Mount Sinai *Center for Computational Psychiatry*

NYCPW - 2025

Tutorial: Bayesian inference

Vincenzo G. Fiore, PhD

Assistant professor
Center for Computational Psychiatry
Icahn School of Medicine at Mount Sinai

Suggested reading

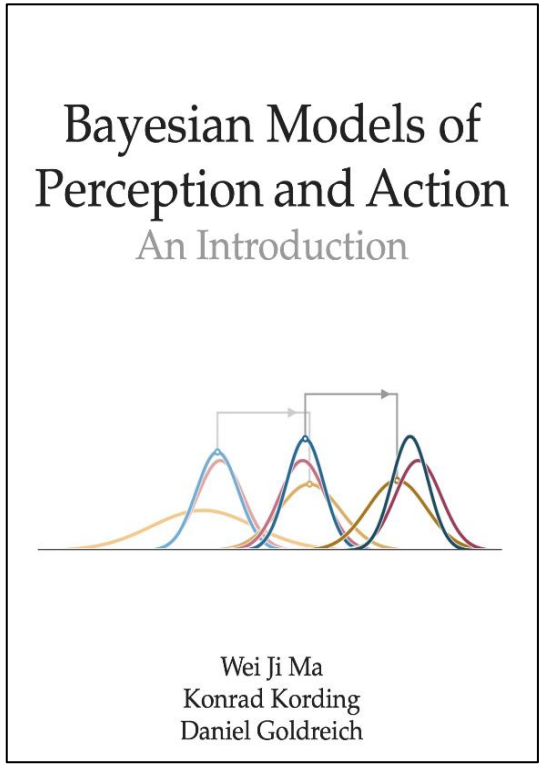

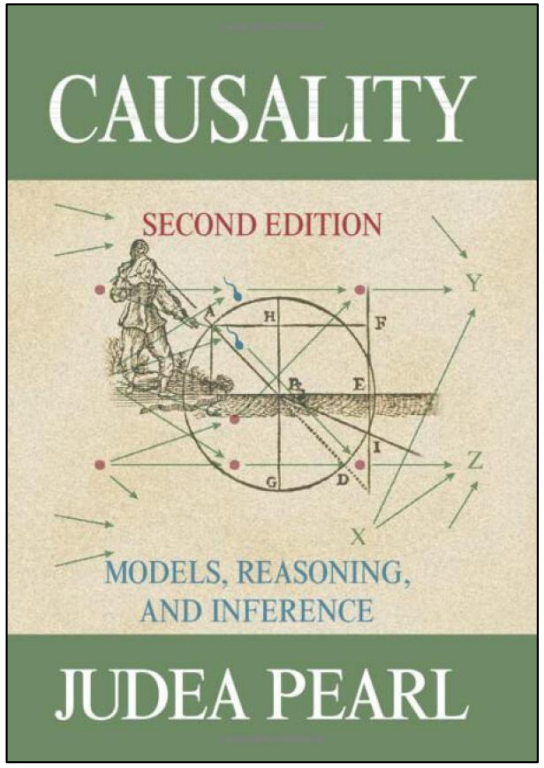

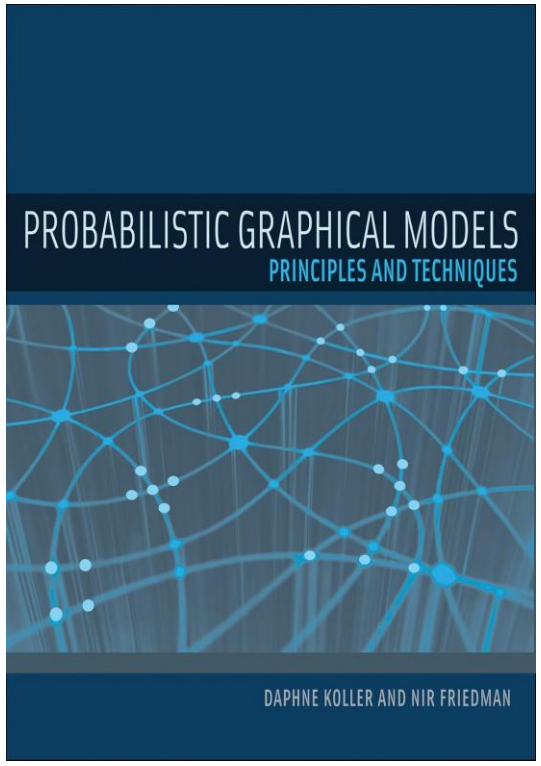

Suggested reading

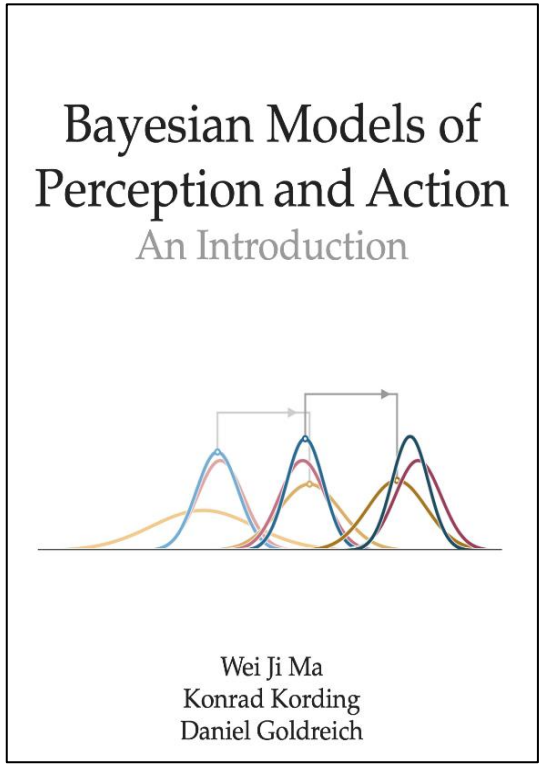

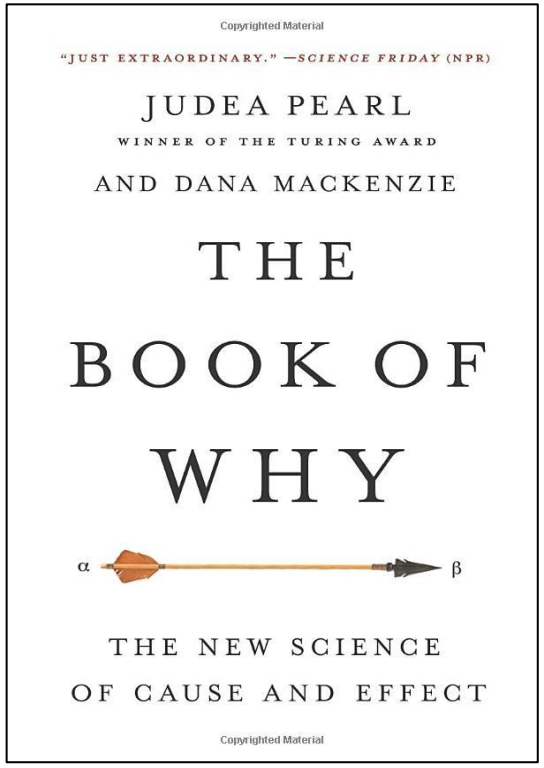

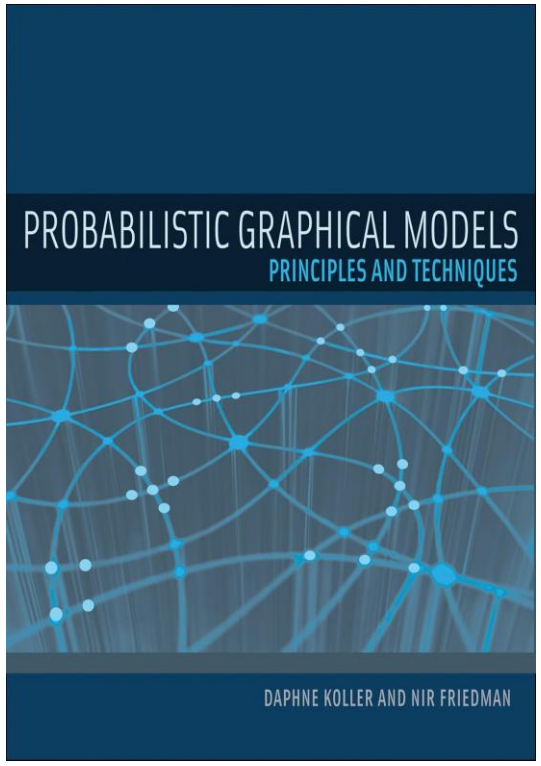

Agenda

- A change of perspective: reward prediction errors and errors of prediction without rewards.
- Basic elements of Bayesian inference (A gentle introduction)
- Phenotypes and hierarchies.

Next

- A change of perspective: reward prediction errors and errors of prediction without rewards.
    1. Beliefs beyond values.
    2. Competing theories.
    3. Dopamine signals of reward prediction error (or maybe not).
- Basic elements of Bayesian inference (A gentle introduction)
- Phenotypes and hierarchies.

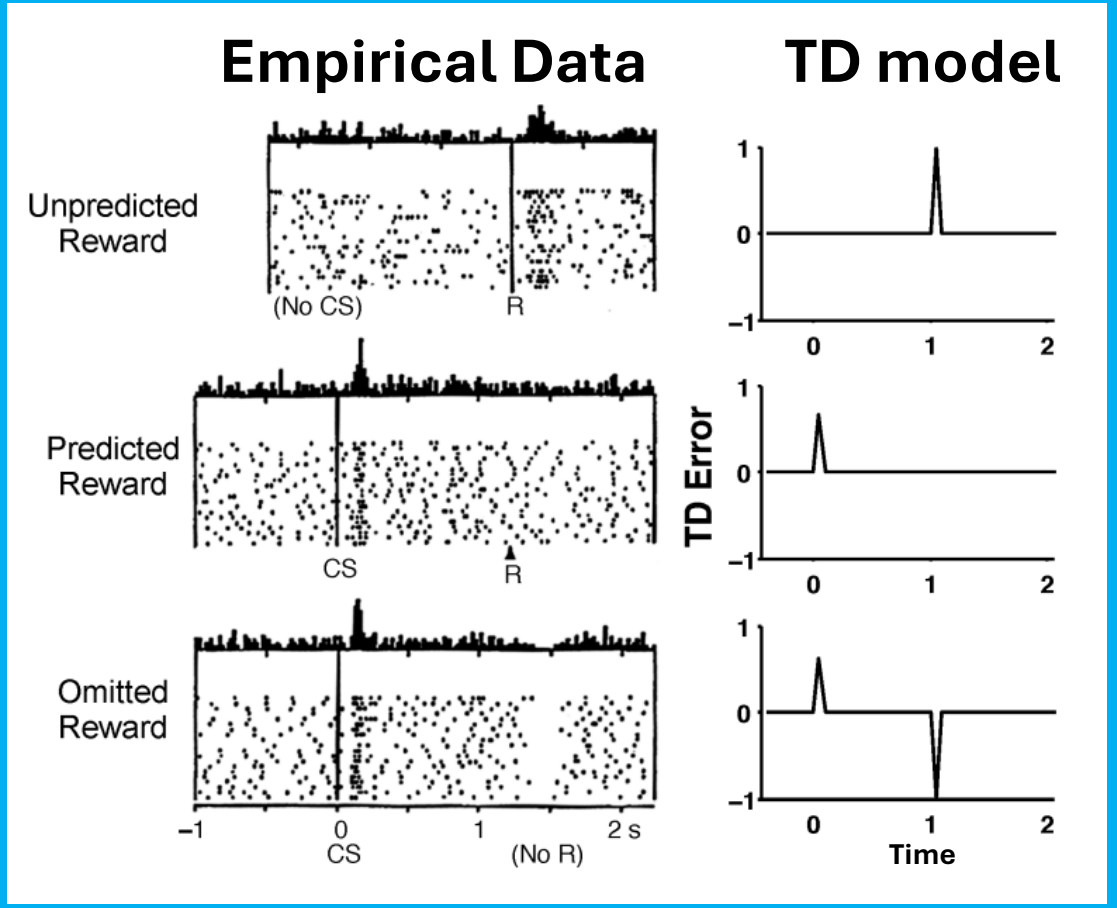

Prior-based Perception

Perception is informed by existing beliefs.

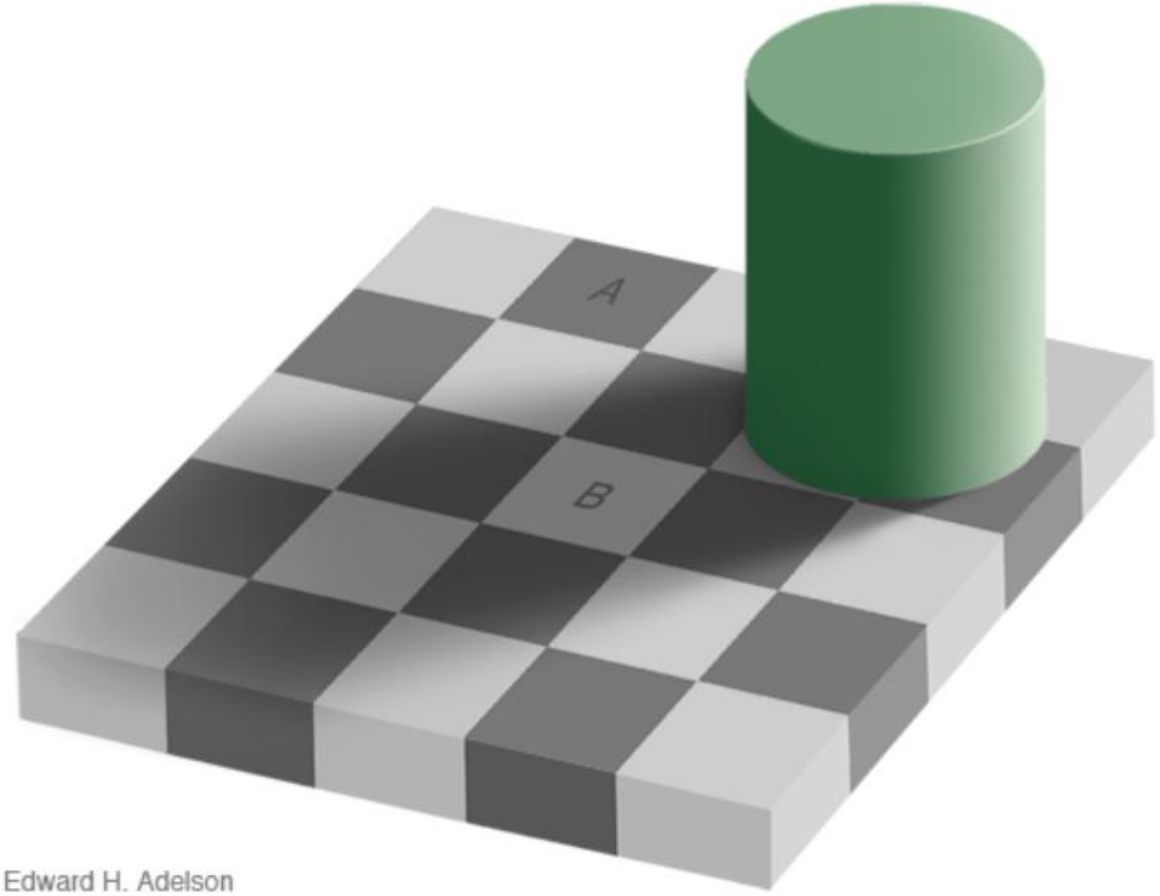

Edward H. Adelson

Prior-based Perception

Perception is informed by existing beliefs.

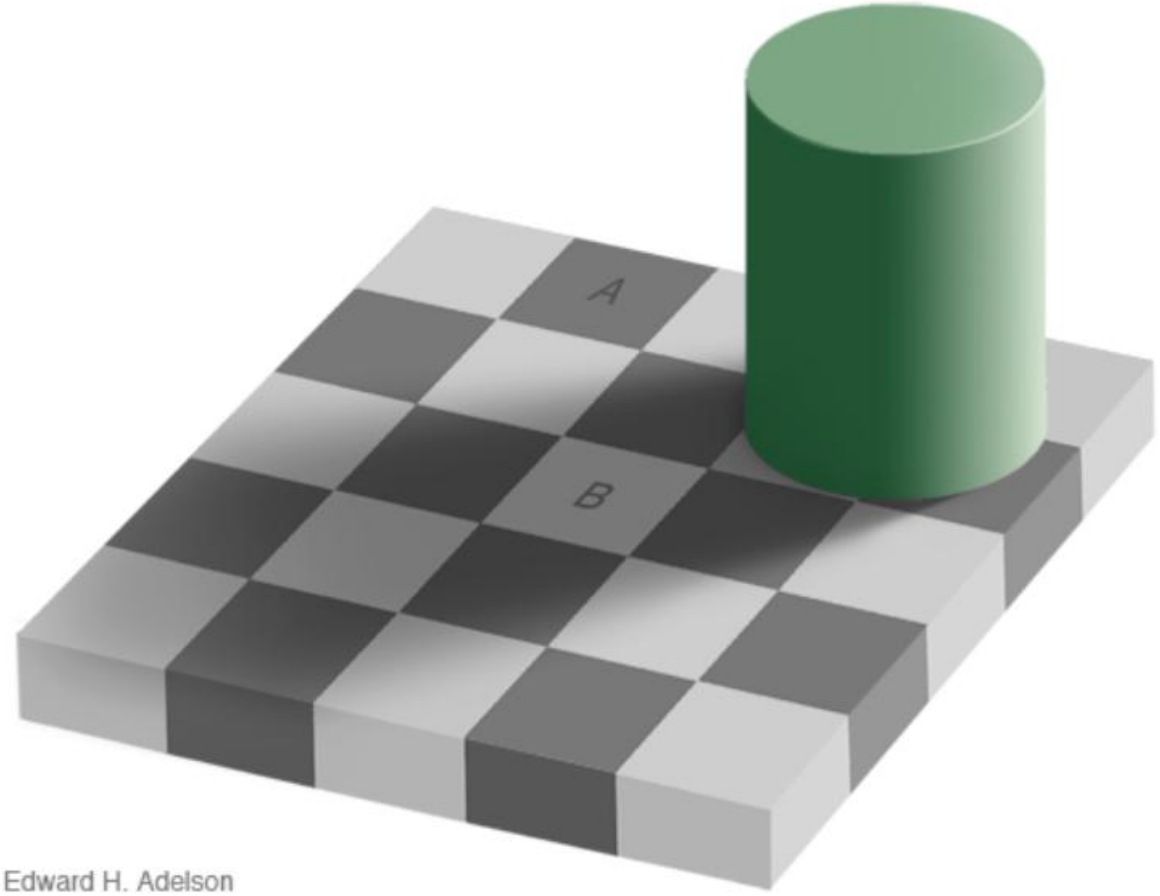

Edward H. Adelson

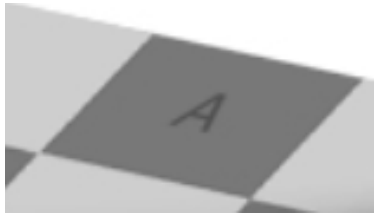

Prior-based Perception

Perception is informed by existing beliefs.

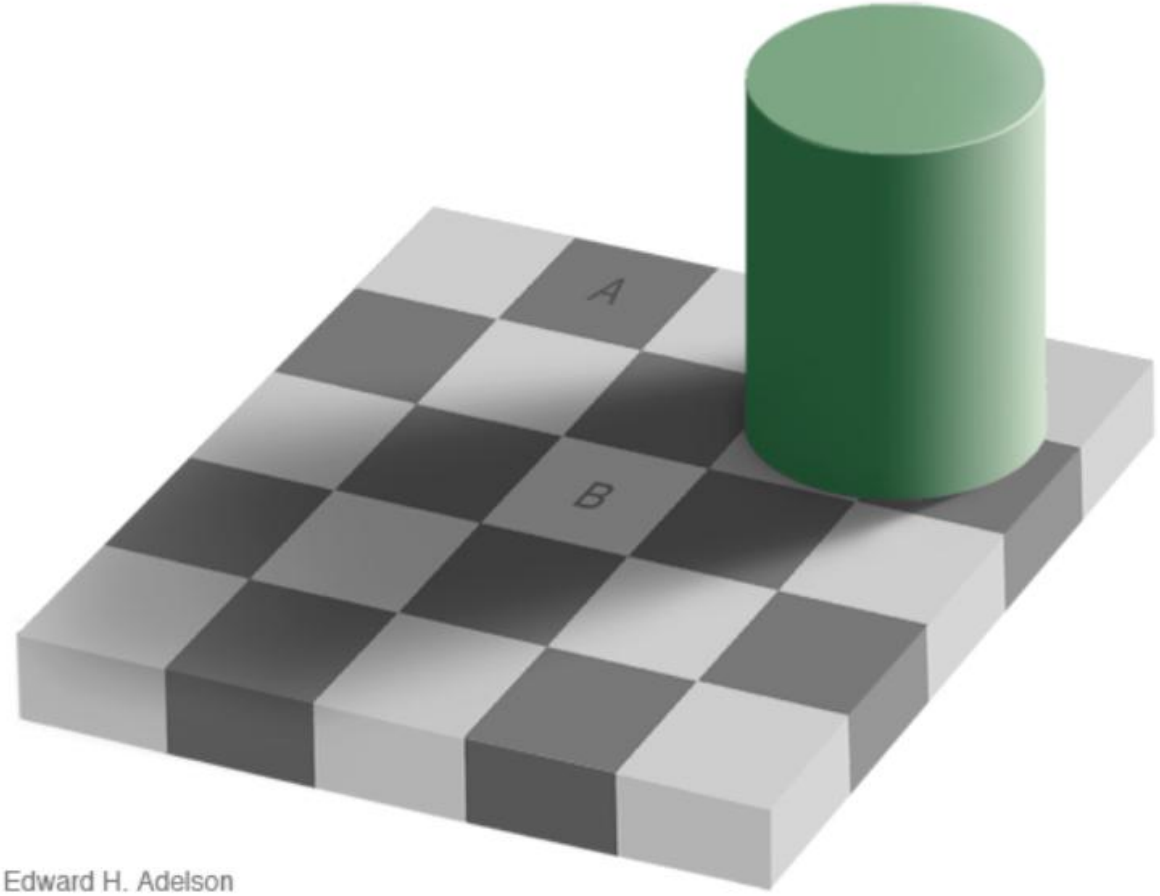

Edward H. Adelson

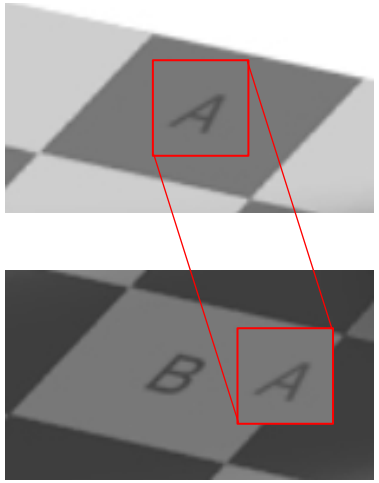

Prior-based Behavior Interpretation

Crow helps hedgehog to cross the street

Prior-based Behavior Interpretation

Crow pecks hedgehog trying to eat it

Prior-based Beliefs (i.e., priors)


The New York Times.

LATE CITY EDITION

GOVERNOR VOTED DOWN ON LIQUOR IN A MAJOR LEGISLATIVE SETBACK; SAYS HE WILL CALL SPECIAL SESSION

Lawmakers to Reconvene On Reform Bills April 15

37 Who Saw Murder Didn't Call the Police

In 1964, Kitty Genovese, a 28-year-old woman, is killed outside her apartment building in New York City. Two weeks after the murder, The New York Times published an article reporting that 37 witnesses saw or heard the attack, and that *none of them called the police or came to her aid*.

The behavior was explained with the ***bystander effect***.

Prior-based Beliefs (i.e., priors)


The New York Times.

LATE CITY EDITION

GOVERNOR VOTED DOWN ON LIQUOR IN A MAJOR LEGISLATIVE SETBACK; SAYS HE WILL CALL SPECIAL SESSION

Lawmakers to Reconvene On Reform Bills April 15

37 Who Saw Murder Didn't Call the Police

In 1964, Kitty Genovese, a 28-year-old woman, is killed outside her apartment building in New York City. Two weeks after the murder, The New York Times published an article reporting that 37 witnesses saw or heard the attack, and that *none of them called the police or came to her aid*.

The behavior was explained with the ***bystander effect***.

Was the story real? Is the effect real in experimental settings?
Is it ecologically validated?

What kind of information would change your mind?

Models of the world

- **Prior (i.e., pre-existing) beliefs:**

NYC is an unforgiving and hostile place.

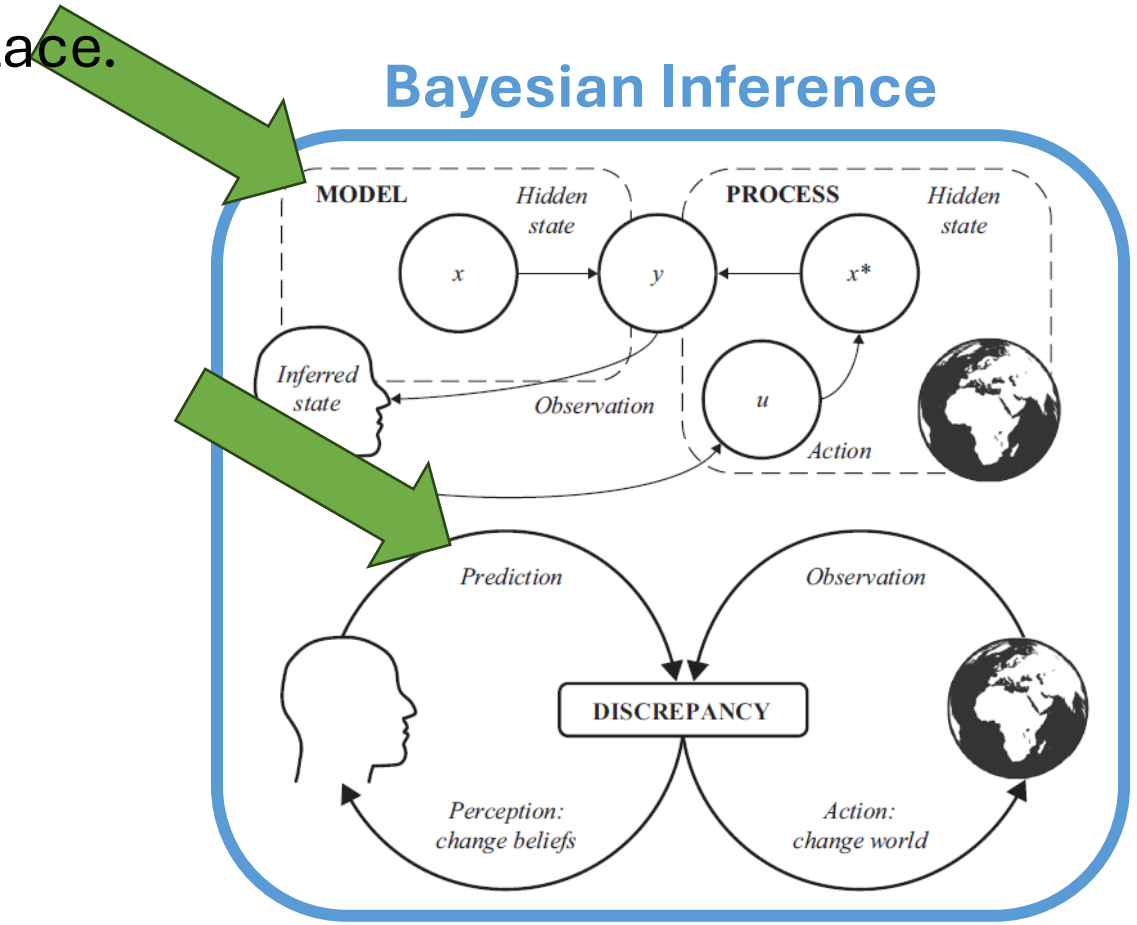

Models of the world

- **Prior (i.e., pre-existing) beliefs:**

NYC is an unforgiving and hostile place.

- **Observation:**

Trusted sources report numbness to violence.

