selection.
- (3) Providing a conception of *active* “mental dispositions”
- (4) The “Default Network” in the brain and its relation to prospection and learning
- (5) What skill might be in self-conscious deliberation and action
- (6) Three families of normative concepts
- (7) Tools for thought

(I) Review

Project

- Recall our project—building up to something recognizable as acting for a reason from parts we independently understand and can provide an *explanation* of how it is possible to *aptly* responsive to reasons,
 - so that we need not treat phenomena such as acting for a reason as an unanalyzed primitive,
 - ... or presuppose it in our very attempt to explain it.

Project

- Recall our project—building up to something recognizable as acting for a reason from parts we independently understand and can provide an *explanation* of how it is possible to *aptly* responsive to reasons,
 - so that we need not treat phenomena such as acting for a reason as an unanalyzed primitive,
 - ... or presuppose it in our very attempt to explain it.
- My morning cereal.

Thus far


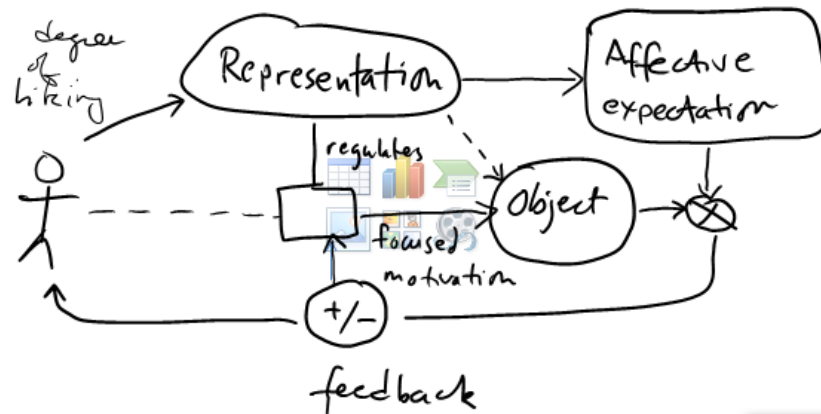
- We have developed notions of desire and belief that have sufficient functional structure to underwrite many of the properties we normally attribute to *acting intentionally*:
 - “under a representation”
 - “desirability characteristic”
 - “answer to question, what is she doing, and why?”
 - “agent’s representation gives satisfaction condition”
 - “bound together over time by agent’s representation”
 - “motivational via representation of end—teleological”
 - “intelligible”—expressive of agent’s concerns, in light of how she sees things—we can see *her* in the action.

We saw a convergence

- ... between this picture and contemporary work in decision theory and in model-based learning and control. Desire and belief *by their nature* incorporate:
 - Decision weights—affective common currency
 - Spontaneous projection—contentful expectations
 - Spontaneous regulation of action-oriented resources—focused motivation and reliance
 - Spontaneous monitoring of action-outcomes—dual paths of feedback
 - Spontaneous learning—updating strength of preference and confidence by discrepancy with outcomes

In that infrastructure ...

- ... desire and belief have a *prospective, regulative* character:



(2) Intelligence and skill in implicit deliberation

There is, of course, ...

- ... nothing to be said for trying to legislate language.
- I partly have in mind here the distinction some make between *acting intentionally* (I intentionally make the verb of the sentence I'm speaking agree with the subject) vs. *acting from an intention* (I formed an intention to make the verb agree—presumably, a rarer event among fluent speakers).

Acting intentionally

- What we needed to explain was how we could be acting intentionally—“guided by a representation”, “for the sake of an end”, “in steps bound together over time”, etc.—without self-consciously forming and deploying a prior intention.
 - As, for example, in fluent speech.

Thanks to this ...

- ... we could offer an explanation of how it *is* possible to self-consciously form and be guided by a prior intention, *as* intentional activity, without needing to self-consciously form and be guided by an intention to do *it*.
 - But those who would be happier to reserve ‘acting intentionally’ for the special case in which there is a propositional object, for example,
 - ... please substitute “purposive behavior” or, as we saw last time, “acting intelligently”.

Our ultimate goal ...

- ... is to characterize an infrastructure that is aptly responsive to reasons so that reasoning, forming an intention, etc., could be:
 - ... more than narrowly rational—that is, more than merely internally coherent
 - ... and capable of being broadly rational—that is, *itself* aptly responsive to reasons in the ways *self-conscious reasoning* distinctively can.
- Today we will discuss some of those ways, and try to say how the infrastructure gives us the *skill with reasons and reasoning* needed for this.

Skill with reasons: some failures

- Some examples of *lack of skill* with reasons and reasoning:
 - Williams’ “one thought too many”
 - An athlete’s self-conscious, deliberative interference with her performance
 - Kant’s would-be reasoner who demands a rule for applying any rule
 - Buridan’s Ass
 - Aristotle’s would-be reasoner who demands a demonstration of all premises before accepting an inference
 - Carroll’s Achilles talking with Tortoise

Skill with reasons: some failures

- Hume's self-defeating skeptic, who undermines even the possibility of his own reasoning
- Ditherers who let opportunities pass from inability to let go of opportunities
- The various forms of practical regress we've discussed
- These are individuals who render themselves *less aptly* responsive to reasons, and may even place reasoning itself outside their grasp, by lacking a certain intelligence and competence in deploying self-conscious deliberation and decision “at the right time, in the right way, toward the right objects, with the right end”, so to speak.

For example: Default, defeasible self-trust

- Suppose that someone proposes not to rely upon his senses and faculties without deciding whether they're reliable.
 - If the decision is not to be mere fiat, it will need evidence and reasoning.
 - But one cannot gain evidence without using one's senses and faculties.
- If instead we have some measure of *default, defeasible trust* in our senses and faculties—without requiring evidence for this—we can begin to learn and discriminate
 - ... through the formation and reliance upon expectations
 - ... permitting feedback and adjustment going forward

