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 3:

"Intelligent Action without Regress"

Peter Railton

(University of Michigan)

Oxford, May 2018

Seminar discussion

- Tomorrow, Thursday, at 9:00 am
- Ryle Room
- Radcliffe Humanities Building (next door)

- All are welcome!

P.A.M. Dirac (1902-1984)

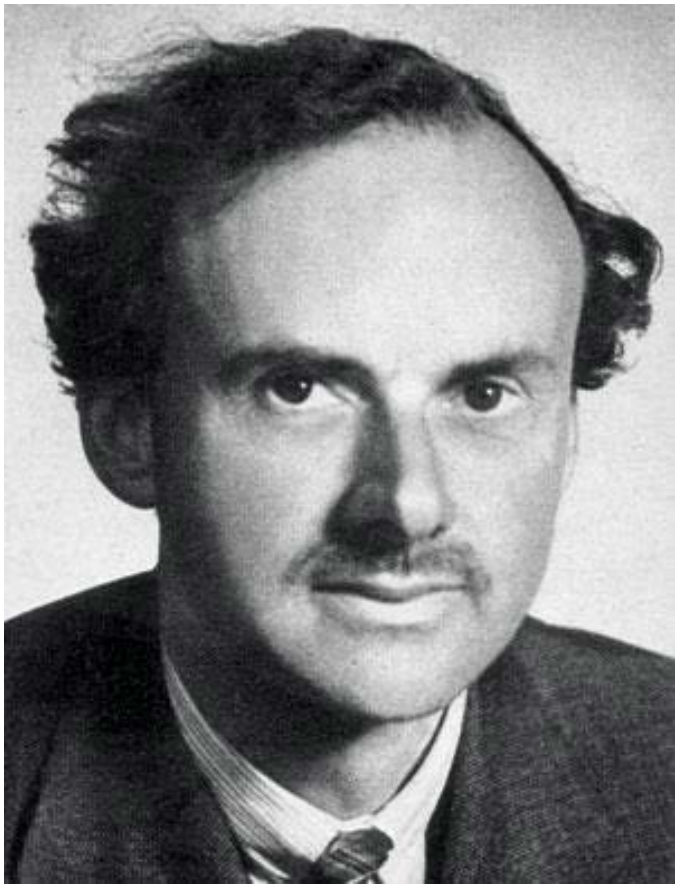

- "I thought I was entering the orderly abode of reason …"

P.A.M. Dirac (1902-1984)

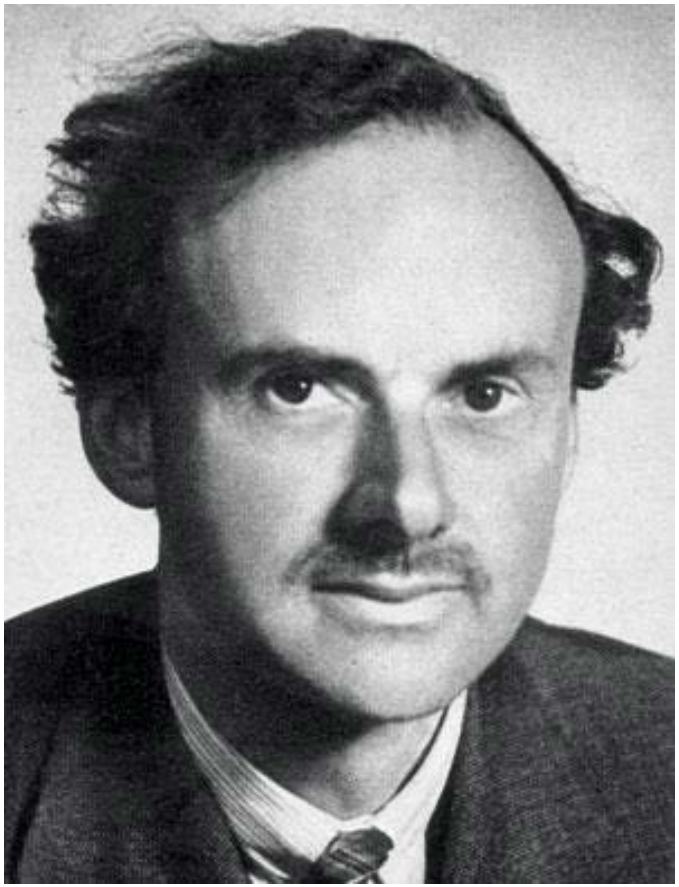

- "I thought I was entering the orderly abode of reason …"
- "… but instead I was entering a factory."

Well, …

- … this is a factory.

Well, …

- … this is a factory.

- But with luck, a reason factory.

Starting point: the orthodox belief-desire model of action, and a puzzle

- **belief + desire → action**
- *representational* *non-representational*
- *inert* *motivating*
- *mind-to-world* *world-to-mind*
- *T/F* *not T/F*
- *cognitive* *non-cognitive*

- *potentially rational + ????? →*

Starting point: the orthodox belief-desire model of action, and a puzzle

- **belief + desire → action**
- *representational* *non-representational*
- *inert* *motivating*
- *mind-to-world* *world-to-mind*
- *T/F* *not T/F*
- *cognitive* *non-cognitive*

- *potentially rational + ????? → potentially rational*

You could give up on the belief-desire model. Instead we've built new models of desire and belief.

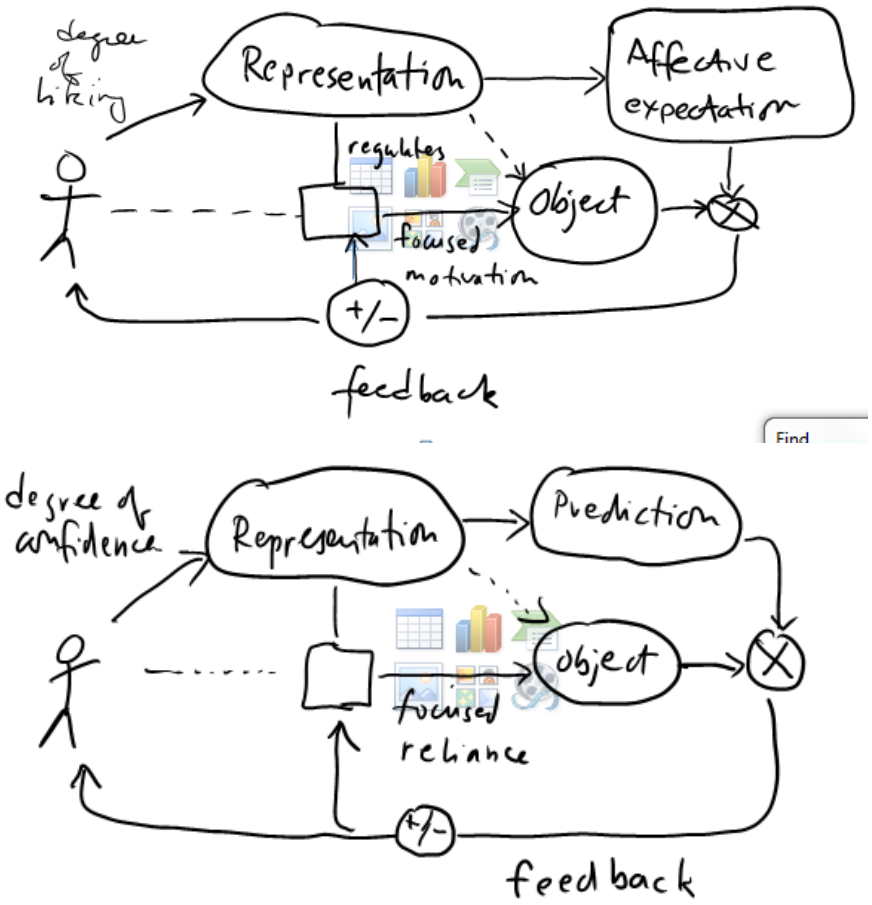

Desire and belief: *de se* representation of intentional object—acting through mode of presentation, opaque context

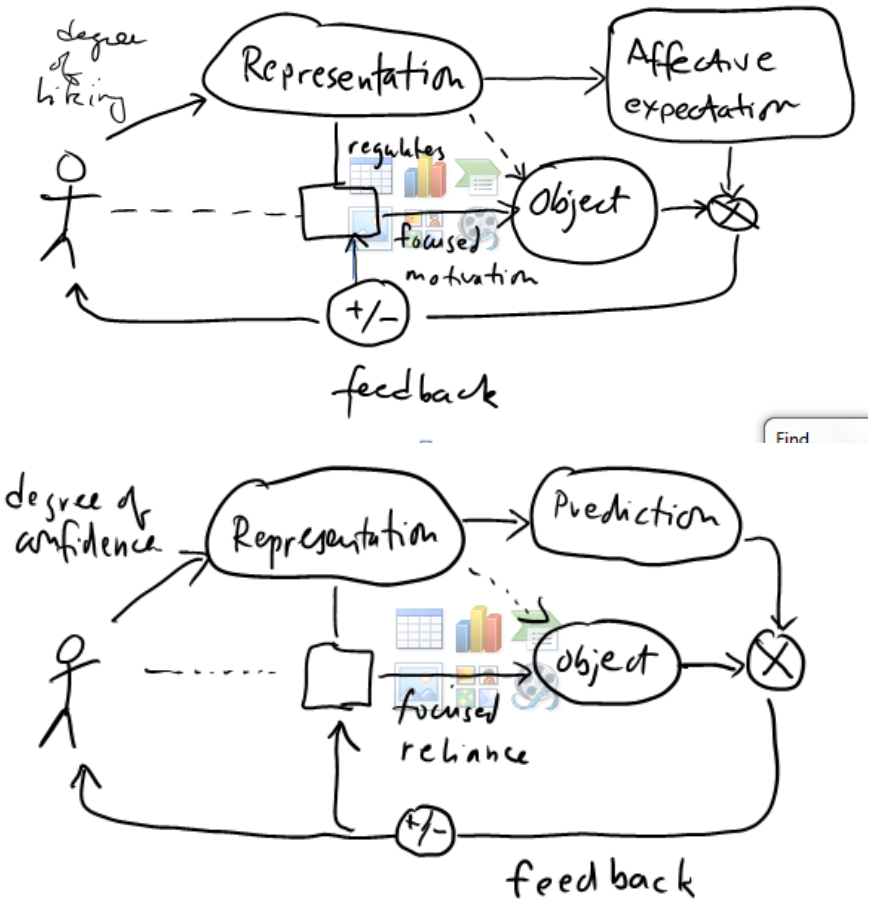

Desire and belief: sustained action guidance "under an idea"

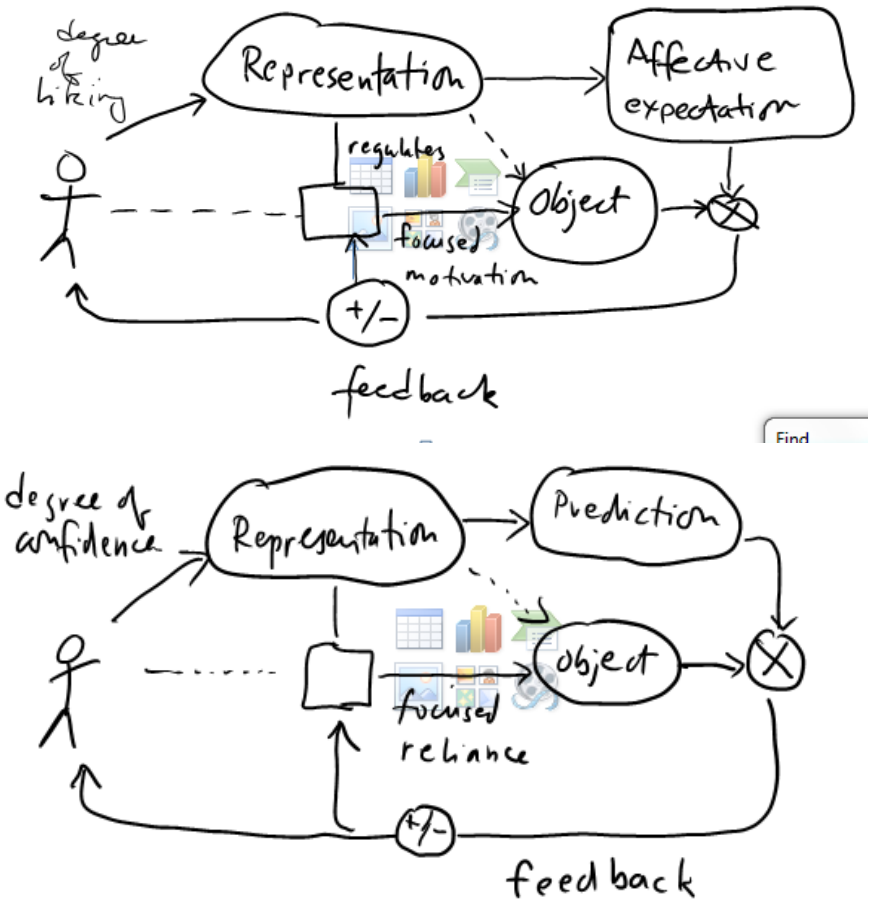

Desire and belief: contribute *decision-weights*, to guide choice, regulate behavior, and make it intelligible