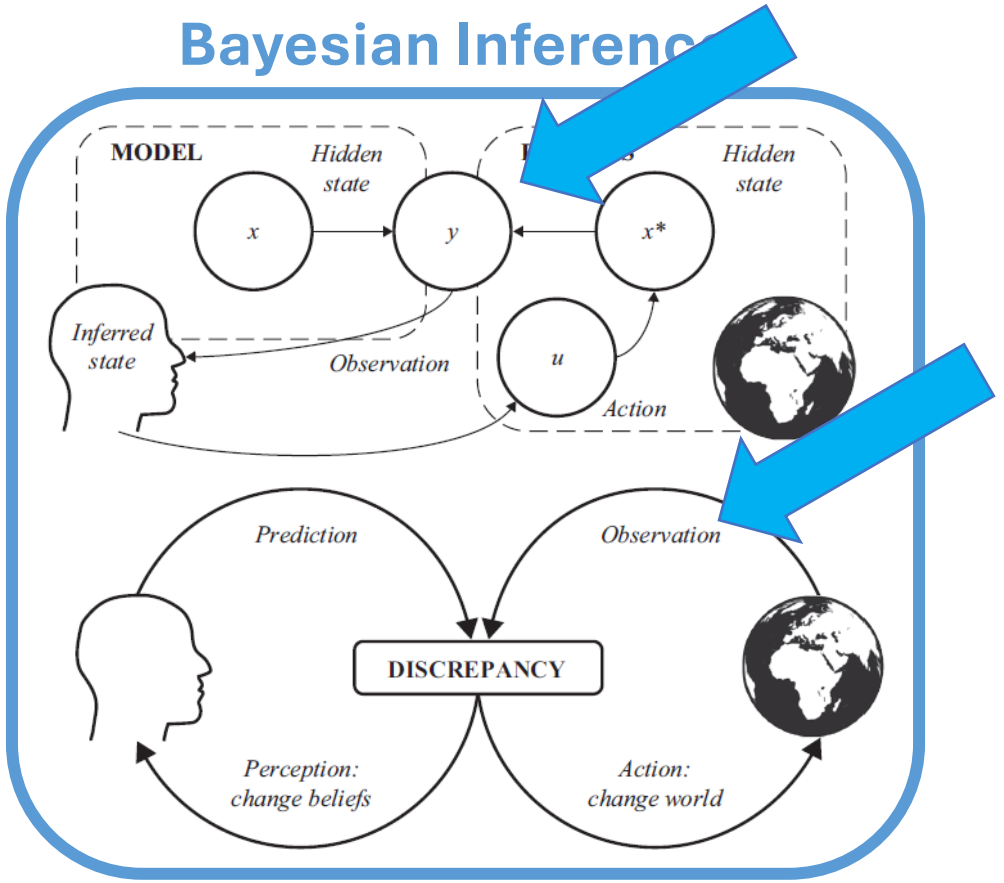

Models of the world

- **Prior (i.e., pre-existing) beliefs:**

NYC is an unforgiving and hostile place.

- **Observation:**

Trusted sources report numbness to violence inability to help fellow human.

- **Posterior (inferred state):**

NYC is a model for all (you are on your own).

Bayesian Inference

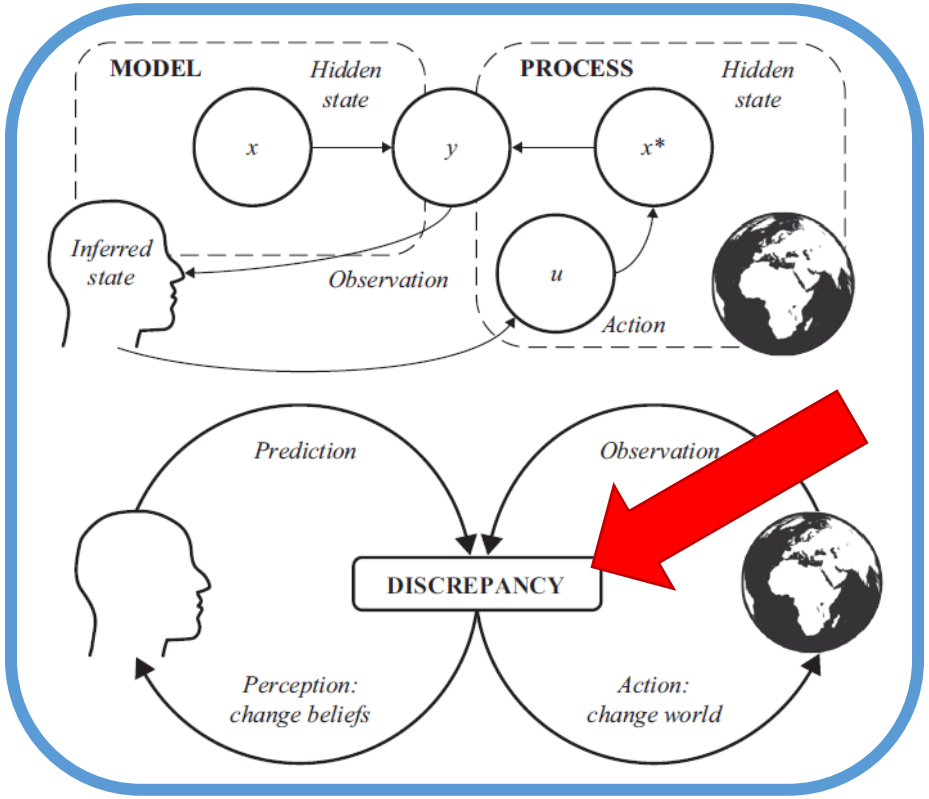

Models of the world

Mismatch between *prior* and *observation* results in *prediction error* and belief updating, generating the new *posterior*.

Bayesian Inference

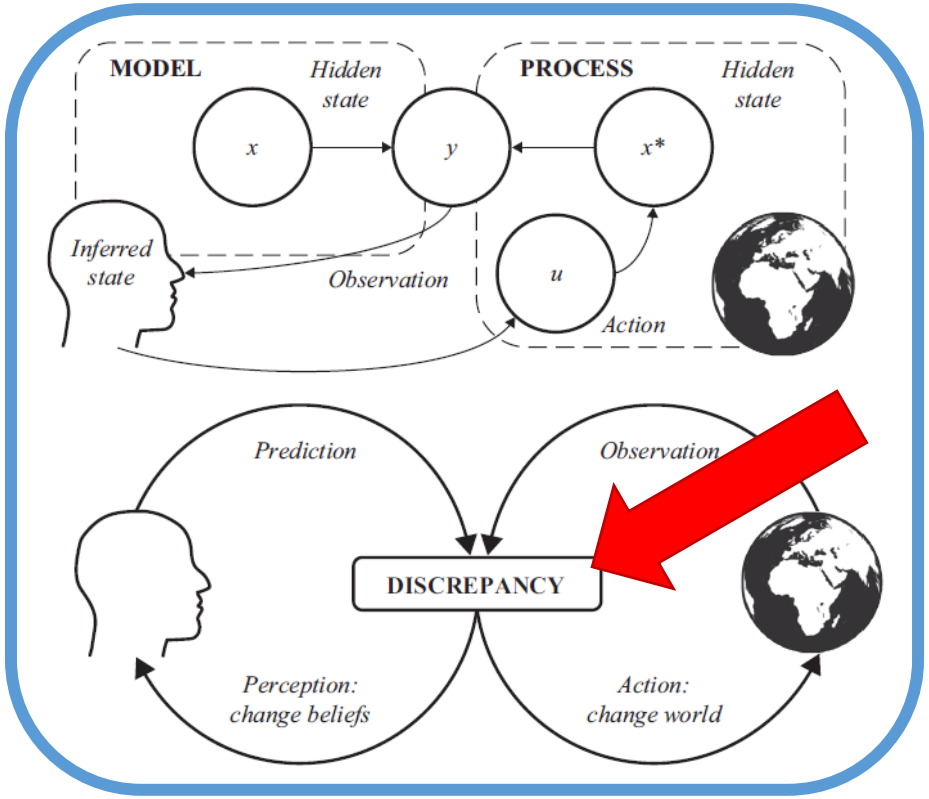

Competing theories

Reinforcement Learning

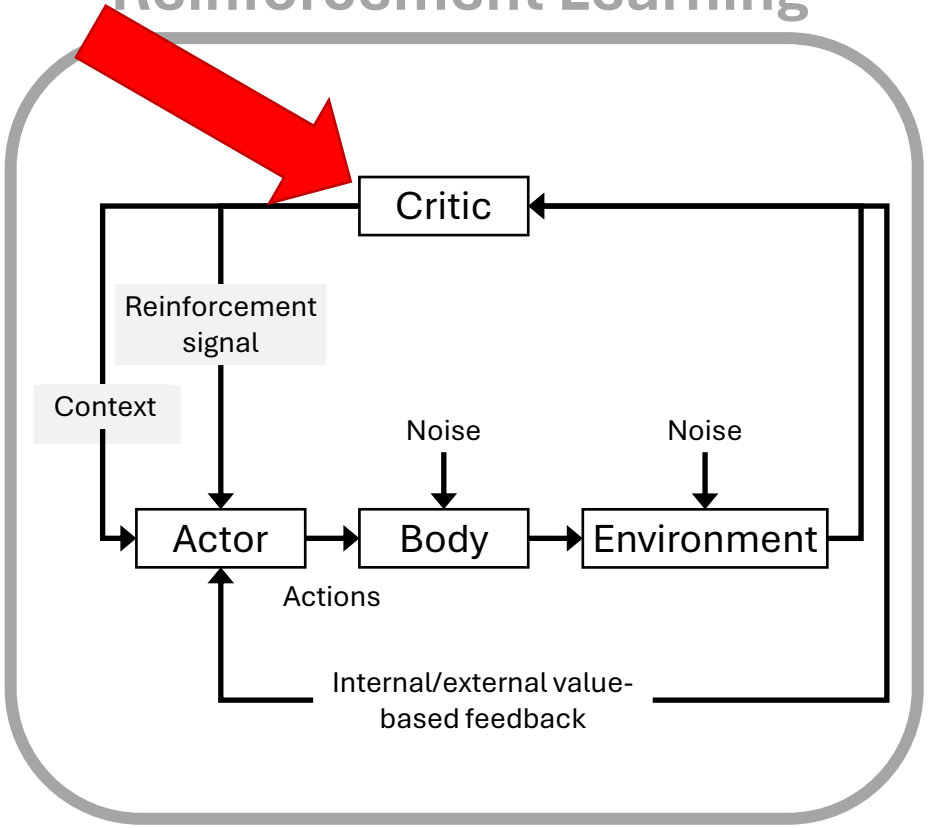

Bayesian Inference

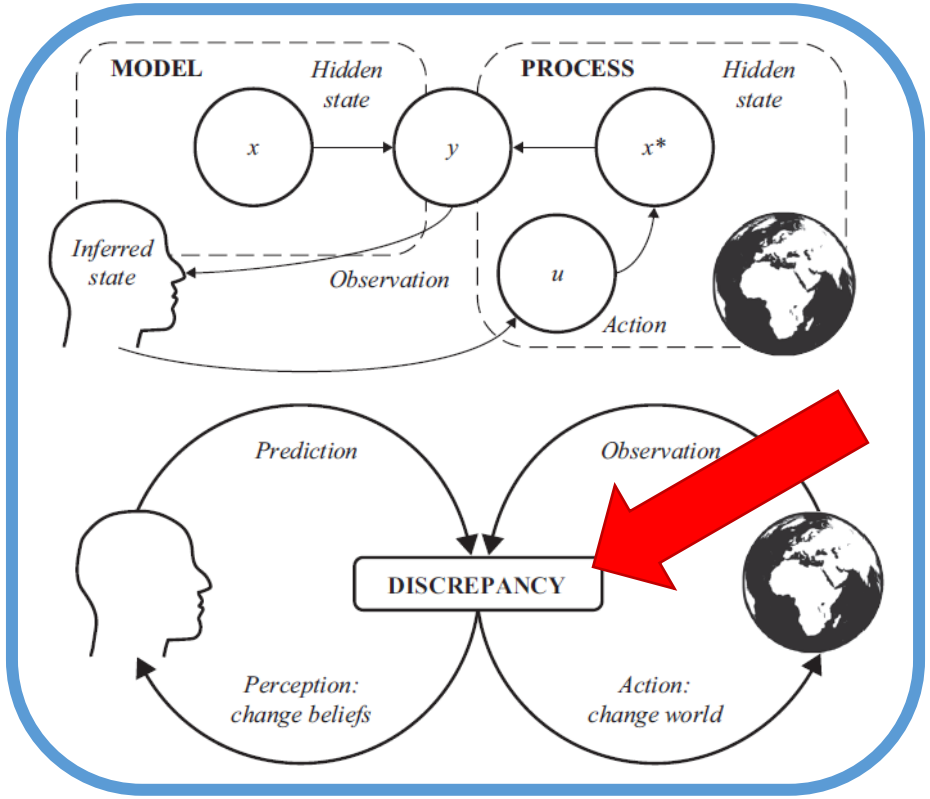

Prediction errors

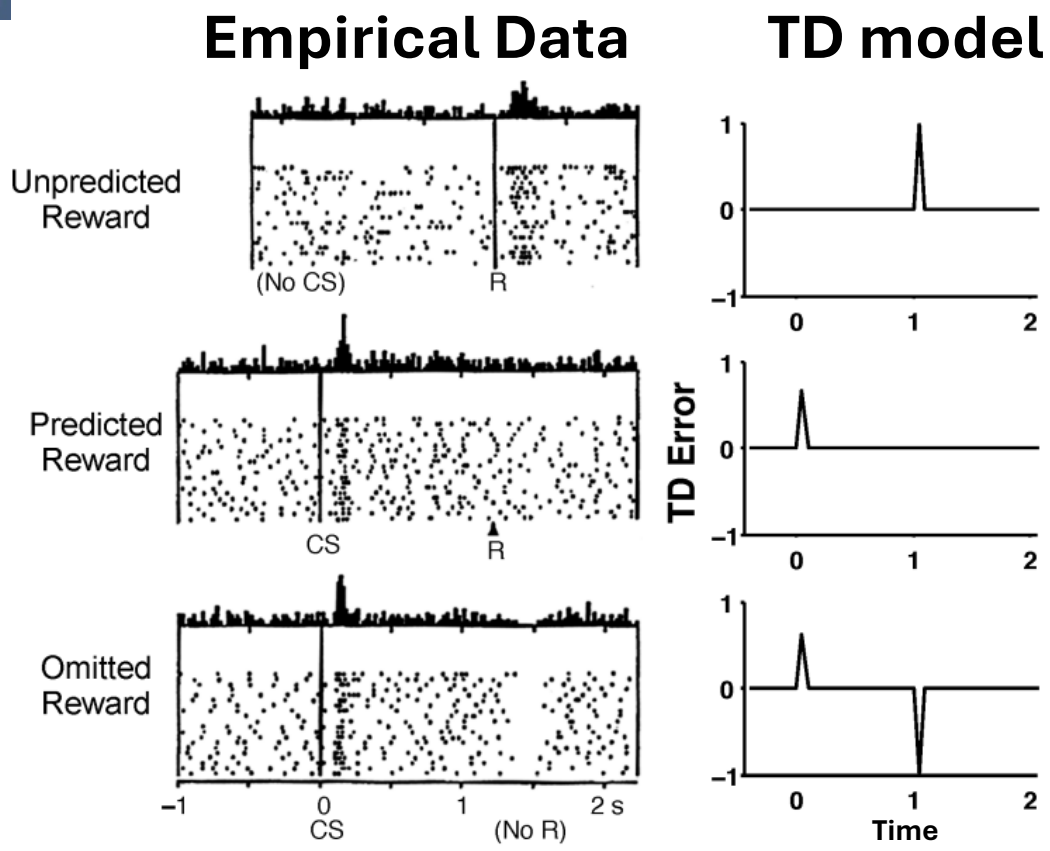

Schultz W, et al. 1997 A neural substrate of prediction and reward. Science. 14;275(5306):1593-9.

Prediction errors

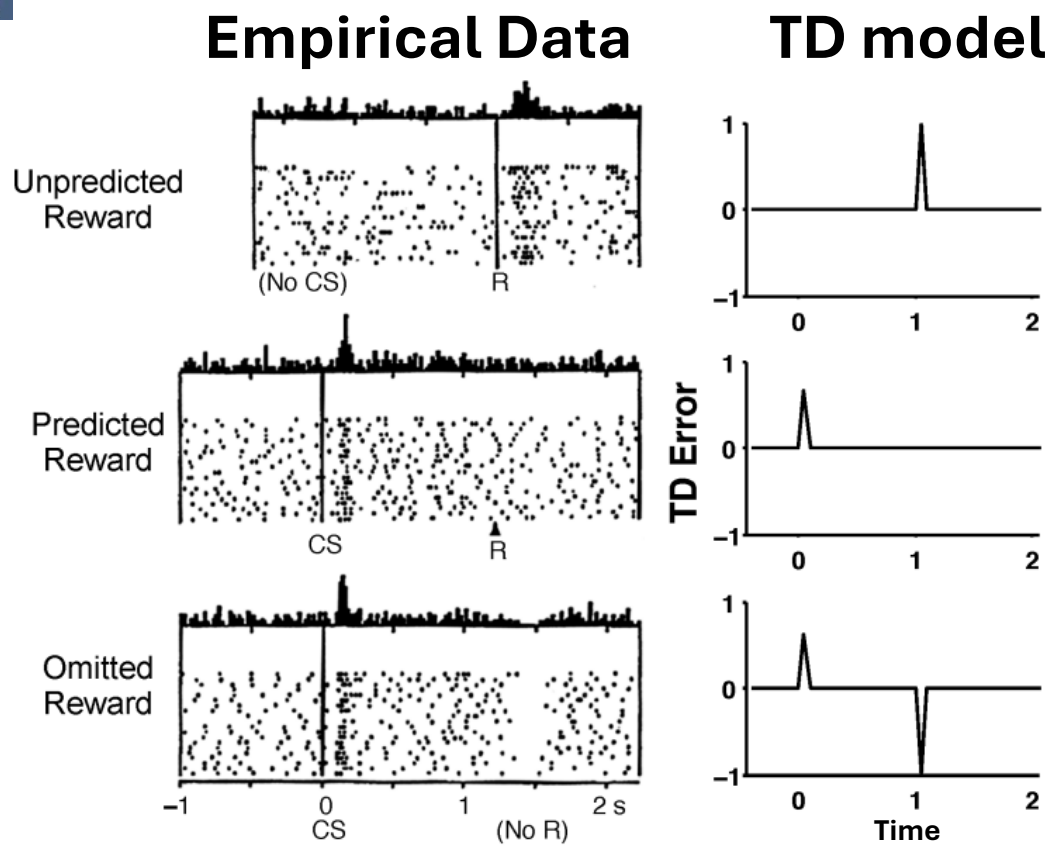

Schultz W, et al. 1997 A neural substrate of prediction and reward. Science. 14;275(5306):1593-9.

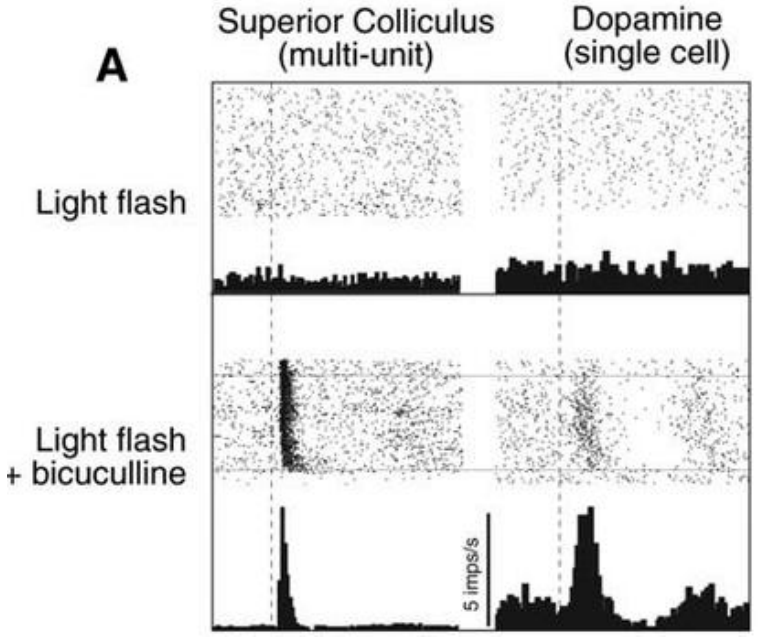

Dommett E, et al. 2005 How visual stimuli activate dopaminergic neurons at short latency. Science. 2005 Mar 4;307(5714):1476-9.

Prediction errors

Is a flash of light rewarding?

Is dopamine (also) responding to salient or novel stimuli?

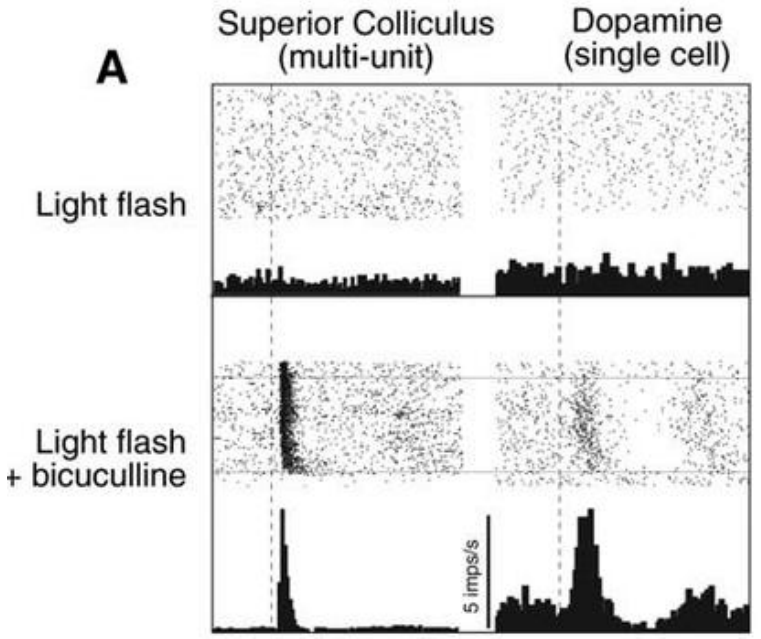

Dommett E, et al. 2005 How visual stimuli activate dopaminergic neurons at short latency. Science. 2005 Mar 4;307(5714):1476-9.

Prediction errors

Neuron
Review

Dopamine in Motivational Control: Rewarding, Aversive, and Alerting

Ethan S. Bromberg-Martin,[1] Masayuki Matsumoto,[1,2] and Okihide Hikosaka[1,*]
[1]Laboratory of Sensorimotor Research, National Eye Institute, National Institutes of Health, Bethesda, Maryland 20892, USA
[2]Primate Research Institute, Kyoto University, Inuyama, Aichi 484-8506, Japan
*Correspondence: oh@lsr.nei.nih.gov
DOI 10.1016/j.neuron.2010.11.022

Midbrain dopamine neurons are well known for their strong responses to rewards and their critical role in positive motivation. It has become increasingly clear, however, that dopamine neurons also transmit signals related to salient but nonrewarding experiences such as aversive and alerting events. Here we review recent advances in understanding the reward and nonreward functions of dopamine. Based on this data, we propose that dopamine neurons come in multiple types that are connected with distinct brain networks and have distinct roles in motivational control. Some dopamine neurons encode motivational value, supporting brain networks for seeking, evaluation, and value learning. Others encode motivational salience, supporting brain networks for orienting, cognition, and general motivation. Both types of dopamine neurons are augmented by an alerting signal involved in rapid detection of potentially important sensory cues. We hypothesize that these dopaminergic pathways for value, salience, and alerting cooperate to support adaptive behavior.

ng to

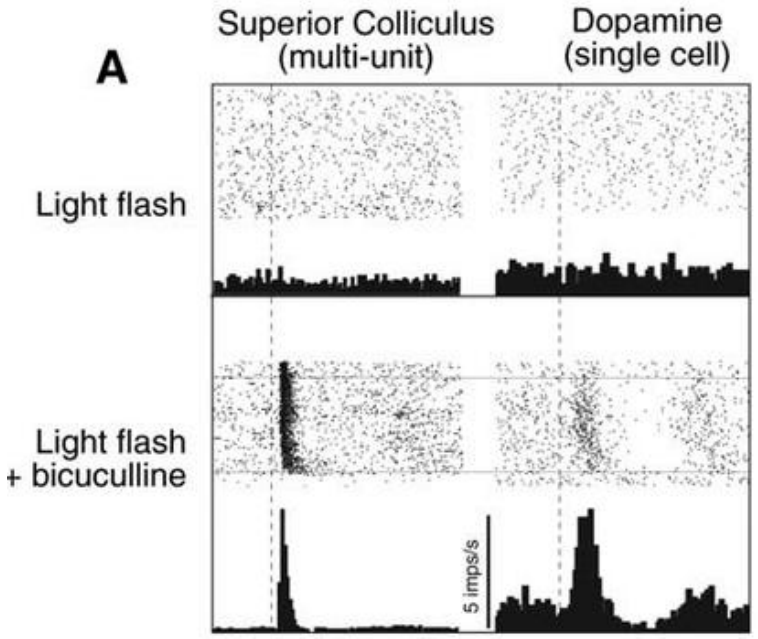

Dommett E, et al. 2005 How visual stimuli activate dopaminergic neurons at short latency. Science. 2005 Mar 4;307(5714):1476-9.

Prediction errors

Neuron
Review

Dopamine in Motivational Control: Rewarding, Aversive, and Alerting

Ethan S.
[1]Laborato
[2]Primate R
*Correspo
DOI 10.10

Midbrai
in posit
signals
recent a
propose
have di
brain ne
brain n
augmen
hypothe
adaptiv

eLIFE
elifesciences.org

RESEARCH ARTICLE

Reward prediction error does not explain movement selectivity in DMS-projecting dopamine neurons

Rachel S Lee, Marcelo G Mattar, Nathan F Parker, Ilana B Witten*, Nathaniel D Daw*

Department of Psychology, Princeton Neuroscience Institute, Princeton University, New Jersey, United States

Abstract Although midbrain dopamine (DA) neurons have been thought to primarily encode reward prediction error (RPE), recent studies have also found movement-related DAergic signals. For example, we recently reported that DA neurons in mice projecting to dorsomedial striatum are modulated by choices contralateral to the recording side. Here, we introduce, and ultimately reject, a candidate resolution for the puzzling RPE vs movement dichotomy, by showing how seemingly movement-related activity might be explained by an action-specific RPE. By considering both choice and RPE on a trial-by-trial basis, we find that DA signals are modulated by contralateral choice in a manner that is distinct from RPE, implying that choice encoding is better explained by movement direction. This fundamental separation between RPE and movement encoding may help shed light on the diversity of functions and dysfunctions of the DA system.
DOI: https://doi.org/10.7554/eLife.42992.001

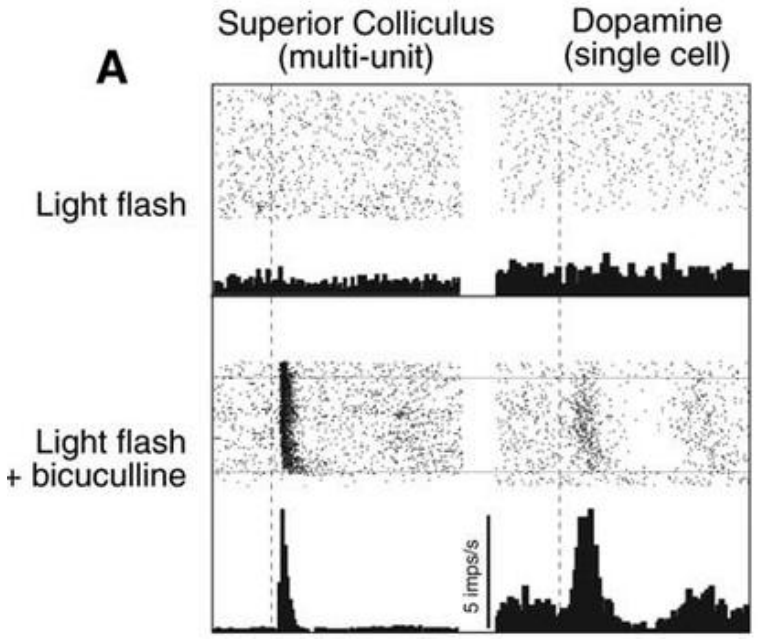

Dommett E, et al. 2005 How visual stimuli activate dopaminergic neurons at short latency. Science. 2005 Mar 4;307(5714):1476-9.