Being skilled with epistemic reasons

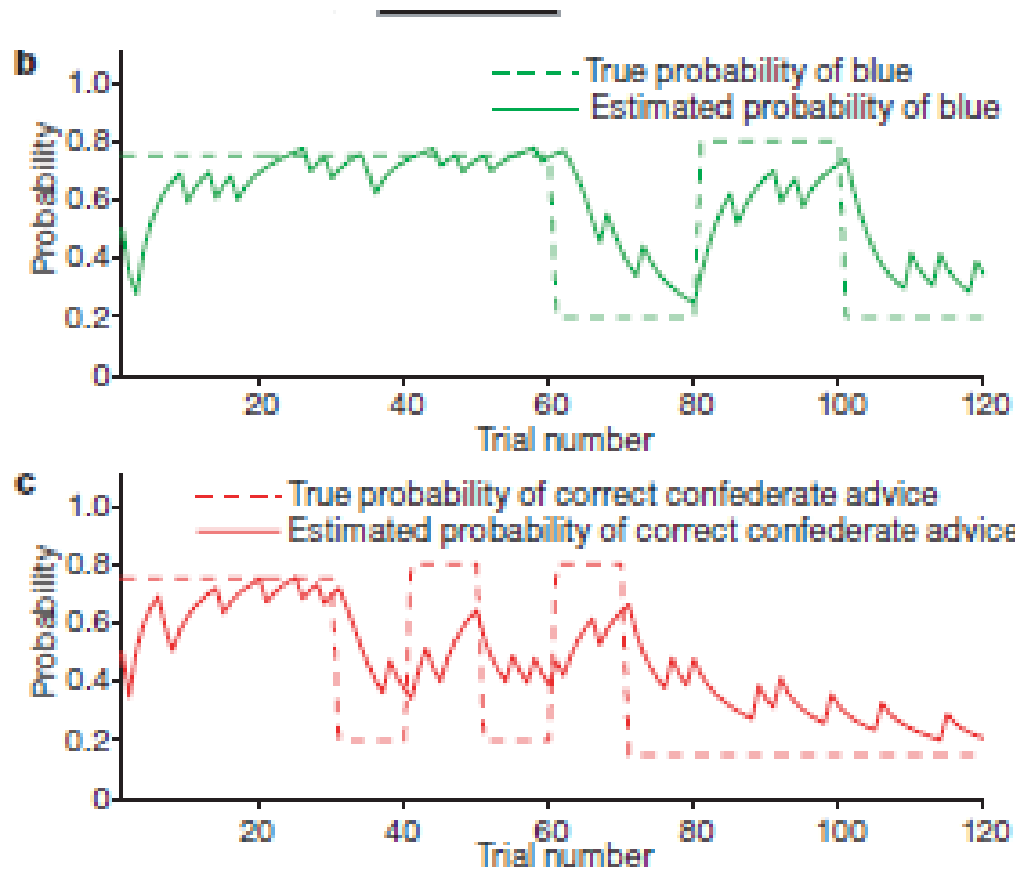
- That is just what belief and desire do.
 - Projection, trust, expectation, and reliance are inherent in their operation.
 - Because they commit action-oriented resources, they set us up for failure in a way tied to what matters to us,
 - ... so that we are motivated to notice and respond to feedback, and thereby learn.
- As we saw, infants display this kind of active, expectation-based learning from very early on.
- The desires and beliefs need not be seen as *themselves* the ground of our reasons—the ground can lie in the inner and outer world they help us stay in touch with.

Being skilled with epistemic reasons, “Stronger Together”

- Of course, as thinking about infants makes clear, we never really do this all on our own.
 - If a would-be learner extends some default, defeasible degree of trust to others, such that she relies upon information from them prospectively, she can again gain experience-based feedback and learn to be more discriminating.
 - Infants do this with respect to speech, food preference, advice, etc. over the course of the first few years (refs.).

Social evaluation of reliability

(Behrens *et al.*, 2011)



(3) Active, intelligent, prospective “action tendencies”

The “Good Regulator Theorem” and beyond

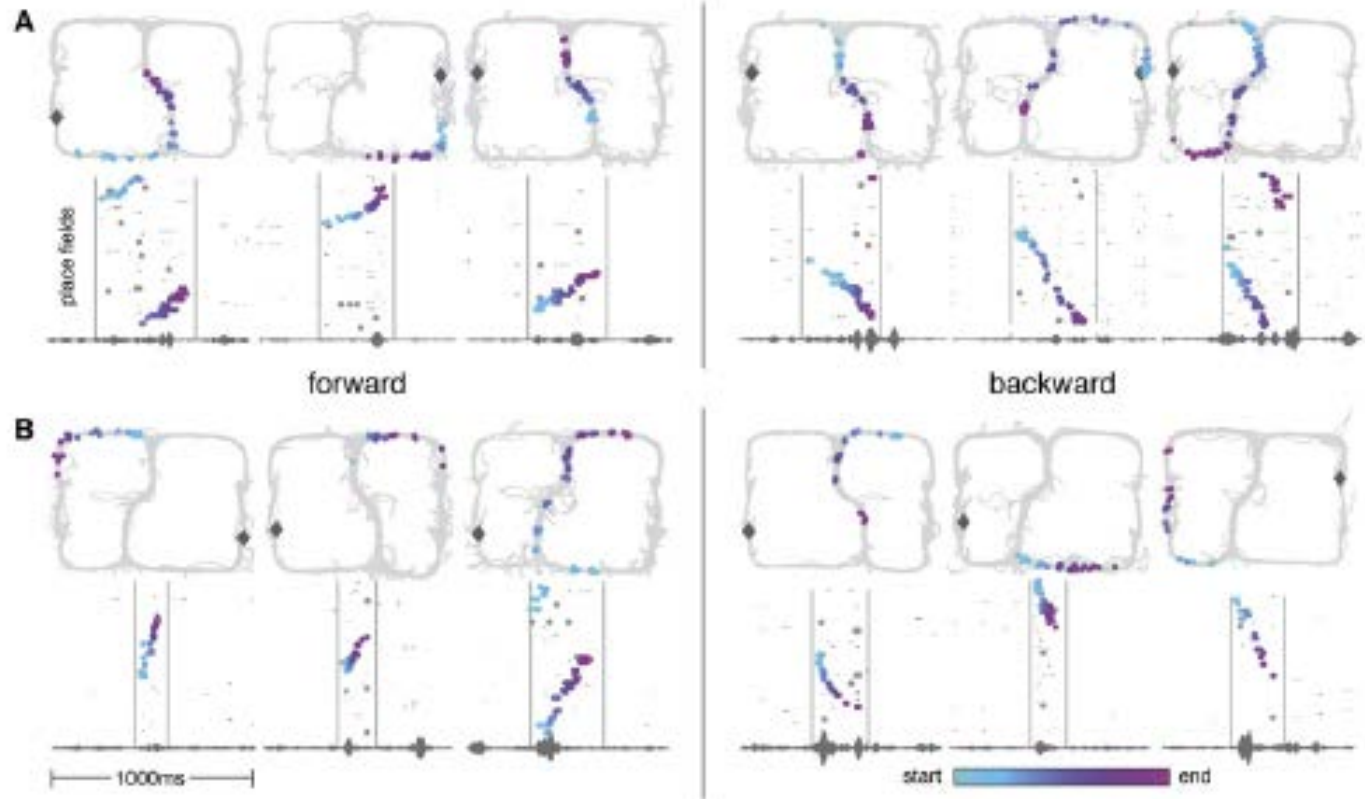
- Last time we discussed briefly Conant and Ashby’s “Good Regulator Theorem” (1970), as bridge from the regulative role of desire and belief to the building of prospective models.
- They themselves speculate:
 - “... the living brain, so far as it is to be successful and efficient as a regulator for survival, must proceed, in learning, by the formation of a model (or models).”
- Rather than place too much weight on this result: In control theory today there is considerable consensus that intelligent systems must *have* or *build* causal-evaluative models, *use* them in simulation and active control, and *learn* through feedback.