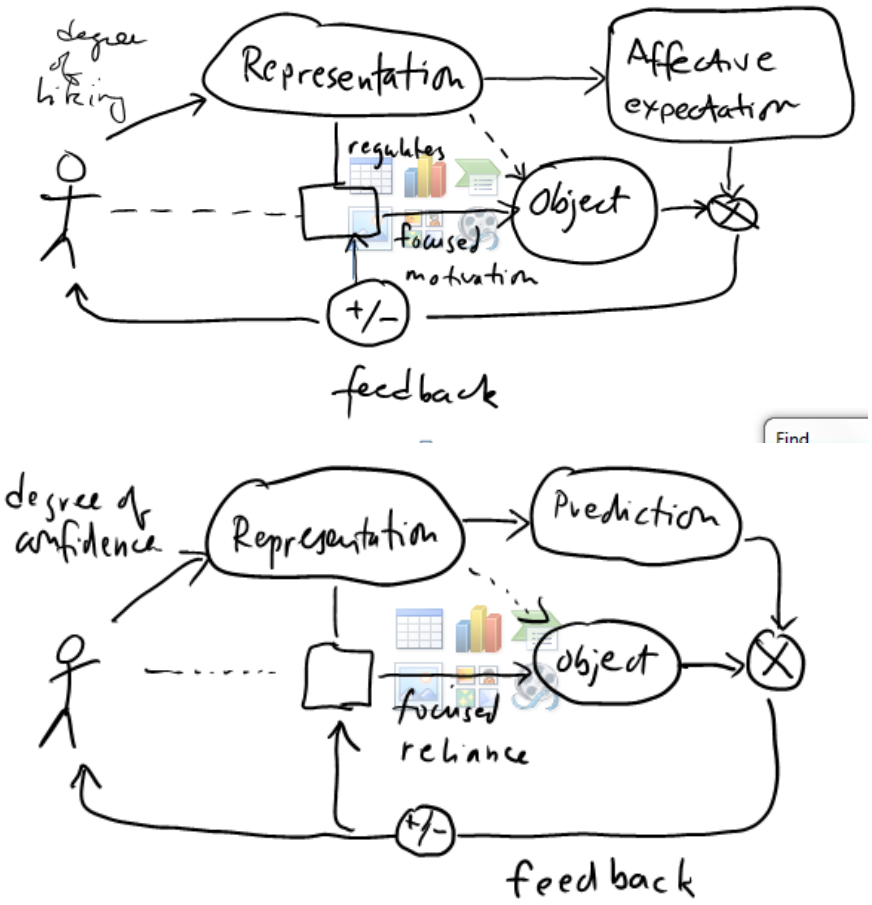

Desire and belief: spontaneously *projective*

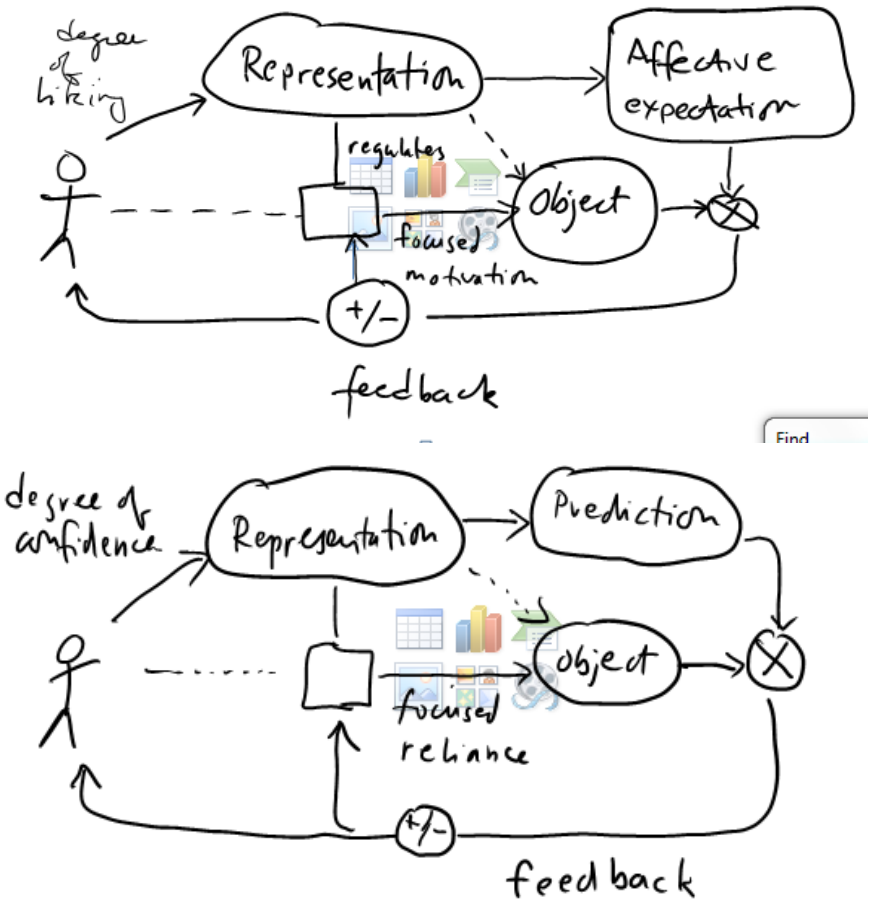

Desire and belief: spontaneous experience sensitivity

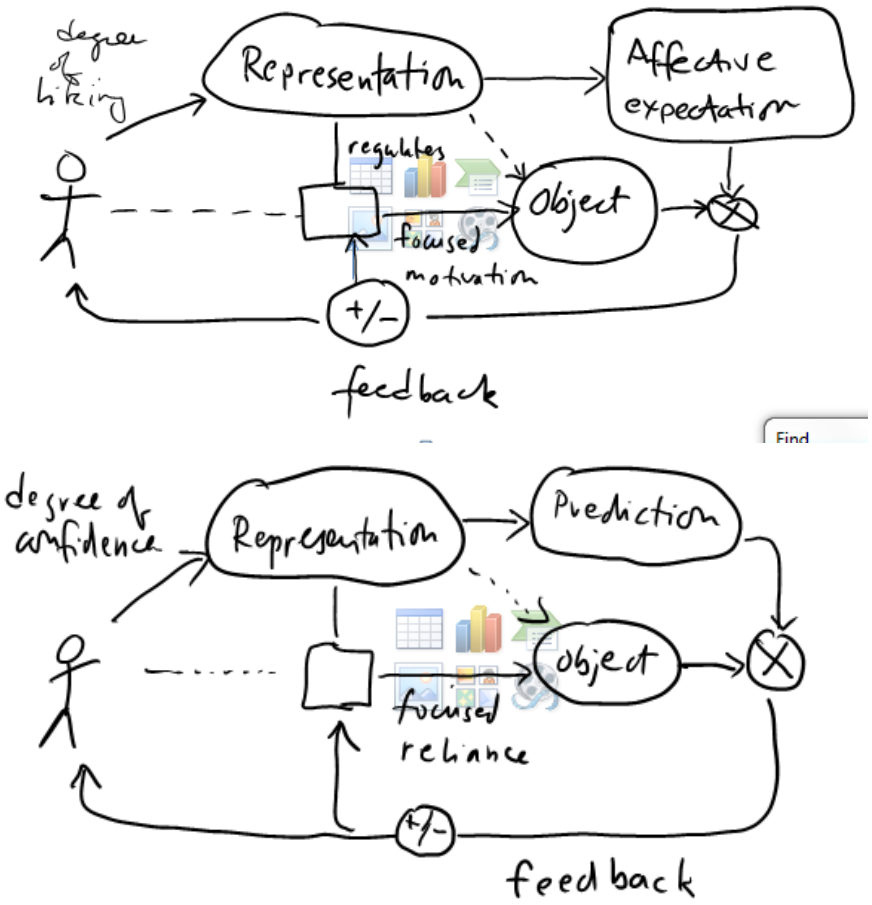

Desire and belief: … and transfer thought- and action-guiding force among them

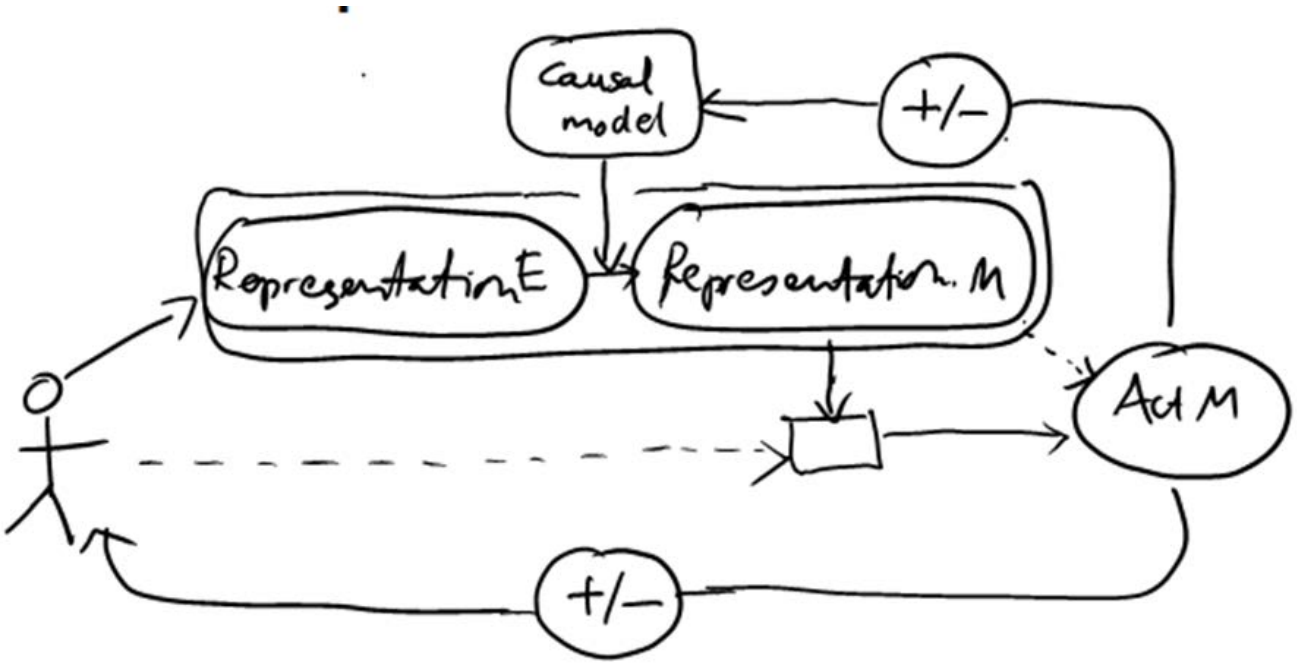

Desire and belief: two kinds of strength

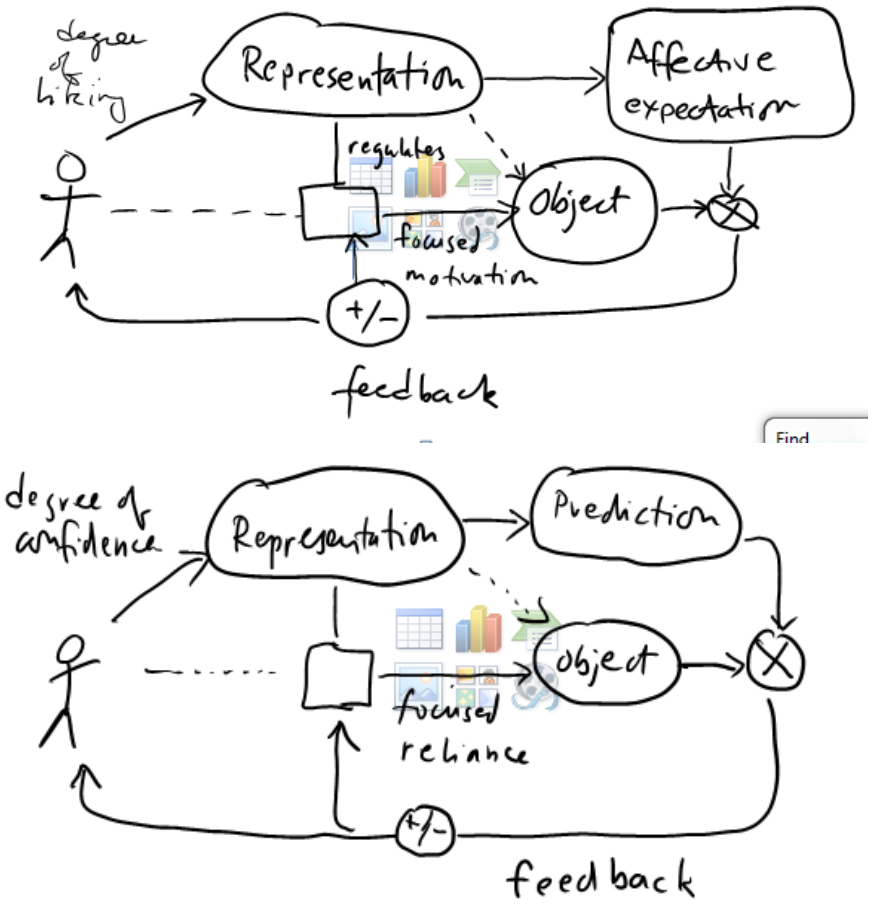

Desire and belief: regulation of motivation and reliance by evaluation—and failures in *akrasia*, *addiction*, and *phobia*

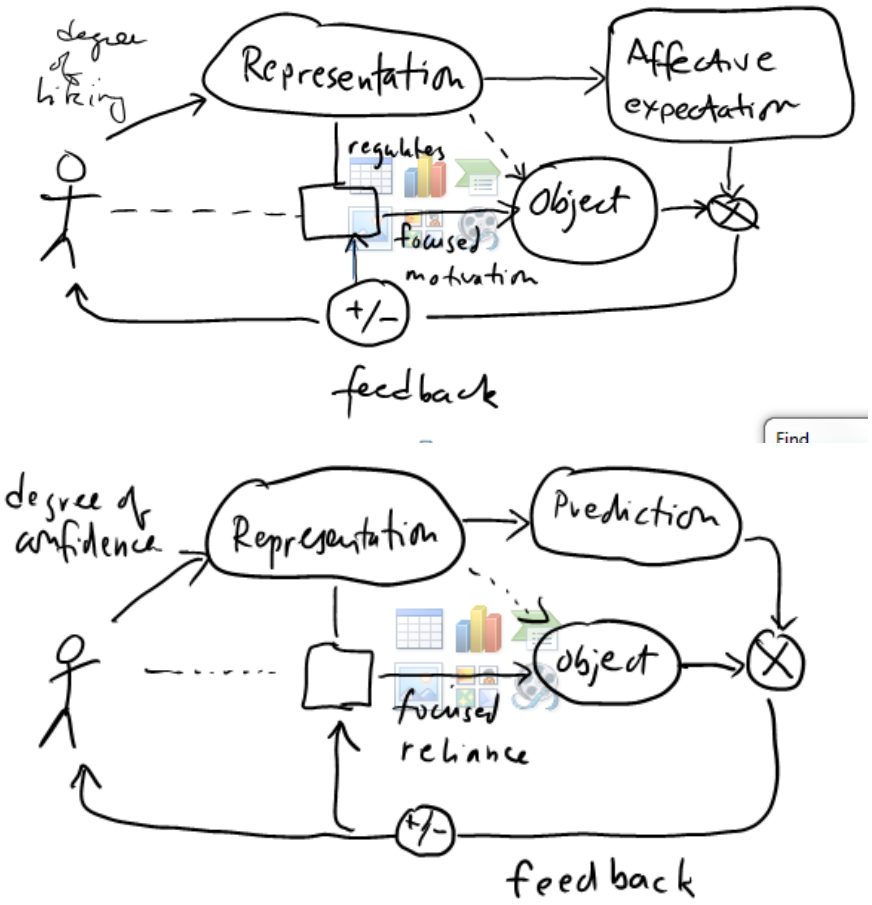

Desire and belief: spontaneous *de se* co-ordination of *active action-tendencies* around shared intentional content

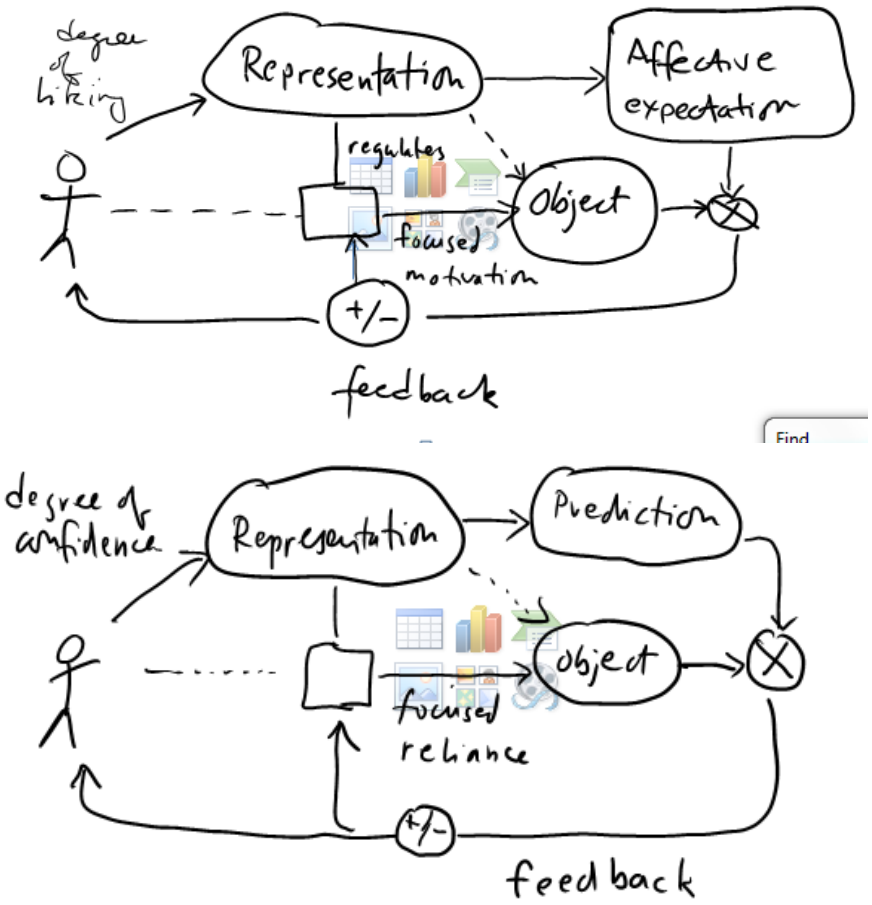

Resulting revisionist belief-desire model

- **belief + desire → action**
- *representational representational*
- *affective+act-guiding affective+act-guiding*
- *mind-to-world m-t-w as well as w-t-m*
- *accuracy of predict. accuracy of evaluative predict.*
- *representation representation*
  regulates reliance regulates motivation
  with learning with learning
  from discrepancy from discrepancy

- *potentially rational + potentially rational →*

We introduced the idea of multiple dimensions of mind-to-world fit

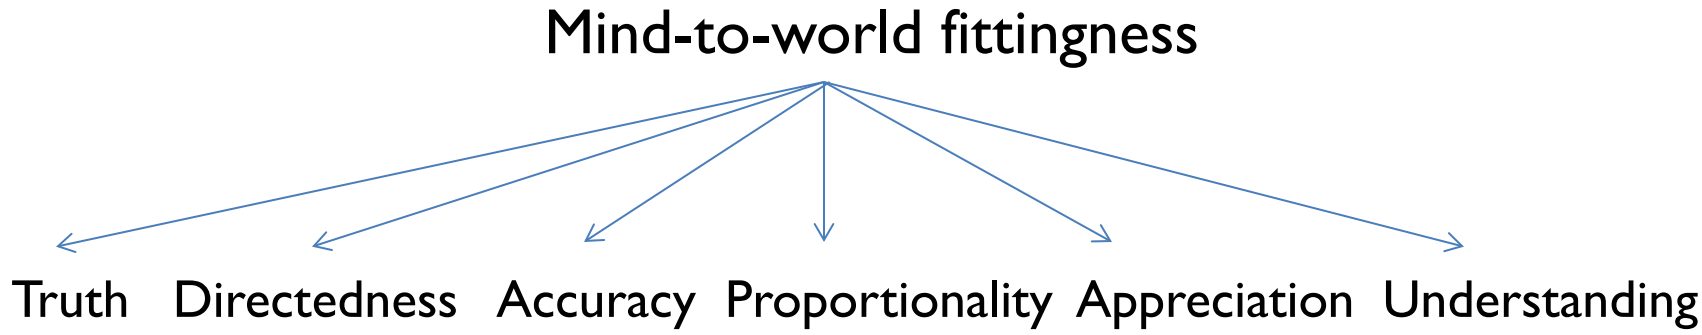

Resulting revisionist belief-desire model

- **belief + desire → action**
- *representational representational*
- *affective+act-guiding affective+act-guiding*
- *mind-to-world m-t-w as well as w-t-m*
- *accuracy of predict. accuracy of evaluative predict.*
- *representation representation*
  regulates reliance regulates motivation
  with learning with learning
  from discrepancy from discrepancy

- *potentially rational + potentially rational → potentially rational*

The claim is not …

- … put together any two rational things, and you'll get some third, rational thing.
- Rather, there can be specific ways of putting together rational parts to make a rational whole.
  - These ways are very restrictive, and reflect the distinctive character of rationality in the *broad sense*:
    - the idea of capacities for, and exercise of, apt responsiveness to reasons.
- Moreover, the parts must in some sense be "made" such that this kind of combination is possible, and will work to yield apt responsiveness to reasons. Combination alone won't do it.

Is there a recognizable approach to rationality that could play this dual role in both belief and desire?

- For both belief and desire, it could be *dynamic rationality* of the kind found in contemporary formal epistemology and rational decision theory.
  - On such views, rationality is about how one responds to experience, not about one's starting point.
- A generic version of such responsivieness or *updating*, (*):

- *prior expectation* → *experience* → *discrepancy-detection* → *discrepancy-reducing posterior adjustment in expectation* → *new prior expectation* → *new experience* → …

Updating and attunement

- On the revised account of desire and belief given here, they both have this (*)-like generic, updating character.
  - While any (*)-like process is dependent upon its starting point, under some broad assumptions, such updating will, given sufficiently extensive and diverse evidence, have the result that individuals starting with different prior expectations will tend toward expectations "attuned" to what actually happens—in terms of frequency or reward value. (We saw examples of this "attunement" in reinforcement learning and real and simulated foraging tasks.)