Prediction errors

Neuron
Review

Dopamine in Motivational Control: Rewarding, Aversive, and Alerting

Ethan S.
[1]Laborato
[2]Primate R
*Correspo
DOI 10.10

Midbrai
in posit
signals
recent a
propose
have dis
brain ne
brain n
augmen
hypothe
adaptiv

eLIFE
elifesciences.org

RESEARCH ARTICLE

Reward prediction error does not explain movement selectivity in DMS-projecting dopamine neurons

Rachel S Lee, Marcelo G Mattar, Nathan F Parker, Ilana B Witten*, Nathaniel D Daw*

LETTER

https://doi.org/10.1038/s41586-019-1261-9

Specialized coding of sensory, motor and cognitive variables in VTA dopamine neurons

Ben Engelhard[1,2], Joel Finkelstein[1,3], Julia Cox[1], Weston Fleming[1], Hee Jae Jang[1], Sharon Ornelas[1], Sue Ann Koay[1], Stephan Y. Thiberge[1,2], Nathaniel D. Daw[1,3], David W. Tank[1,2] & Ilana B. Witten[1,2,3]*

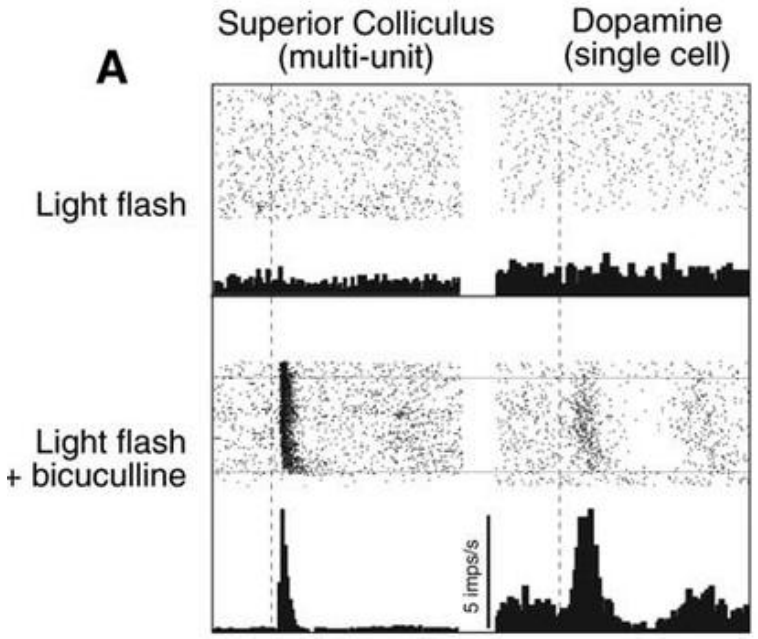

Dommett E, et al. 2005 How visual stimuli activate dopaminergic neurons at short latency. Science. 2005 Mar 4;307(5714):1476-9.

Prediction errors

Neuron
Review

Dopamine in Motivational Control: Rewarding, Aversive, and Alerting

Ethan S.
[1]Laborato
[2]Primate R
*Correspo
DOI 10.10

Midbrai
in posit
signals
recent a
propose
have dis
brain ne
brain n
augmen
hypothe
adaptiv

eLIFE
elifesciences.org

RESEARCH ARTICLE

Reward prediction error does not explain movement selectivity in DMS-projecting dopamine neurons

Rachel S Lee, Marcelo G Mattar, Nathan F Parker, Ilana B Witten*, Nathaniel D Daw*

LETTER

https://doi.org/10.1038/s41586-019-1261-9

Specialized coding of sensory, motor and cognitive variables in VTA dopamine neurons

Ben Engelhard[1,2], Joel Finkelstein[1,3], Julia Cox[1], Weston Fleming[1], Hee Jae Jang[1], Sharon Ornelas[1], Sue Ann Koay[1], Stephan Y. Thiberge[1,2], Nathaniel D. Daw[1,3], David W. Tank[1,2] & Ilana B. Witten[1,2,3]*

PERSPECTIVES

OPINION

The short-latency dopamine signal: a role in discovering novel actions?

Peter Redgrave and Kevin Gurney

Abstract | An influential concept in contemporary computational neuroscience is the reward prediction error hypothesis of phasic dopaminergic function. It maintains that midbrain dopaminergic neurons signal the occurrence of unpredicted reward, which is used in appetitive learning to reinforce existing actions that most often lead to reward. However, the availability of limited afferent sensory processing and the precise timing of dopaminergic signals suggest that they might instead have a central role in identifying which aspects of context and behavioural output are crucial in causing unpredicted events.

to what is being computed and why[19,28]). A notable exception is the reward prediction error hypothesis proposed by Montague *et al.*[8,7] and by Schultz and colleagues[2,8–10]. These investigators suggest that the short-latency, sensory-evoked DA responses signal reward prediction errors, which are used by reinforcement learning mechanisms in the basal ganglia, and elsewhere, to select actions that will maximize the future acquisition of reward. The reward prediction error hypothesis has received much empirical support[11–17] and is now widely accepted by many biological[9,18–30] and computational neuroscientists[7,31–35]. In this article, however, we wish to question this view and make an alternative suggestion. To do this, we first need to outline certain important aspects of phasic DA signalling.

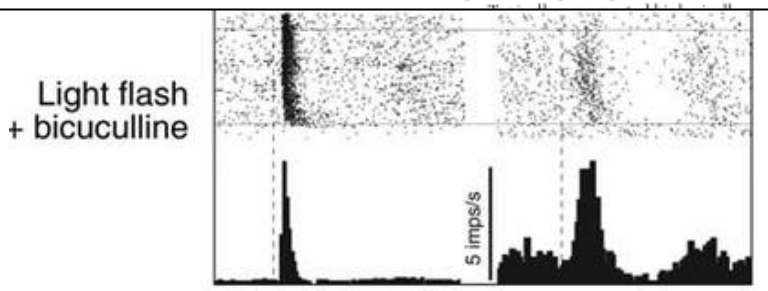

Dommett E, et al. 2005 How visual stimuli activate dopaminergic neurons at short latency. Science. 2005 Mar 4;307(5714):1476-9.

Prediction errors

Neuron
Review

Dopamine in Motivational Control:
Rewarding, Aversive, and Alerting

Ethan S.
[1]Laborato
[2]Primate R
*Correspo
DOI 10.10

Midbrai
in posit
signals
recent a
propose
have dis
brain ne
brain n
augmen
hypothe
adaptiv

eLIFE
elifesciences.org

RESEARCH ARTICL

Reward prediction error does
movement selectivity in DMS
dopamine neurons

Rachel S Lee, Marcelo G Mattar, Nathan F Parker, Ilana
Nathaniel D Daw*

LETTER

Specialized coding of senso
variables in VTA dopamine

Ben Engelhard[1,2], Joel Finkelstein[1,3], Julia Cox[1], Weston Fleming[1],
Stephan Y. Thiberge[1,2], Nathaniel D. Daw[1,3], David W. Tank[1,2] & Ila

Current Biology

Article

Dopamine neurons drive spatiotemporally heterogeneous striatal dopamine signals during learning

Highlights

- VTA DA drives learning and spatiotemporally heterogeneous DA signals in striatum
- Cue-evoked DA emerges in the NAc core and DMS, but not shell or DLS
- NAc core and DMS DA reflect movement invigoration and prediction errors
- SNc DA drives learning but does not engage prediction-related striatal DA encoding

Authors

Liv Engel, Amy R. Wolff, Madelyn Blake, Val L. Collins, Sonal Sinha, Benjamin T. Saunders

Correspondence

bts@umn.edu

In brief

Engel and Wolff et al. examine how dopamine (DA)-neuron-mediated Pavlovian learning engages DA signals across the striatum. They find a spatiotemporally heterogeneous landscape of DA encoding that is qualitatively distinct for VTA versus SNc DA neurons. These results offer broad insight into the functional architecture of the striatal network.

PERSPECTIVES

hat is being computed and why[10,28]). A
able exception is the reward prediction
r hypothesis proposed by Montague
[l,8,7] and by Schultz and colleagues[2,8–10].
se investigators suggest that the short-
ncy, sensory-evoked DA responses
al reward prediction errors, which are
d by reinforcement learning mechanisms
he basal ganglia, and elsewhere, to select
ons that will maximize the future acqui-
on of reward. The reward prediction
r hypothesis has received much empiri-
support[11–17] and is now widely accepted
many biological[9,18–20] and computational
roscientists[7,10–15]. In this article, however,
wish to question this view and make an
rnative suggestion. To do this, we first
d to outline certain important aspects of
sic DA signalling.

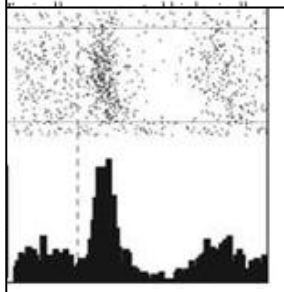

Dommett E, et al. 2005 How visual stimuli activate dopaminergic neurons at short latency. Science. 2005 Mar 4;307(5714):1476-9.

Prediction errors

Neuron
Review

Dopamine in Motivational Control:
Rewarding, Aversive, and Ale

Ethan S.
[1]Laborato
[2]Primate
*Correspo
DOI 10.10

Midbrai
in posit
signals
recent a
propose
have di
brain ne
brain n
augmen
hypothe
adaptiv

eLIFE
elifesciences.org

Reward
moveme
dopamin

Rachel S Lee, M
Nathaniel D Da

LET

Specia
variab

Ben Engelhard
Stephan Y. Thi

nature neuroscience

Article

https://doi.org/10.1038/s41593-024-01689-1

A feature-specific prediction error model explains dopaminergic heterogeneity

Received: 8 February 2022

Accepted: 22 May 2024

Published online: 03 July 2024

Check for updates

Rachel S. Lee[1], Yotam Sagiv[1], Ben Engelhard[1], Ilana B. Witten[1] & Nathaniel D. Daw[1,2]

The hypothesis that midbrain dopamine (DA) neurons broadcast a reward prediction error (RPE) is among the great successes of computational neuroscience. However, recent results contradict a core aspect of this theory: specifically that the neurons convey a scalar, homogeneous signal. While the predominant family of extensions to the RPE model replicates the classic model in multiple parallel circuits, we argue that these models are ill suited to explain reports of heterogeneity in task variable encoding across DA neurons. Instead, we introduce a complementary 'feature-specific RPE' model, positing that individual ventral tegmental area DA neurons report RPEs for different aspects of an animal's moment-to-moment situation. Further, we show how our framework can be extended to explain patterns of heterogeneity in action responses reported among substantia nigra pars compacta DA neurons. This theory reconciles new observations of DA heterogeneity with classic ideas about RPE coding while also providing a new perspective of how the brain performs reinforcement learning in high-dimensional environments.

PERSPECTIVES

Article

orally
nals during

Amy R. Wolff, Madelyn Blake,
lins, Sonal Sinha,
T. Saunders

ondence
.edu

l Wolff et al. examine how
(DA)-neuron-mediated
learning engages DA signals
e striatum. They find a
porally heterogeneous
e of DA encoding that is
ely distinct for VTA versus SNc
ns. These results offer broad
o the functional architecture of
l network.

what is being computed and why[19,28]). A
able exception is the reward prediction
r hypothesis proposed by Montague
[8,7] and by Schultz and colleagues[2,8–10].
se investigators suggest that the short-
ncy, sensory-evoked DA responses
al reward prediction errors, which are
d by reinforcement learning mechanisms
he basal ganglia, and elsewhere, to select
ons that will maximize the future acqui-
on of reward. The reward prediction
r hypothesis has received much empiri-
support[31–37] and is now widely accepted
many biological[9,28–30] and computational
roscientists[7,10–15]. In this article, however,
wish to question this view and make an
rnative suggestion. To do this, we first
d to outline certain important aspects of
sic DA signalling.

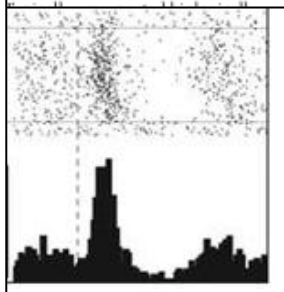

Dommett E, et al. 2005 How visual stimuli activate dopaminergic neurons at short latency. Science. 2005 Mar 4;307(5714):1476-9.

Prediction errors

Neuron
Review

Dopamine in Motiv
Rewarding, Aversi

Ethan S.
[1]Laborato
[2]Primate R
*Correspo
DOI 10.10

Midbrai
in posit
signals
recent a
propose
have dis
brain ne
brain n
augmen
hypothe
adaptiv

eLIFE
elifesciences.org

Cell
PRESS

nature neuroscience

PERSPECTIVE
https://doi.org/10.1038/s41593-018-0152-y

What does dopamine mean?

Joshua D. Berke

Dopamine is a critical modulator of both learning and motivation. This presents a problem: how can target cells know whether increased dopamine is a signal to learn or to move? It is often presumed that motivation involves slow ('tonic') dopamine changes, while fast ('phasic') dopamine fluctuations convey reward prediction errors for learning. Yet recent studies have shown that dopamine conveys motivational value and promotes movement even on subsecond timescales. Here I describe an alternative account of how dopamine regulates ongoing behavior. Dopamine release related to motivation is rapidly and locally sculpted by receptors on dopamine terminals, independently from dopamine cell firing. Target neurons abruptly switch between learning and performance modes, with striatal cholinergic interneurons providing one candidate switch mechanism. The behavioral impact of dopamine varies by subregion, but in each case dopamine provides a dynamic estimate of whether it is worth expending a limited internal resource, such as energy, attention, or time.

variab

Ben Engelhard
Stephan Y. Thil

suited to explain reports of heterogeneity in task variable encoding across DA neurons. Instead, we introduce a complementary 'feature-specific RPE' model, positing that individual ventral tegmental area DA neurons report RPEs for different aspects of an animal's moment-to-moment situation. Further, we show how our framework can be extended to explain patterns of heterogeneity in action responses reported among substantia nigra pars compacta DA neurons. This theory reconciles new observations of DA heterogeneity with classic ideas about RPE coding while also providing a new perspective of how the brain performs reinforcement learning in high-dimensional environments.

PERSPECTIVES

what is being computed and why[19,28]). A
able exception is the reward prediction
r hypothesis proposed by Montague
[8,7] and by Schultz and colleagues[2,8–10].
se investigators suggest that the short-
ncy, sensory-evoked DA responses
al reward prediction errors, which are
d by reinforcement learning mechanisms
he basal ganglia, and elsewhere, to select
ons that will maximize the future acqui-
on of reward. The reward prediction
r hypothesis has received much empiri-
support[31–37] and is now widely accepted
many biological[9,28–30] and computational
roscientists[7,10–15]. In this article, however,
wish to question this view and make an
rnative suggestion. To do this, we first
d to outline certain important aspects of
sic DA signalling.

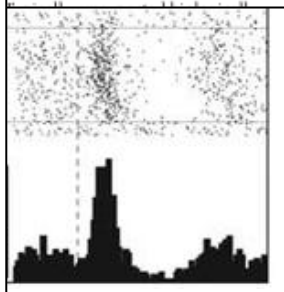

ns. These results offer broad
o the functional architecture of
l network.

Dommett E, et al. 2005 How visual stimuli activate dopaminergic neurons at short latency. Science. 2005 Mar 4;307(5714):1476-9.

Prediction errors

Is a flash

Is dopar
salient c

~~Reward~~ prediction error

nature neuroscience

PERSPECTIVE

https://doi.org/10.1038/s41593-018-0152-y

What does dopamine mean?

Joshua D. Berke

Dopamine is a critical modulator of both learning and motivation. This presents a problem: how can target cells know whether increased dopamine is a signal to learn or to move? It is often presumed that motivation involves slow ('tonic') dopamine changes, while fast ('phasic') dopamine fluctuations convey reward prediction errors for learning. Yet recent studies have shown that dopamine conveys motivational value and promotes movement even on subsecond timescales. Here I describe an alternative account of how dopamine regulates ongoing behavior. Dopamine release related to motivation is rapidly and locally sculpted by receptors on dopamine terminals, independently from dopamine cell firing. Target neurons abruptly switch between learning and performance modes, with striatal cholinergic interneurons providing one candidate switch mechanism. The behavioral impact of dopamine varies by subregion, but in each case dopamine provides a dynamic estimate of whether it is worth expending a limited internal resource, such as energy, attention, or time.

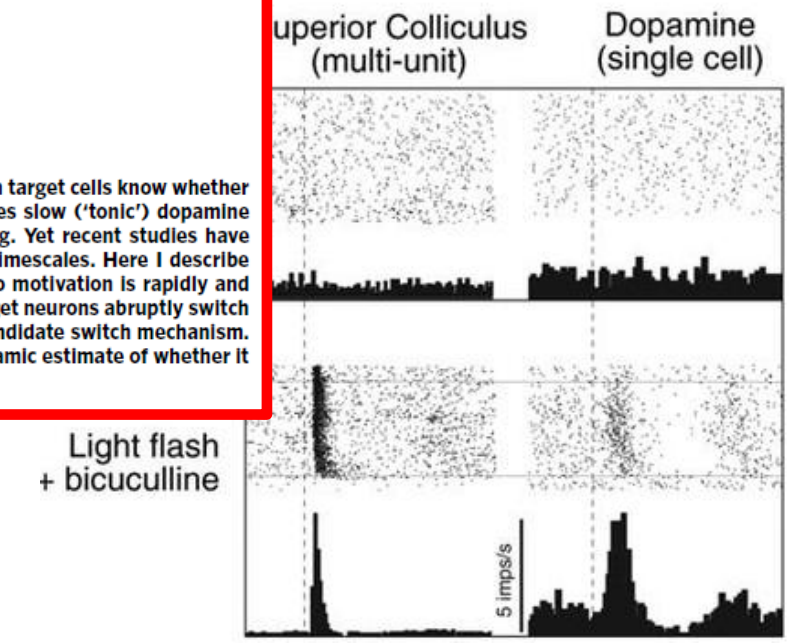

Dommett E, et al. 2005 How visual stimuli activate dopaminergic neurons at short latency. Science. 2005 Mar 4;307(5714):1476-9.

Prediction errors

Is a flash of light rewarding?

Is dopamine (also) responding to salient or novel stimuli?

~~Reward~~ prediction error

Or novelty detection.

Or onset/offset detection.

Or extrinsic/intrinsic reward signals.

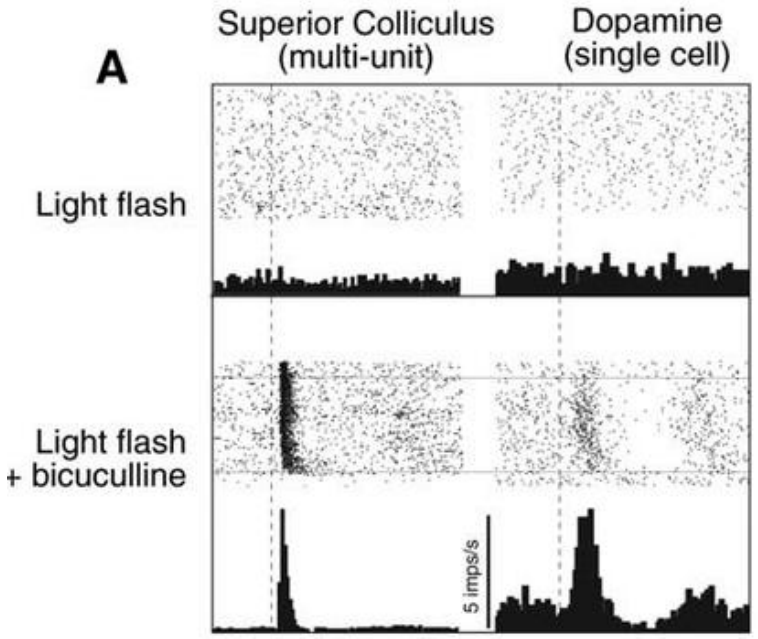

Dommett E, et al. 2005 How visual stimuli activate dopaminergic neurons at short latency. Science. 2005 Mar 4;307(5714):1476-9.

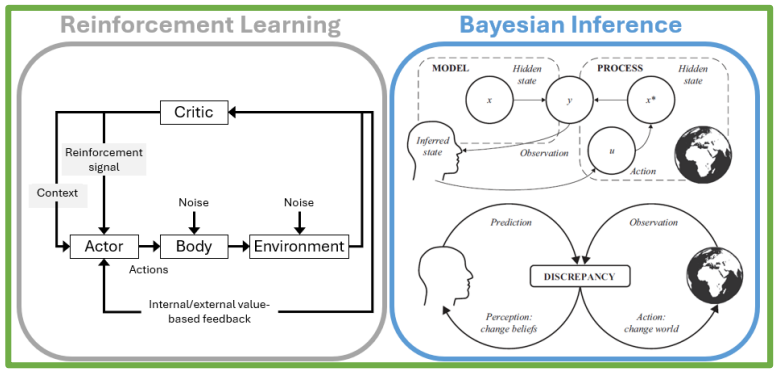

Competing theories

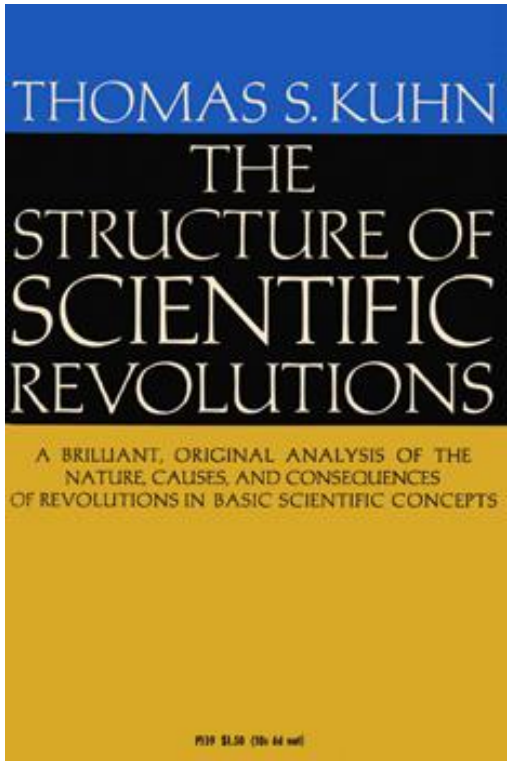

Philosophy of science (mildly scary version)

- **Thomas Kuhn**

" [...] past scientific revolutions were not the unambiguous expression of sheerly logical and experimental factors, rationally played out according to a well-defined methodology. Rather, they were the expression of a variety of nonlogical factors as well: social, psychological, metaphysical, technological, aesthetic, and personal."

The Structure of Scientific Revolutions

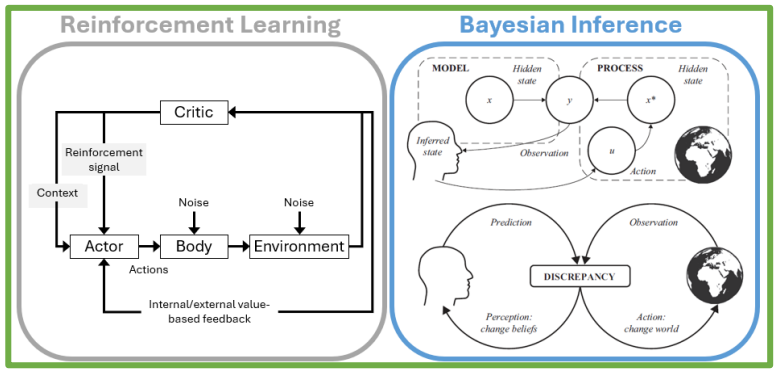

Competing theories

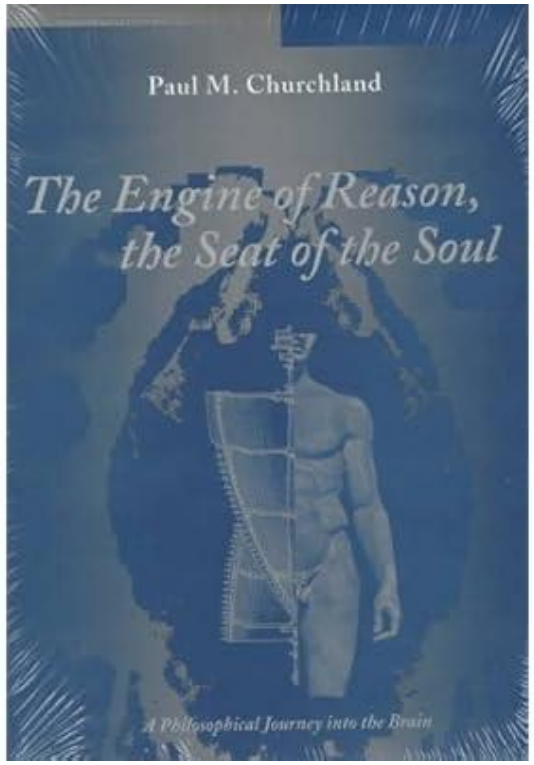

Philosophy of science (mildly scary version)

- ~~**Thomas Kuhn**~~
- **Paul Churchland**

" [...] past scientific revolutions were not the unambiguous expression of sheerly logical and experimental factors, rationally played out according to a well-defined methodology. Rather, they were the expression of a variety of nonlogical factors as well: social, psychological, metaphysical, technological, aesthetic, and personal."

~~**The Structure of Scientific Revolutions**~~

The Engine of Reason

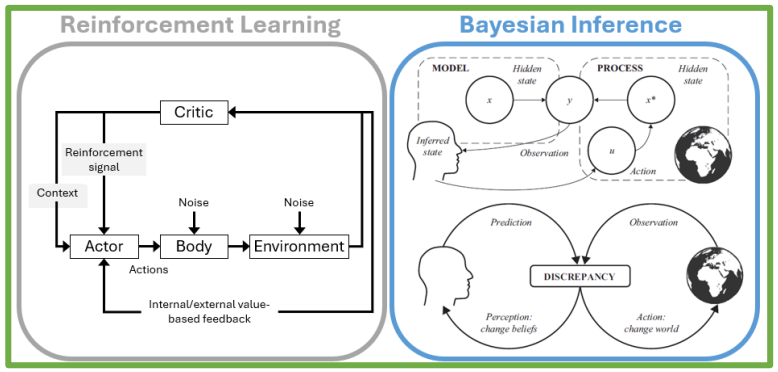

Competing theories

Philosophy of science (really scary version)

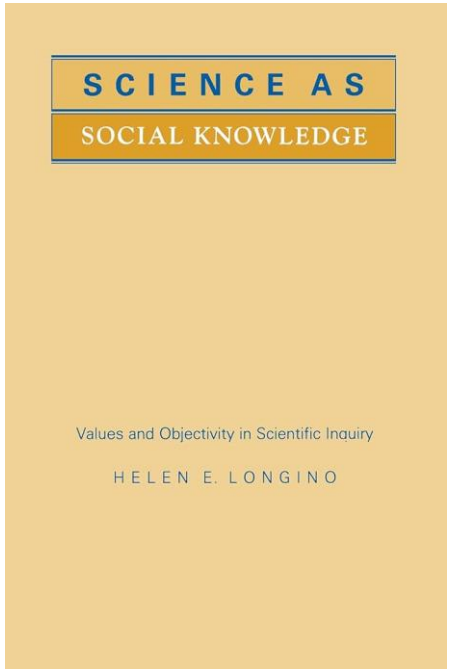

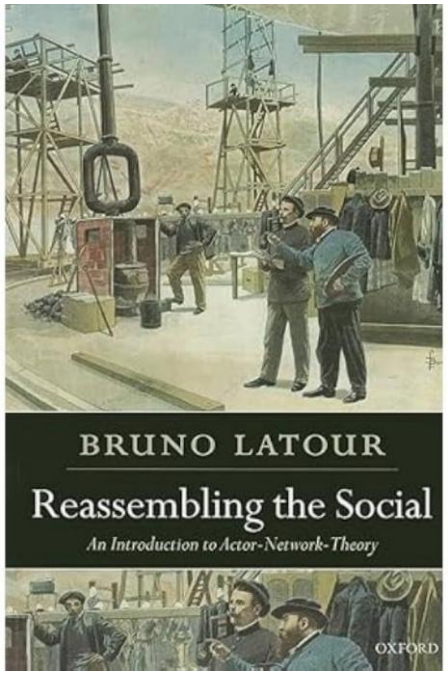

Key words for future reference:

- Social epistemology of science or social empiricism.
- Actor-Network Theory (ANT).
- Cultural Attractor Theory (CAT) or cultural darwinism.

Questions?

- A change of perspective: reward prediction errors and errors of prediction without rewards.