Natural intelligence

- Last time we reviewed some evidence from cognitive and affective neuroscience that intelligent animals like rats and monkeys do indeed construct spatial and causal-evaluative models through experience,
 - ... use such models to simulate possibilities in advance, including the generation of novel possibilities
 - ... and are actively guided by them in real-time decision and action.


Hippocampal construction of novel paths in sleep

(Gupta *et al.*, 2010)



Examples of Forward and Backward Replay

A rat using a causal-evaluative model in real-time decision-making (Johnson & Redish, 2007)




Prospection as a way of life

- Action extends into the future, so we cannot simply rely upon perception and memory—we must *generate* possible futures, explore them, evaluate them, and act in light of them.
 - This creates pressure to develop models with *generality*, *projectibility*, and *abstraction* to *types* of situations, actions, outcomes, etc.


Personal/subpersonal

- It has been tempting to think in terms of two kinds of mental processes:
 - Personal, deliberative – “System 2”
 - Subpersonal, automatic – “System 1”
- But the subpersonal can be, and is, *highly* deliberative—just subpersonally.
 - Some might call these *mere casual dispositions*, since they are not under direct control of self-conscious deliberation,
 - ... but they approximate the *form and principles* of rational choice theory and Bayesian belief-revision.


The brain I was brought up on:




The brain I was brought up on:




The brain I was brought up on:



The brain I was brought up on:



Today:



Passive dispositions vs. action-tendencies

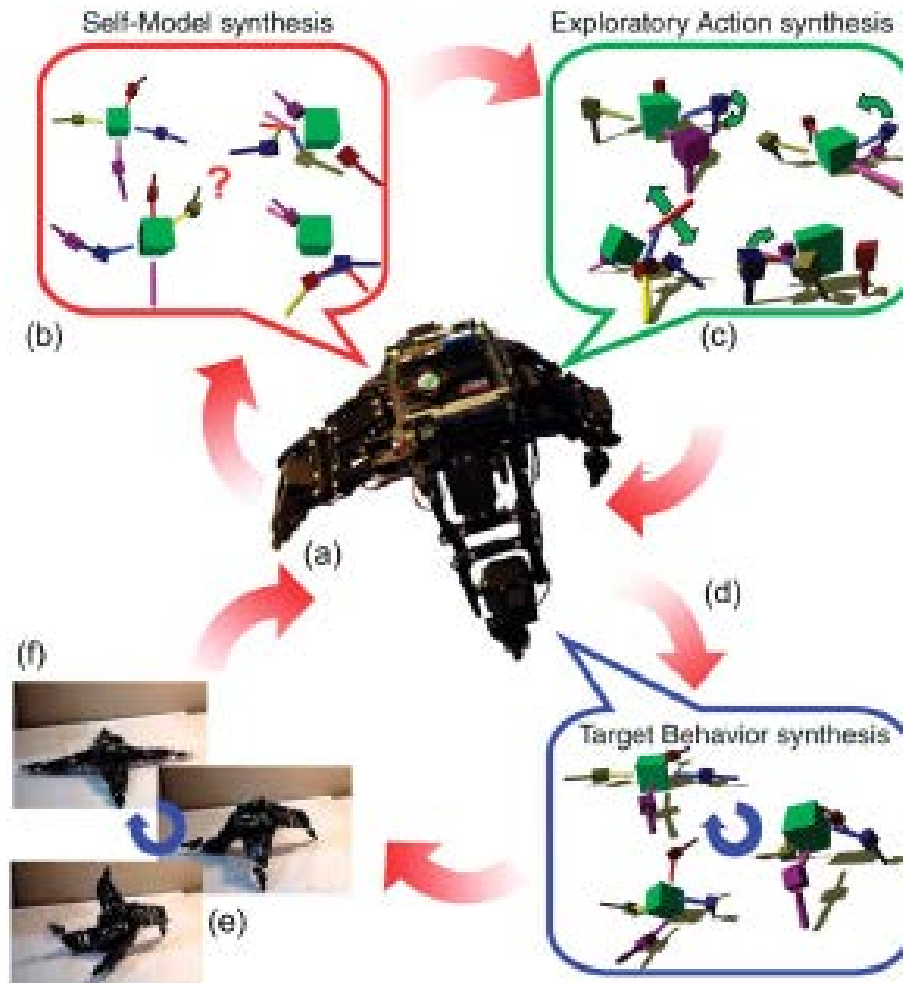
- The prospective brain does not wait for a stimulus in order to act or learn.
 - As we saw in the case of the rat, rats actively discriminate among stimuli the most predictive cues (Rescorla, 1988).
 - Representationally-based off-line simulation plays an important role in learning, focusing on possibilities that have the greatest information value (Li & Wilson, 2007)—not ingrained habit.
- Simulation also plays an active role in real-time behavior, as the brain “runs ahead” of the situation, projecting and assessing possibilities to guide action selection.

Passive dispositions vs. action-tendencies

- The individual *engages* the world actively, *explores* for information, *imagines* possibilities, uses these processes to *build* action-guiding models, *notices and seeks out* opportunities to realize desires or meet needs.
 - As core affective states, desire and belief actively regulate the deployment of these capacities for attention, perception, memory, inference, motivation, and behavior in a way congruent with our goals and viewpoint.
- That is, they are not a form of passivity on the part of the agent. Rather, they are forms of agency—intelligent, implicit, goal-directed agency.