Effective learning

- (*)-type learning processes are therefore found throughout the animal world,
- … in behavior and in brain systems,
- … in human infants' learning patterns (Wellman, 2014),
- … and increasingly are in building inquirers "from scratch"—e.g., in artificial intelligence and scientific investigation.

Some limitations

- Such learning of course always is limited in many dimensions, and is highly dependent upon the kinds of experiences and concepts available to the individual.
  - As we'll see in subsequent lectures, this gives an important and distinguished role to more self-conscious thought.
- But it gives us an entry point into potential rationality.

Dynamic rationality of this kind is rationality in the broad sense …

- … —that is, apt responsiveness to reasons—rather than narrow: although these are rule-like processes, the learning typically takes place implicitly and does not involve self-conscious applications of the rule by animals or humans.

  – But: Let's not fight over a word: instead, let's call this kind of apt responsiveness to reason-constituting considerations *intelligence*—on the model of animal intelligence, artificial intelligence, intelligent system, etc.

The point, however, is not to lower the bar

- … from strict, narrow rationality.
  - Indeed while the following picture is attractive:

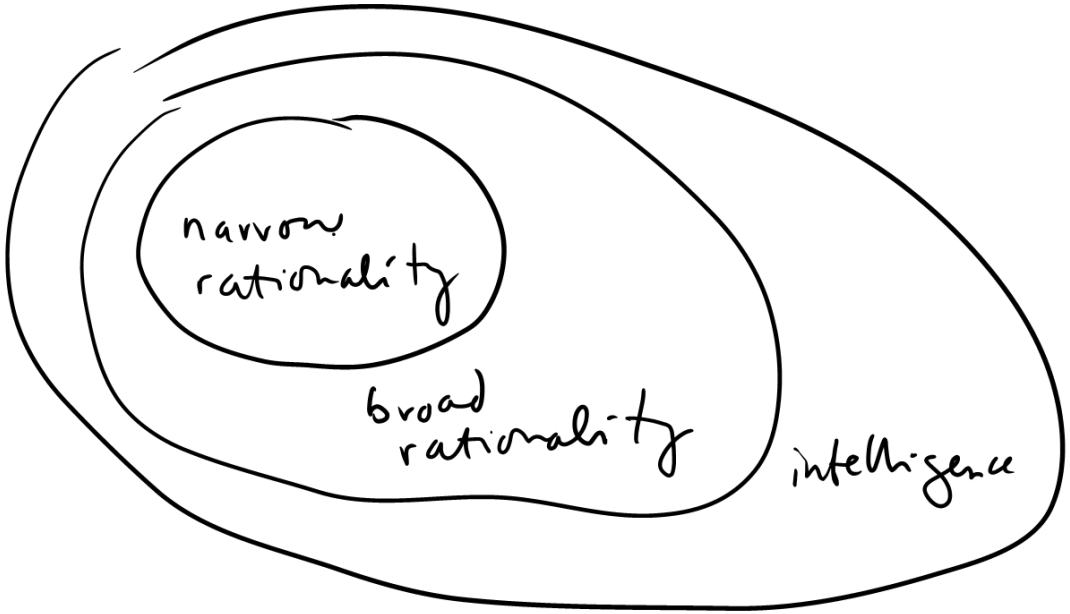

  - … it can't be right.

Narrowly rational agents …

- … are guided by reasoning and consistency constraints—consistency in their beliefs, means-end consistency in action, and "judgment-sensitivity" to such reasoning in what they think and feel.

But …

- … if the agent's conclusions about what reasons she has arise from a relatively limited information or parochial concerns, then she might be "perfectly rational" in the narrow sense, but unresponsive to important considerations that lie outside her purview—and that might only be able to influence her if she were *less* coherent.
  - This is not merely a notional problem—all of us are like this sometimes, and in such cases it is not clear that we'd rate being narrowly rational as either broadly rational or intelligent.
- Which would you prefer in a friend? Or parent? Or child?

Desire and belief are regulative processes—what might make such a process intelligent?

- Consider for example a regulative system like a thermostat, which controls the activity of a heating system.
  - If given a fixed set point, it will tend to regulate the allocation of heating power to hold room temperature at or near that point.
  - However, if the set point is changed, it will typically overshoot and resettle slowly.
- If the thermostat could predict the behavior of the heating system, and the change you make in temperature setting, it could increase or decrease heat more smoothly, improving both energy efficiency and efficacy in having the room at temperature desired. An *intelligent* thermostat learns to predict by (*)-like learning of your *pattern* of use.

The "Good Regulator Theorem"

- That is, the intelligent thermostat learns a *model* of your behavior and that of the heating system, and guides its activation accordingly.
- According to the "Good Regulator Theorem" of Conant and Ashby (1970; cf. Eykoff, 1994):
  - "… under very broad conditions, … any regulator that is maximally both successful and simple must be isomorphic with the system being regulated."
- That is, a good regulator of a system will contain a *model* the system.
  - "… Making a model is thus necessary."

Since the mind is the regulator …

- … of the organism's interactions with the environment and its own body, intelligent animals will learn models of potential states of these systems, potential actions available, likely outcomes, effects on needs or goal-attainment, etc.
  - An intelligent being, therefore, is a model-building creature.
  - Such models are *projective*, projecting the model's relations forward to guide subsequent action
  - … and they are *hierarchical*, since they extract more general patterns and relations, which improve predictability and efficacy.

Thus *model-based regulative control*, …

- … like model-based learning, is ubiquitous in animals, brain subsystems, and, increasingly, artificially intelligent agents.
  - Intelligent action, then, we will understand to be action that is guided by learning and acting through such causal-evaluative models, updating dynamically in light of the outcomes.

Really?—Intelligent animals build causal-evaluative models, engage in forward and inverse inference, etc.?

Warning:

NOT PHILOSOPHY

We've already seen some of the evidence

Recall: separate representation of expected value vs. risk

(Fiorillo *et al.*, 2003)

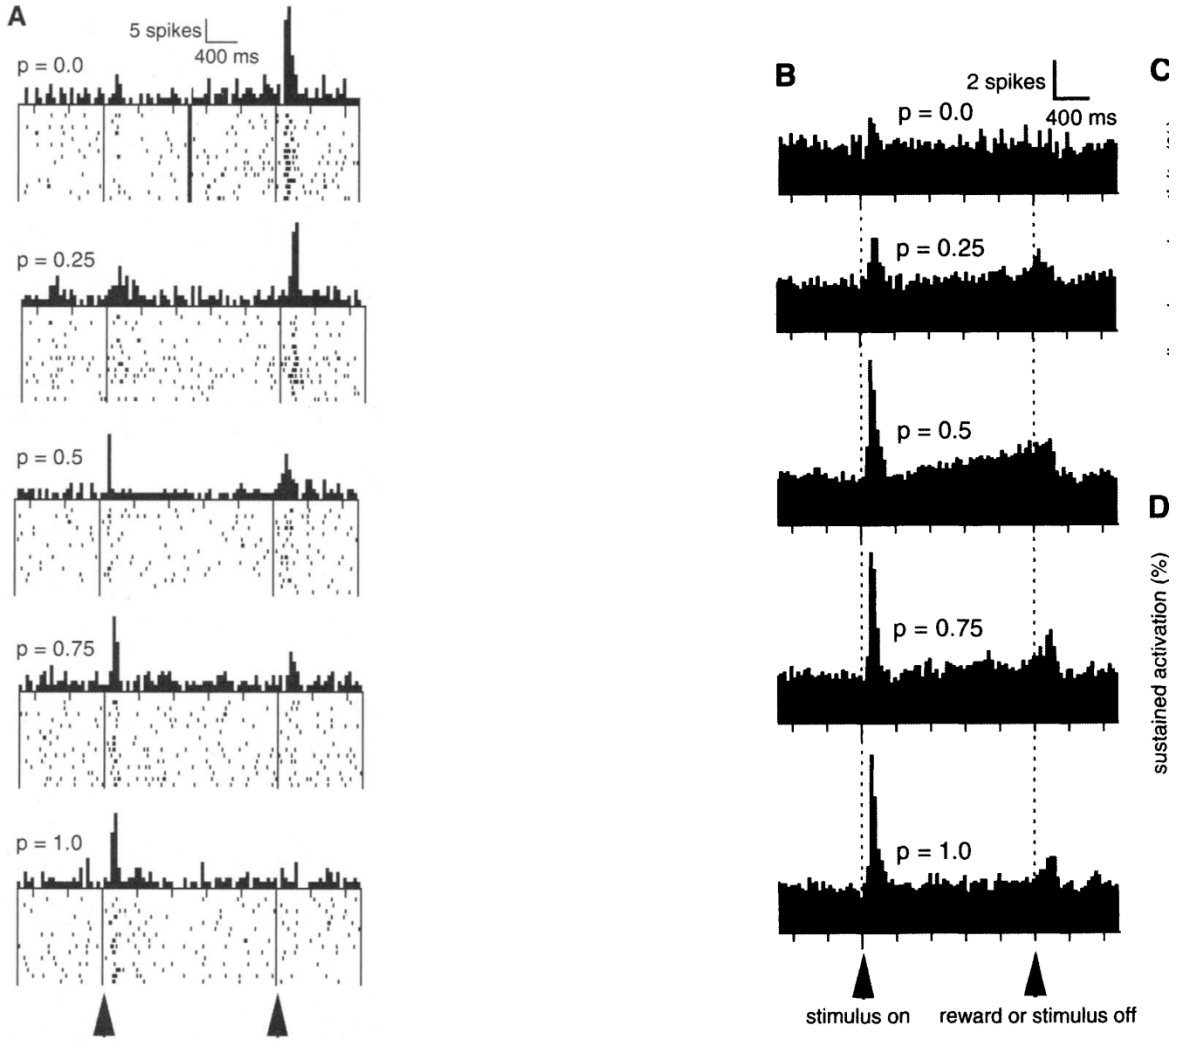

38

Separate representation of expected value vs. risk—activation in the human ventral striatum (Quartz, 2009)

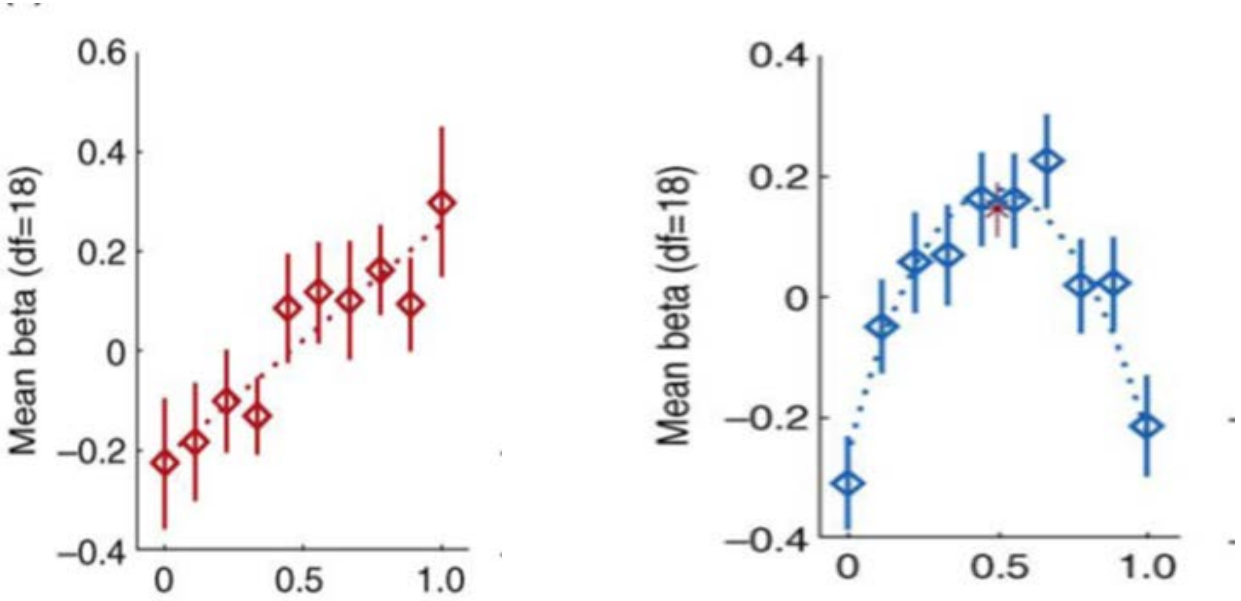

Constructing predictive utility functions from gambles

(Stauffer *et al.*, 2014)

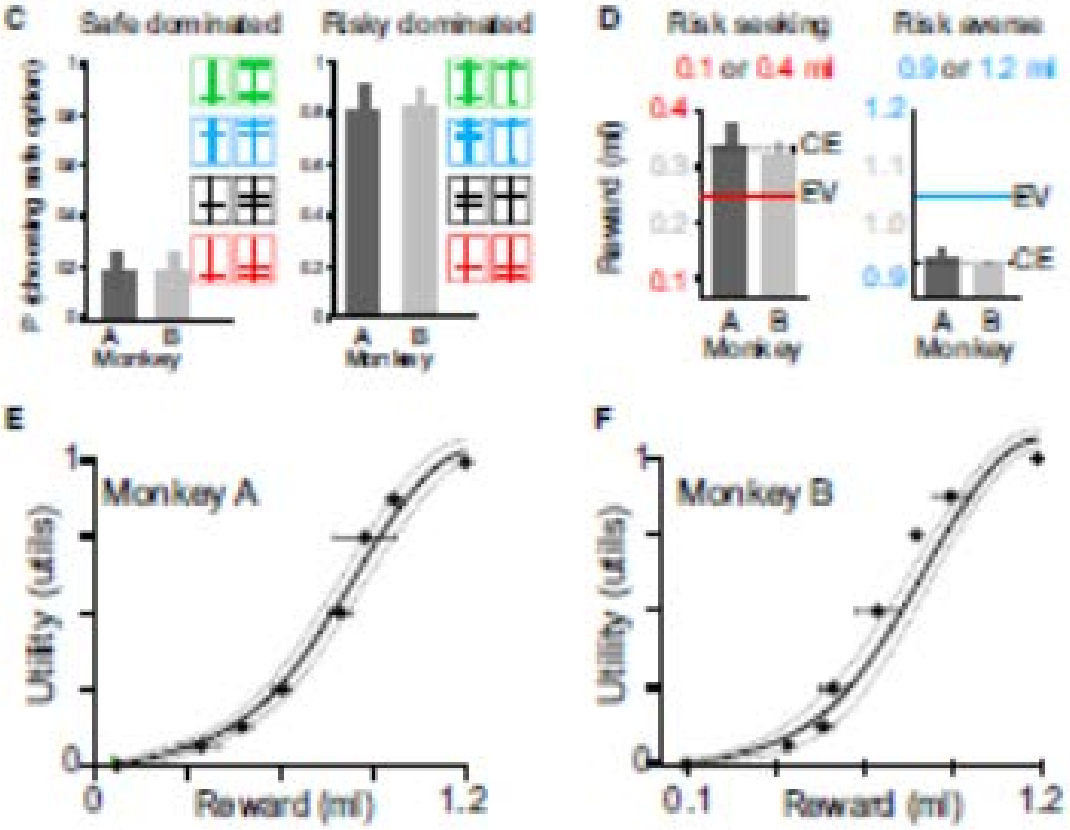

Modeling space: place and grid cells in the rat hippocampus and entorhinal cortex (Moser & Moser 2014)