    1. Beliefs beyond values.
    2. Competing theories.
    3. Dopamine signals of reward prediction error (or maybe not).

- Basic elements of Bayesian inference (A gentle introduction)

- Phenotypes and hierarchies.

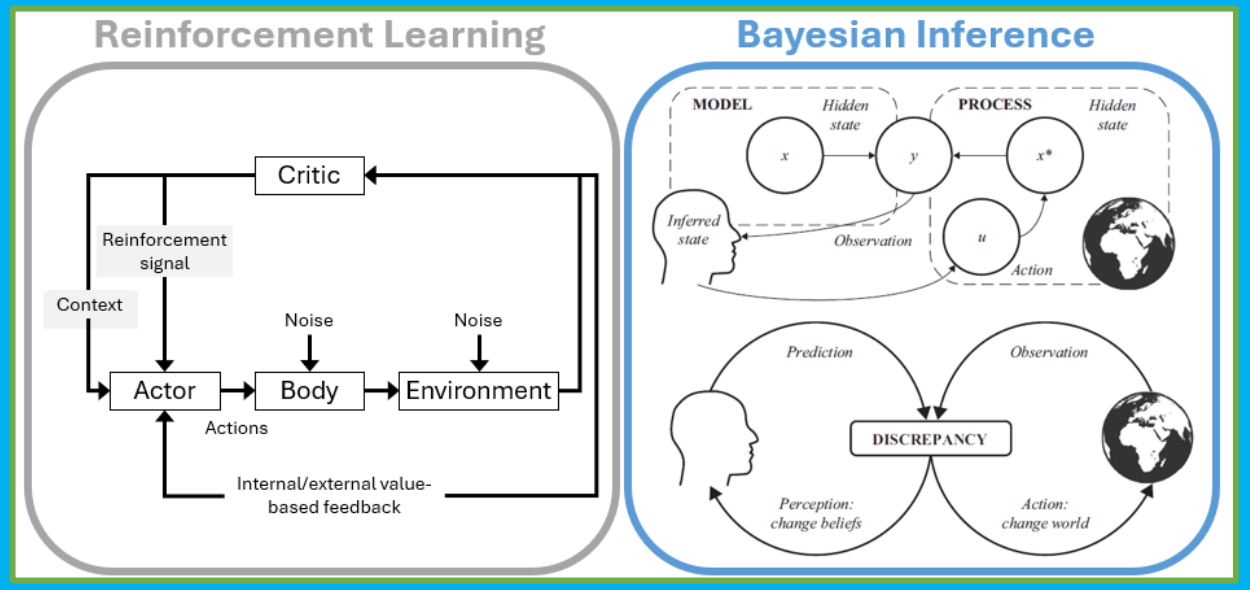

Next

- A change of perspective: reward prediction errors and errors of prediction without rewards.
- Basic elements of Bayesian inference (A ***gentle*** introduction)
  1. Elements of probability theory
  2. Belief updating with discrete evidence
- Phenotypes and hierarchies.

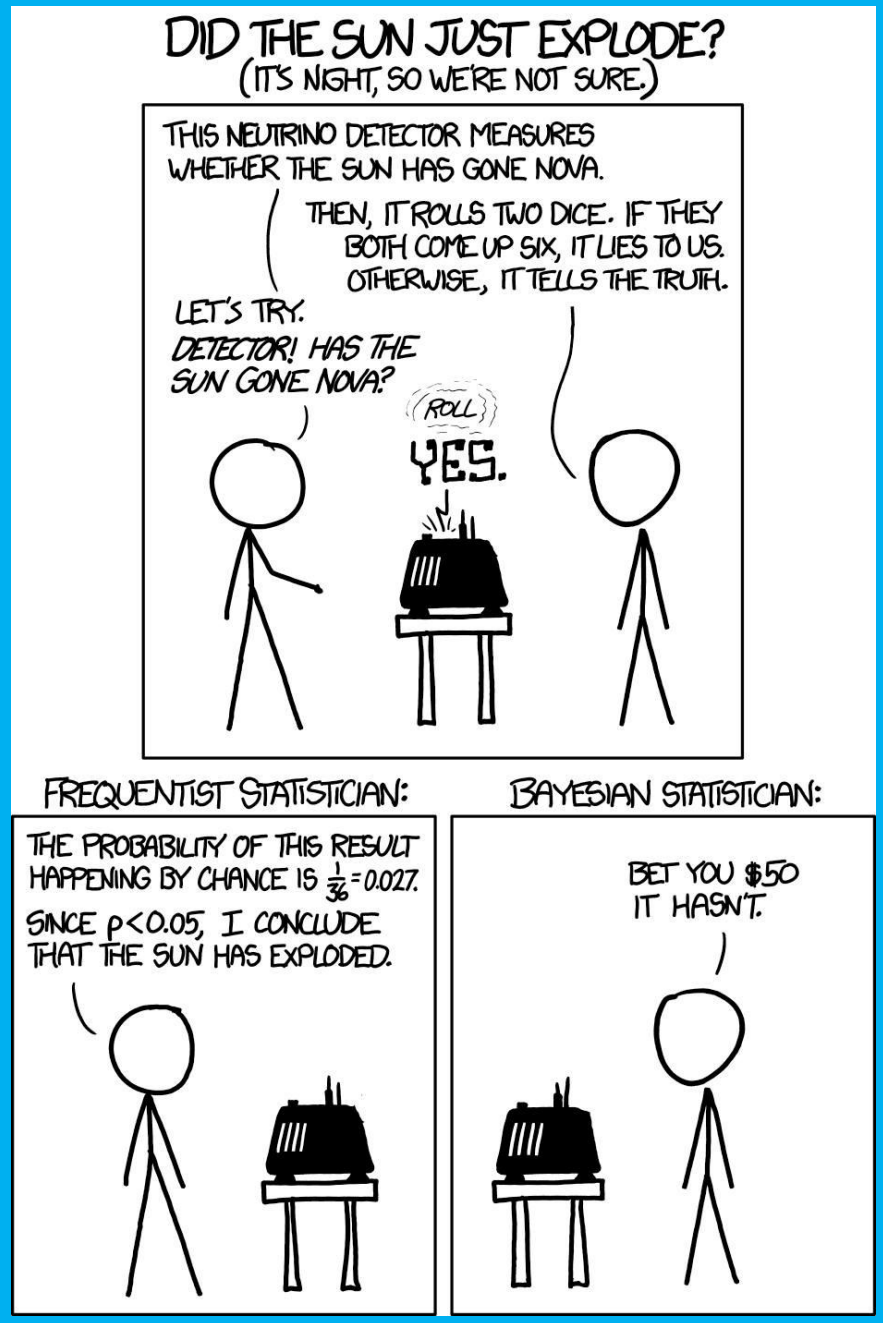

Independent probability

Probability of one variable or *likelihood* of an event.

It quantifies how likely it is for something to happen: the probability of a specific outcome for a random variable.

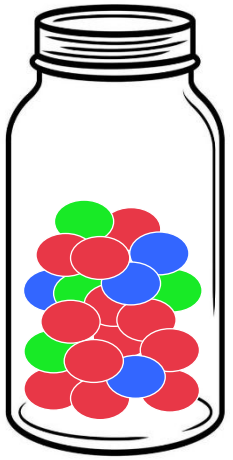

Independent probability

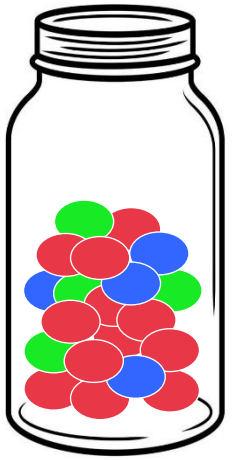

Probability of one variable or *likelihood* of an event.

It quantifies how likely it is for something to happen: the probability of a specific outcome for a random variable.

Probability of event A = p(A)

$$\text{p(red marble)}=\frac{\text{number of desired outcomes}}{\text{total possible outcomes (if equally probable)}}=\frac{12}{20}=0.60$$

Assuming a jar with 12 red marble, 4 blue and 4 green

Independent probability

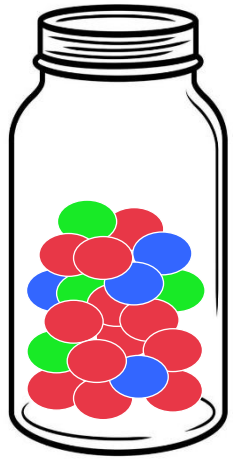

Probability of one variable or *likelihood* of an event.

It quantifies how likely it is for something to happen: the probability of a specific outcome for a random variable.

Probability of event A = p(A)

$$p(\text{red marble}) = \frac{\text{number of desired outcomes}}{\text{total possible outcomes (if equally probable)}} = \frac{12}{20} = 0.60$$

Assuming a jar with 12 red marble, 4 blue and 4 green

Sum of the probabilities for all outcomes = 1

Probability of impossible outcome = 0

Probability of certain outcome = 1

Joint probability

The probability of two (or more) simultaneous events, e.g. the outcomes of two (or more) different random variables.

$p(A,B) = p(A \cap B) = p(A \text{ ^ } B) = p(A \text{ and } B)$

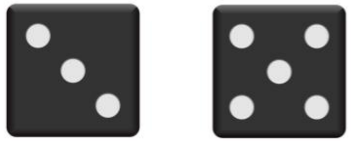

Joint probability

The probability of two (or more) simultaneous events, e.g. the outcomes of two (or more) different random variables.

p(A,B) = p(A ∩ B) = p(A ^ B) = p(A and B)

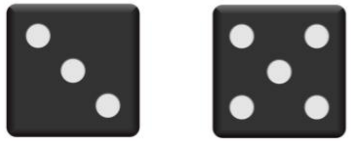

If A and B are statistically independent variables, it is calculated as:

p(A , B) = p(A)*p(B) e.g., $\frac{1}{6} * \frac{1}{6} = \frac{1}{36}$

Assuming two standard 6 face dices, rolled at once.

Joint probability

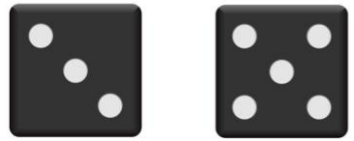

If A and B are statistically independent variables, it is calculated as:

p(A , B) = p(A)*p(B) e.g., $\frac{1}{6} * \frac{1}{6} = \frac{1}{36}$

Assuming two standard 6 face dices, rolled at once.

With dependent variables, it is equivalent to :

p(A , B) = p(B, A) = p(A given B) * p(B) = p(B given A) * p(A)

A=p(odd umber)= $\frac{1}{2}$; B=p(blue face)= $\frac{1}{2}$

p(A , B) = p(B, A) = $\frac{2}{3} * \frac{1}{2} = \frac{2}{6} = \frac{1}{3}$

Assuming a standard 6 face dice, with 3 blue and 3 red faces

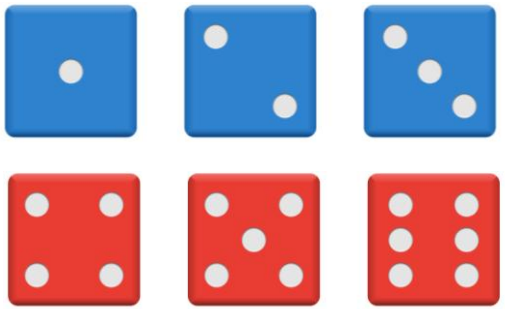

Conditional probability

The probability of an event given the occurrence of another event. For instance, the probability of event A occurring after or in the presence of event B (p(B=1)), for a given trial.

p(A given B) = p(A | B) = $\frac{p(A \text{ and } B)}{p(B)}$

Where p(B)>0, i.e., the event B is possible (but it does not need to be certain).

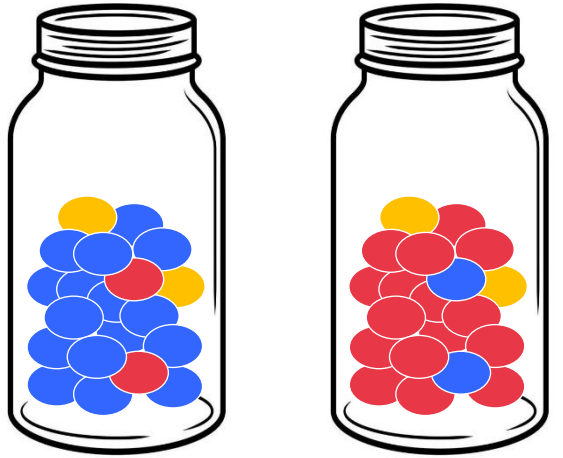

Conditional probability

$$\text{p(red marble)}=\frac{\text{number of desired outcomes}}{\text{total possible (equally probable) outcomes}}=\frac{18}{40}=0.45$$

However,

$$\text{p(red marble| blue extraction jar)} = \frac{p(red\ marble, Blue\ jar)}{p(Blue\ jar)} = \frac{0.05}{0.5} = 0.1$$

Assuming two jars:

“Blue jar”: 16 blue marbles, 2 red and 2 yellow

“Red jar”: 16 red marbles, 2 blue and 2 yellow

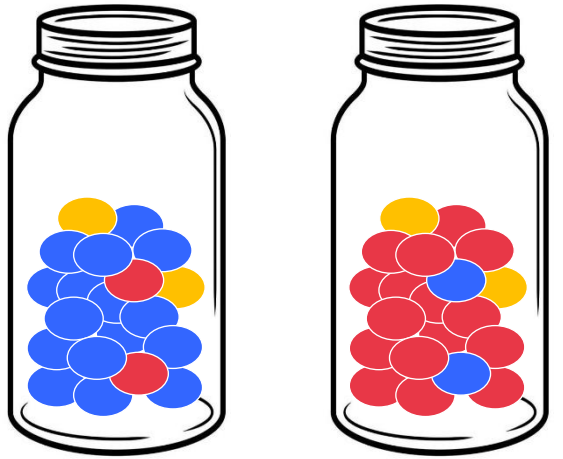

Bayes' theorem

$$p(A , B) = p(B , A) = p(A \mid B) * p(B) = p(B \mid A) * p(A)$$

Bayes' theorem

$p(A , B) = p(B , A) = p(A \mid B) * p(B) = p(B \mid A) * p(A)$

$$p(A \mid B) = \frac{p(B|A) * p(A)}{p(B)}$$

Bayes' theorem

$p(A , B) = p(B , A) = p(A \mid B) * p(B) = p(B \mid A) * p(A)$

$$p(A \mid B) = \frac{p(B|A) * p(A)}{p(B)}$$

Conditional Probability of the data (D) to occur, if the hypothesis (H) is correct.

$$p(H|D) = \frac{p(D|H)\ p(H)}{p(D)}$$

Conditional Probability of the hypothesis (H) to be true, given the data (D).

Bayes' theorem

Conditional Probability of the data (D) to occur, if the hypothesis (H) is correct.

$$p(H|D) = \frac{p(D|H)\ \mathrm{p}(H)}{p(D)}$$

Independent Probability for the hypothesis (H) to occur, in the environment.

Conditional Probability of the hypothesis (H) to be true, given the data (D).

Independent Probability to observe the data (D), in the environment.

Bayesian Inference

A common question:
What is the probability to find wet grass in the park, today?
P(wet grass=1)

- Is the street wet?
- Are sprinkles usually working at this time of the day?
- Is rain likely in this time of the year?

Bayesian Inference

Data

Hypotheses

100%?

Inversion:

I see wet grass in the park, did it rain today?

p(rain | wet grass=1)

- Is the street wet?
- Are sprinkles usually working at this time of the day?
- Is rain likely in this time of the year?

Bayesian Inference

Data

Hypotheses

50%?

50%?

Inversion:

I see wet grass in the park, did it rain today?

p(rain | wet grass=1)

- Is the street wet?
- Are sprinkles usually working at this time of the day?
- Is rain likely in this time of the year?

Bayesian Inference

Data

Hypotheses

P(hypothesis | data)

NOTE: the frequentist approach assigns probabilities to data, not to hypotheses.

The Bayesian approach assigns probabilities to hypotheses, incorporating prior knowledge.

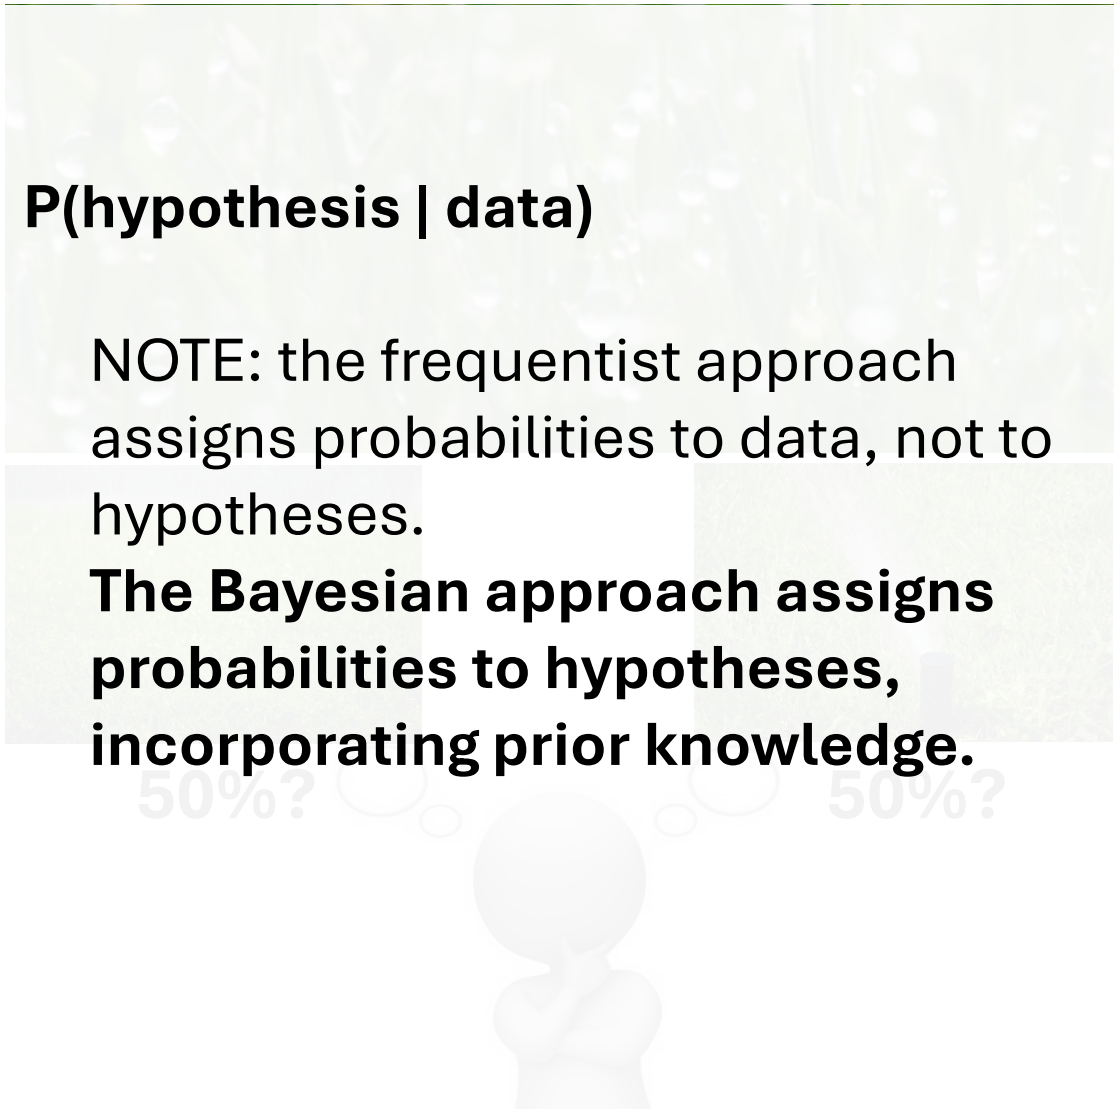

Inversion:

I see wet grass in the park, did it rain today?

p(rain | wet grass=1)

- Is the street wet?
- Are sprinkles usually working at this time of the day?
- Is rain likely in this time of the year?
- What probability was assigned to the possibility of rain, before I collected the data?

Interim Summary

P(hypothesis | data)

NOTE: the frequentist approach assigns probabilities to data, not to hypotheses.

The Bayesian approach assigns probabilities to hypotheses, incorporating prior knowledge.

It took me ~20 minutes to explain this joke.

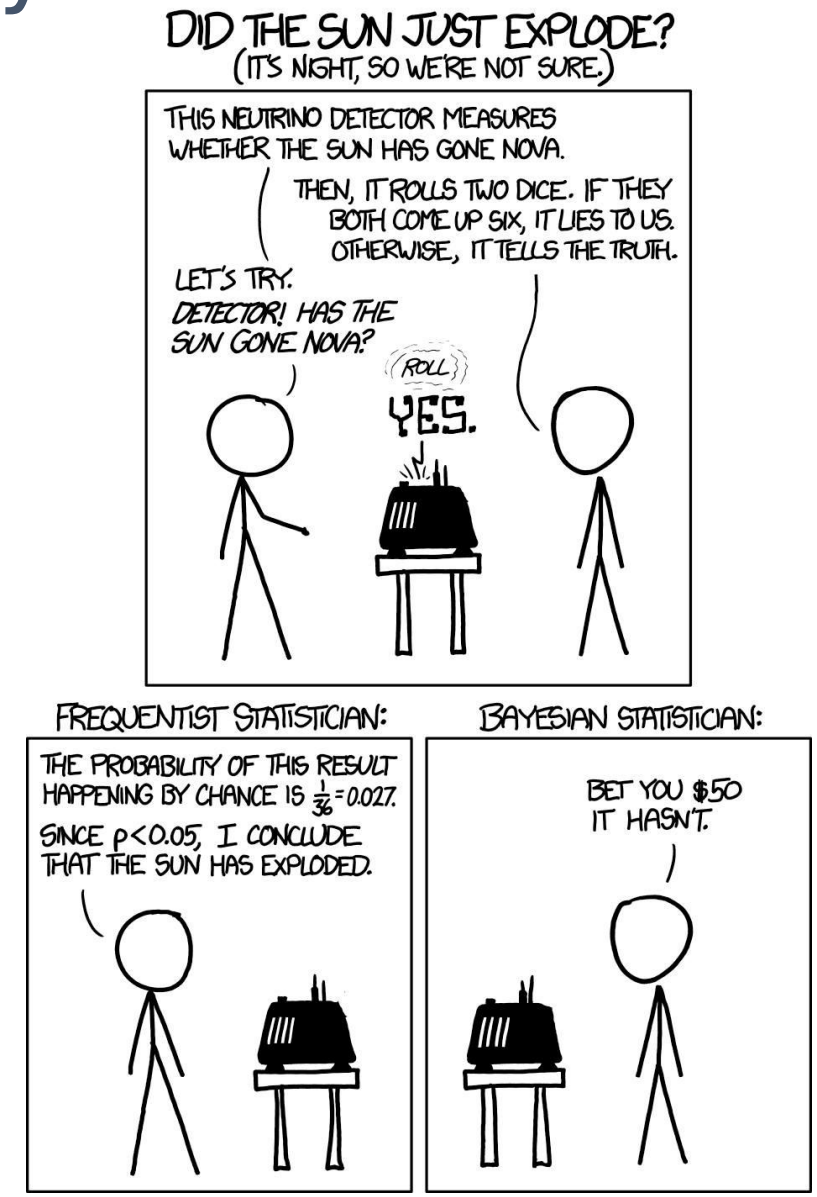

Next

- A change of perspective: reward prediction errors and errors of prediction without rewards.
- Basic elements of Bayesian inference (A *gentle* introduction)
    1. Elements of probability theory
    2. Belief updating with discrete evidence
- Phenotypes and hierarchies.

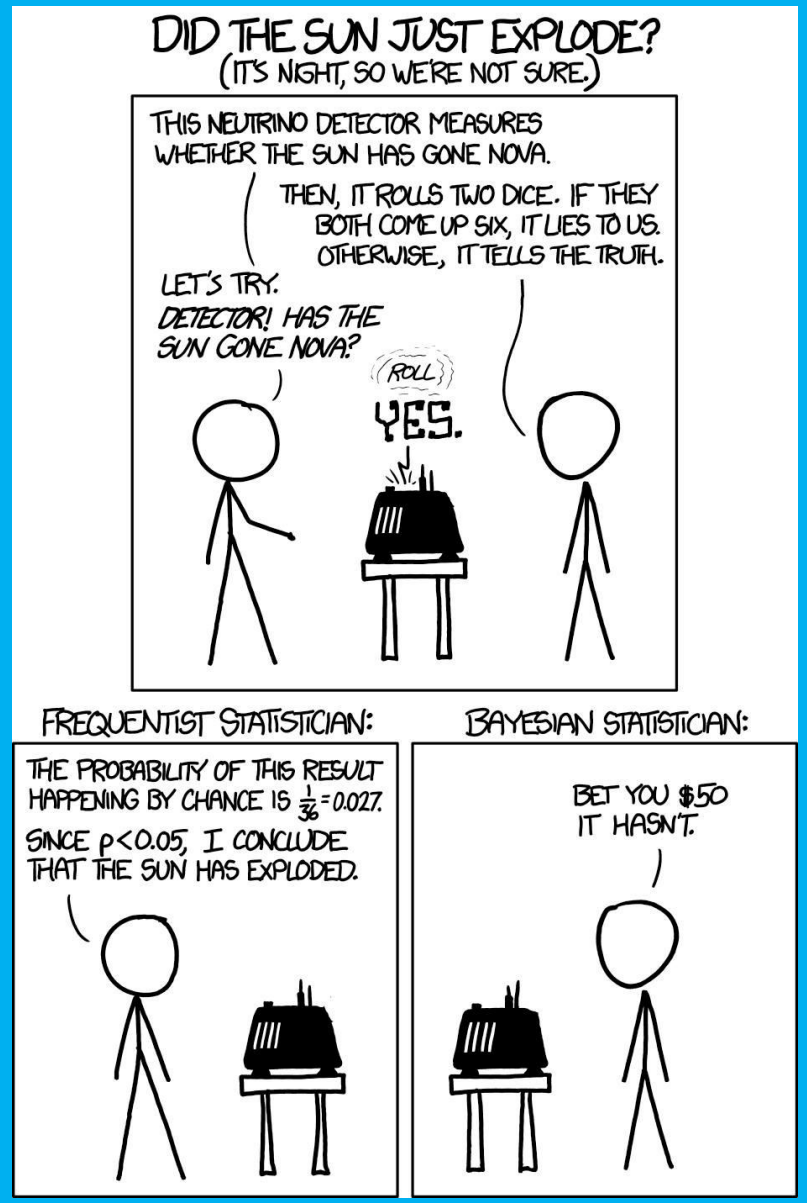

Belief updating with discrete evidence, M1

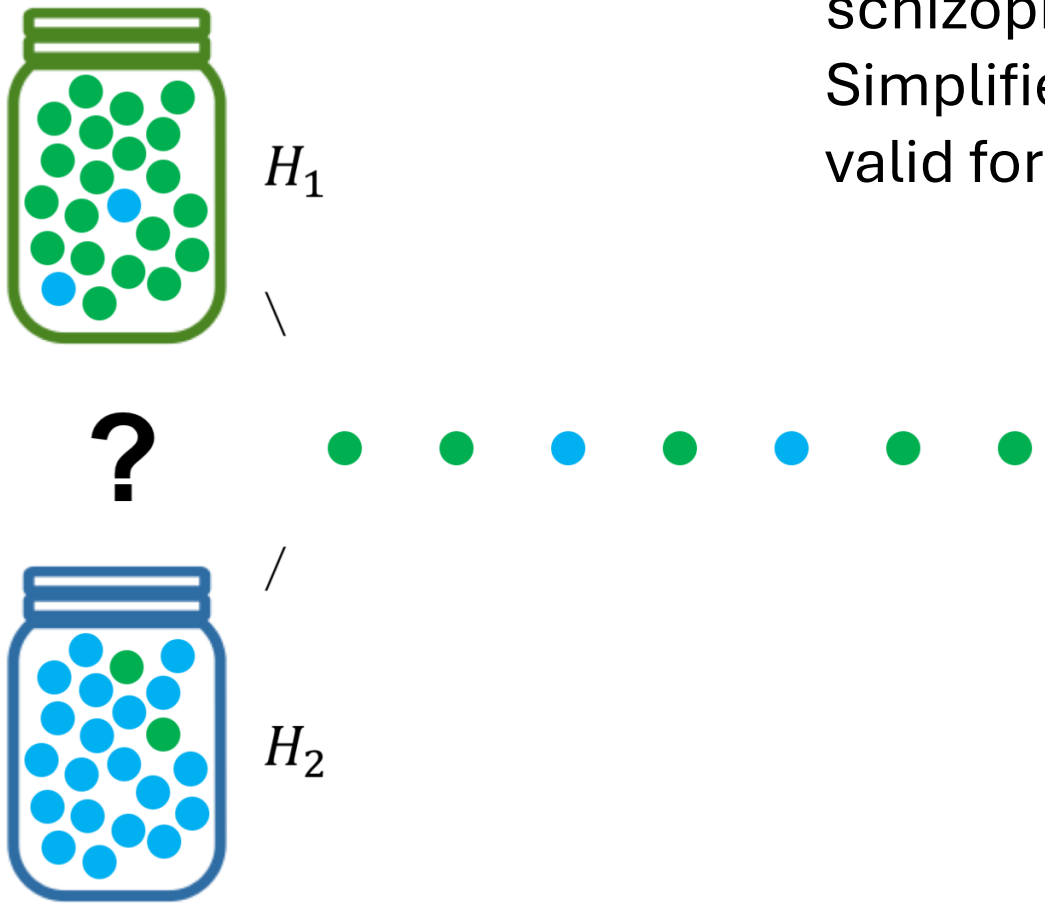

E.g., see Guillermo Horga's work on schizophrenia and delusions.
Simplified Bayesian inference updates, valid for environments with 2 options.