A (robot) spider learning to move


(Bongard & Lipson, 2014)





How important are these prospective processes?

Distribution of cone cells on the retina

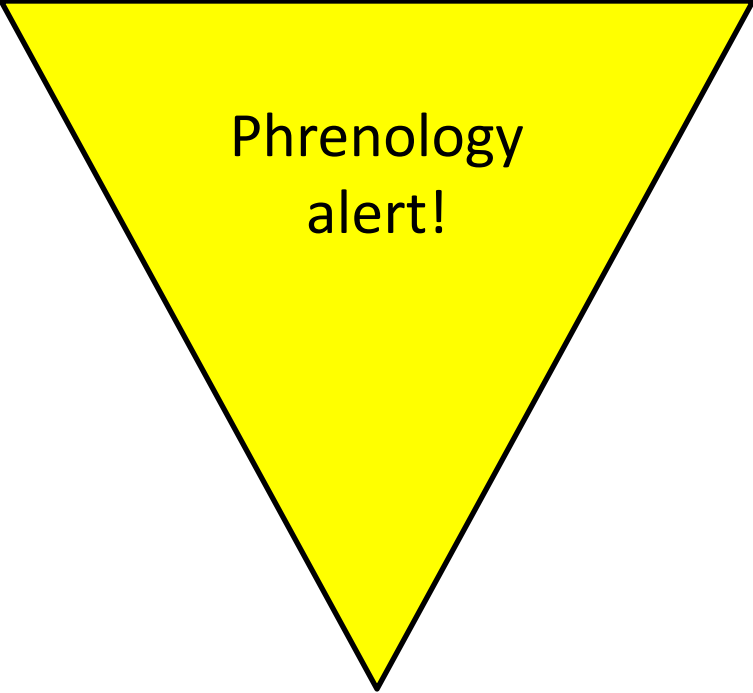
(Johnson, 2010)



The importance in vision of general patterns and pattern completion

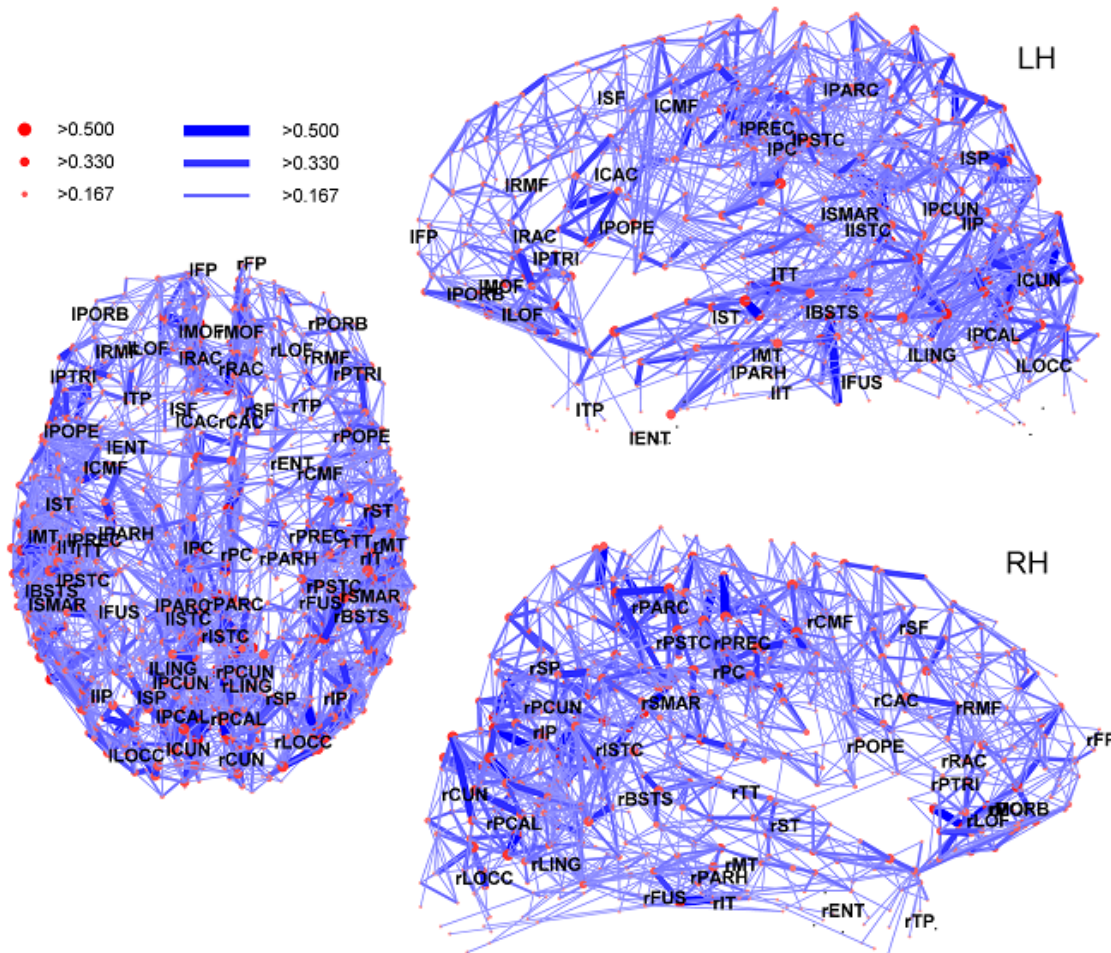


(4) The “default network” and prospection




Phrenology
alert!

Well, not *quite* phrenology, Connectomics: A view of the mind in terms of functionally integrated, persistent networks
(Hagman *et al.*, 2008)



Default network

(Buckner *et al.*, 2008)





Metabolically and developmentally ...

- ... such a brain is tremendously expensive.
- Even in this “resting” state, it consumes 15-20% of your body’s oxygen and calories, despite constituting about 2% of your body weight (Raichle & Gusnard, 2005).
 - And the level of metabolic brain activity is remarkably constant, whether the brain is “on task” or “in default mode”.
 - Moreover, the brain regularly shifts back and forth from task-focused to default mode over the course of time.
- So what’s default brain up to? Is it “resting”?