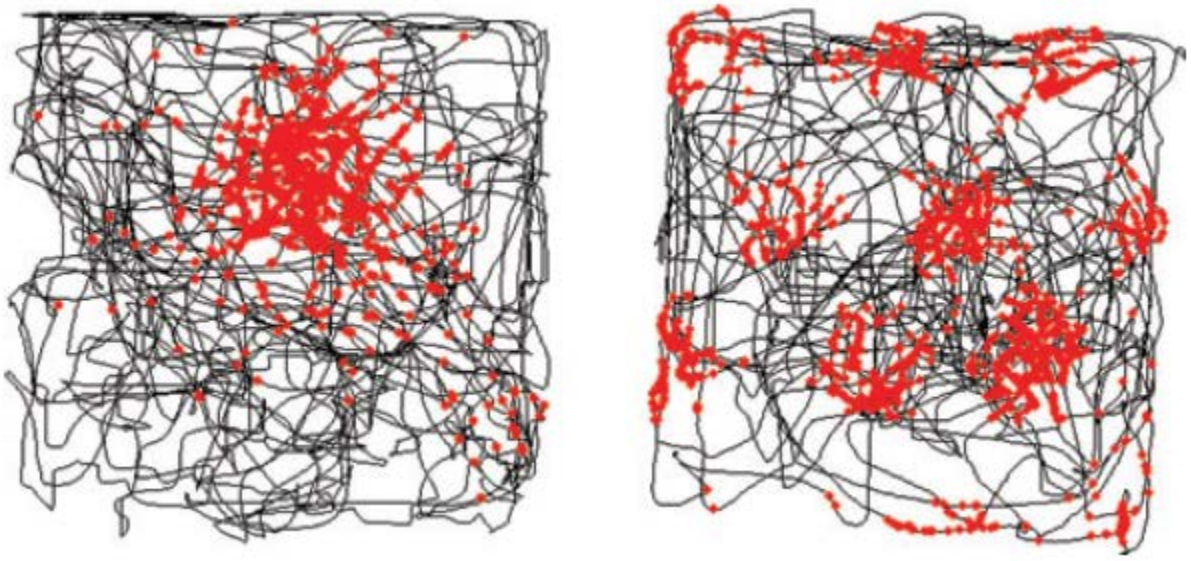

T-maze with return rails

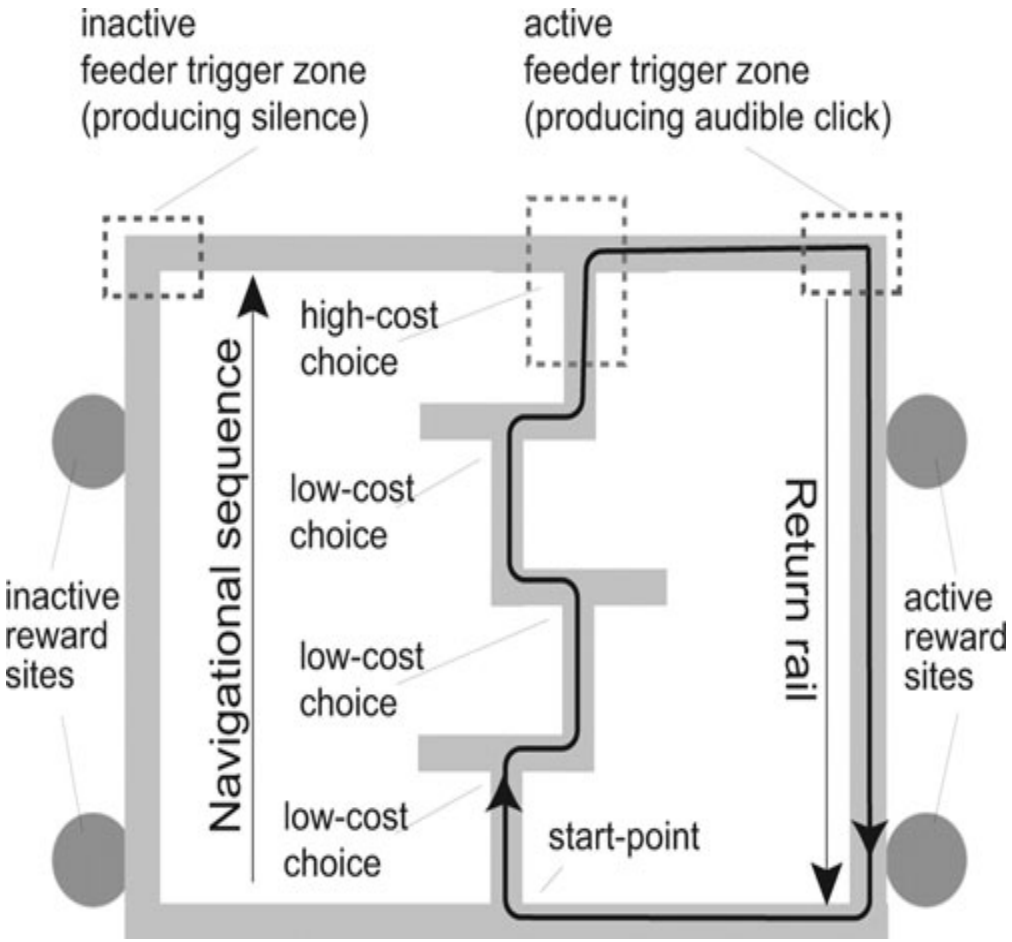

Acquired cognitive map—hippocampal place cells

(Johnson & Redish 2007)

Co-ordinated replay of mapping in the rat hippocampus during sleep (Ji & Wilson, 2007)

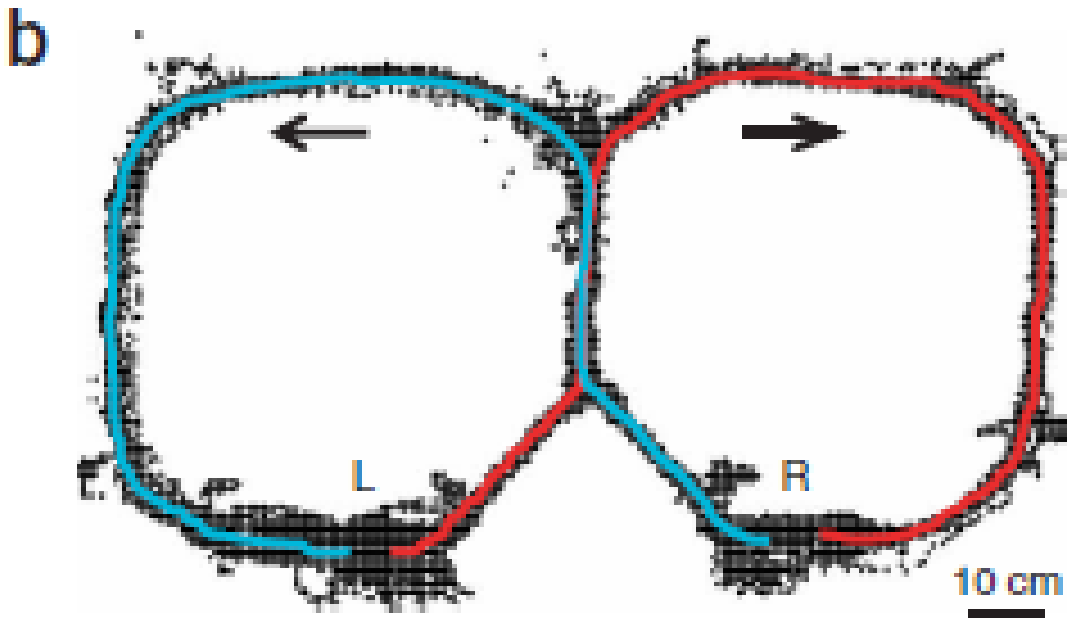

Hippocampal construction of novel paths in sleep

(Gupta *et al.*, 2010)

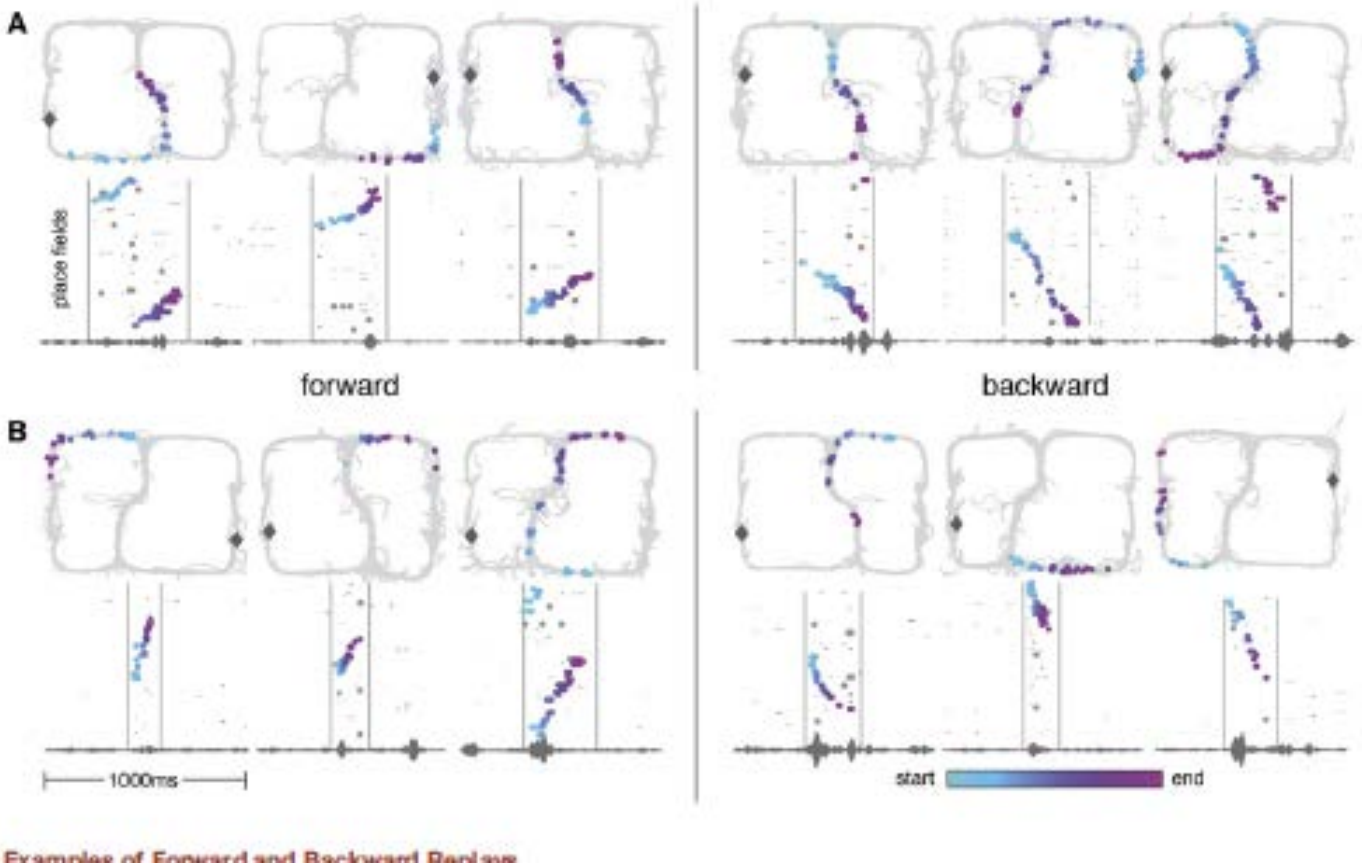

A rat using a causal-evaluative model in real-time decision-making (Johnson & Redish, 2007)

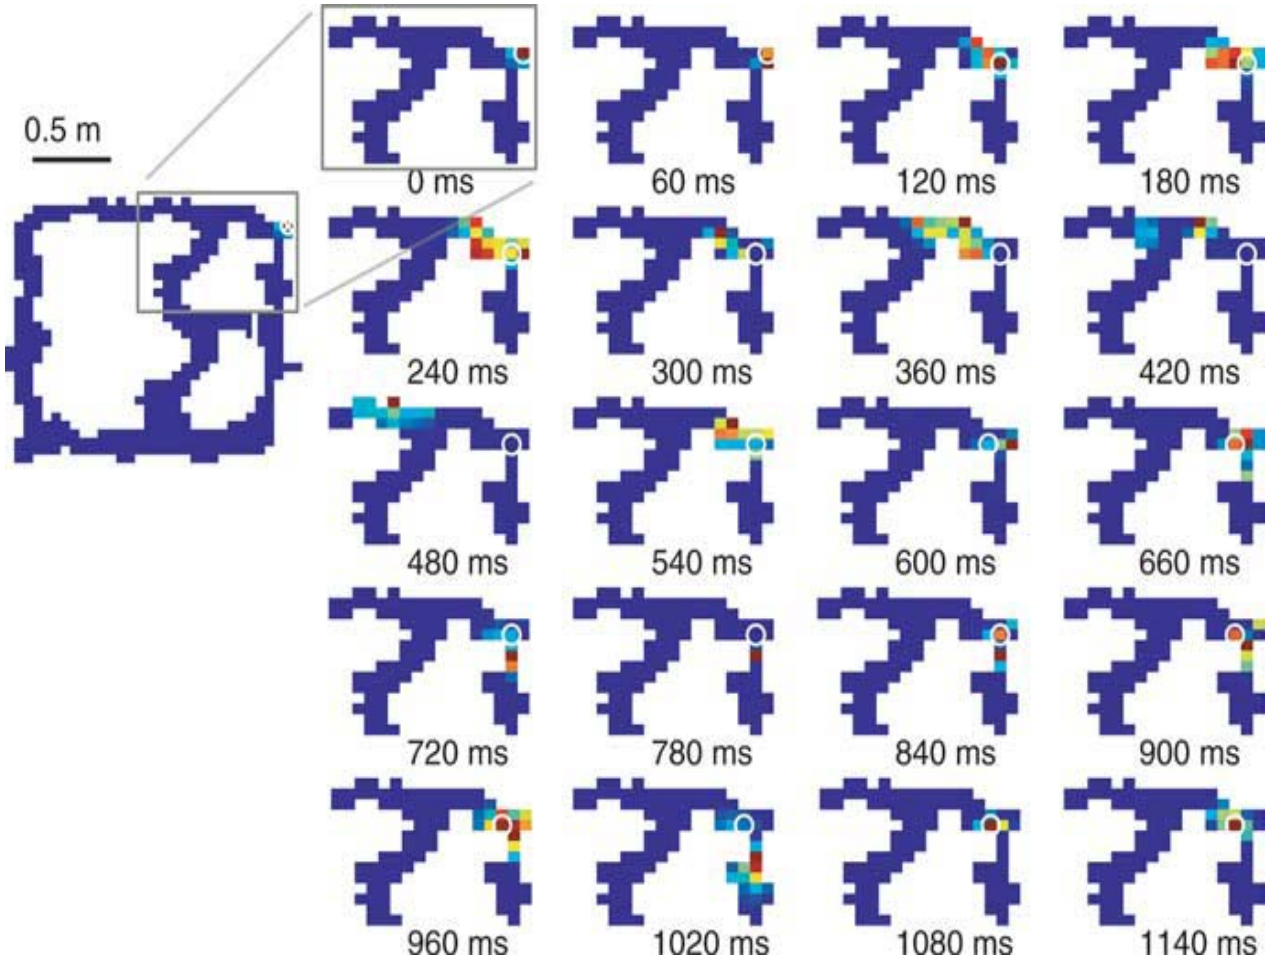

46

Metabolically and developmentally …

- … such a model-building, projective brain is expensive.
- Even in a "resting" state, the brain of intelligent mammals continuously consumes up to 15-20% of the body's oxygen and calories, despite constituting 2% or less of body weight (Raichle & Gusnard, 2005).
  - At the level of metabolic brain activity is remarkably constant, whether the brain is "on task" or "resting".
  - What is it doing in the rest or "default" state?

Simulation and learning

- Evidence suggests that brain is occupied in consolidating, organizing, and anticipatory tasks, simulating possible futures and updating representations and evaluations in light of these simulations and recent experience (Buckner *et al.*, 2008; Lewis, Baldassarre, Committeri, Romani, & Corbetta, 2009; Bollinger, Rubens, Zanto, and Gazzaley, 2010).
  - Prospective, model-based simulation and control gives us an explanation of optimal foraging in animals (Dugatkin, 2004) and humans (Kolling, *et al.*, 2012).
- Animals and human infants display patterns of causal inference that suggests inverse use of causal models (Blaisdell *et al.*, 2006; Gopnik *et al.*, 2004).

What sort of brain did prospection favor?

- The rat brain is estimated to have 200 million neurons, and 5 x 10^{11} synapses.
  - Each neuron is capable of firing from once a second, to many times a second.
  - And many areas of activity can operate in parallel.
- The human brain is estimated to have 80-100 billion neurons, and 10^{14}–10^{15} synapses.
  - It is estimated to have a "cycle time" in the range of 80 billion to 15 trillion action potentials *per second*.
- The architecture of the neo-cortex is remarkably similar throughout—it is the architecture needed for neural net deep learning via general-purpose processes like (*).

Back to where we were …

Back to where we were …

- Intelligent action, we were saying, will be understand here to be action that is guided by such causal-evaluative models and learning.
  - The mental states humans would need to engage in such intelligent action thus would have to be:
    - Representational and content-sensitive
    - Capable of spontaneously organizing and coordinating thought- and action-guidance
    - Projective and responsive to feedback from discrepancy with expectation

This is a lot to ask of mental states …

- … without introducing some additional, intelligent controlling agent to oversee and control these processes.
  - Have we seen any models of human mental states with these spontaneous capacities for intelligent regulation of behavior and learning?

Intelligent, model-based learning and control in humans

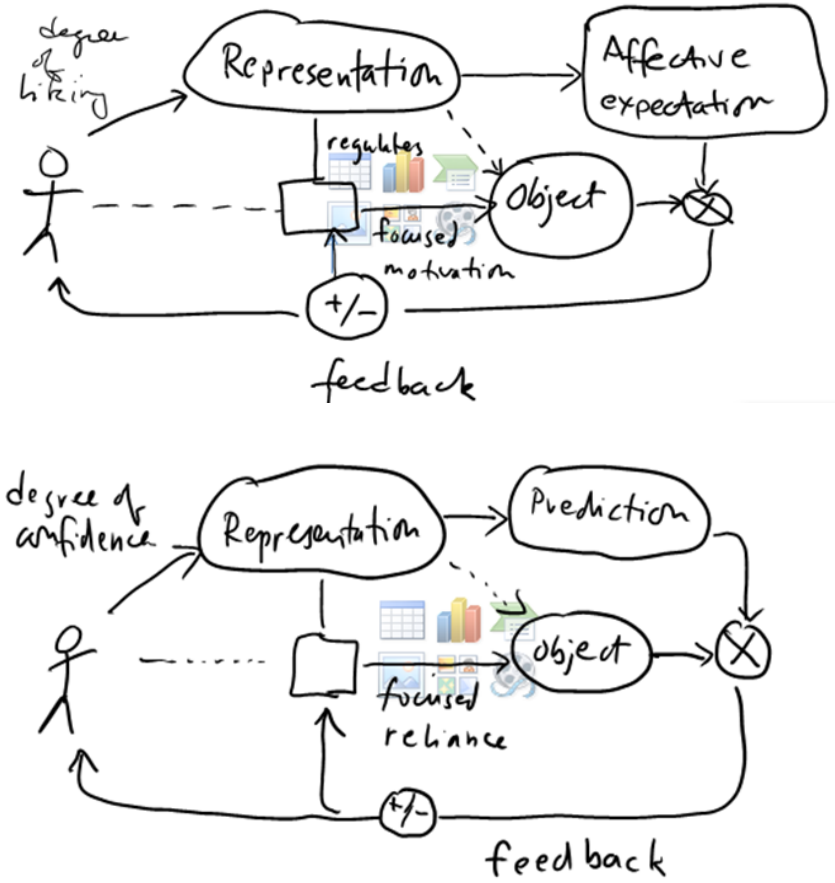

The revisionist revisionist belief-desire model

- **belief + desire → action**
- *representational* *representational*
- *affective+act-guiding* *affective+act-guiding*
- *mind-to-world* *m-t-w as well as w-t-m*
- *accuracy of predict.* *accuracy of evaluative predict.*
- *representation regulates reliance by learning from discrepancy* *representation regulating motivation by learning from discrepancy*

- *potentially intelligent + potentially intell. → potentially intell.*

But still …

- … intelligent as these states might be, this is still just a "causal sequence"—what would put its components together to make anything with the overall telic structure of intentional action, or action for a reason? Where is the unity and intentionality of agency?