Belief updating with discrete evidence, M1

H_1

\

$$\frac{P(H_1 \mid E)}{P(H_2 \mid E)} = \frac{P(H_1)}{P(H_2)} \cdot \frac{P(E \mid H_1)}{P(E \mid H_2)}$$

/

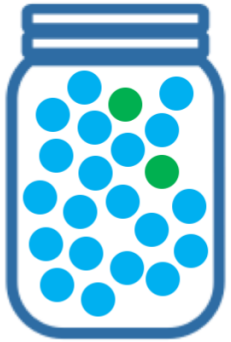

H_2

Belief updating with discrete evidence, M1

\

$$\frac{P(\text{🫙} \mid \text{●})}{P(\text{🫙} \mid \text{●})} = \frac{P(\text{🫙})}{P(\text{🫙})} \cdot \frac{P(\text{●} \mid \text{🫙})}{P(\text{●} \mid \text{🫙})}$$

/

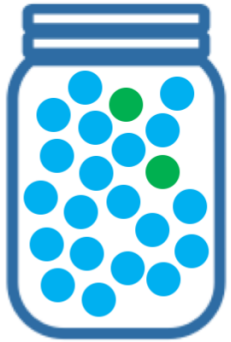

Belief updating with discrete evidence, M1

$$logit(posterior) = logit(prior) + logit(likelihood)$$

$$logit(p) = ln\left(\frac{1-p}{p}\right)$$

Belief updating with discrete evidence, M1

$$logit(posterior) = \omega_p \cdot logit(prior) + \omega_l \cdot logit(likelihood)$$

Belief updating with discrete evidence, M1

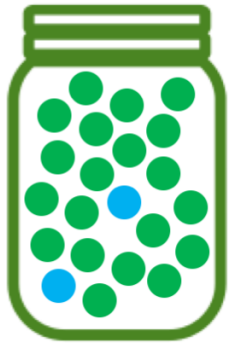

$$logit(posterior)_d = \omega_p \cdot logit(posterior)_{d-1} + \omega_l \cdot LLR$$

Log Likelihood Ratio

Belief updating with discrete evidence, M1

\

$$logit(posterior)_d = \omega_p \cdot logit(posterior)_{d-1} + \omega_l \cdot LLR$$

/

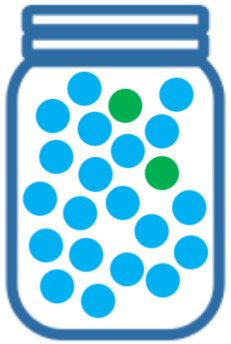

Leaky integration (neuroscience)
Base-rate neglect (economics)

Uscher and McClelland, Psychol Rev 2001

Belief updating with discrete evidence, M2

Generalization: complete Bayesian inference updates, valid for any number of choices.

Belief updating with discrete evidence, M2

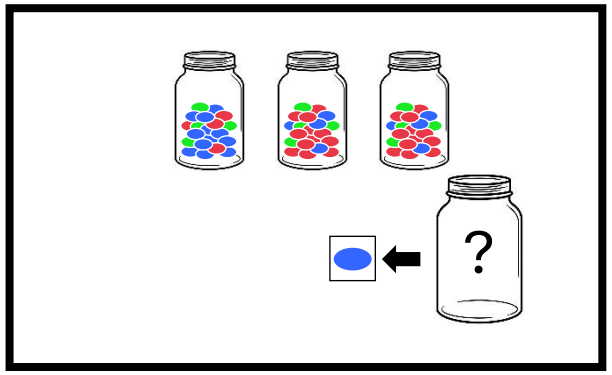

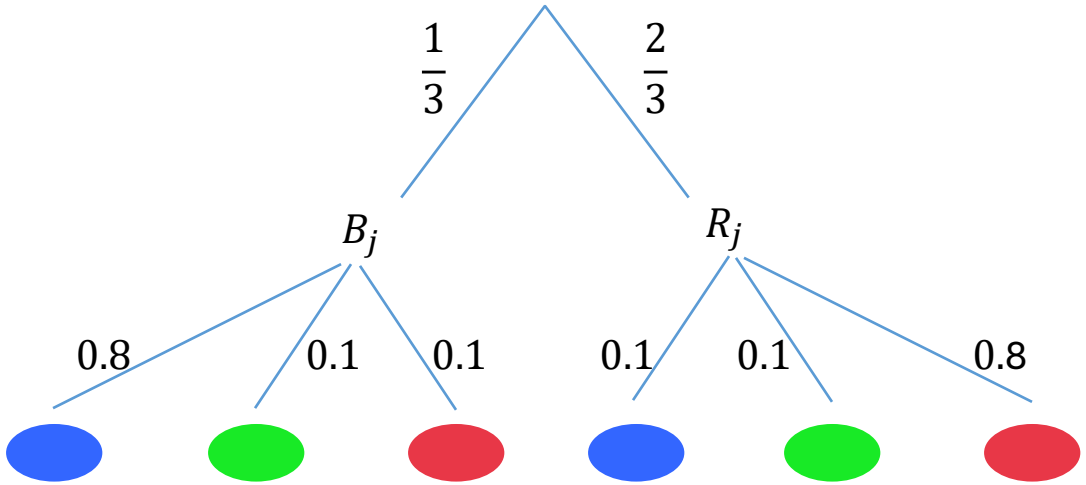

4. Prior probability: the distribution of probabilities, prior to collecting evidence.

$$P(B_j) = \frac{1}{3} \qquad P(R_j) = \frac{2}{3}$$

5. Likelihood: the probability of the data (marble extracted=blue), assuming each hypothesis is correct.

$$P(b|B_j)=0.8 \qquad P(b|R_j)=0.1$$

6. Posterior probability: the probability of each hypothesis to be true, given the data (marble extracted=blue).

$$P(B_j|b)=? \qquad P(R_j|b)=?$$

From priors to posteriors

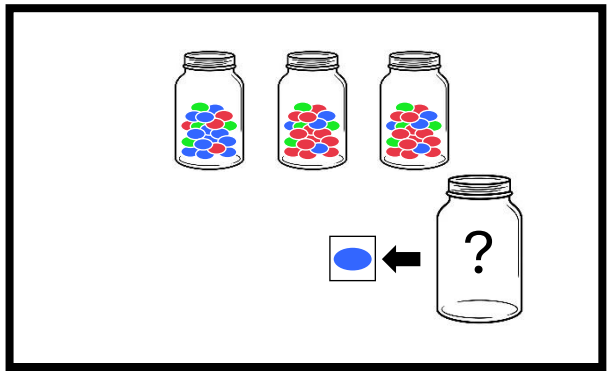

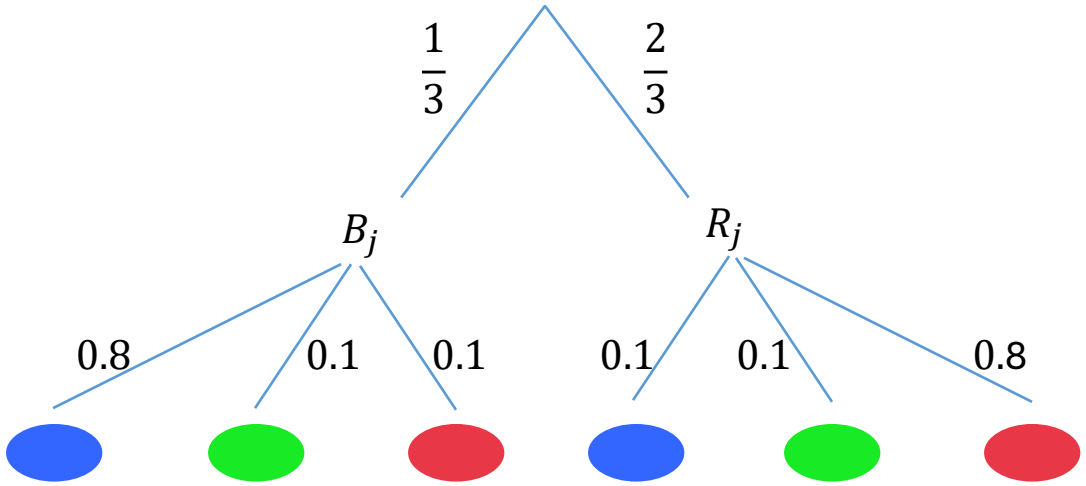

6. Posterior probability: the probability of each hypothesis, given the data (marble extracted=blue).

$$P(B_j|b) = \frac{P(b|B_j)\,P(B_j)}{P(b)} = \frac{0.8 \cdot \frac{1}{3}}{(0.8 \cdot \frac{1}{3} + 0.1 \cdot \frac{2}{3})} = 0.8$$

$$P(R_j|b) = \frac{P(b|R_j)\,P(R_j)}{P(b)} = \frac{0.1 \cdot \frac{2}{3}}{(0.8 \cdot \frac{1}{3} + 0.1 \cdot \frac{2}{3})} = 0.2$$

From priors to posteriors

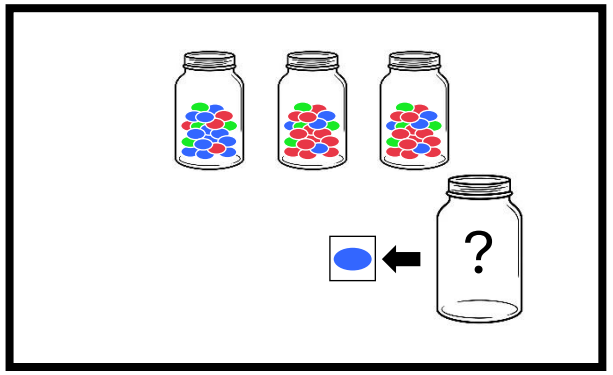

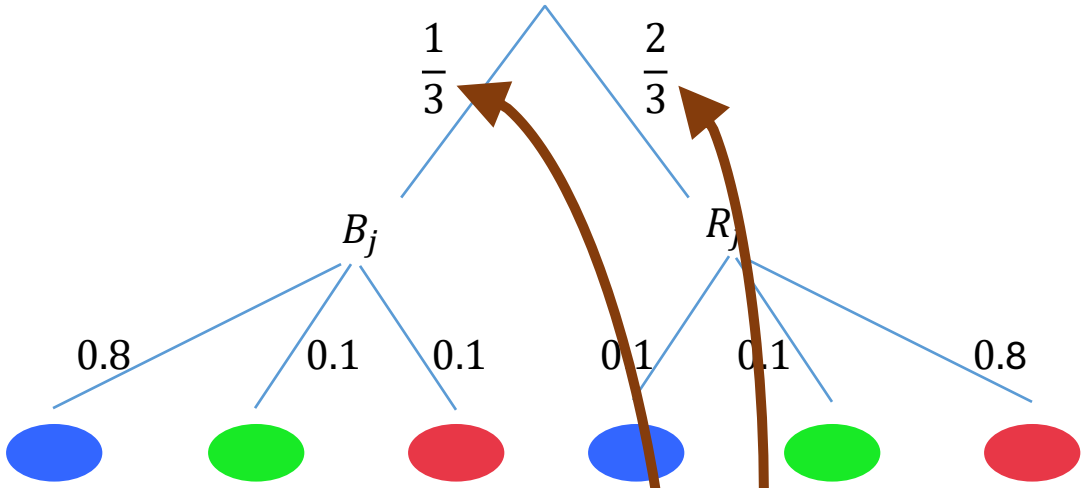

6. Posterior probability: the probability of each hypothesis, given the data (marble extracted=blue).

$$P(B_j|b) = \frac{P(b|B_j)\,P(B_j)}{P(b)} = \frac{0.8 \cdot \frac{1}{3}}{(0.8 \cdot \frac{1}{3} + 0.1 \cdot \frac{2}{3})} = 0.8$$

$$P(R_j|b) = \frac{P(b|R_j)\,P(R_j)}{P(b)} = \frac{0.1 \cdot \frac{2}{3}}{(0.8 \cdot \frac{1}{3} + 0.1 \cdot \frac{2}{3})} = 0.2$$

From priors to posteriors

New priors, new estimations to be updated…

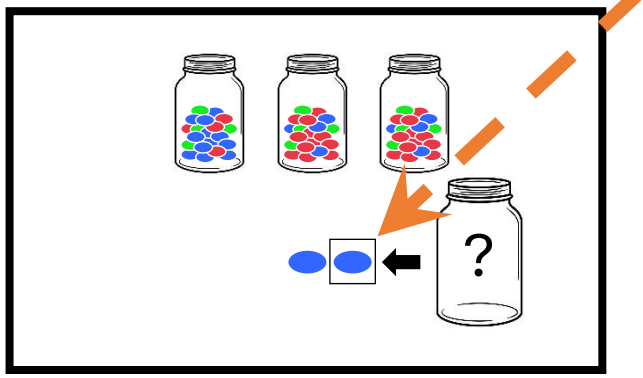

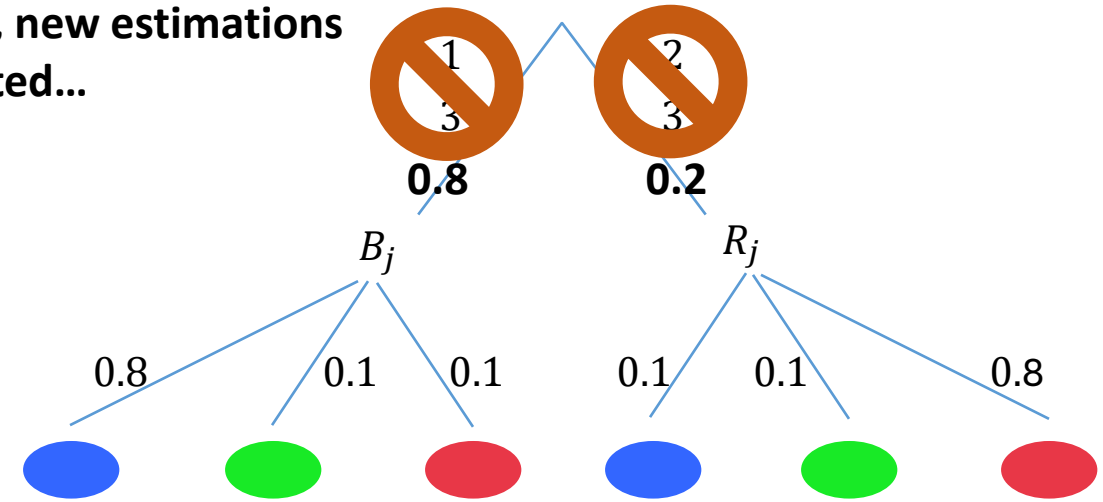

6. Posterior probability: the probability of each hypothesis, given the data (marbles extracted=blue + blue).

$$P(B_j|b) = \frac{P(b|B_j)\,P(B_j)}{P(b)} = \frac{0.8 \cdot \mathbf{0.8}}{(0.8 \cdot \mathbf{0.8} + 0.1 \cdot \mathbf{0.2})} = 0.97$$

$$P(R_j|b) = \frac{P(b|R_j)\,P(R_j)}{P(b)} = \frac{0.1 \cdot \mathbf{0.2}}{(0.8 \cdot \mathbf{0.8} + 0.1 \cdot \mathbf{0.2})} = 0.03$$

…priors have been updated into posteriors!

From priors to posteriors

New priors, new estimations to be updated…

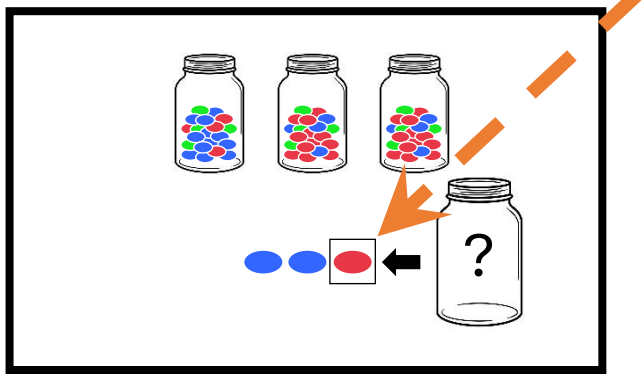

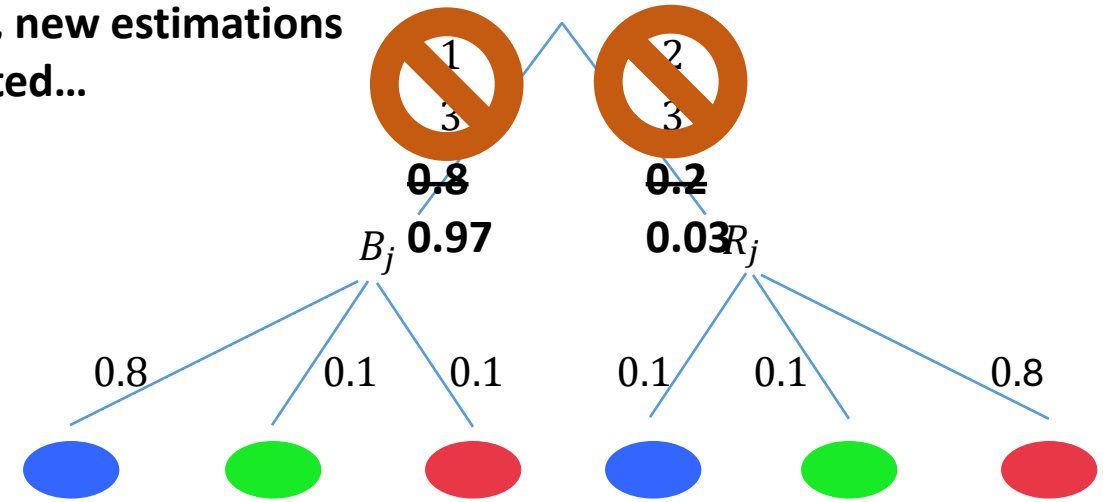

6. Posterior probability: the probability of each hypothesis, given the data (marbles extracted=blue + blue + red).

$$P(B_j|r) = \frac{P(r|B_j)\,P(B_j)}{P(r)} = \frac{0.1 \cdot \mathbf{0.97}}{(0.1 \cdot \mathbf{0.97} + 0.8 \cdot \mathbf{0.03})} = 0.8$$

$$P(R_j|r) = \frac{P(r|R_j)\,P(R_j)}{P(r)} = \frac{0.8 \cdot \mathbf{0.03}}{(0.1 \cdot \mathbf{0.97} + 0.8 \cdot \mathbf{0.03})} = 0.2$$

…priors have been updated into posteriors!

From priors to posteriors

New priors, new estimations to be updated…

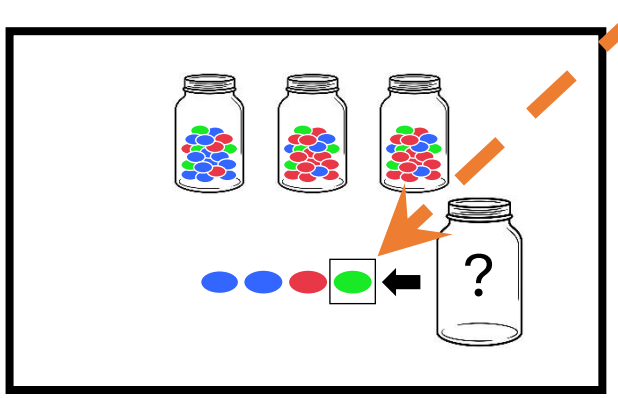

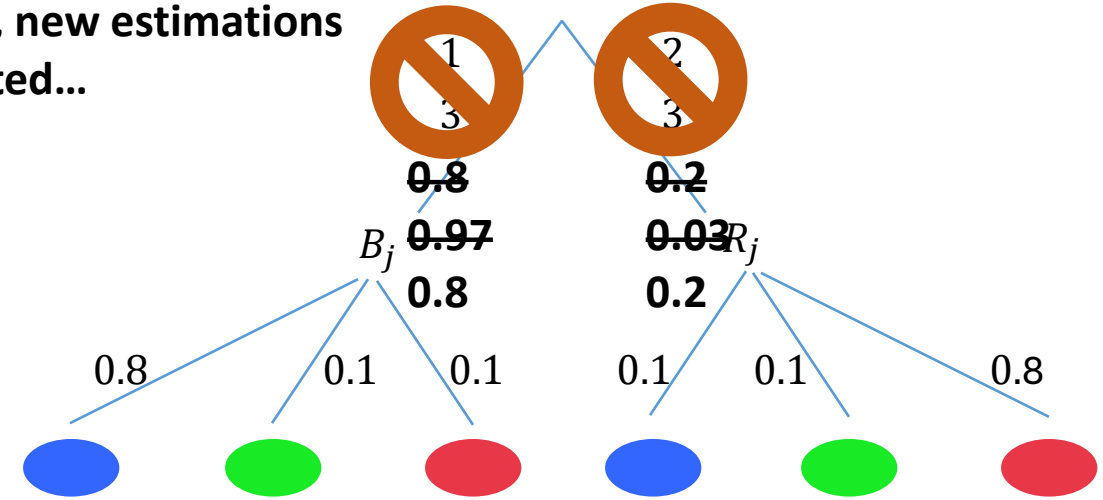

6. Posterior probability: the probability of each hypothesis, given the data (marbles extracted=blue + blue + red + green).

$$P(B_j \mid g) = \frac{P(g \mid B_j)\,P(B_j)}{P(g)} = \frac{0.1 \cdot \mathbf{0.8}}{(0.1 \cdot \mathbf{0.8} + 0.1 \cdot \mathbf{0.2})} = 0.8$$

$$P(R_j \mid g) = \frac{P(r \mid R_j)\,P(R_j)}{P(g)} = \frac{0.1 \cdot \mathbf{0.2}}{(0.1 \cdot \mathbf{0.8} + 0.1 \cdot \mathbf{0.2})} = 0.2$$

…priors have been updated into posteriors!

From priors to posteriors

New priors, new estimations to be updated…

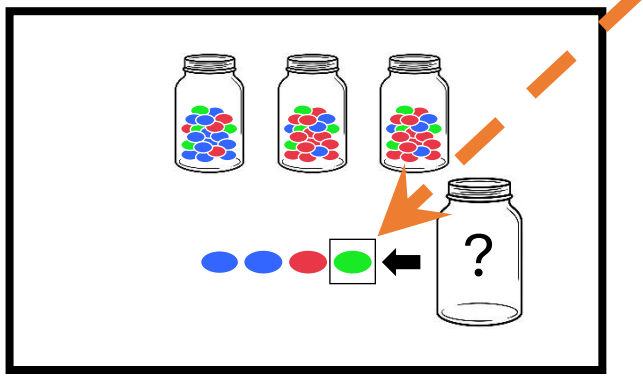

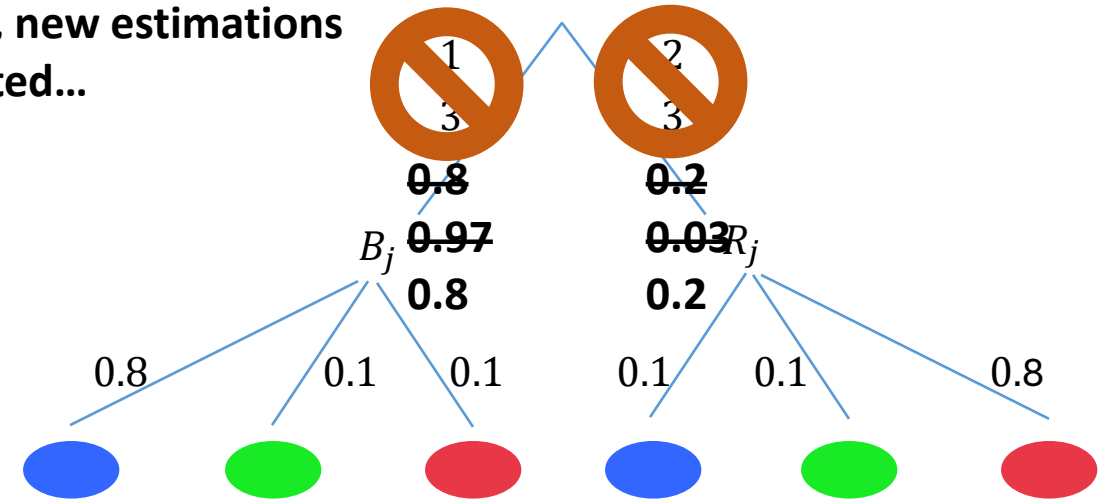

6. Posterior probability: the probability of each hypothesis, given the data (marbles extracted=blue + blue + red + green).

$$P(B_j \mid g) = \frac{P(g \mid B_j)\,P(B_j)}{P(g)} = \frac{0.1 \cdot \mathbf{0.8}}{(0.1 \cdot \mathbf{0.8} + 0.1 \cdot \mathbf{0.2})} = 0.8$$

$$P(R_j \mid g) = \frac{P(r \mid R_j)\,P(R_j)}{P(g)} = \frac{0.1 \cdot \mathbf{0.2}}{(0.1 \cdot \mathbf{0.8} + 0.1 \cdot \mathbf{0.2})} = 0.2$$

Priors=Posteriors

Data can be meaningless!

…priors have been updated into posteriors!