Default network

(Buckner et al., 2008)




Prospection and memory


- Memory, we know, is remarkably unreliable and malleable.
- New information, or new thoughts processed while recalling a memory, are then *stored* back in memory, seamlessly.
 - This looks like a serious defect if memory is about storing information.
 - But if memory is “thinking for doing”, then updating with new information and thought makes sense.
 - As does the ability to select and recombine episodes from memory to use them to fit new contexts.
 - Binder & Desai (2011) found:

Default network


(Buckner *et al.*, 2008)



AUTOBIOGRAPHICAL MEMORY



ENVISIONING THE FUTURE




From infancy onward ...


- ... the most important context for us to understand is often the *social* context.
 - To understand the people around us and “size up” social contexts
 - ... to see the prospects and perils of the situation
 - ... to observe even third-party interactions to find the interrelations present, potential companions, sources of information, people in need, people in anger, potential allies, and so on (compare Melis *et al.*, 2006).

Default network


(Buckner et al., 2008)




AUTOBIOGRAPHICAL MEMORY



ENVISIONING THE FUTURE



THEORY OF MIND




Skill combines intelligence with fluency

- Intelligent, implicit agency can improve through practice *and* simulation, resulting in greater fluency and skill.
 - Think of how we learn and use language


General semantic network vs. default state network

(Binder et al., 2008—150 studies)

General Semantic Network



"Default State" Network



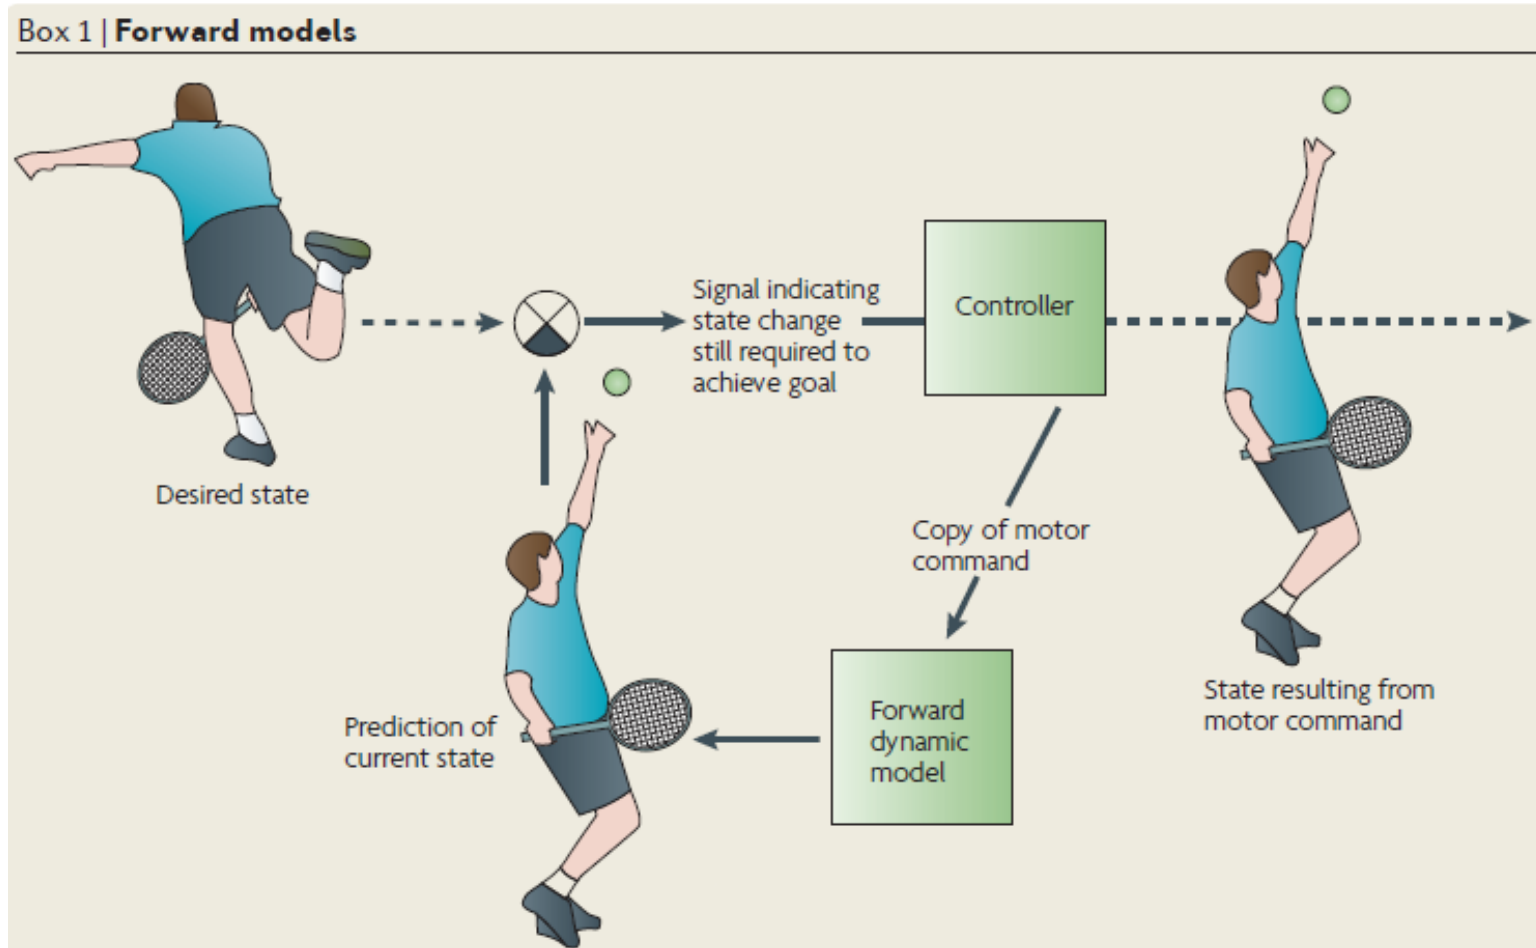
(5) Developing skill in self-conscious deliberation and action

Athletics and imagination

- Elite athletes have typically practiced thousands of hours more than good athletes—actively, perceptually, and imaginatively.
 - Studies of elite athletes indicate that they mentally rehearse and revisit much more extensively than merely good athletes.
 - They attend to different features a game situation, gather more information from them, are able to predict the next event more rapidly and accurately, and can adjust their own response “on the fly”—their movements are *less* predictable and *more* dynamically variable and exploratory than very good players (Yarrow et al., 2009).