Going to the dentist

1. She desires: Get rid of toothache.
2. She believes: Going to the dentist gets rid of toothaches.

<*jointly cause*>

3. She initiates: Going to the dentist.

Joint causation is not distinguishable from madness

(Korsgaard 1997, 221-222)

[A]s Nagel points out in *The Possibility of Altruism*, the specifically rational character of going to the dentist to avert an unwanted toothache depends upon *how* the belief [that dentists cure toothaches] and the desire [to be rid of one's toothache] are 'combined'. It is certainly not enough to say that they jointly *cause* the action, or that their bare co-presence effects a motive, for a person might be conditioned so that he responds in totally crazy ways to the co-presence of certain beliefs and desires. . . . If the belief and desire . . . operate on the person merely by having a certain causal efficacy when co-present, the rational action is only accidentally or externally different from the mad one. . . . For the person to act rationally, she must be motivated by her own *recognition* of the appropriate conceptual connection between the belief and the desire.*

Bringing in recognition of a conceptual relation

1. She desires: Get rid of toothache.
2. She believes: Going to the dentist gets rid of toothaches.
3. She recognizes: (1) + (2) have a conceptual connection such that they constitute a reason to go to the dentist.

<*jointly cause*>

4. She desires: Going to the dentist.

<*causes*>

5. She initiates: Going to the dentist.[10]

- But couldn't (1)-(5) also just come about through some bizarre conditioning? So we'd have to move to:

Bringing in recognition of a conceptual relation

1. She desires: Get rid of toothache.
2. She believes: Going to the dentist gets rid of toothaches.
3. She recognizes: (1) + (2) have a conceptual connection such that they constitute a reason to go to the dentist.
4. She recognizes: (1) + (2) + (3) have a conceptual structure such that they constitute a reason to want to go to the dentist.

<*jointly cause*>

5. She desires: Going to the dentist.

<*jointly cause*>

6. She initiates: Going to the dentist.

(Note that the problem does not arise because of the employment of 'cause'—even if we think the recognition of a reason *constitutes* a desire, which then *partially constitutes* an intention, we still need first 2 desires, then 3, …)

Theoretical and practical tortoises

- This suggests that the lesson from Lewis Carroll's "Tortoise and Achilles" needs to be appreciated more widely.
  - It is often said that Carroll's argument shows that we cannot think of rules of inference as additional premises of the inference, on pain of regress.
  - Yes, but the problem is worse—we also cannot think of *applying the rules* as operations that intervene explicitly or implicitly between every step in an inference.
- So what *is* an inference, if not applying rules explicitly or implicitly to go from one step to the next?
  - A series of *non-inferential* steps, *not agentially-mediated*, the taking of which *constitutes an inference by the agent*.

Theoretical and practical tortoises

- Of course, a fully rational agent can perfectly well "step back" from a given inference, and take the series of inferential steps by self-consciously invoking and applying rules.
  - But even this involves non-inferential steps of noticing that a rule might apply, invoking a rule, applying it to the premise at hand, etc. Such steps cannot be even *implicitly* agential and deliberative.
  - So an agent's self-conscious rule-following rests upon inferences carried out by a causal sequence of non-inferential steps, not unified by agential intervention.

"Blind" dispositions?

- Wittgenstein (*Investigations* §219) suggested that what we need are *blind dispositions* to make these immediate transitions in thought, "without reasons".
  - This makes it hard to see how the agent could take herself to be explicitly or implicitly *following a line of thought* or *inferring for a reason*—forms of acting for a reason.
  - Or how *we* could see her as inferring logically or rationally. Here Korsgaard's worry resurfaces.
- Such dispositions would work only if they *weren't* really blind—they would have to be responsive to the occurrence and content of certain thoughts, relations of semantic or logical relevance, strength of belief, etc.

"Intelligent dispositions"

- What is needed, it would seem, are dispositions to make direct transitions in thought (whether the thought is conscious or unconscious) that are:
  - *Alive* to the occurrence of certain thoughts.
  - *Sensitive* to the form or contents of the thoughts, and to their relevance and intentionality, etc.
  - *Regulative*, non-accidentally guiding what mental transitions one makes in light of the thoughts' contents and context, one's degree of confidence in them, etc.
  - And involve *monitoring and feedback*—whether the transition results in a thought that seems irrelevant, etc.
- These are ways dispositions might be *intelligent*.

Once again, …

- Have we seen any examples of how we could have intelligent dispositions to believe or act inferentially in the sense just described?
  - Spontaneously projective and thought- and action-guiding
  - Content-sensitive
  - Regulative
  - Feedback

Once again, …

- Have we seen any examples of how we could have intelligent dispositions to believe or act inferentially in the sense just described?
  - Spontaneously thought-guiding and projective
  - Content-sensitive
  - Regulative
  - Feedback

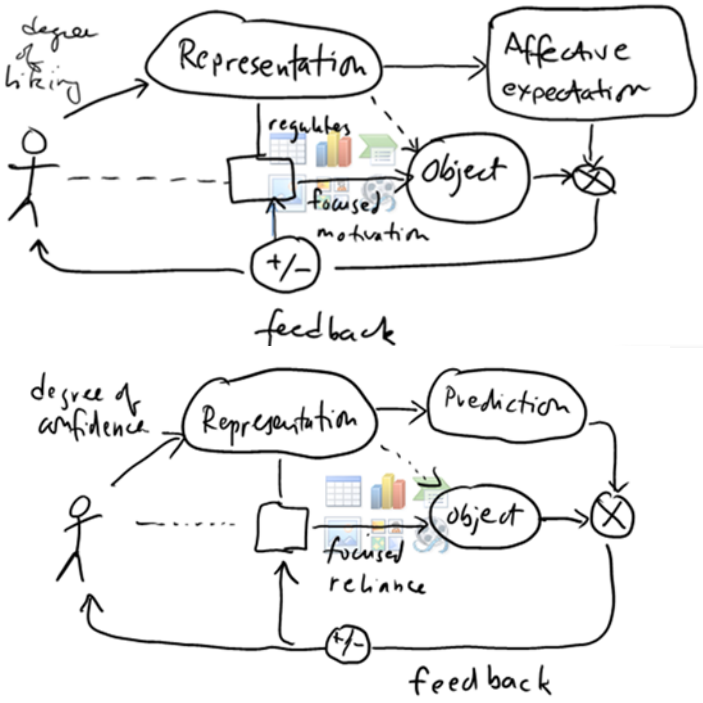

In model-based regulation and control

- The relations relevant to the behavior of the system—spatial, causal, evaluative, semantic, correlative, etc.—are *encoded* in the transition-matrix of the model structure of the regulator, such that inputs yield outputs without deliberation or agential intervention.
  - These relations in turn are revised depending upon whether outcomes fit output expectations.

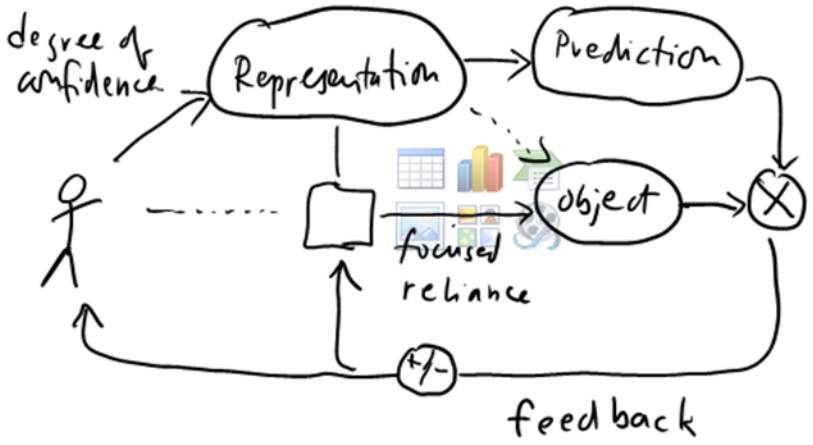

Let's look at some further features of the revised account of desire and belief …

- … to see how far they might go in their tandem operation to yield the features of intentional action.
- Field notes: What are some of the features people have taken to be vital for intentional action?

Desire and belief: *de se* representation of intentional object—acting through mode of presentation, opaque context

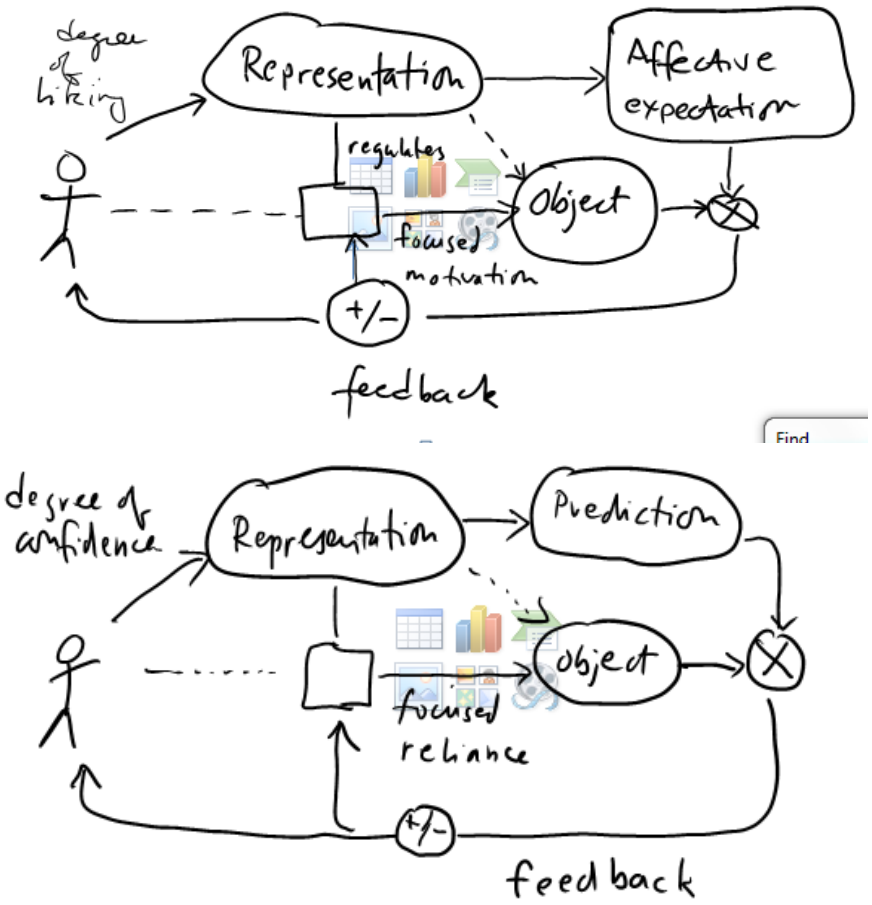

Desire and belief: spontaneous *de se* co-ordination of *active action-tendencies* around shared intentional content

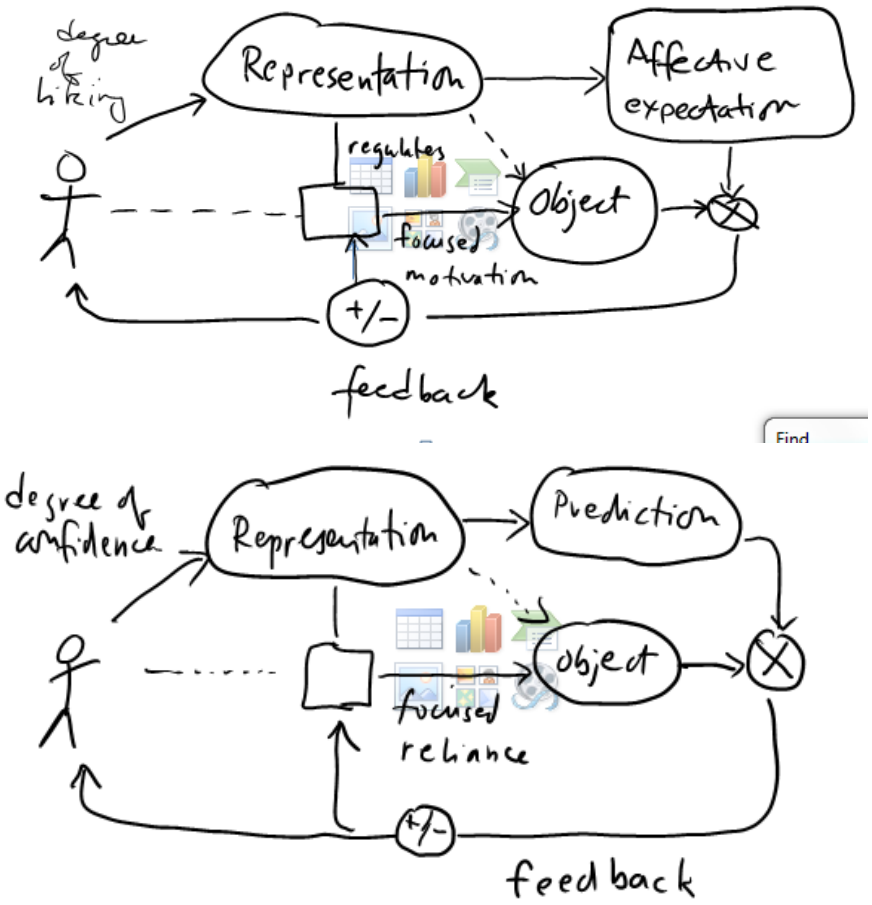

Desire and belief: sustained action guidance "under an idea"

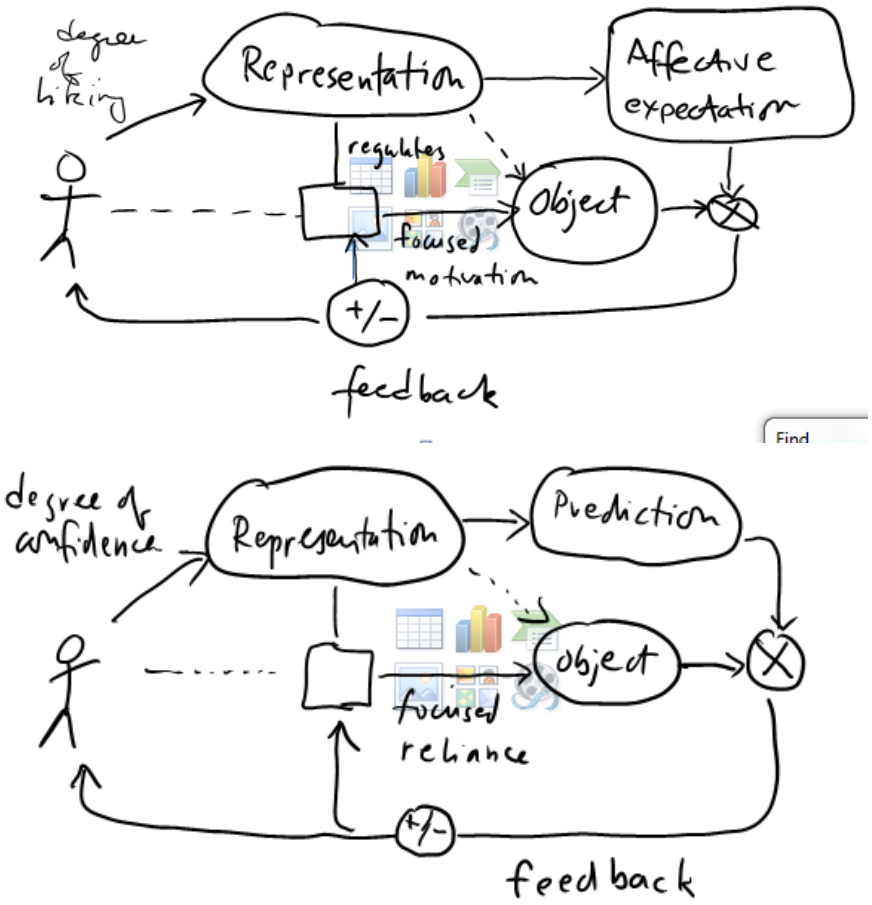

Desire and belief: "desirability characteristic", satisfaction condition that is also an answer to what one is doing

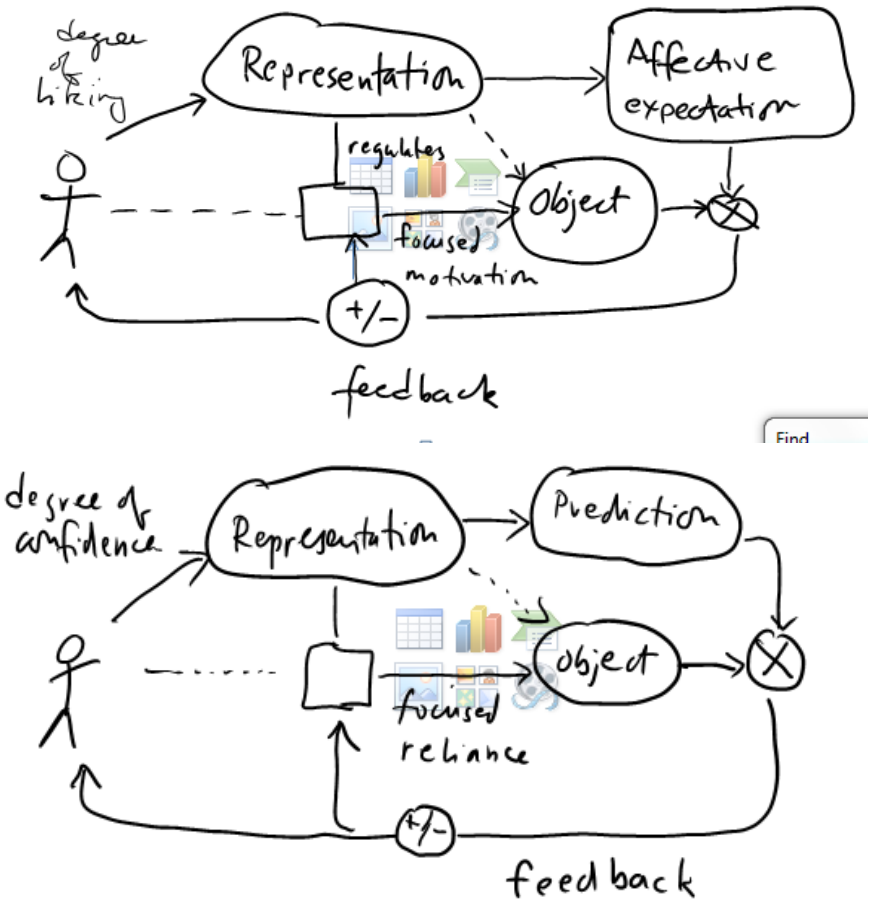

Desire and belief: contribute *decision-weights*, to guide choice, regulate behavior, and make it intelligible