To be noted

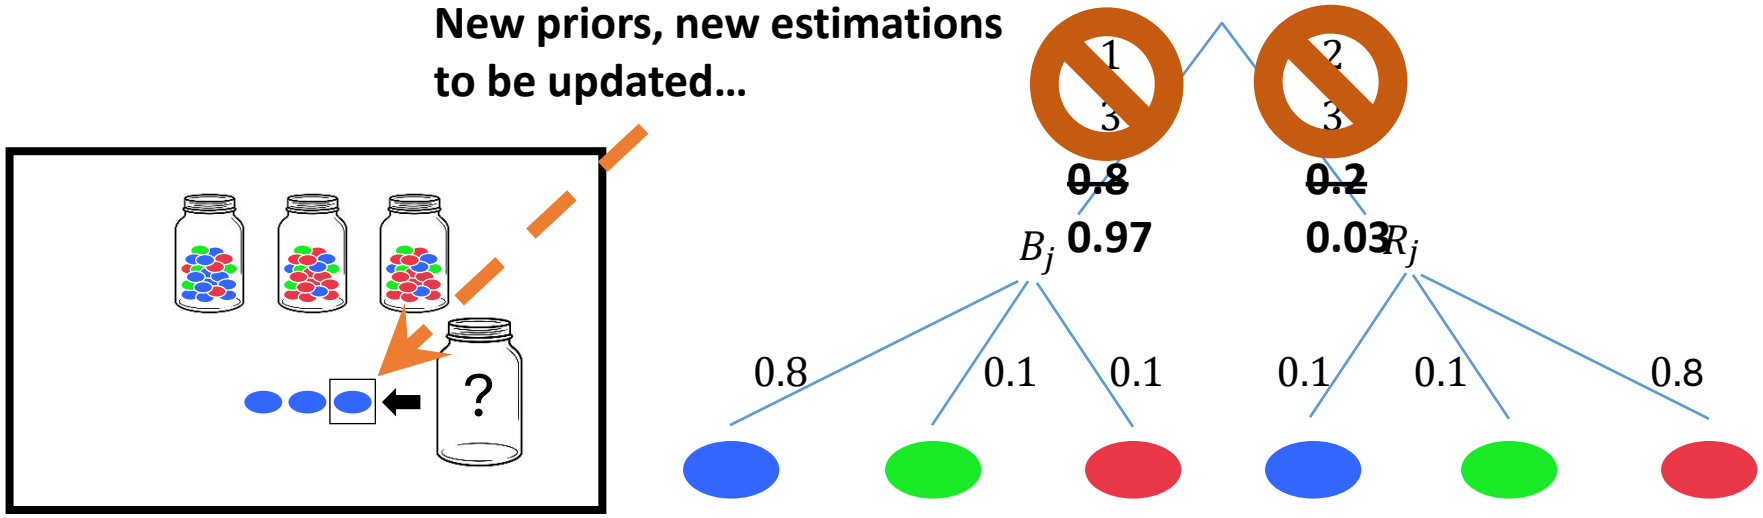

6. Posterior probability: the probability of each hypothesis, given the data (marbles extracted=blue + blue + blue).

$$P(B_j \mid b) = \frac{P(b \mid B_j)\,P(B_j)}{P(b)} = \frac{0.8 \cdot \mathbf{0.97}}{(0.8 \cdot \mathbf{0.97} + 0.1 \cdot \mathbf{0.03})} = 0.9961$$

$$P(R_j \mid b) = \frac{P(b \mid R_j)\,P(R_j)}{P(b)} = \frac{0.1 \cdot \mathbf{0.03}}{(0.8 \cdot \mathbf{0.97} + 0.1 \cdot \mathbf{0.03})} = 0.0039$$

To be noted

New priors, new estimations to be updated…

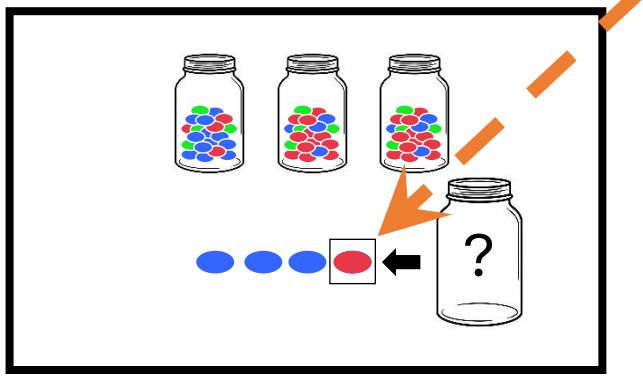

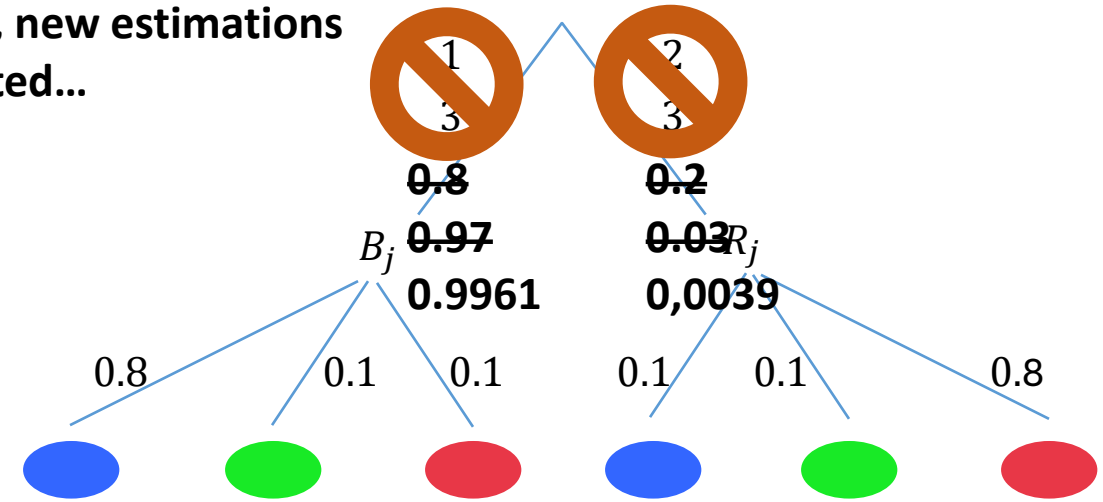

6. Posterior probability: the probability of each hypothesis, given the data (marbles extracted=blue + blue + blue + red).

$$P(B_j \mid r) = \frac{P(r \mid B_j)\,P(B_j)}{P(r)} = \frac{0.1 \cdot \mathbf{0.9961}}{(0.1 \cdot \mathbf{0.9961} + 0.8 \cdot \mathbf{0.0039})} = 0.9696$$

$$P(R_j \mid r) = \frac{P(r \mid R_j)\,P(R_j)}{P(r)} = \frac{0.8 \cdot \mathbf{0.0039}}{(0.1 \cdot \mathbf{0.9961} + 0.8 \cdot \mathbf{0.0039})} = 0.0304$$

High precision of prior beliefs leads to downplaying incoming sensory input.

Questions?

- A change of perspective: reward prediction errors and errors of prediction without rewards.
- Basic elements of Bayesian inference
  (A *gentle* introduction)
  1. Elements of probability theory
  2. Belief updating with discrete evidence
- Phenotypes and hierarchies.

Next…

- A change of perspective: reward prediction errors and errors of prediction without rewards.
- Basic elements of Bayesian inference (A *gentle* introduction)
    1. Elements of probability theory
    2. Belief updating with discrete evidence
- Phenotypes and hierarchies (A gentle introduction?).

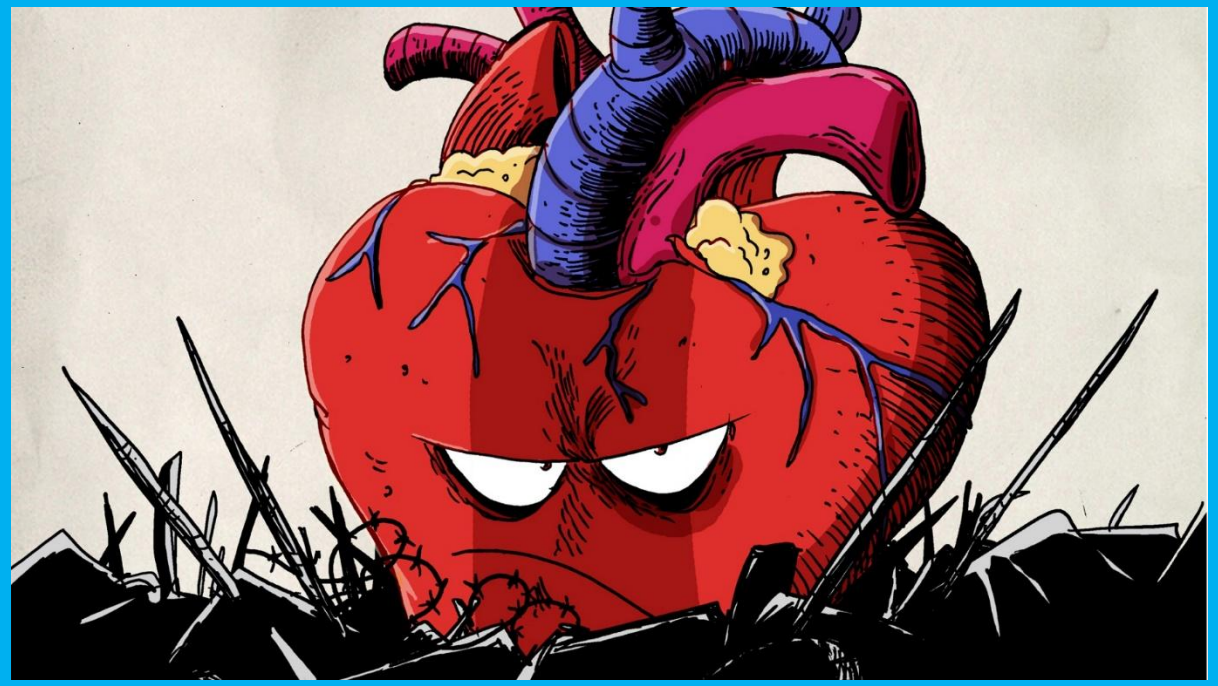

Modelling principles in action: phenotypes

$$logit(posterior)_d = \omega_p \cdot logit(posterior)_{d-1} + \omega_l \cdot LLR$$

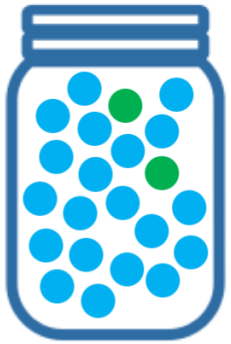

Leaky integration (neuroscience)
Base-rate neglect (economics)

Uscher and McClelland, Psychol Rev 2001

Modelling principles in action: phenotypes

$$logit(posterior)_d = \omega_p \cdot logit(posterior)_{d-1} + \omega_l \cdot LLR$$

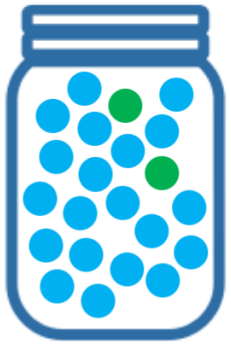

Subjects (e.g., clinical population vs healthy control) can differ in terms of the **relevance (weight) assigned to the prior** (ω_p) **or the likelihood** (ω_l).

Uscher and McClelland, Psychol Rev 2001

Modelling principles in action: phenotypes

Posterior

Prior

$$P(H|D) = \frac{P(D|H)P(H)}{P(D)}$$

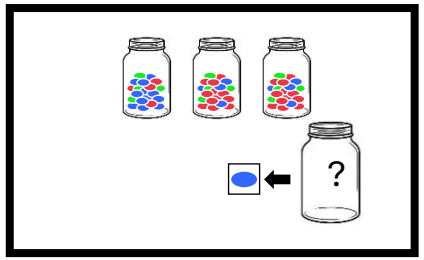

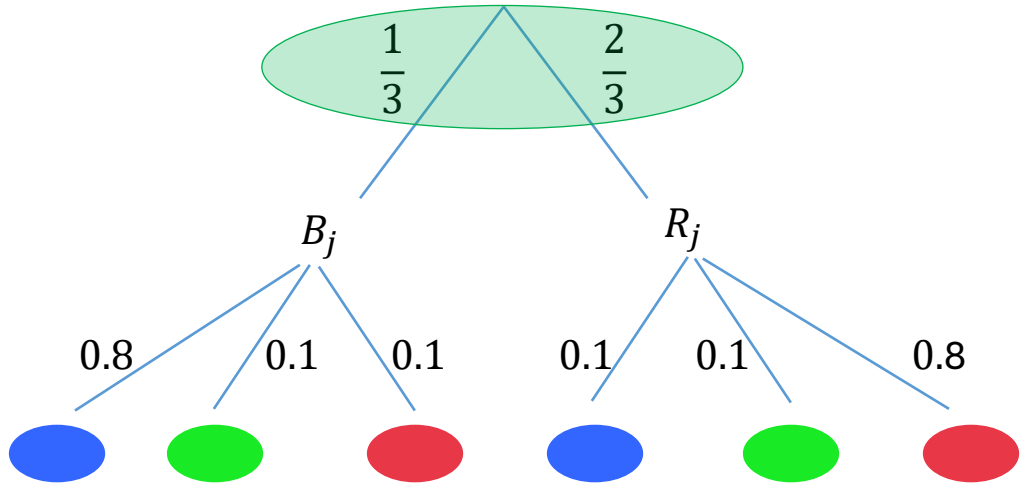

Modelling principles in action: phenotypes

Likelihood Prior

Posterior

$$P(H|D) = \frac{P(D|H)P(H)}{P(D)}$$

Total probability of the data to occur

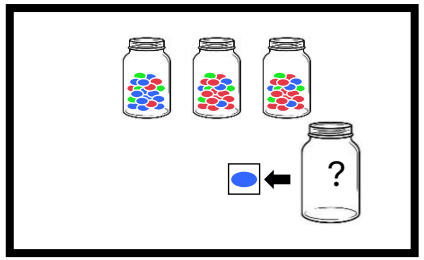

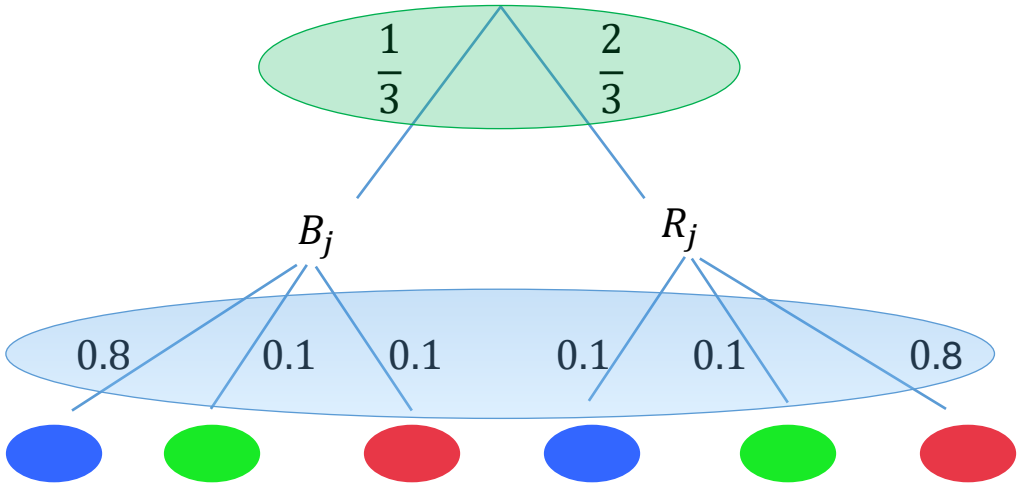

What if the likelihood of the data, given the hypothesis, is also unknown?

Modelling principles in action: phenotypes

Bayesian modelling usually assumes that the beliefs of biological agents are normally distributed.
Thus, subjects differ depending on their assumptions about how the events are (normally) distributed.

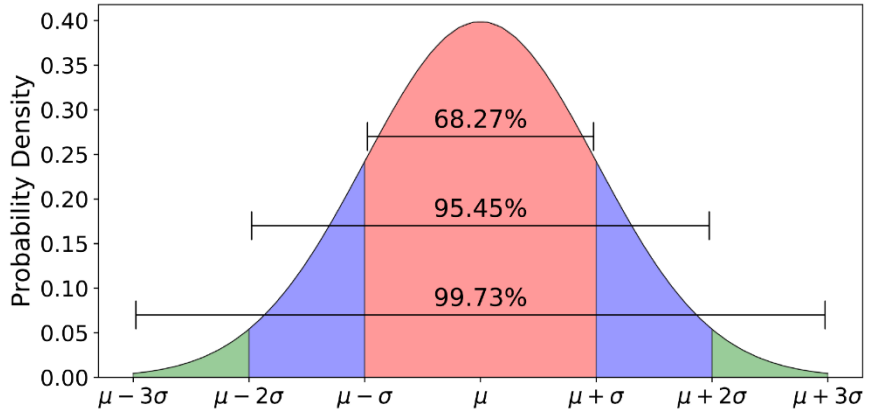

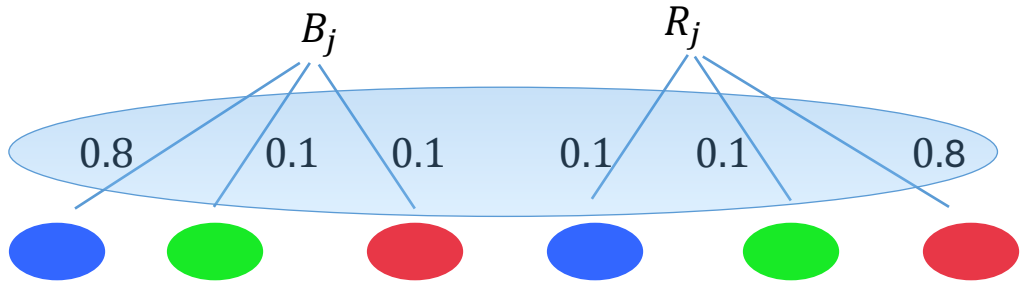

What if the likelihood of the data,
given the hypothesis, is also unknown?

Modelling principles in action: phenotypes

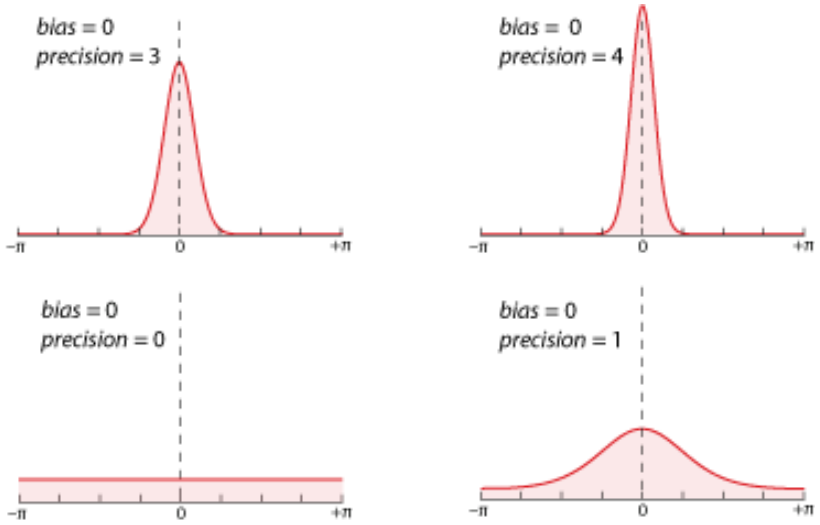

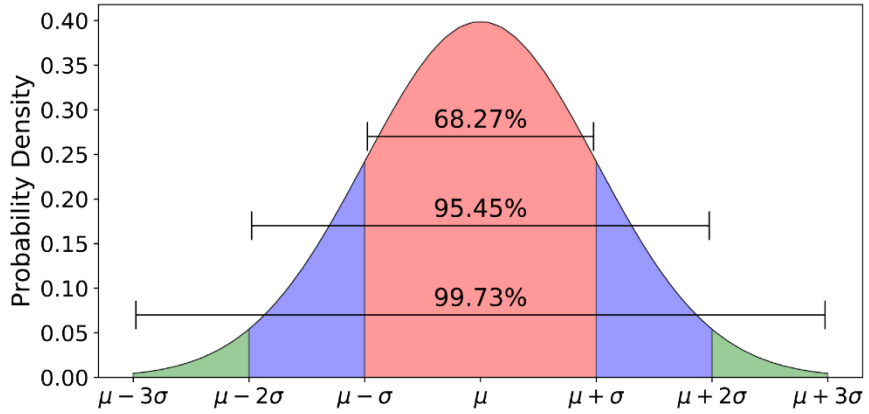

Different assumptions concerning these distributions of events result in different paces for the belief update.

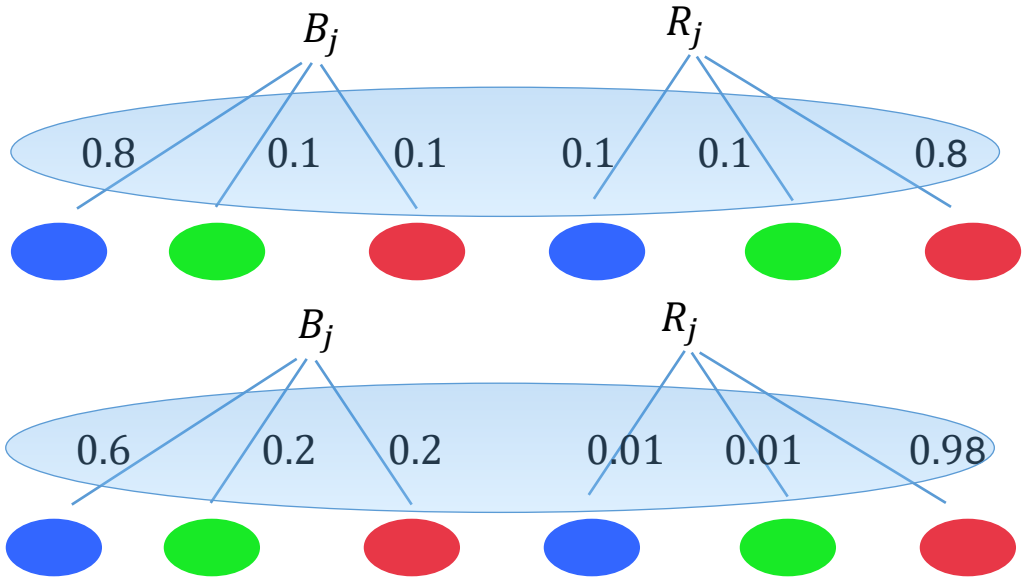

Modelling principles in action: phenotypes

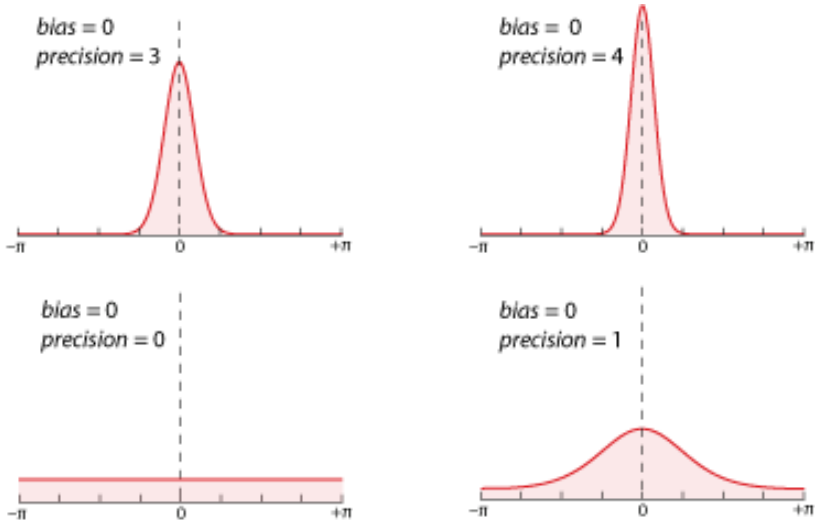

Cumulative distribution function

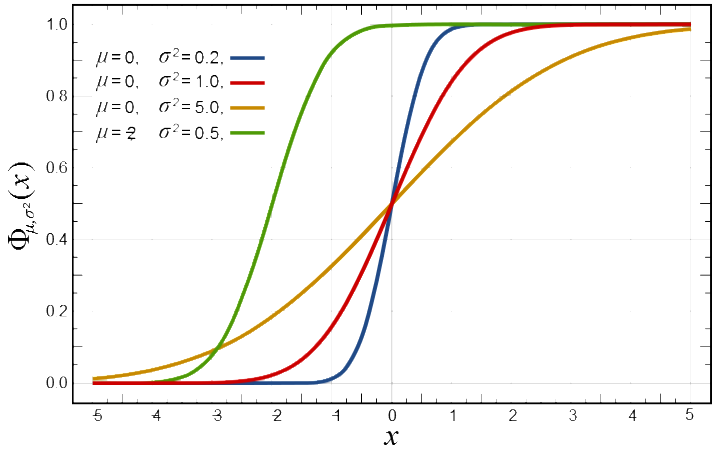

Different assumptions concerning these distributions of events result in different paces for the belief update.

Quasi-deterministic assumptions (high precision) entail fast updates, whereas assuming high stochasticity (low precision) entails slow updates.

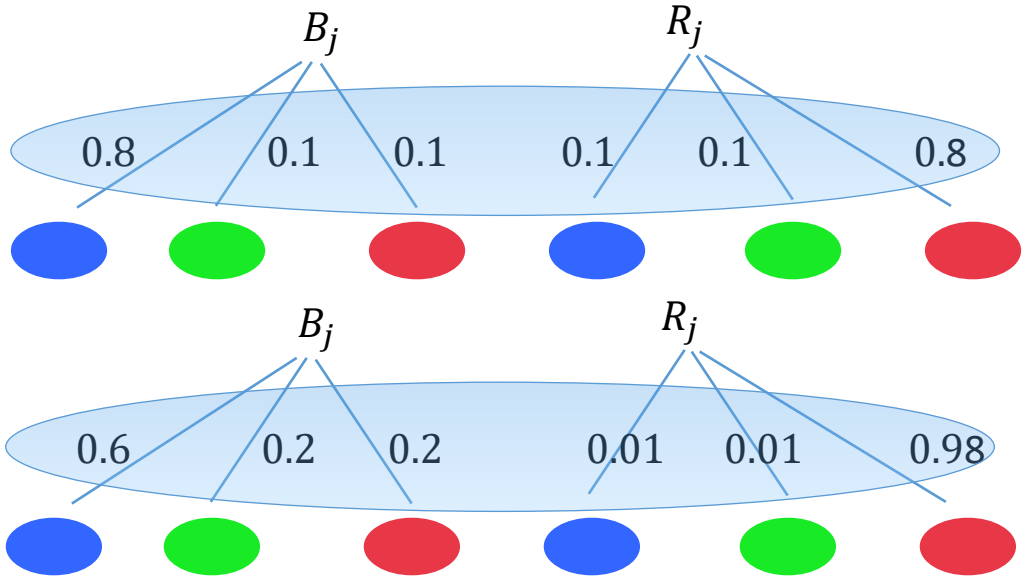

Modelling principles in action: phenotypes

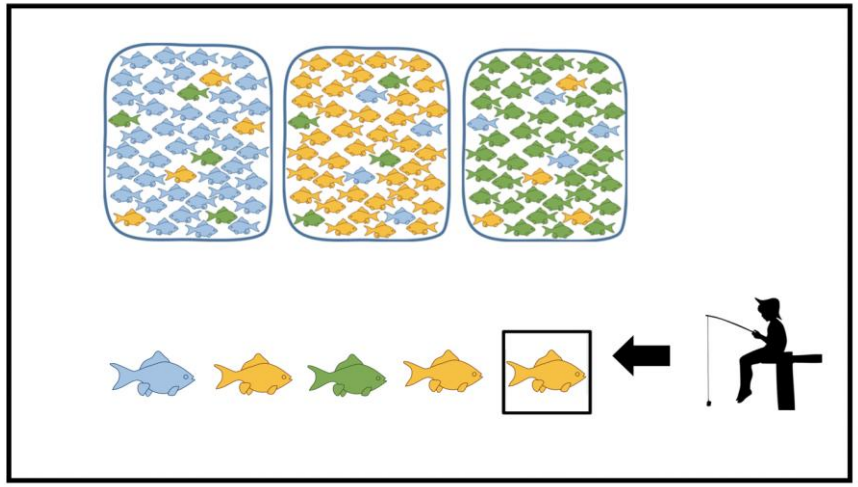

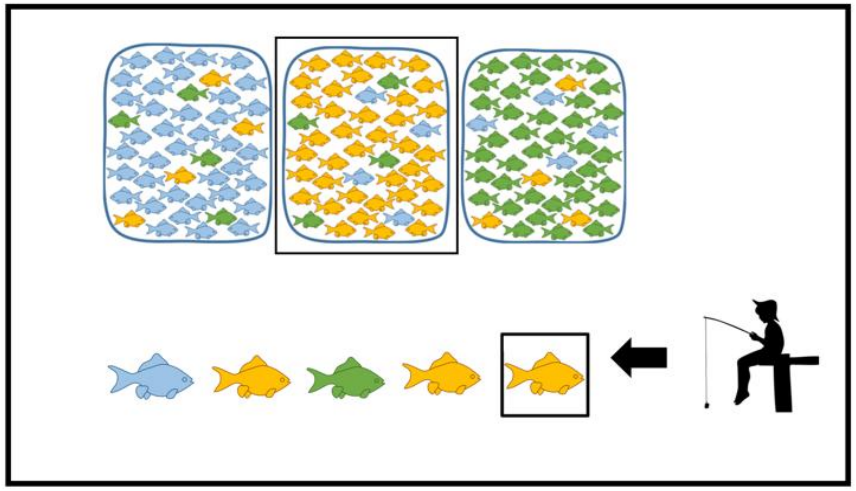

Probabilistic perceptual task

- No feedback provided after choice selections.
- 10 blocks.
- 1 or 2 reversals per block. 15 trials per block.