Skill in humans

(Yarrow et al., 2009)



Skill combines intelligence and fluency in weighing multiple factors—that is, deliberation


- Typically there are multiple variables at stake in a game-situation—coordination with others, various dimensions of possible advantage or loss, associated degrees of confidence or risk, utilization of scarce resources, and so on.
 - The highest level of accomplishment requires reliable estimates of these variables for plausible available courses of action, and these values must be compared in a common pathway to attention and action.
- To be effective, this *implicit weighing* must be done *fluently*, without extensive conscious intervention.

Skilled action selection

- Implicit weighing of alternative actions, where the weights come from appropriately-learned and relevant information, represented in terms of options and associated expectations and evaluations, *is* deliberation. (Recall the evidence of sweeping ahead to compare alternative paths in maze choices.)
- It is a causal process, but not a “mere causal process”—it is at the same time a weighing of decision-appropriate information.
 - Skill in this form is not “muscle memory” or “non-representational”.
- How does implicit deliberation issue in action selection?
There are various models, applicable in different contexts of choice. Here’s one:

Real-time deliberation and action selection

(Grabenhorst & Rolls, 2011)



Thinking and doing

- *How do acts come from thought?* We need an intelligent *implicit* competence or skill in considering and weighing options, making decisions, mobilizing in accord with valuation, launches action, and monitoring it over time (including such an action as self-conscious deliberation), so that we act “at the right time, in the right way, toward the right end”, even when self-conscious deliberation would be self-defeating.
 - ‘Intelligence’ comes from **-lect** (to gather)
- The intelligent animal is continuously, actively collecting and weighing information about expected value of competing courses of action (say, to stay on patch or explore), and this guides the allocation of action for optimal foraging.

Buridan's Ass

- Isn't a deliberator of this kind vulnerable to the problem of Buridan's Ass? Mustn't we introduce higher-order reasoning to break the tie by forming an intention?
 - In actual neural systems, and artificial control systems, *noise* is spontaneously generated to prevent such “overfitting”.
- Note that we need cannot offer a *deliberative* route to a solution via the notion of forming an intention alone—for there would be a tie between two candidate intentions. Higher-order deliberation, as much as implicit deliberation, needs a degree of randomness, and some speculate that there has been selection for this in neural evolution (Grabenhorst & Rolls, 2011)

Skill and intentionality

- Skilled activity—in sports, speech, music, artisanship, social life, self-conscious deliberation—is the agent *acting through* a body of knowledge, capacities, and competencies to achieve ends she cares about in light of how she sees things.
 - Such activity is guided by an evaluative representation of a situation that reflects the agent's concerns, understanding, and goals.
- It is not blind, but fluently *attentive* to the situation and *responsive* to relevant features ...
 - ... and we can *learn* and *understand* that many reasons—athletic, social, practical, deliberative—may require fluency if we are to respond aptly to them.

We have, then, a model of *implicit* deliberation and choice ...

- ... and can see how it can be intelligent and fluent thanks to experience and prospection.
- We now need to assemble some field notes about self-conscious deliberation, to see whether such skill with implicit deliberation and decision could underwrite skill in self-conscious deliberation and decision,
 - ... that is, skill with reasons through reasoning.

Ada and Bruce

Notice the importance ...

- ... of what Ada notices
- ... of which thoughts occur to her, and which don't
- ... of how plausible these various thoughts initially seem
- ... of the possible actions that come to mind
- ... of what causes her to stop and rethink
- ... of how she settles her mind
- ... and of how her model of the situation, Bruce, and potential futures guides her eventual action
- These are all parts of her deliberation, but not *products* of deliberation as such. They're products of her underlying knowledge and values, the information she acquires from the situation, and her imaginativeness and competence.

Moreover ...

- In these ways, Ada's deliberation can be aptly responsive to reasons.
- Moreover, we see in all these elements an expression of who Ada is, what matters to her, and how she sees the world. The deliberation, and the eventual action, are expressive of *her* as an agent.
 - So, she shows skill in deliberation and intention formation—responsiveness to reasons for forming the intention.
 - Without forming a prior intention to form that intention, and without mere automaticity that would make her thoughts and action opaque to her (recall Radioman).

Causal deviance

- Let's change the story slightly.

Given our account of the causal structure of her behavior ...

- ... we can see why this *isn't* a case in which we credit her with acting intentionally.
 - Although her behavior achieved the goal she sought, it did so in a way that *surprised* her, that was not represented by her as available, not regulated by her representation, and does not accord with her evaluative assessments.
- Thus, while her beliefs and desires *caused* her to deliberate, and her deliberation *caused* her to take an inattentive step that *did* achieve her goal, we do not see this behavior as *guided* by her prospective model in a way that makes it intentional.
 - Davidson's climber, Oedipus on the road to

This does not eliminate the problem of deviant chains ...

- ... but it carries us further down the road in understanding why *these* examples are examples of such chains,
 - ... and what counts in such cases as *causation in the right way*.
- In these examples, the agent's prospective model is either not controlling the action (e.g., the climber), or controlling it under a different representation (e.g., Oedipus).
 - A signature of this is that the agent will immediately be *surprised* that his or her goal has been met in the way it has.

We might compare Ada


- ... with Clarence, who fails to notice his cat's new lounging place, with the predictable result.
 - Or who has an unacknowledged desire to punish Bruce.
 - His act is not wholly unintentional, we'd say—to some degree, he aimed at just this result. His implicit deliberation and action-guidance is not aligned with his self-representation of the situation.
 - What we *don't* want to say is: “It is in no way Clarence's act—it is an act only of some subsystem within him.”
- Contrast Desiree, who has a nervous tick, and steps on her cat's tail in consequence (even if the result fits her aims).

Bigger picture?

- Davidson, who originally sought to defend the idea that belief and desire could provide the structure of acting intentionally, quit psychology and decision theory when he came to the view that there is no fact of the matter about how to distribute belief vs. desire, etc. on the basis of the evidence.
 - This pushed him toward interpretivism.
- But we now have a great deal more evidence, which suggests principled ways of making and testing these attributions.

Recall: separate representation of expected value vs. risk

(Fiorillo *et al.*, 2003)



Bigger picture?

- Of course, we can always reinterpret *these* representations as well.
 - But what is wanted is sufficient, principled empirical grounds for making the attributions, not for refuting generic kinds of skepticism.
- We may not need to take rule-following as a primitive, if we can see it as an instance of model-based control. What would that look like?
- Let's follow Wittgenstein's lead, and head to a construction site.