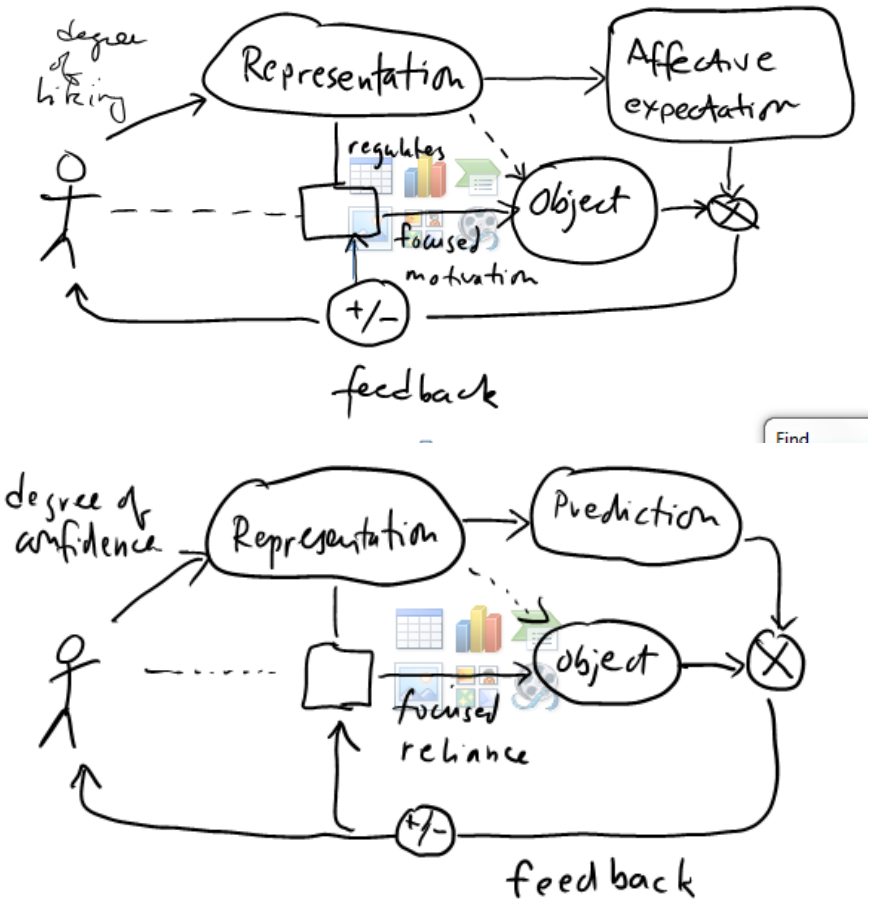

Desire and belief: take multiplicity of objects—persons, processes, actions, abilities, propositions, etc. …

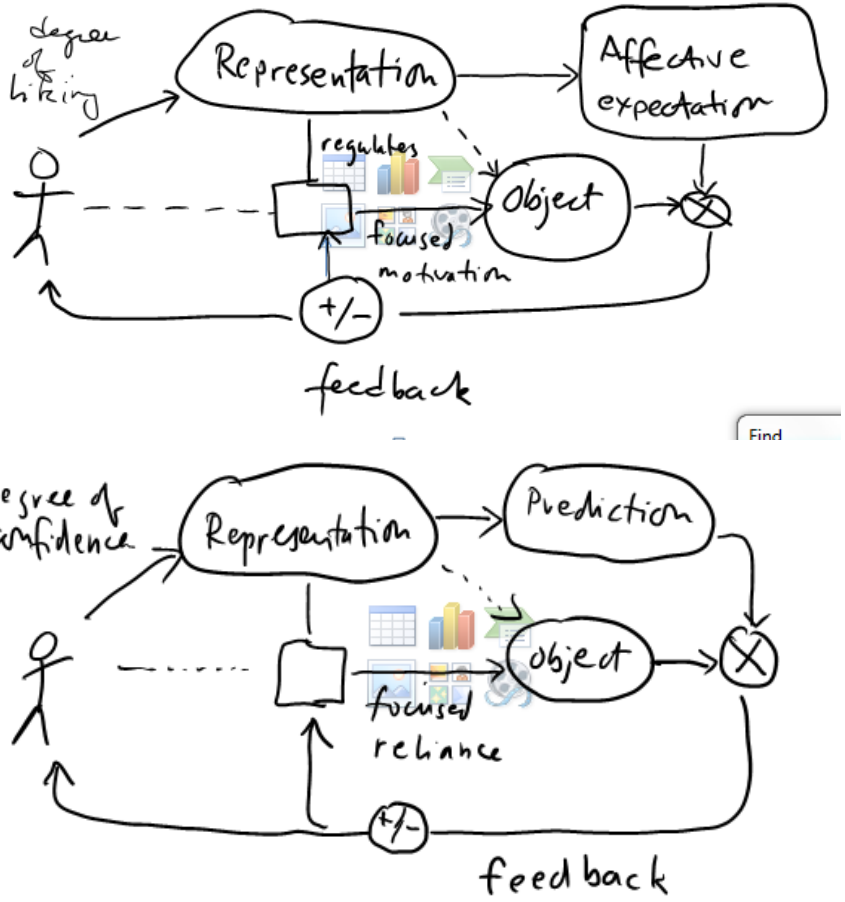

Desire and belief: part of a "common pathway" from the agent's point of view to action

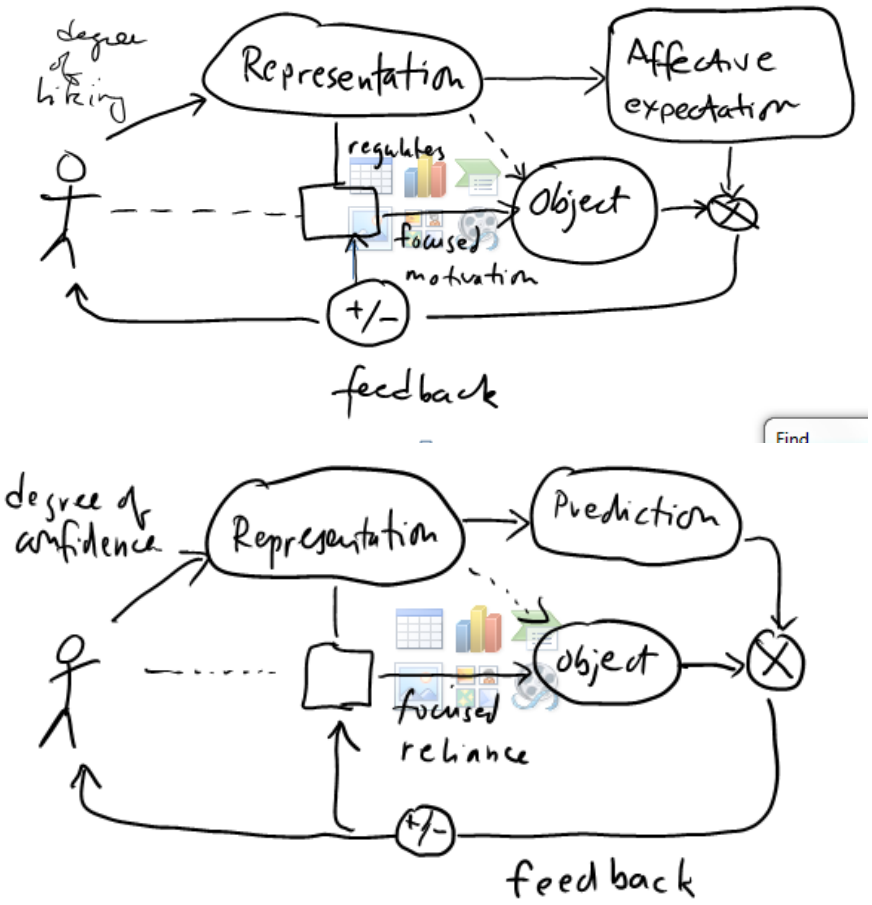

The broader mental economy of affect—these all contribute valence and magnitude—common path

- *Desire*
  - *Liking*
  - *Affection*
  - *Care*
  - *Interest*
  - *Attraction*
  - *Admiring*
  - *Disliking*
  - *Hating*
  - *Indifference*
- *Belief*
  - *Confidence*
  - *Trust*
  - *Assurance*
  - *Conviction*
  - *Hunch*
  - *Certainty*
  - *Uncertainty*
  - *Doubt*
  - *Suspicion*

Desire and belief: spontaneously *projective*

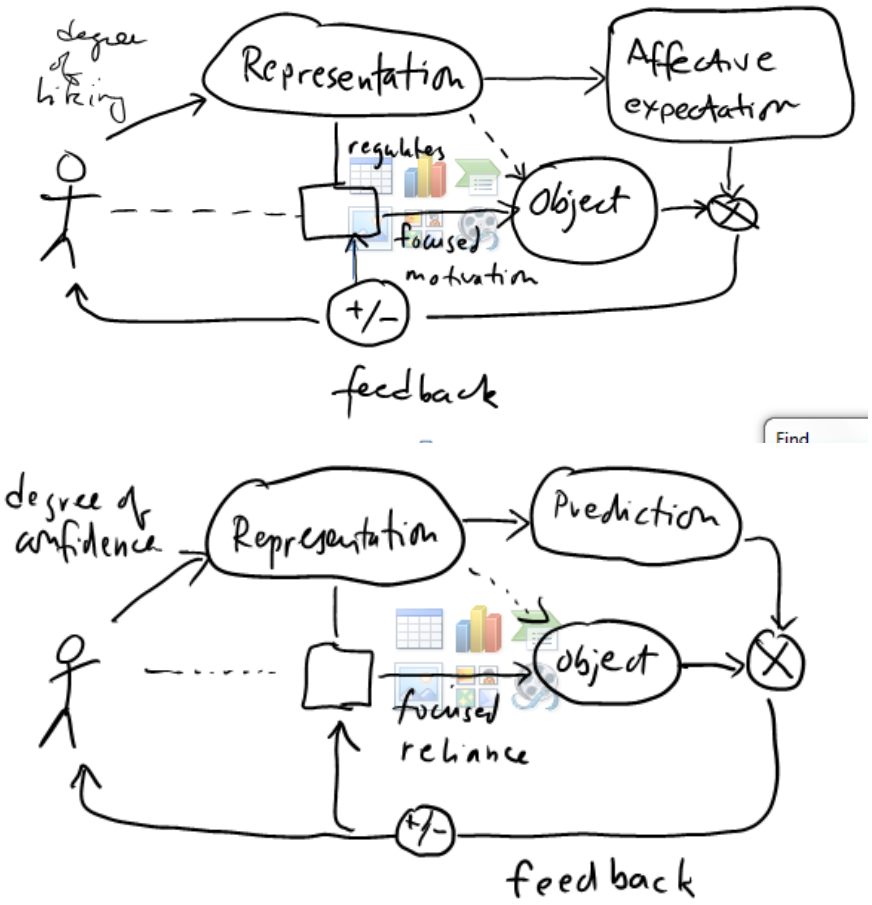

Desire and belief: spontaneous experience sensitivity

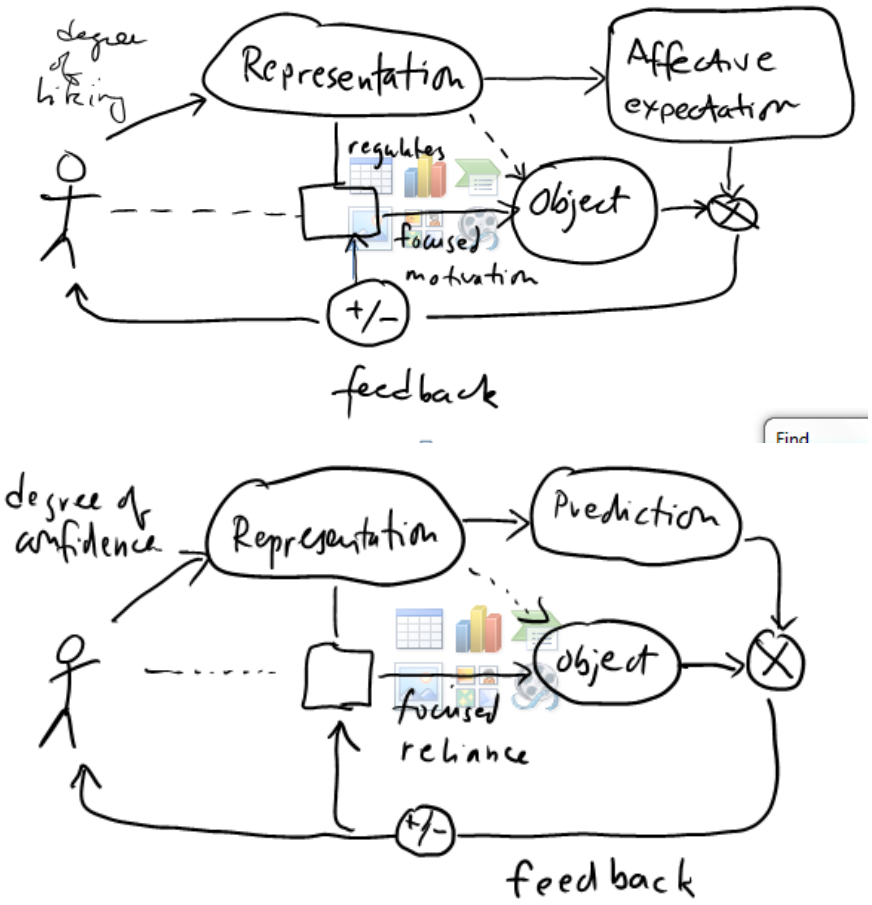

Desire and belief: non-voluntary and resistant to mere instrumentalization—so is intention