Modelling principles in action: phenotypes

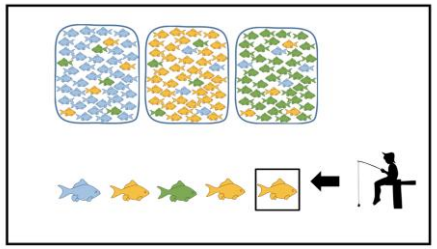

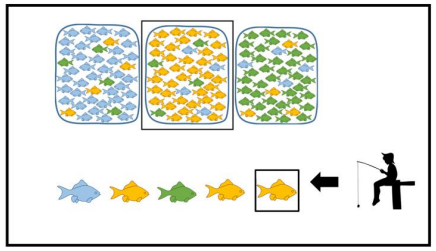

Likelihood Prior

Posterior

$$P(H|D) = \frac{P(D|H)P(H)}{P(D)}$$

Total probability of the data to occur

N=100

Likelihood: $0.9<\lambda<1$

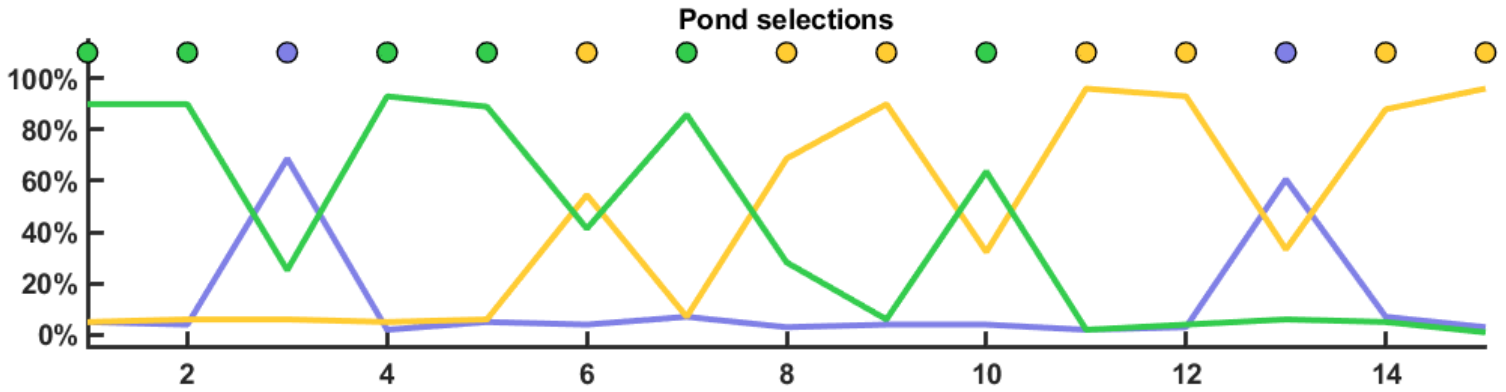

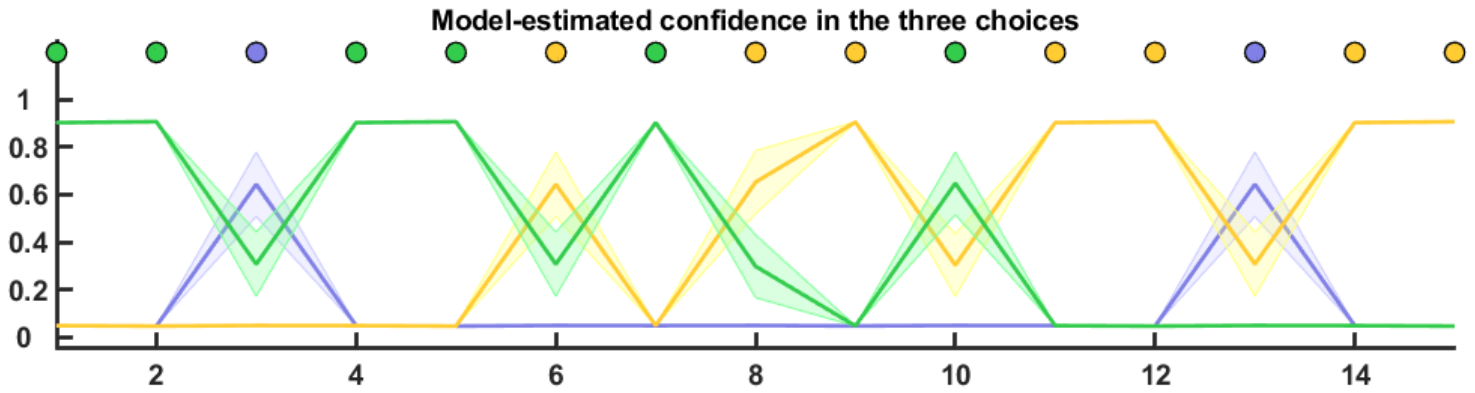

Modelling principles in action: phenotypes

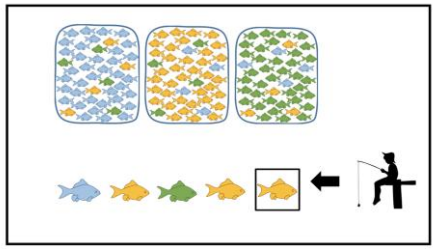

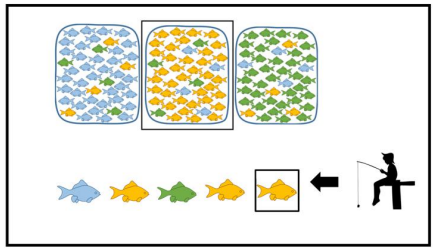

Posterior Likelihood Prior

$$P(H|D) = \frac{P(D|H)P(H)}{P(D)}$$

Total probability of the data to occur

N=100

Likelihood: $0.8<\lambda<0.9$

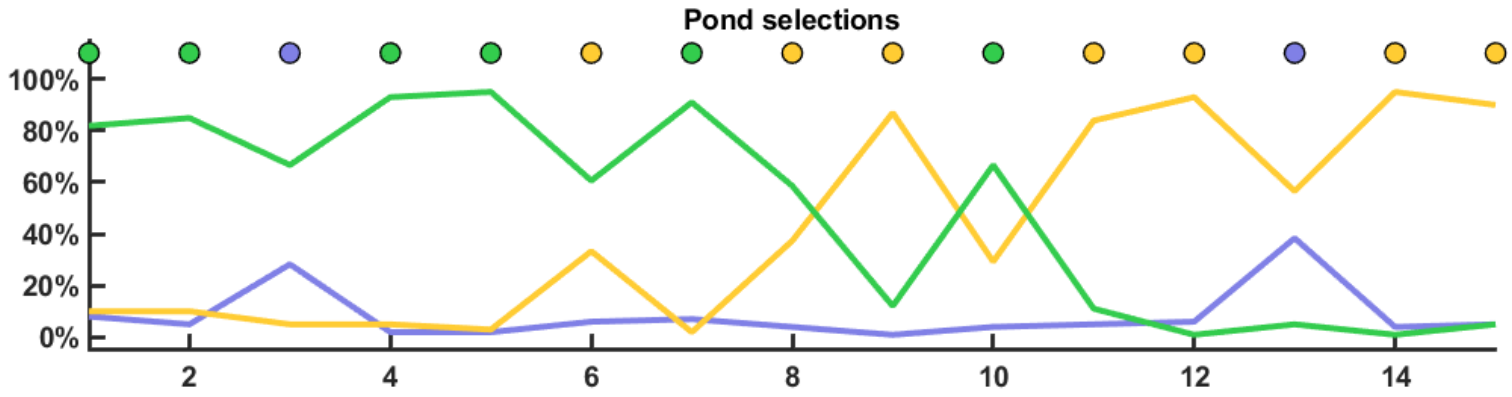

Modelling principles in action: phenotypes

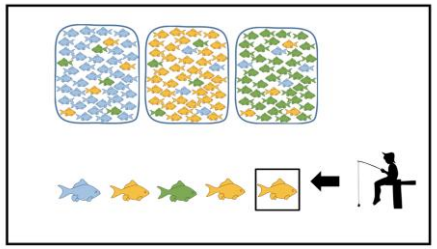

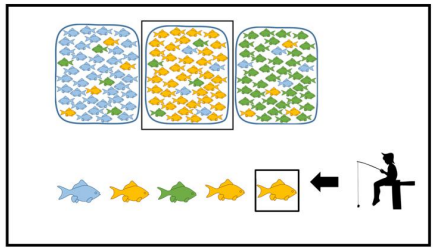

Posterior

$$P(H|D) = \frac{P(D|H)\,P(H)}{P(D)}$$

Likelihood Prior

Total probability of the data to occur

N=100

Likelihood: $0.7<\lambda<0.8$

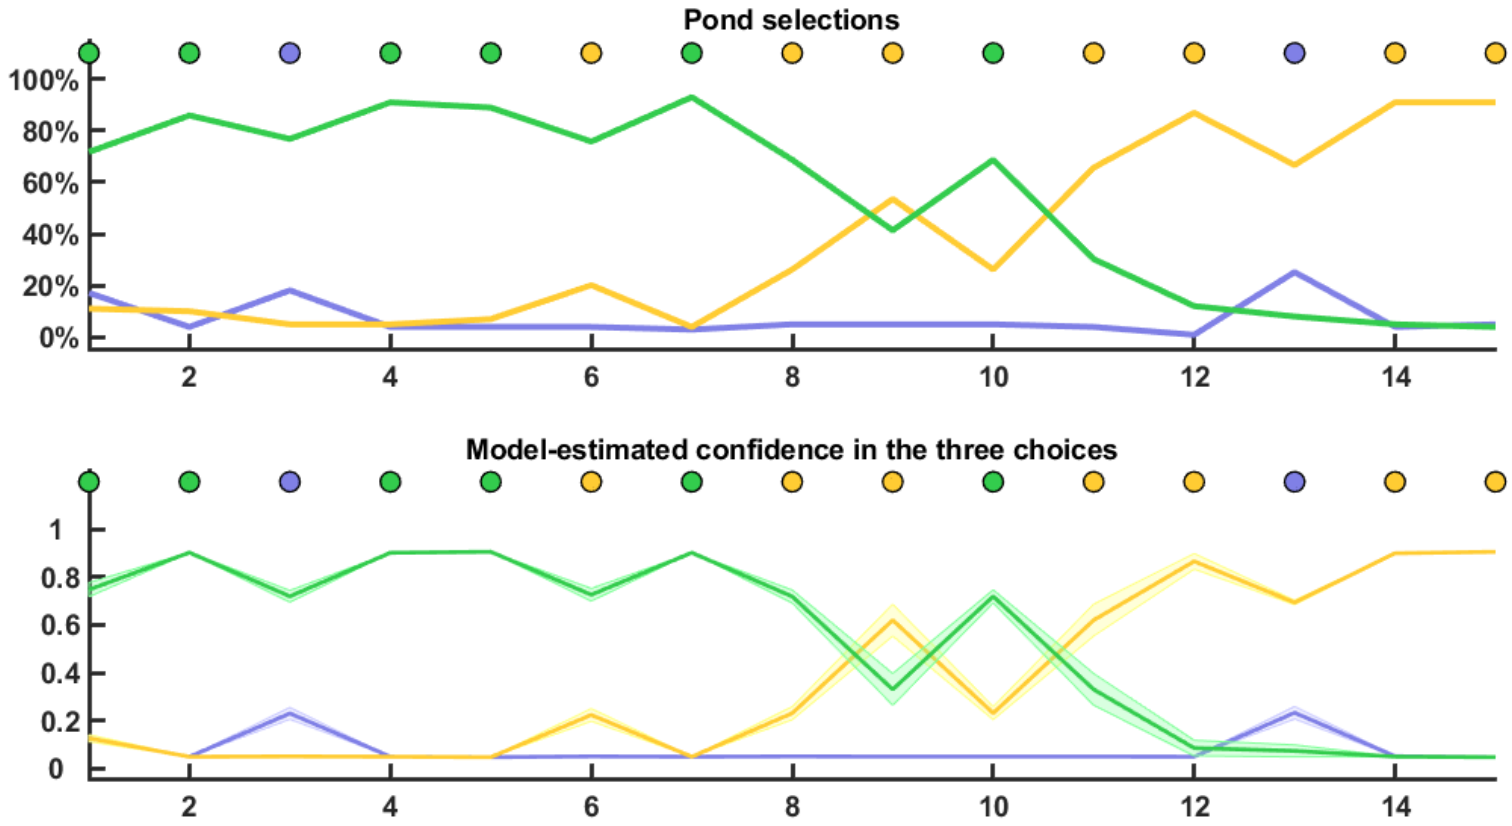

Modelling principles in action: phenotypes

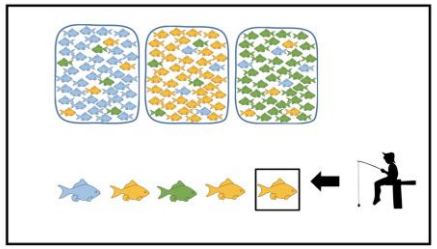

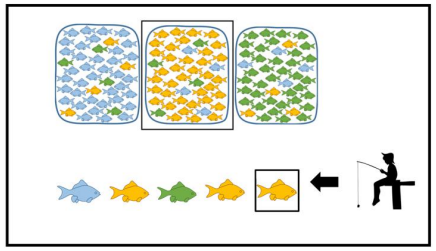

Posterior Likelihood Prior

$$P(H|D) = \frac{P(D|H)P(H)}{P(D)}$$

Total probability of the data to occur

N=100

Likelihood: $0.6<\lambda<0.7$

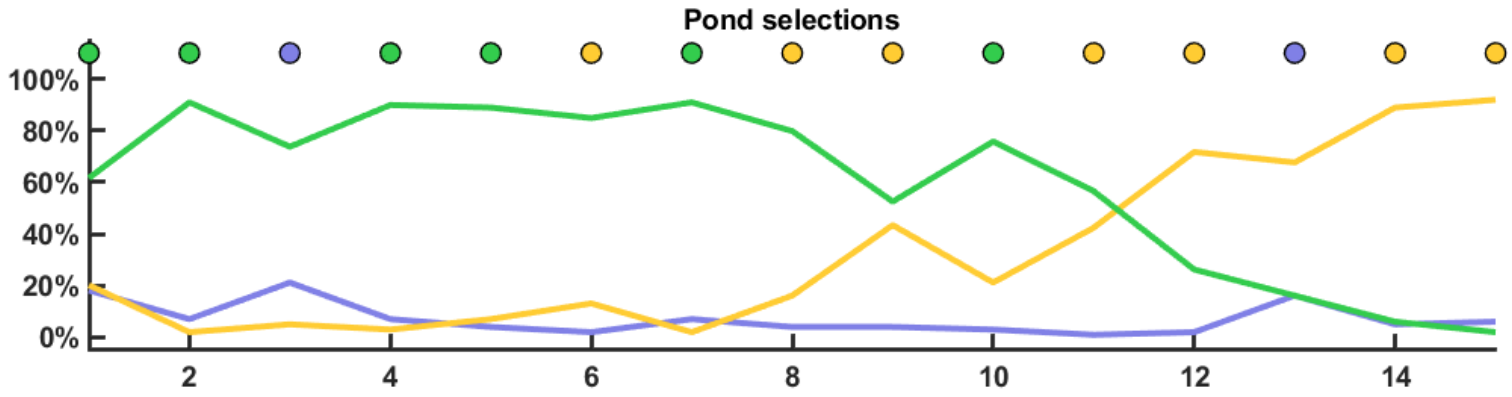

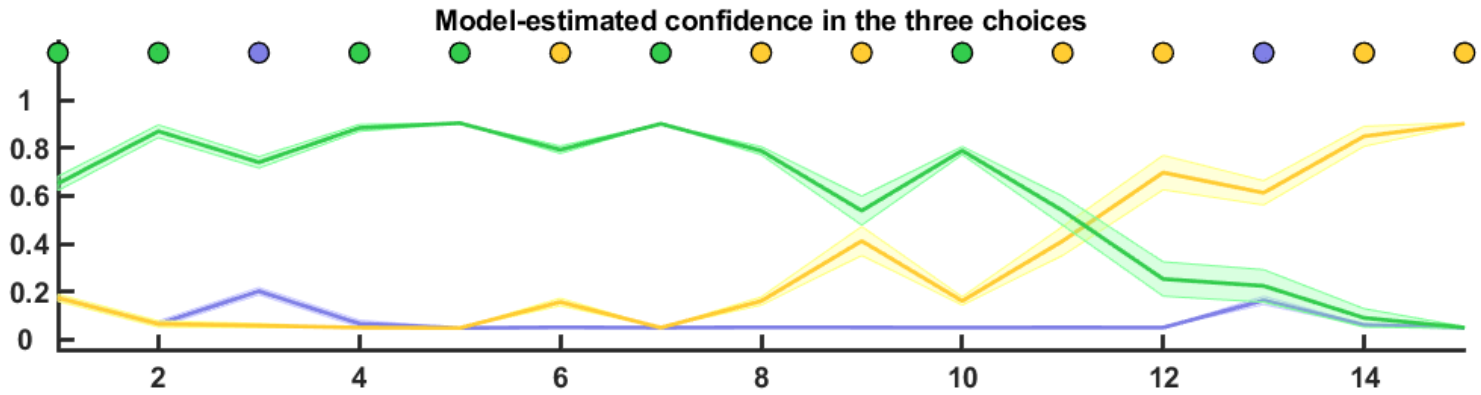

Modelling principles in action: phenotypes

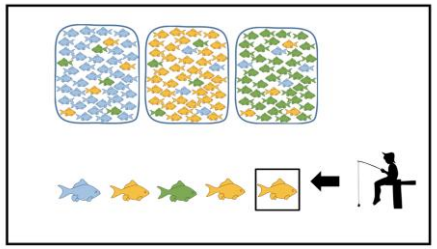

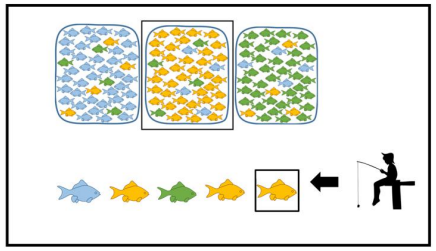

Posterior Likelihood Prior

$$P(H|D) = \frac{P(D|H)P(H)}{P(D)}$$

Total probability of the data to occur

N=100

Likelihood: $0.5<\lambda<0.6$

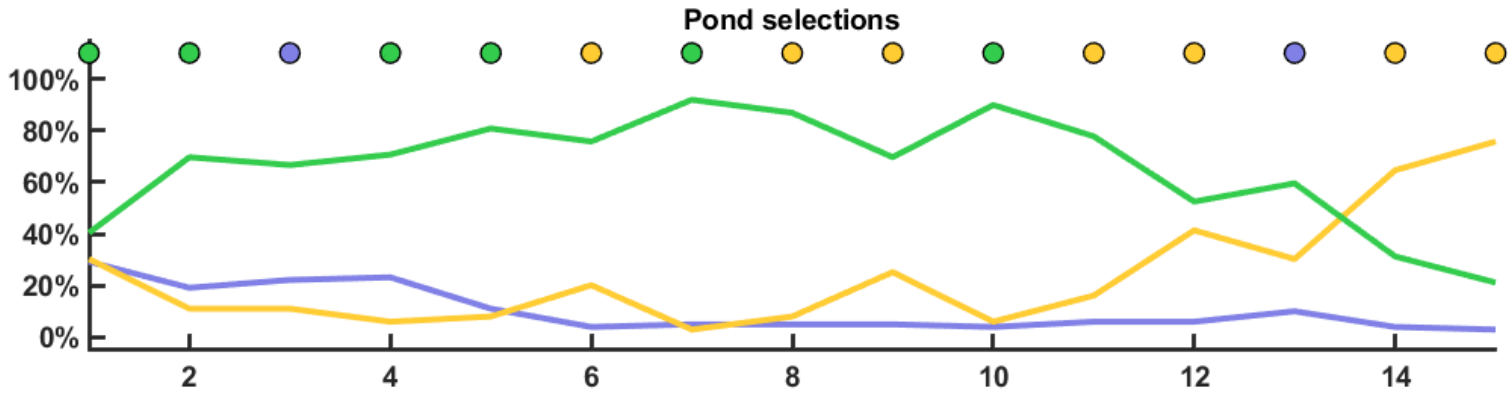

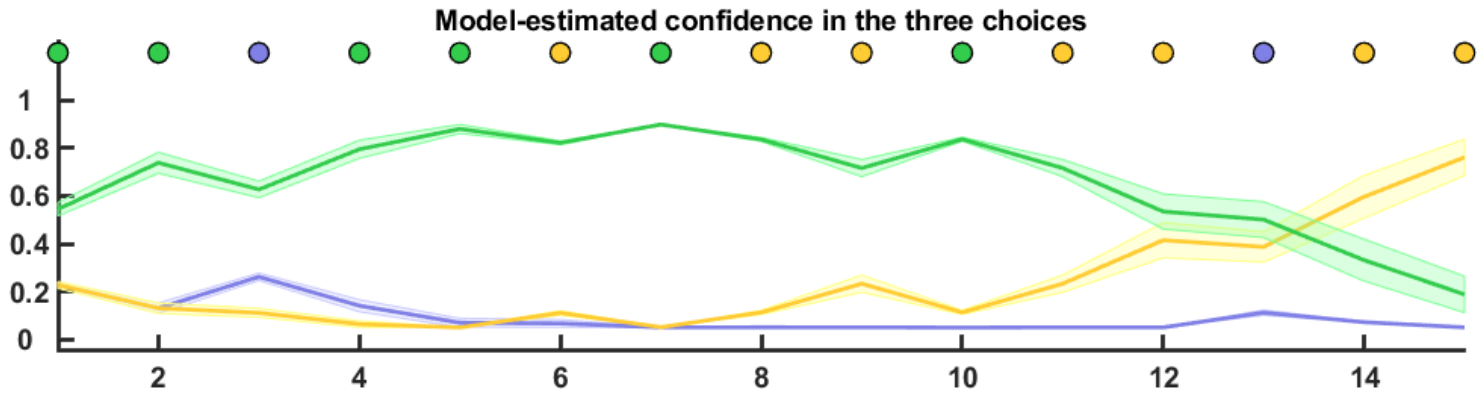

Modelling principles in action: phenotypes

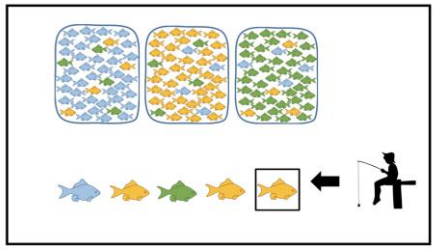

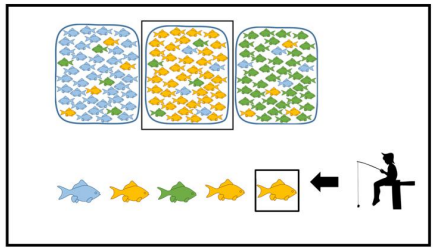

Likelihood Prior

Posterior

$$P(H|D) = \frac{P(D|H)P(H)}{P(D)}$$

Total probability of the data to occur

N=100

Likelihood: $0.4<\lambda<0.5$

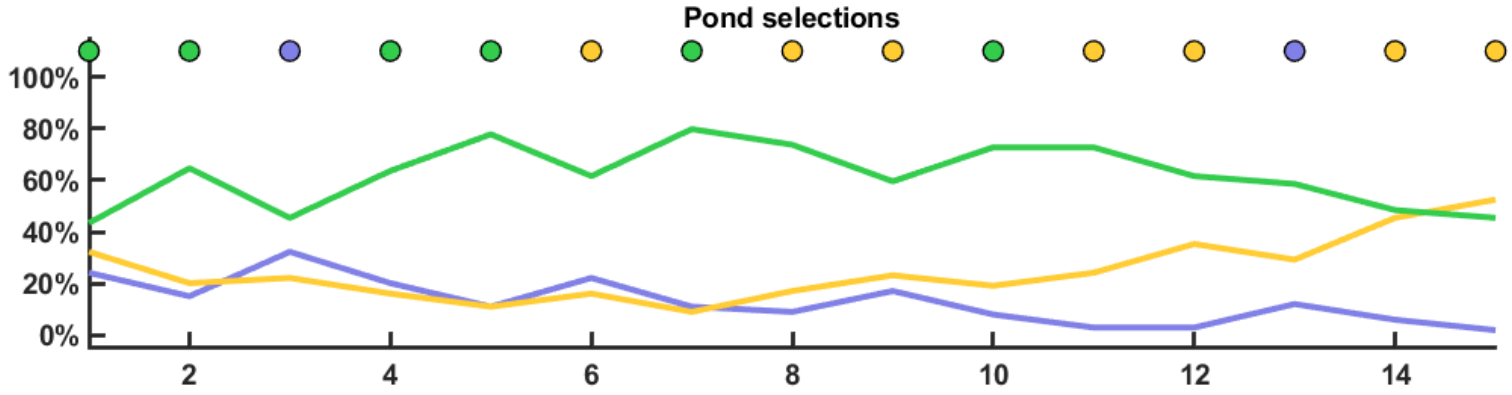

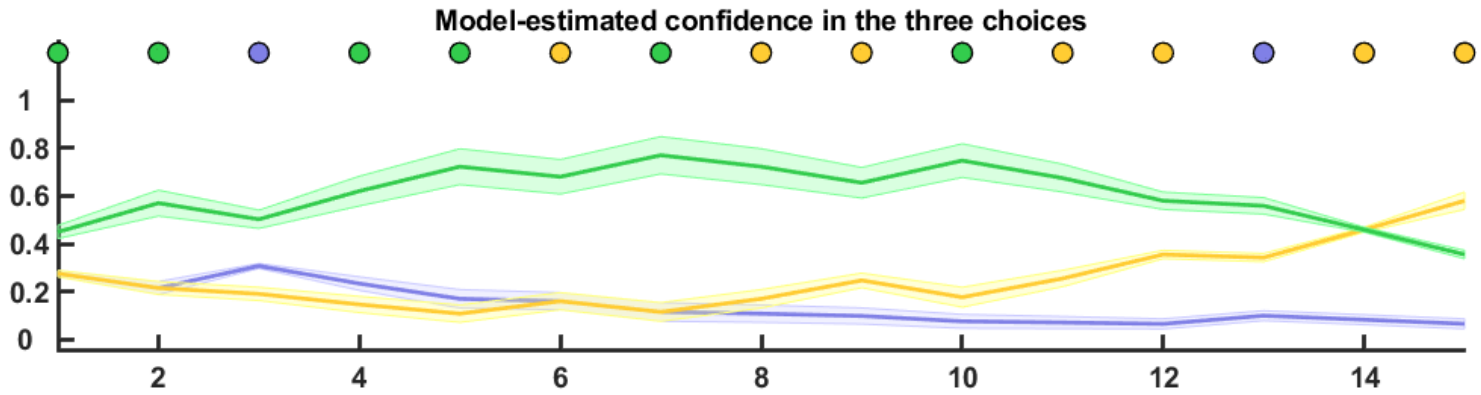

Modelling principles in action: hierarchies

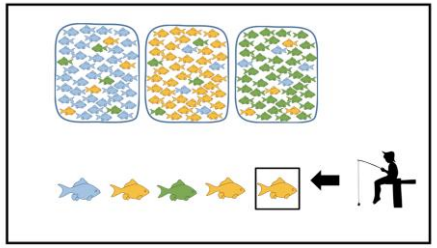

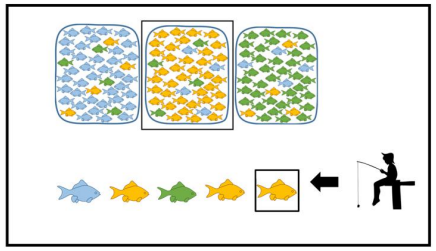

Posterior $P(H|D)$ = Likelihood $P(D|H)$ Prior $P(H)$ / Total probability of the data to occur $P(D)$

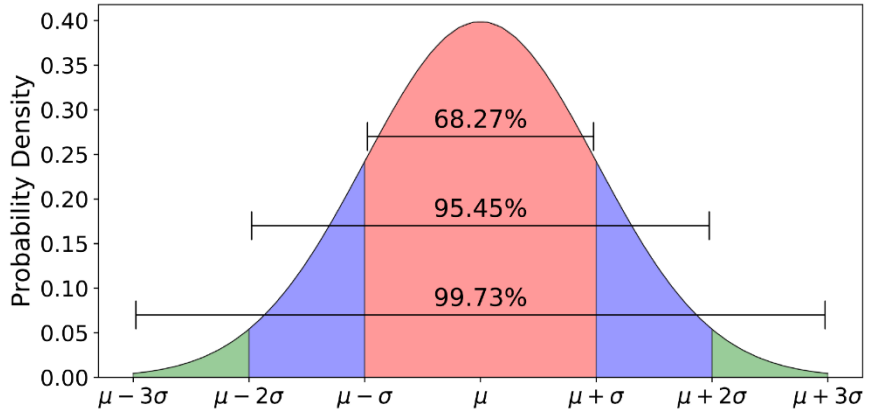

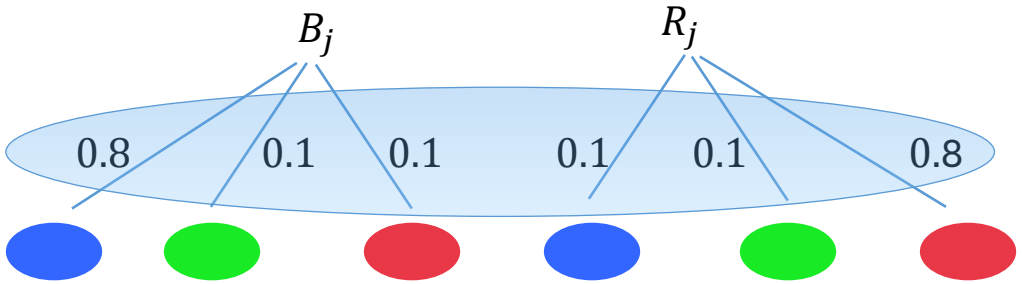

What if the likelihood of the data,
given the hypothesis, is both *unknown* and *varies over time*?