What carpentry can tell us about the structure of the normative realm

(6) Three families of normative concepts

Three families of normative concepts

- **Regulatives:**
 - *Norm, rule, law, right, wrong, correct, incorrect, standard, canon, criterion, law*
- **Evaluatives:**
 - *Value, good, poor, virtuous, vicious, desirable, fine, base, noble, important, true, trustworthy, intelligent, beautiful, sublime*
- **Deliberatives:**
 - *Ought, must, reason, rational, irrational, fitting, unfitting, warrant, merit, deserve, deliberate, conclude, decide*


Norma



TOOLSTATION

Regulus





Regulatives, word origins

Norm – builder's square

Rule – ruler, straight-edge

Right – straight, straight angle

Wrong – twisted, bent

Correct – co-aligned

Standard – standing straight, regimental sign

Consistent – standing together

Ortho- – straight, upright

Criterion – measuring rod

Canon – measuring rod

Regulatives, in function

- Involve a standard or rule relative to which an assessment can be made with greater or lesser definiteness, often as a matter of *form*.
 - Syntactic well-formedness of a sentence
 - Logical validity of an argument
 - Standard of measurement
- Typically, we can voluntarily hold ourselves to such a standard—it is within the normal power of the will.

Regulatives, in function

- *If we can hold ourselves to these standards, we can achieve:*
 - Joint regulation of a shared building project running to plan
 - Communication via common language with shared meanings
 - Scientific , legal, or organizational regimentation and standardization that enables us to cooperate over time through an extensive division of labor with a degree of ambition and accuracy otherwise impossible—in social forms, knowledge and technology, or production.

Regulatives, in function

- This is *not* to say that regulatives tell us what premises to start from, or what kinds of buildings are good to live in, or what is worth saying, or what is a worthwhile common project.
 - Validity is not soundness or importance
 - Square angles might not be best for an airplane wing
 - Grammatical sentences can be uninformative or false or pointless or awkward
 - A well-defined statistical test might perform poorly in making the distinctions we need in hypothesis testing
- We must bring them to life by disciplining of ourselves (or others) to them.

Regulatives, in function

- Rather, if we have good thoughts, plans, designs, projects, etc., regulatives help us to set these out clearly, deploy them effectively, communicate them reliably, reproduce them accurately, and do so in a coordinated, effective way.
 - Note: we are forming and acting on what are, in effect, prior intentions, to create a self-enforced normative practice—with an eye to what it can achieve.
- But we need skill at applying and following tools—*that* isn't a tool, it's a skill. And it's value-guided, not rule-like
 - We need the same kind of goal-sensitive skill at applying and following norms and standards

Evaluatives

- *Value, good, poor, virtuous, vicious, desirable, fine, base, noble, important, true, trustworthy, intelligent, credible, beautiful, sublime, lousy, fair, fortunate, disgusting*

ghedh- (good)



wal- (value)



wī-ro- (virtue)



Evaluatives, word origins

Value – strength

Virtue – virility

Credible – heart

True, trustworthy – firm

Poor – few, little

Sublime – transcending limits

Fair – lovely, pleasant

Good – together, united

Important – having causal power

Beauty – favored

False – fail

Noble – knowing

The positive evaluative terms concern characteristics one wants for one's own life, or for one's kith or kin, or society.

A *good* friend isn't one who is “always correct”, but one who shows firmness, heart, solidarity, favor, understanding, etc.

A *fair* marriage or arrangement, or a *sublime* sonata, is more than a matter of strict following of rules.

Evaluatives, in use

- Application typically cannot be assessed in purely formal ways—a more substantive notion of *fit* is involved.
 - Is this house plan a good fit for the needs of those who'll live there?
 - Is this work of art one that is fitting to admire—can we expect it to survive the test of time?
 - Is zoo life a fitting existence for a wild animal, if this is the only alternative to extinction—how might it be better?
- Typically, there is no set of rules or criteria that suffice to determine whether an idea, or remark, or house, or life is good, worthwhile, important, or beautiful.

Evaluatives, in use

- Also typically, we cannot attain these evaluative states simply by an act of will.
 - Often open-ended and difficult to settle
 - Usually a matter of *degree* rather than bivalent
- Fitting attitudes toward these values typically are affective or motivating: *liking, admiring, appreciating, hating, etc.*
- We can through our efforts typically become better at discerning or embodying these values.
 - And this can make possible, and ground, personal or shared goals or ideals, or a shared appreciative practice.

Evaluatives, in function

- Evaluatives can also characterize the *meaning* of what we say or do, giving substantive guidance in thought and action in ways distinct from regulatives.
 - Mere legality vs. “the rule of law” as a value and ideal.
 - Kant’s Doctrine of Right (obedience to law—*observantia*) vs. Doctrine of Virtue (respect for the law—*reverentia*) — only the latter, not the former, has unconditional worth.


Evaluatives, in function

- So this *isn't* about being a consequentialist or a deontologist
 - Consequentialists need regulatives to represent *how* value is to guide action, e.g., a criterion of right action.
 - Deontologists need evaluatives to explain the importance and value made possible by following rules, and to avoid a regress in the application of rules.

Deliberatives

- *Ought, must, reason, rational, irrational, fitting, unfitting, warrant, merit, deserve, deliberate, conclude, decide*
- Given evaluative and regulative concepts and capacities, we are equipped to navigate the world in a certain manner—setting a goal and plotting and measuring a course aligned with it.
- But what if there are multiple possible and worthy goals, or paths that vary in risk or promise? We need some way of assessing what to do in relation to such questions. And for that, we need deliberatives.

libra- (deliberate)




libra- (deliberate)



med- (must)



ar- (rational)



Deliberatives, word origins

Ought – to possess, be master of

Must – to take appropriate measure

Deliberate – to weigh

Reason, Rational – to fit together

Merit – to receive a share

Decide – to cut off

Conclude – to close off

Deliberatives, in use and function

- Deliberation is not just about weighing values and risks, but also about *how* to weigh them, and how to use self-control to suspend action, reach a conclusion, and implement it.
 - These are functions distinct from regulatives or evaluatives.

(7) Tools for thought

Implicit normative guidance

- Norms figure in the shaping of our conduct in various ways, for example, by constraining the range of options we consider, or the way in which we feel about situations.
 - Consider for example implicit guidance by norms of conversational distance.
 - We've all had the experience of entering a new society and being made uncomfortable by local norms of appropriate conversational distance.
 - Perhaps for the first time we discover that we have acquired such norms in our home society.