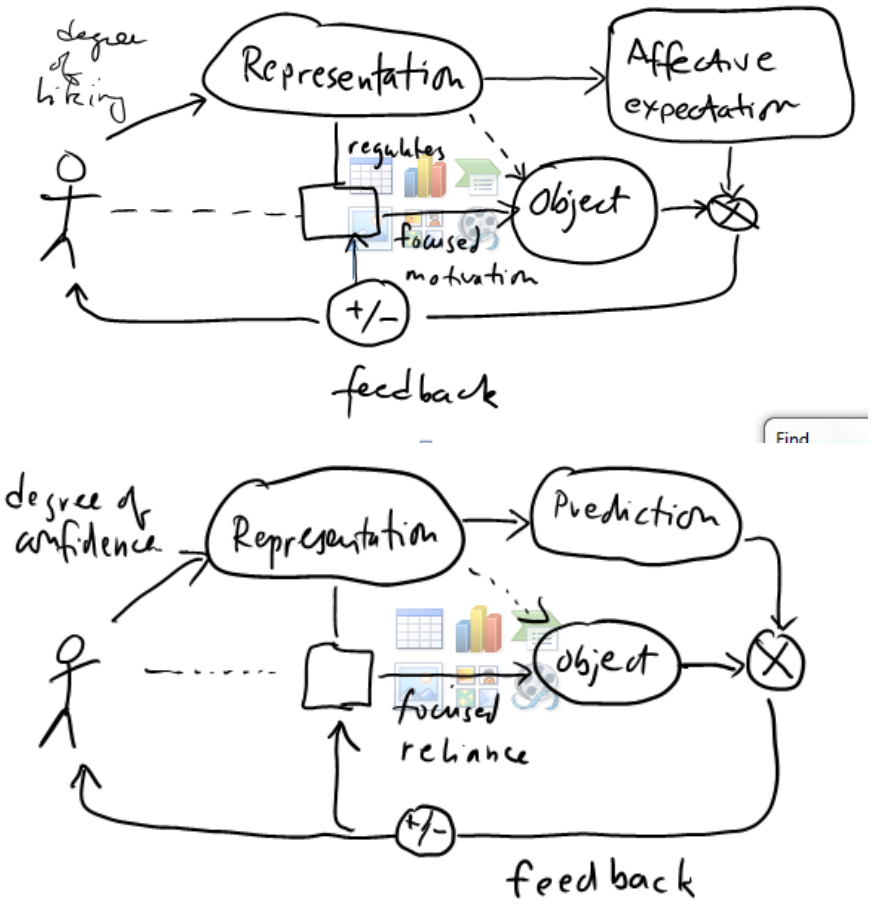

Toxin puzzle

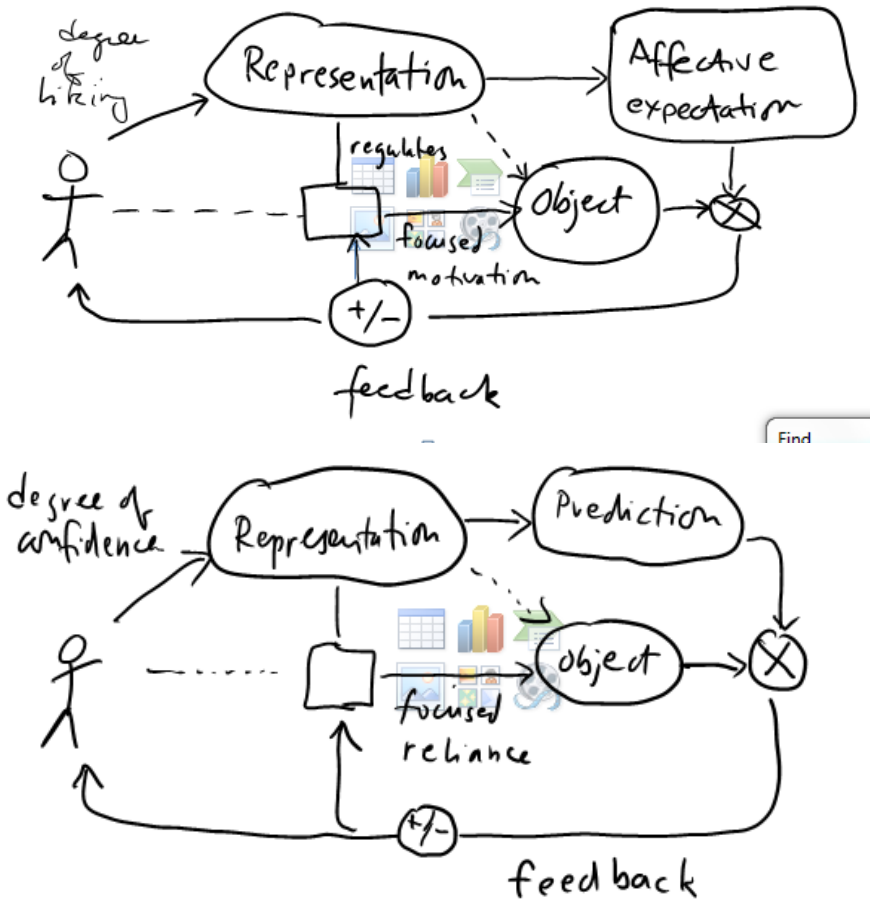

Desire and belief: strong mind-to-world direction of fit

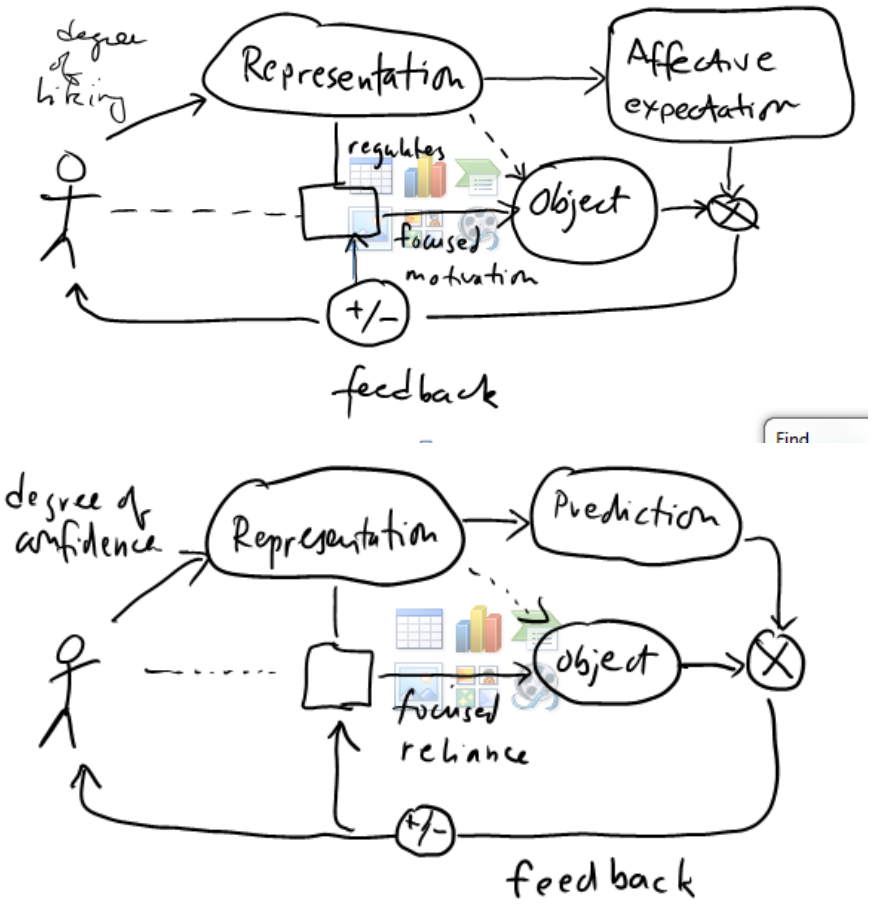

Simple vs. strong mind-to-world fit

- For example: Recall that we could not appeal to such "simple" dimensions of mind-to-world fit as *representational character* or *truth-aptness* to explain how *belief that p* is different from *supposition that p* or *imagining that p*.
  - *Supposition* or *imagining that p* are not out of place, or other than they should be, even in the face of conclusive evidence that *not-p* or perceptual experience as of *not-p*. These do not give us reasons not to suppose or imagine that *p*.
  - That is not so with *belief that p*, which tends, appropriately, to be weakened by evidence that *not-p* or perceptual experience as of *not-p* is a reason to reject it.

Fitting attitudes

- For example, In Lecture 2, we claimed that beliefs have *strong* mind-to-world fit—the inherent dynamic of belief is to be spontaneously responsive to evidence of fit with the world—responsive *by the very state it is*, without requiring *agents* to hold their beliefs to this. But also orients to world practically—*reliance* as an element of practical knowledge.

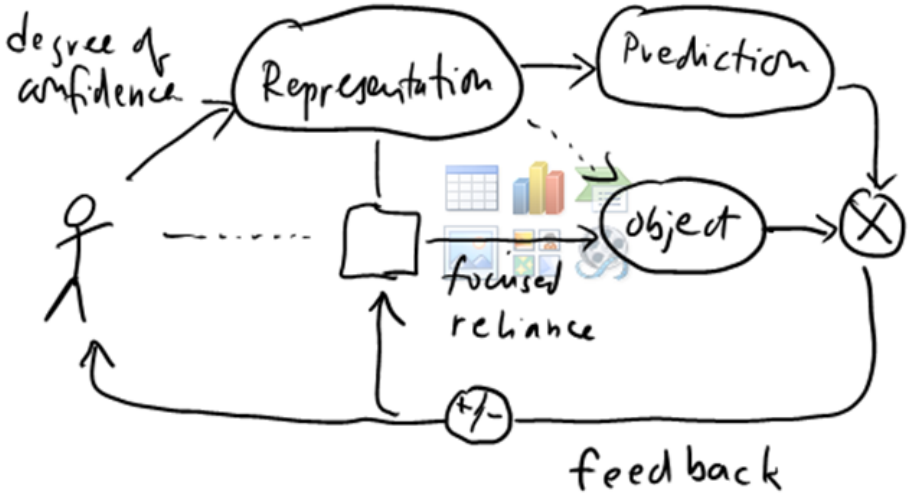

And desire?

- In Lecture 1, we argued that desire is a fitting attitude toward value, with *strong* mind-to-world fit. But what does desire hold us to? Consider the two components of desire, one belonging to the family of positive affective appraisal, the other to the family of focused motivation and pursuit.
- Value, we argued, "calls for" or "merits" both an appreciative representation and focused effort to bring it about.

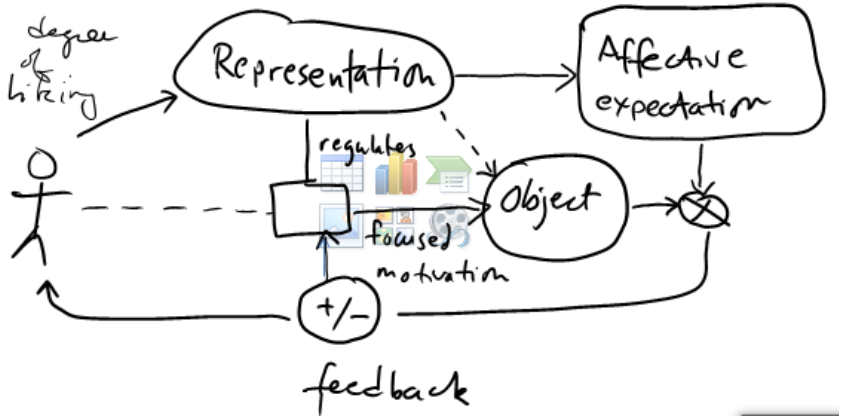

Desire and strong direction of fit

- In Lecture 1, argued that desire captures both attractive, evaluative aspect and effortful mattering. Genuinely *practical* thought—focused attention and motivation, frustration and practical pleasure, sustaining of unsatisfied representation. The football coach.

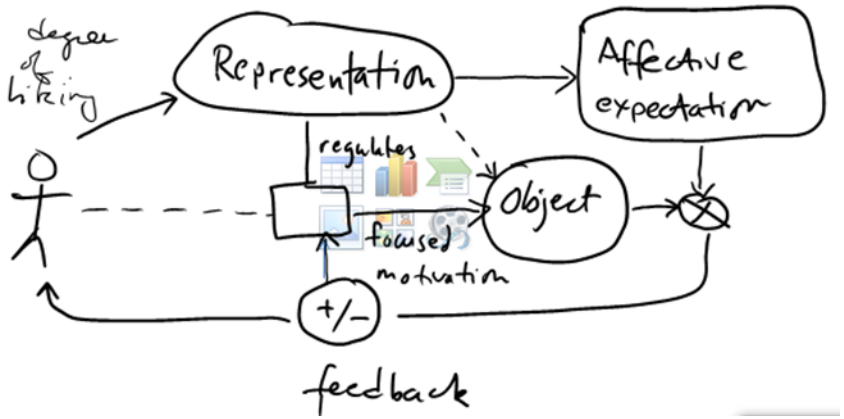

Desire and belief: spontaneous *de se* co-ordination of *active action-tendencies* around shared intentional content

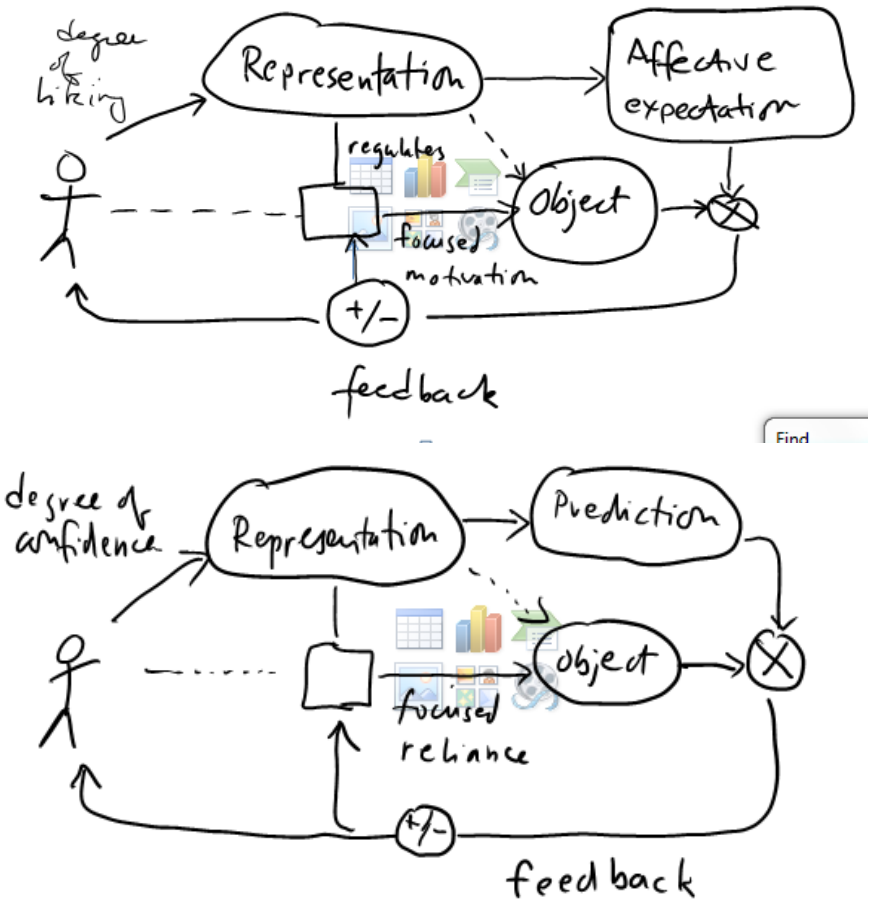

Practical intellect

- Belief and desire function to *orient agents toward the world* and *hold them to an aim*—as non-voluntary states that commit thought- and action-guiding resources and sustains expectation, reliance, and feedback, they have an inherent *practical* dynamic of *tying our representations and actions to the world rather than our will or fancy.*
  - Practical thought can begin in motivation and end in intention (Aristotle, ref.).
  - Practical intellect can anchor us where reason cannot (Hume, ref.)
- Affect makes it matter and makes it work—we *trust and rely upon*, have *confidence and are committed*, *are vindicated or surprised or disappointed*, *frustrated* or *satisfied,* etc.

Desire and belief: … and transfer thought- and action-guiding force among them

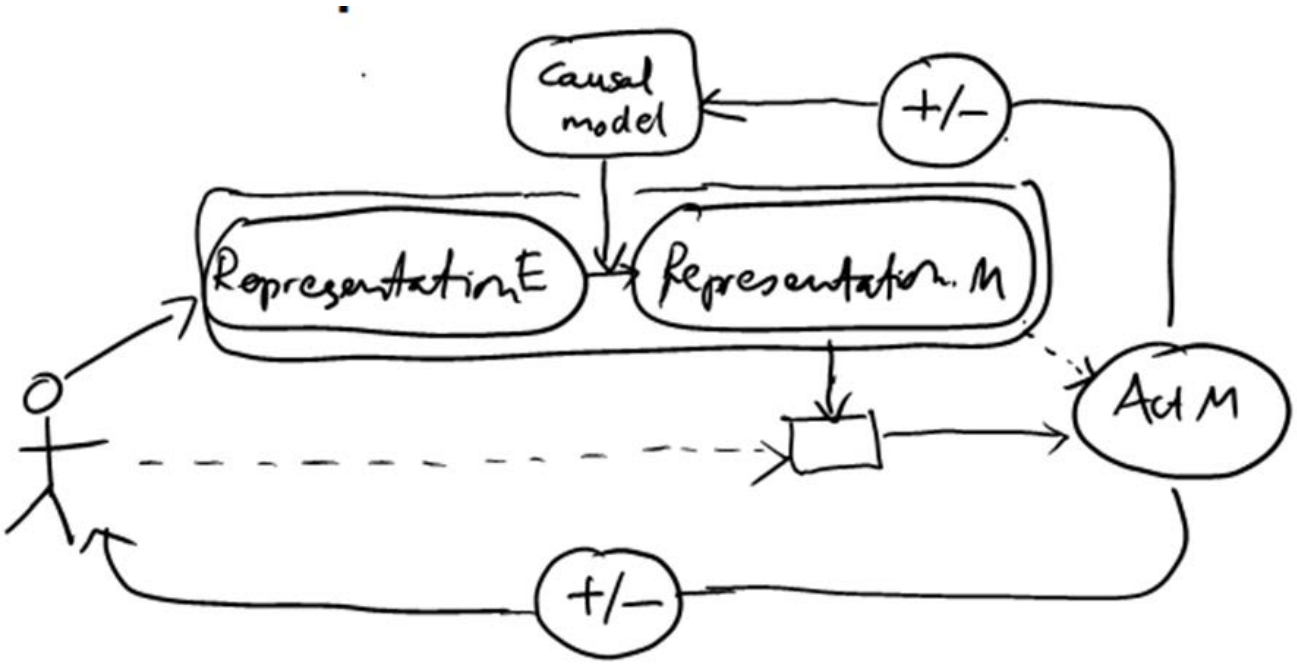

Desire and belief: … and transfer thought- and action-guiding force among them … but not "unthinking"

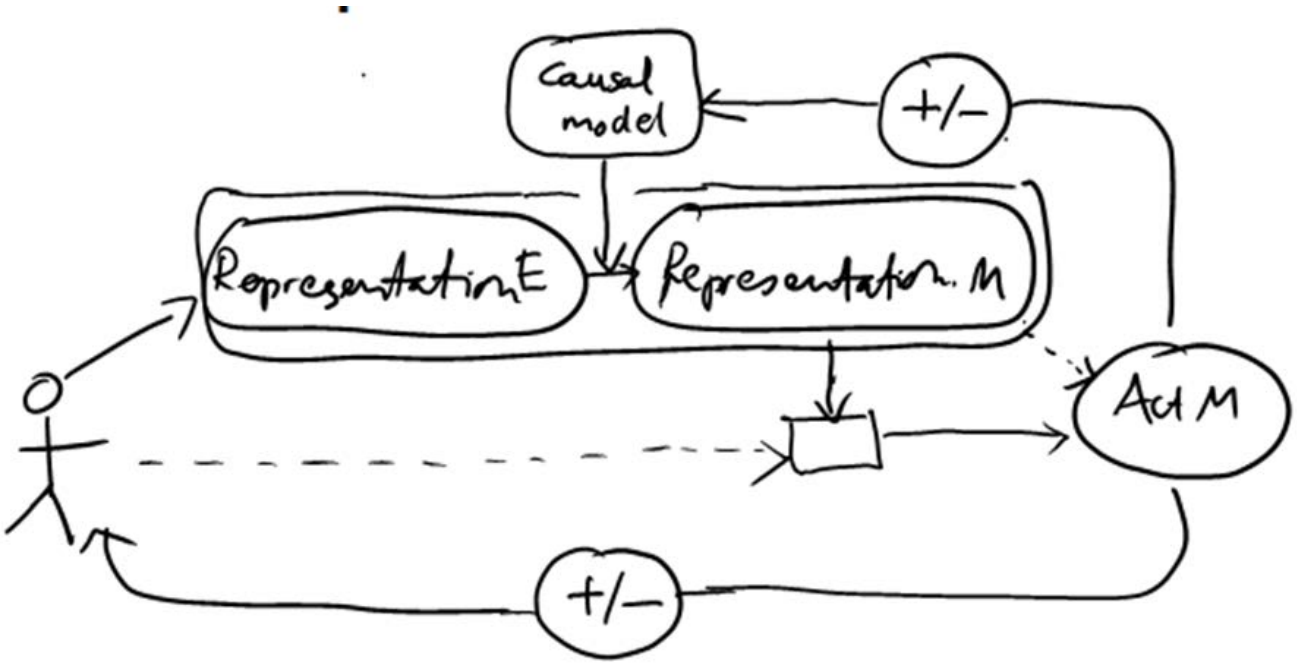

Desire and belief: two kinds of strength

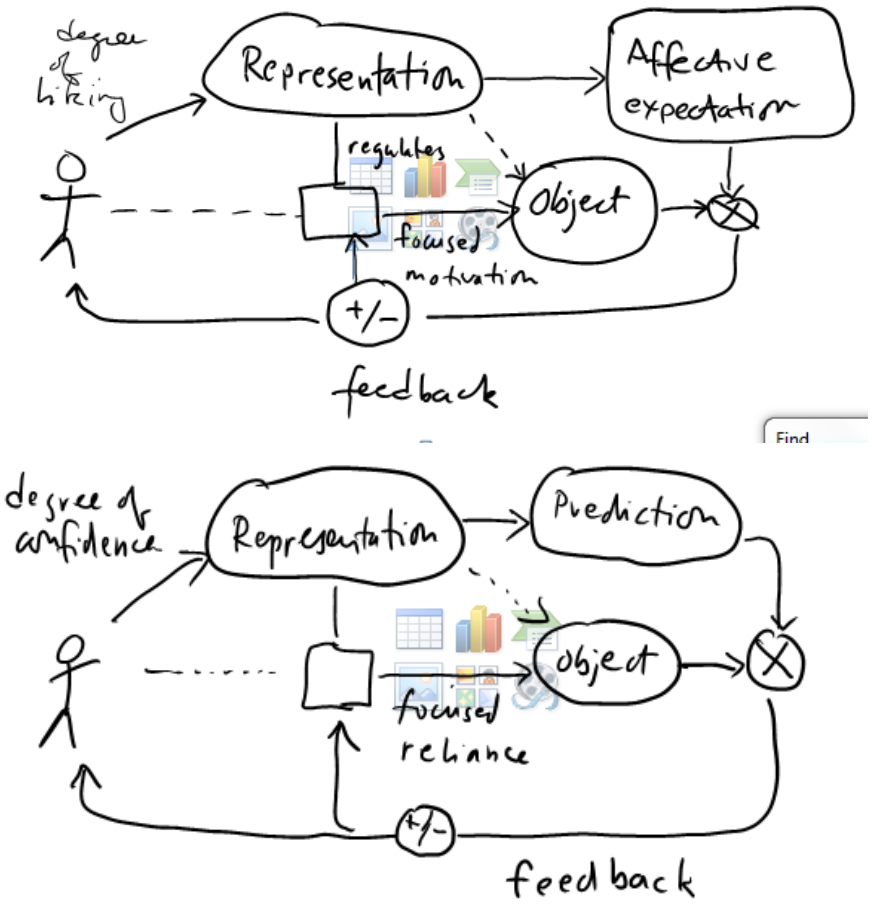

Desire and belief: regulation of motivation and reliance by evaluation—and failures in *akrasia*, *addiction*, and *phobia*

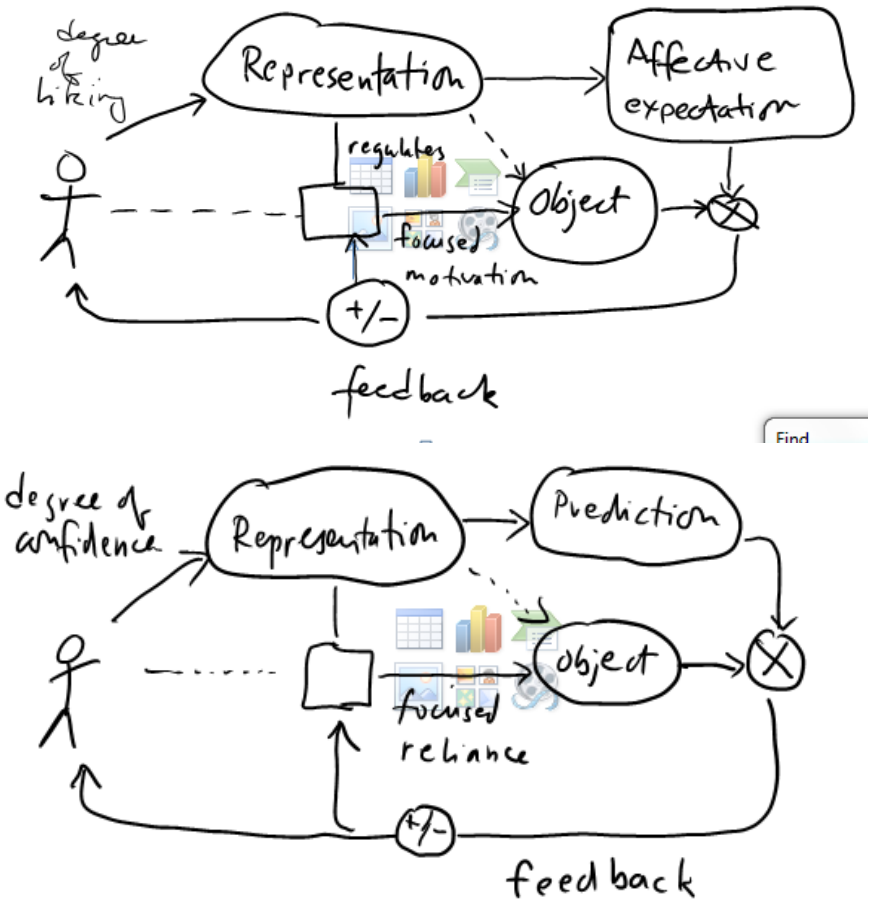

Control and the will

- These models are desire and belief when functioning well—parts can be present without others and we'll likely speak of "irrational desire" or "irrational belief" or … . We need a taxonomy of these that is informative and explanatory—not just odds and sods.

Normal regulation and control

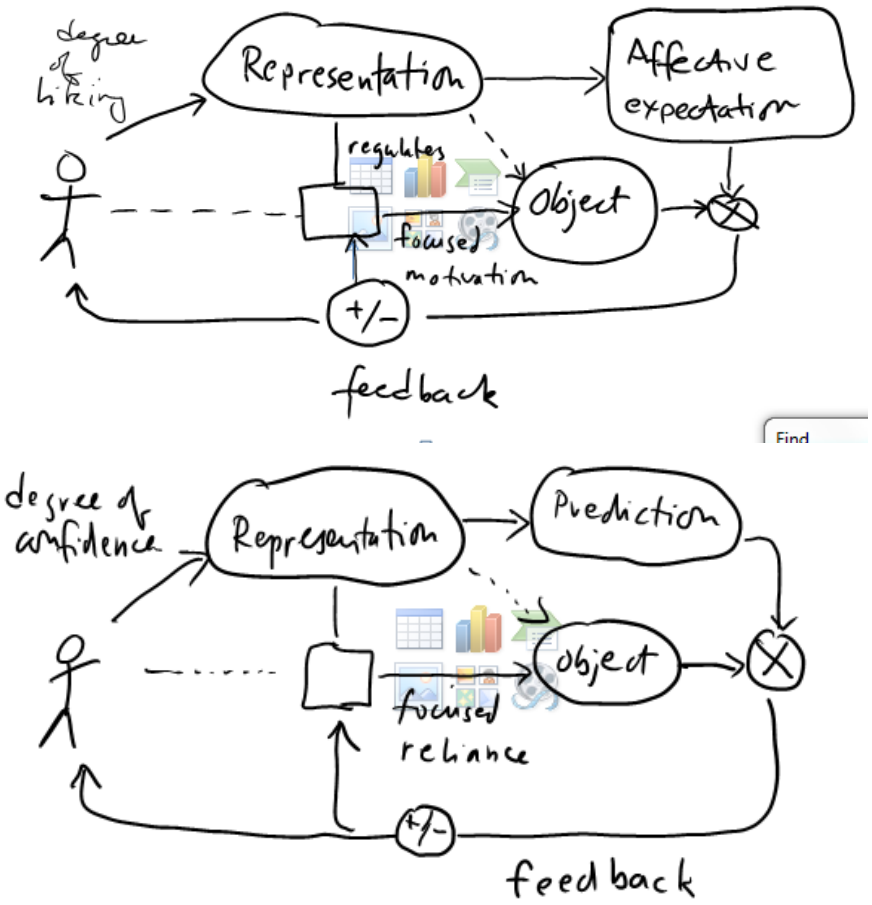

Agency, control, and higher-order desires

(cf. Frankfurt 1971)

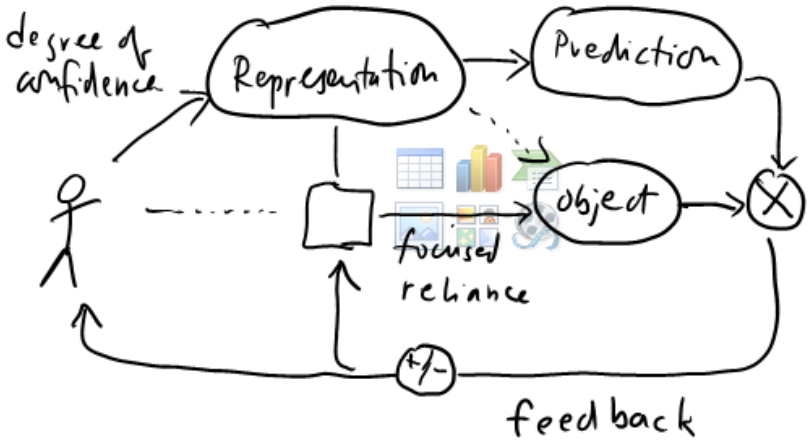

Agency, control, and higher-order evaluation

(cf. Watson 1975)

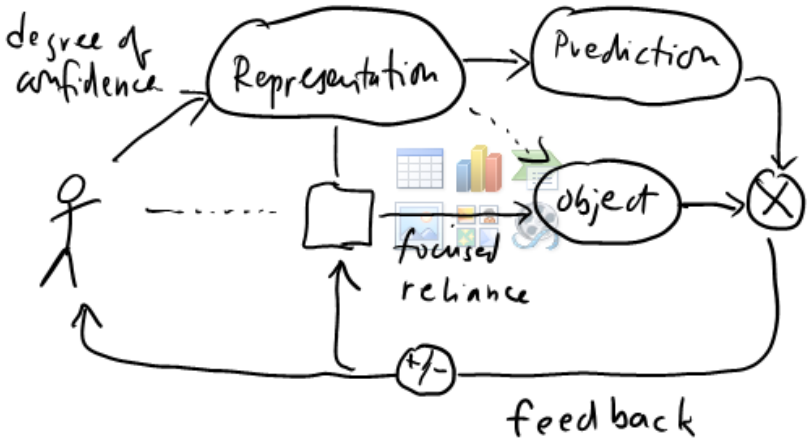

Practical intellect

- Belief and desire function to *orient agents toward the world* and *hold them to an aim*—as non-voluntary states that commit thought- and action-guiding resources and sustains expectation, reliance, and feedback, they have an inherent *practical* dynamic of *tying our representations and actions to the world rather than our will or fancy.*
  - Practical thought can begin in motivation and end in intention (Aristotle, ref.).
  - Practical intellect can anchor us where reason cannot (Hume, ref.)
- Affect makes it matter and makes it work—we *trust and rely upon*, have *confidence and are committed, are vindicated or surprised or disappointed, frustrated* or *satisfied,* etc.

The structures needed for *acting intentionally* …

- … thus can be present without requiring that an agent "intervene in the causal process", "stand back", "form an intention", or "fill gaps" necessary for control and monitoring of behavior.
  - That's very good news for *acting on the basis of an intention* as a self-conscious process.
  - It means we aren't asking the self-conscious agent to do what she can't (form intentions by forming intentions to form intentions …),
  - … and giving her the *practical intelligence* to be able to do what she can (form intentions *intentionally and intelligently*, but *without forming an intention to do so*).

The structures needed for inference …

- … likewise can be in place.
  - Namely, a capacity to make content-sensitive, situationally-relevant, directed transitions in thought, modulated by our degrees of confidence in those thoughts and our own thought process,
  - … without having to *think* all of this,
  - … and use it to *infer* the next step.
- That is, the agent's structure of beliefs and goals can afford a *matrix* within which inference can occur by a sequence of *intelligent non-inferential* steps, which nonetheless are alive to the beliefs and aims of the agent, not blind.

Desire and belief thus impart the kind of …

- … structure, phenomenology, dynamic, intentionality, and integration over time that are thought to be marks of intentional action.
  - While retaining their intelligent character.
  - This is a first element of normativity
  - Fundamental for broad rationality, and for narrow rationality to be a boon.

The regulative model

- … goes beyond the "dispositional" picture of behaviorists.
- Regulation *is* a form of activity, not latency, and it gives to activity control and unity around a represented target or aim.

Desires and beliefs as *default* attitudes with a (generally) thin phenomenology

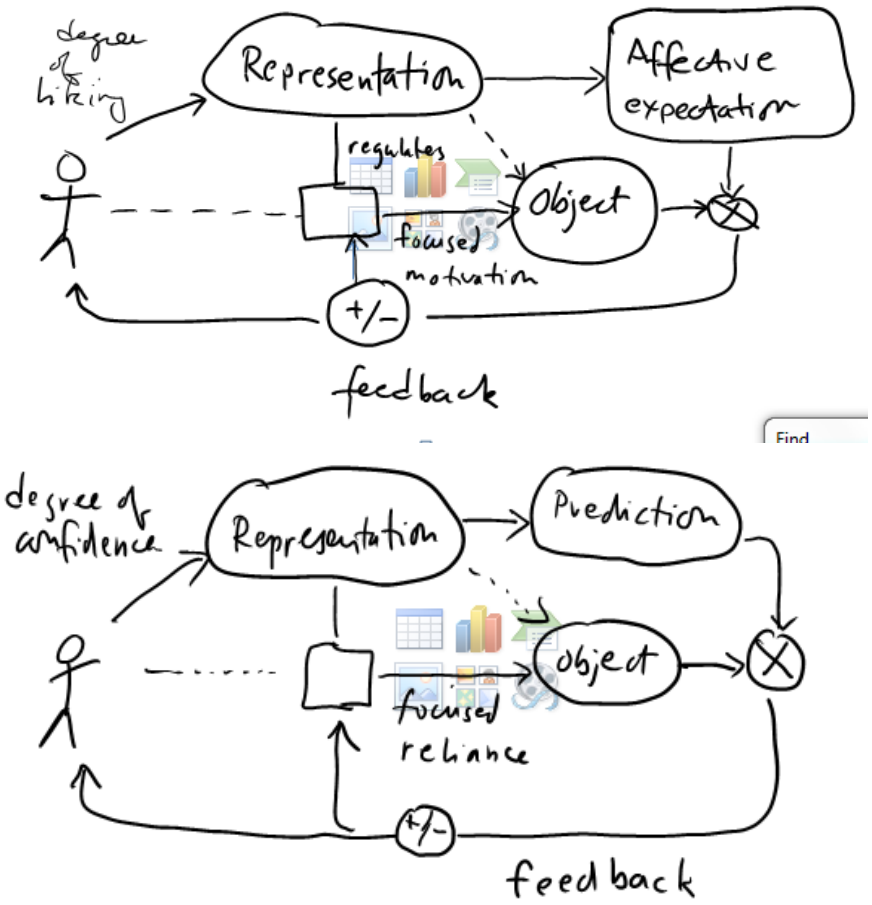

Default affective attitudes

- *Positive* affective attitudes—such as *liking, attraction, trust*, and *confidence*—play a fundamental, default role in learning. These are forms of *approach* affect—e.g., taking information in:
  - A child can get learning underway if she has default trust in her senses and faculties, forming expectations accordingly, and then learning selectively from when these expectations are or are not borne out.
  - A child who who extended no default trust to her senses or faculties, taking nothing in until it is confirmed, would not acquire evidence, or become more discerning.

Default affective attitudes

- Similarly:
  - A child can come to know better what she likes or needs if she has default trust in her initial preferences, forms expectations accordingly, relies upon them in action, and then learns selectively from whether the expectations are borne out.
  - A child who would not rely upon any preferences until she had confirmation that her expectations would be borne out would end up relying upon no preferences, and fail to learn what she likes or needs from experience.
- As *default attitudes*, trust, liking, etc., have a *thin phenomenology*

Of course

- … all manner of fantastical things can be, and have been, believed.
  - But Hume was right about skepticism.
  - And at the level of ordinary experience, where the relation of belief to action, and feedback from outcomes, is clearest, people are able to navigate the everyday world reliably and be remarkably accurate (consider the statistical evidence mentioned last time, e.g., Gallistel *et al*., 2014).

Default affective attitudes

- Infants—human or animal—who are raised in environments where they cannot form trusting relations with the adults around them or confidence in their circumstances suffer profound deficits in intellectual, emotional, and social development, and capacity for self-control (Deaton et *al.*, 2014).

Default affective attitudes

- Infants—human or animal—who are raised in environments where they cannot form trusting relations with the adults around them or confidence in their circumstances suffer profound deficits in intellectual, emotional, and social development, and capacity for self-control (Deaton et *al.*, 2014).
- And political systems—largely human!—that lose reasonable default levels of mutual trust, confidence, or positive concern suffer profound deficits in public intellectual, social, and emotional life, and capacity for democratic self-governance.

With the idea of model-based action …

- … we're on the way to skill.
- And to higher order rationality as an exercise of skill.