Implicit normative guidance

- We need not imagine a rule of some sort, nested in our mind.
 - It is enough if our model of ordinary social interactions assigns heightened significance to being closer than a certain amount, such that we find ourselves unable to avoid experiencing conversational distance less than this amount as needing “correction” when it is discordant with the closeness of our relation to the other person.
- Recall the nature of regulatives—if we accept a norm we sense a need to “correct” ourselves or the situation, not simply to treat a departure from the pattern as *unusual* or *surprising* (as we would a departure from what we view as a bad habit). We will experience some motivation to “correct” even without external incentive. (Railton, 2006)

Implicit normative guidance

- Growing up in a society, most of us acquire skill in applying norms of conversational distance—mode of address, public behavior, etc.—by acquiring an array of sensitivities that make us alive to when these norms are violated, and discomfited and motivated to correct.
 - We acquire the skill at *taking a situation to be in compliance, or not*, without deliberation or knowledge or a rule.
- The existence of these norms and our skills with them does not mean it is *good* to act this way. (That is why ‘correct’ has been in inverted commas.) We can try to unearth these norms, and challenge their value—for example, they may serve to entrench hierarchy or enforce gender distinctions.

Implicit inference

- The important point for our purposes is that first-order, unself-conscious norm-observance need not be explained by the agent applying a rule
 - A causal-evaluative model *encodes* the relevant relations, generating expectations of oneself or others in a given situation (real or imagined), motivation to “correct”, etc. via a direct input/output function or pattern completion (e.g., of the kind found in the perceptual system), with no need to bring in norm—even when the encoded *relation* is itself normative.
- E.g., evidence that spontaneous expectation-updating in infants conforms approximately to Bayesian updating—no epistemic norm is being applied by the infant, even if her expectations are evolving in an epistemically normative pattern and the result is a perceptually-warranted degree of confidence.

Implicit inference

- Model-based projection or simulation thus are effectively forms of *direct inference* that are not deliberate applications of the relevant rules.
 - And that is a good thing, since we will need such capacities to explain how we *can* explicitly follow a rule, and do so in apt response to reasons.
- Recall that what we needed were non-inferential steps that can, taken together, constitute an inference.
- Simulation is a form of content-sensitive, goal-directed transference from one thought to another, via the semantic, causal, and evaluative relations encoded in the model. It is found in animals likely incapable of rule-based inference.

Going explicit: skill in using tools

- How can we leverage this to get self-conscious rule-following inference?
- Perhaps in the same way we can leverage carpenter's *skill with tools* to apply tools without needing another tool for this.
 - The carpenter has a well-developed, model-based skill in using various measuring tools to attain various goals. (For example, he will often quickly see the point, and use, of a new measuring tool.)
- He understands that, by holding his cuts to a given measuring tool, he can, together with others, make tight joints, stable frames, openings to fit an industry-standard door, etc.

Going explicit: skill in using tools

- The skilled carpenter also learns over time which measuring tools to trust, how much, and in what contexts and for what purposes.
 - Recall that the nature of trust or confidence, as an affective state capable of regulating expectation and reliance, enables it to take different objects, and to transfer trust from one object to another.
 - For example, trust in a recipes will yield trust in an outcome, distrust of a measuring tool will yield reluctance to rely upon its result, etc.

This week's shameless appeal to authority



Wittgenstein on logic

- “**81**. F. P. Ramsey once emphasized in conversation with me that logic was a 'normative science'. I do not know exactly what he had in mind, but it was doubtless closely related to what only dawned on me later: namely, that in philosophy we often compare the use of words with games and calculi which have fixed rules” [PI, 38]
- “**131**. For we can avoid ineptness or emptiness in our assertions only by presenting the model as what it is, an object of comparison—as, so to speak, a measuring-rod; not as preconceived idea to which reality must correspond. (The dogmatism into which we fall so easily in doing philosophy.)” [PI, 50-51]

Logic, then,

- ... is not the “necessary laws of thought”, but an artificial regulative model—really, a set of such models—which we can *bring to bear on thought* by holding our use of inference to the transitions permitted by a logic.
 - In so doing, we are treating the rules of the logic as *a priori* and *regulative*, much as carpenters treat squares and levels—”corrections” are to be made to fit rules and rulers, not adjustment of the rules or rulers to what actually is done (*a posteriori* and *descriptive*).

At the same time

- ... mathematicians develop selective trust in various methods or procedures, using multiple forms of validation to check the work.
 - But there will be some fundamental evaluative questions, e.g., about whether to accept non-constructive proofs, for which there is no rule—only the skilled practice of mathematicians, responding to a number of competing values.

Such self-discipline by norms ...

- ... whether resulting from choice or socialization, will enable our patterns of reasoning to possess some of the valuable properties of the logic in question, e.g., truth preservation in the case of deductive logic, or progressive diminution of the influence of priors in Bayesian inference, or protocols that make experimental results shareable in a wide community.
- Of course, such self-discipline is *policed* in the end by ourselves, not by the “laws of thought”, and so we face the challenge of “building more refined and accurate tools with less refined and accurate tools”—an inescapable challenge in this territory if we are ever to have more refined and accurate tools, and the advantages they confer.

Tools

- We must use crude tools to make refined tools. And we must develop a disciplined practice. But we are moved to do this because we can see the *goods* achievable thereby.
 - This is not blind, then, but intelligent tool use, sensitive to content, context, purpose. “Choice of logic” likewise.
- Of course our actual skills, even when combined together, are finite and fallible—we must *find and hold ourselves to* sources of regulation that go beyond our native endowment.
 - That’s why weights and measures are now standardized using the most reliable physical phenomena we know.
 - And why mathematicians work together and develop algorithms and indirect tests to check hugely-long proofs.

Beyond logic

- Similarly, we can have a legal practice that is skilled at applying laws, thereby making legal regulation more nearly possible.
 - Or time-keeping, calendars, engineering standards, experimental methods, historical documentation, ...
- As Aristotle, Hume, Kant, and Wittgenstein have told us—we need a form of *understanding* that is aptly reasons-sensitive but *not* deliberative rule-application. In turn, this can give us a capacity to develop the skills needed for *consciously*-deliberative rule-following that also is aptly reasons-sensitive.
 - And does not simply get in its own way.