Modelling principles in action: hierarchies

nature neuroscience

Learning the value of information in an uncertain world

Timothy E J Behrens[1,2], Mark W Woolrich[1], Mark E Walton[2] & Matthew F S Rushworth[1,2]

Our decisions are guided by outcomes that are associated with decisions made in the past. However, the amount of influence each past outcome has on our next decision remains unclear. To ensure optimal decision-making, the weight given to decision outcomes should reflect their salience in predicting future outcomes, and this salience should be modulated by the volatility of the reward environment. We show that human subjects assess volatility in an optimal manner and adjust decision-making accordingly. This optimal estimate of volatility is reflected in the fMRI signal in the anterior cingulate cortex (ACC) when each trial outcome is observed. When a new piece of information is witnessed, activity levels reflect its salience for predicting future outcomes. Furthermore, variations in this ACC signal across the population predict variations in subject learning rates. Our results provide a formal account of how we weigh our different experiences in guiding our future actions.

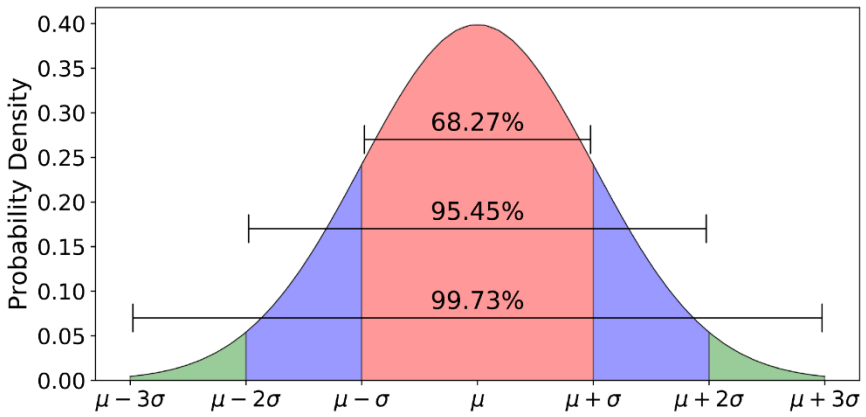

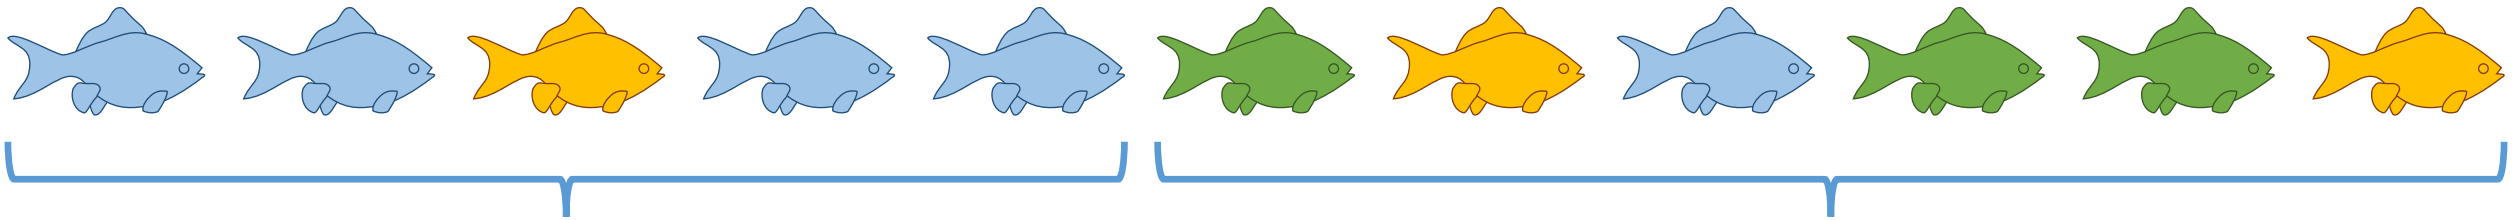

Volatility: 80%-10%-10%

Volatility: 50%-25%-25%

Modelling principles in action: hierarchies

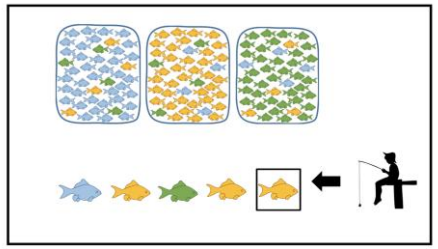

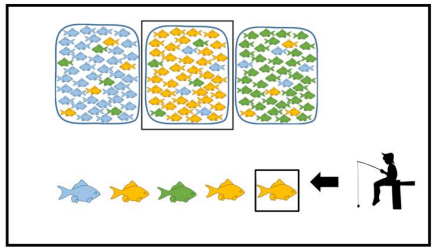

Environment
(generative process)

Inferred state
(generative model)

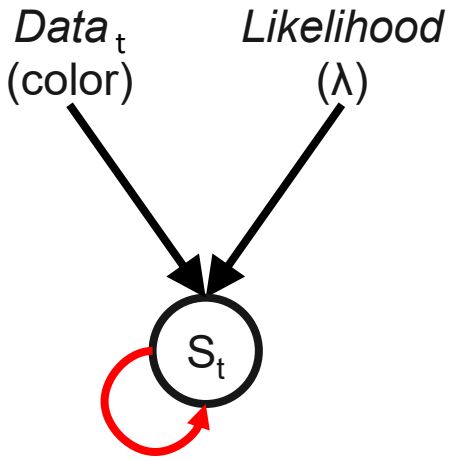

$P(\text{pond} \mid \text{color}_t, \lambda, \text{prior})_t$

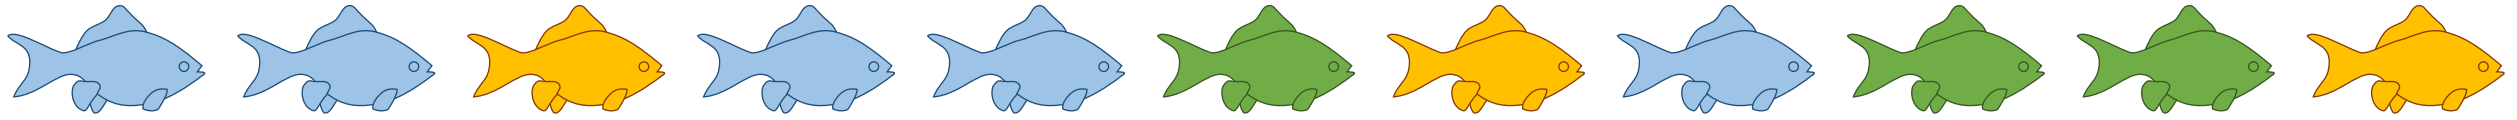

Volatility: 80%-10%-10%

Volatility: 50%-25%-25%

Modelling principles in action: hierarchies

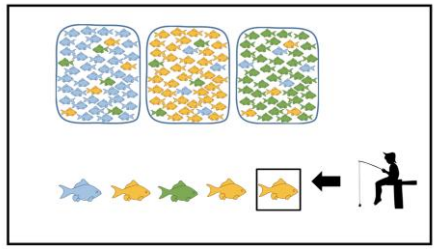

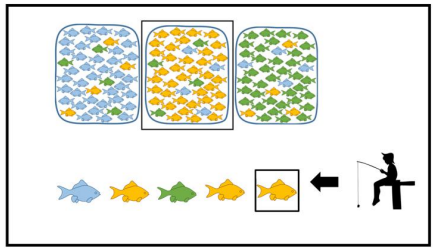

Environment
(generative process)

Inferred state
(generative model)

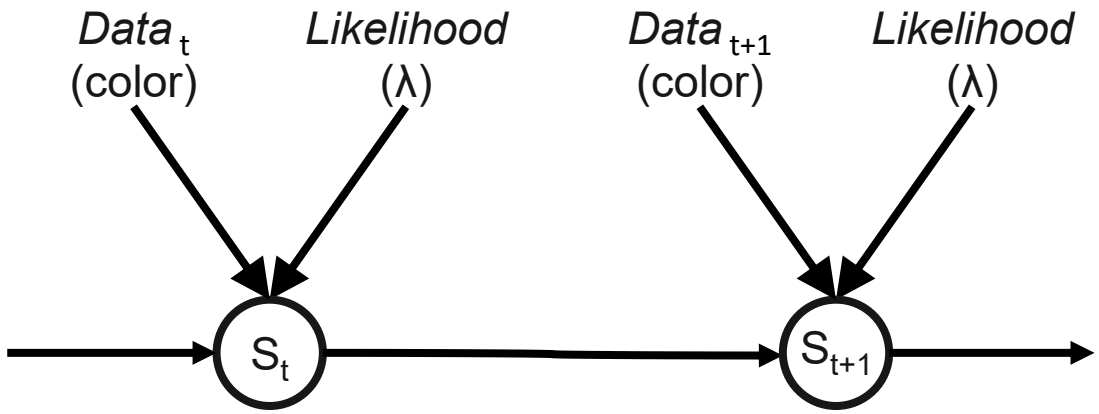

$P(\text{pond}|\text{color}_t, \lambda, \text{prior})_t$ $P(\text{pond}|\text{color}_{t+1}, \lambda, \text{prior})_{t+1}$

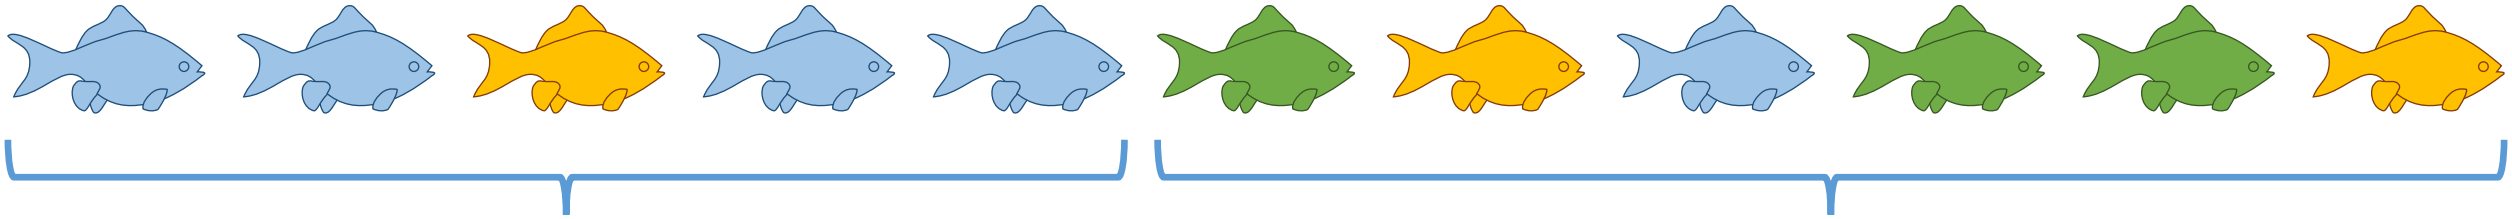

Volatility: 80%-10%-10%

Volatility: 50%-25%-25%

Modelling principles in action: hierarchies

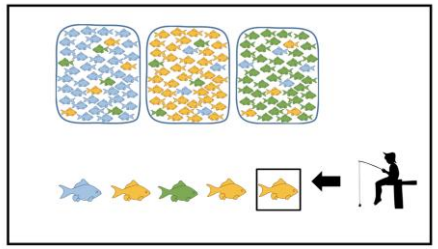

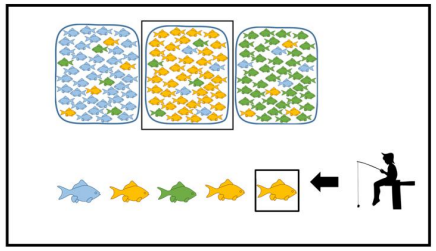

Environment
(generative process)

Inferred state
(generative model)

$\lambda_t = f(\lambda_0, Data_{1:t})$

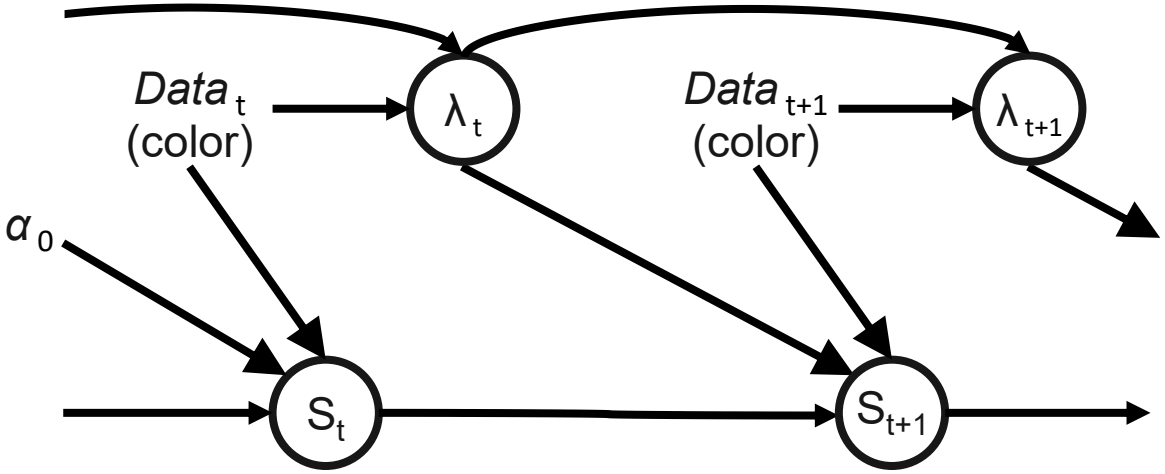

$P(pond|color_t, \lambda_t, prior)_t$ $P(pond|color_{t+1}, \lambda_{t+1}, prior)_{t+1}$

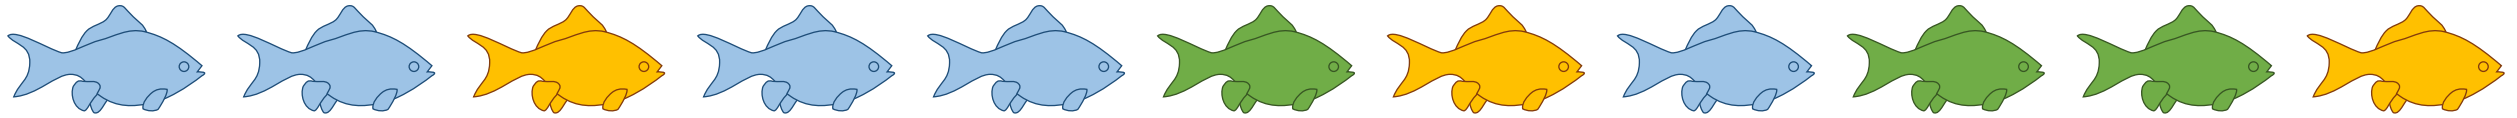

Volatility: 80%-10%-10%

Volatility: 50%-25%-25%

Modelling principles in action: hierarchies

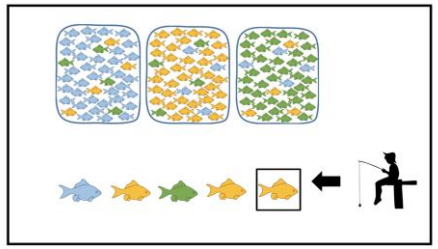

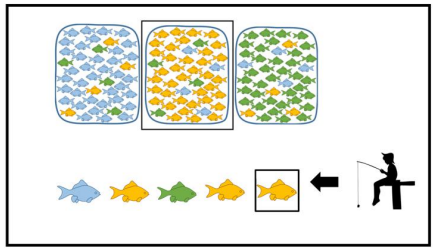

$\lambda_t = f(\alpha_0, k, Data_{1:t})$

Dirichlet update:

$$\lambda_t^{(i)} = \frac{\alpha_0^{(i)} + k\, data_{1:t}^{(i)}}{\sum(\alpha_0^{(i)} + k\, data_{1:t})}$$

e.g.:

$$\alpha_0^{(blue\ pond)} = [8\ 1\ 1]; \qquad data_{1:5} = [4\ 1\ 0]; \qquad k = 0.5;$$

$$\lambda_6^{(blue\ pond)} = \frac{[10\ 1.5\ 1]}{12.5} = [0.8\ 0.12\ 0.08]$$

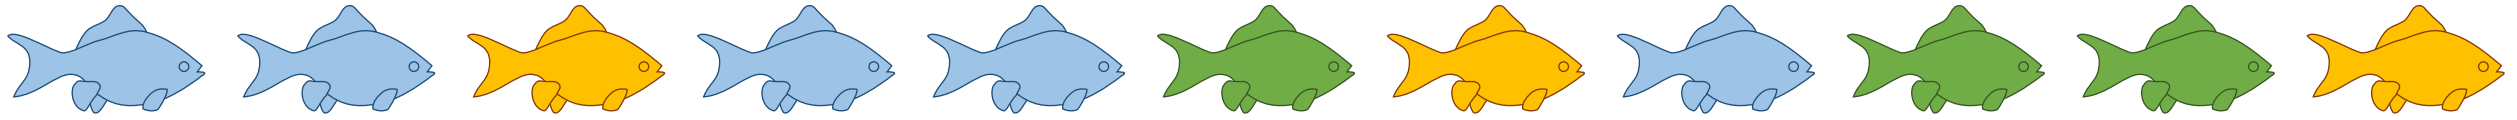

Volatility: 80%-10%-10%

Volatility: 50%-25%-25%

Modelling principles in action: hierarchies

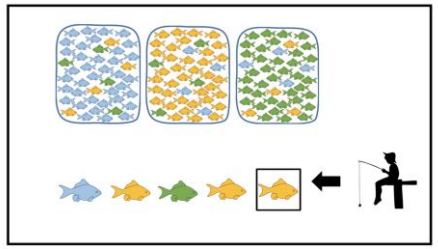

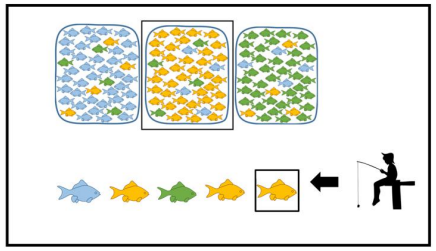

$\lambda_t = f(\alpha_0, k, Data_{1:t})$

Dirichlet update:

$$\lambda_t^{(i)} = \frac{\alpha_0^{(i)} + k\, data_{1:t}^{(i)}}{\sum(\alpha_0^{(i)} + k\, data_{1:t})}$$

e.g.:

$$\alpha_0^{(yellow\ pond)} = [1\ 8\ 1]; \qquad data_{1:5} = [4\ 1\ 0]; \qquad k = 0.5;$$

$$\lambda_6^{(blue)} = \frac{[10\ 1.5\ 1]}{12.5} = [0.8\ 0.12\ 0.08]; \qquad \lambda_6^{(yellow)} = \frac{[3\ 8.5\ 1]}{12.5} = [0.24\ 0.68\ 0.08]$$

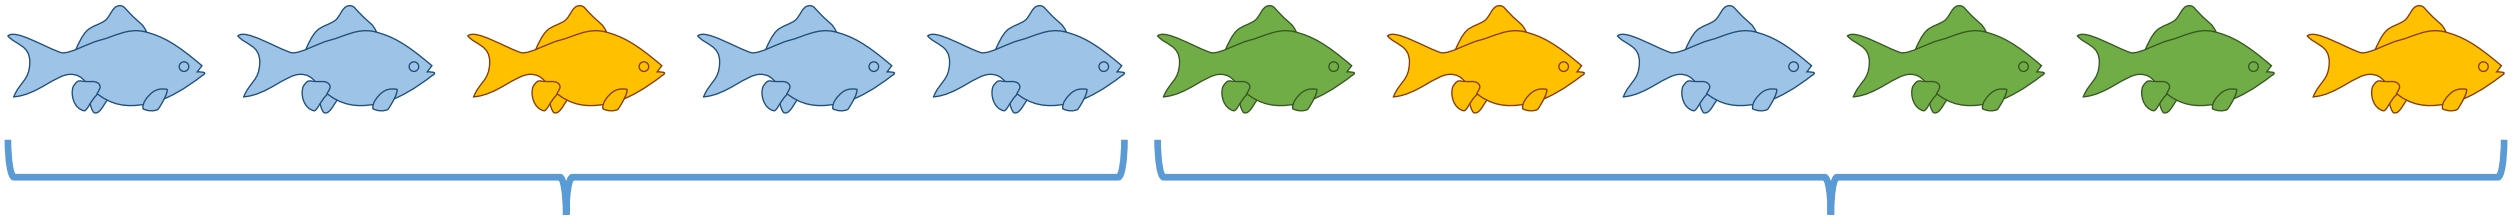

Volatility: 80%-10%-10%

Volatility: 50%-25%-25%

Modelling principles in action: hierarchies

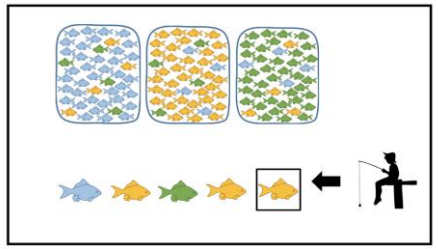

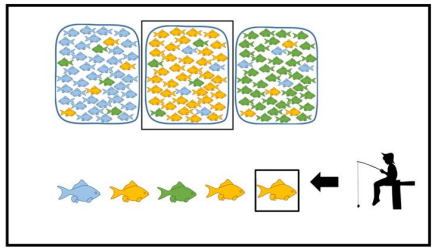

$\lambda_t = f(\alpha_0, k, Data_{1:t})$

Dirichlet update:

$$\lambda_t^{(i)} = \frac{\alpha_0^{(i)} + k\, data_{1:t}^{(i)}}{\sum(\alpha_0^{(i)} + k\, data_{1:t})}$$

e.g.:

$$\alpha_0^{(green\ pond)} = [1\ 1\ 8]; \qquad data_{1:5} = [4\ 1\ 0]; \qquad k = 0.5;$$

$$\lambda_6^{(blue)} = \frac{[10\ 1.5\ 1]}{12.5}; \ \lambda_6^{(yellow)} = \frac{[3\ 8.5\ 1]}{12.5}; \qquad \lambda_6^{(green)} = \frac{[3\ 1.5\ 8]}{12.5} = [0.24\ 0.12\ 0.64]$$

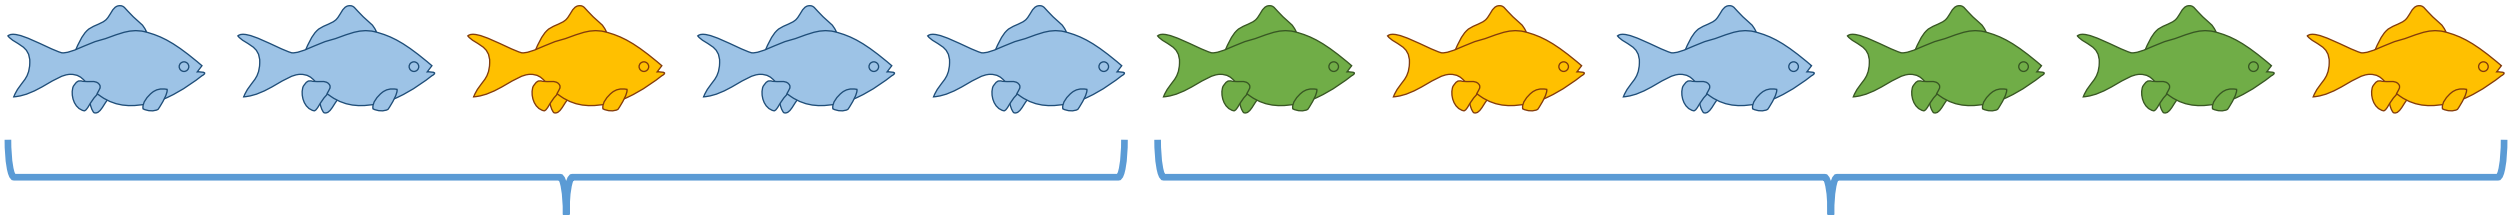

Volatility: 80%-10%-10%

Volatility: 50%-25%-25%

Modelling principles in action: hierarchies

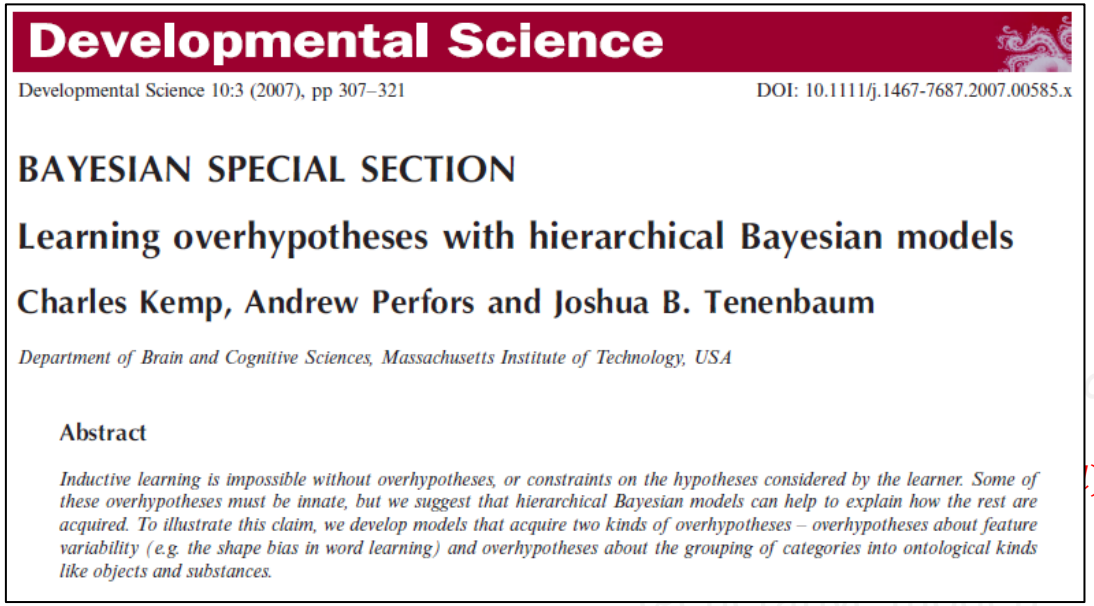

Developmental Science

Developmental Science 10:3 (2007), pp 307–321 DOI: 10.1111/j.1467-7687.2007.00585.x

BAYESIAN SPECIAL SECTION

Learning overhypotheses with hierarchical Bayesian models

Charles Kemp, Andrew Perfors and Joshua B. Tenenbaum

Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, USA

Abstract

Inductive learning is impossible without overhypotheses, or constraints on the hypotheses considered by the learner. Some of these overhypotheses must be innate, but we suggest that hierarchical Bayesian models can help to explain how the rest are acquired. To illustrate this claim, we develop models that acquire two kinds of overhypotheses – overhypotheses about feature variability (e.g. the shape bias in word learning) and overhypotheses about the grouping of categories into ontological kinds like objects and substances.

$\lambda_t = f(\alpha_0, k, Data_{1:t})$

$$\lambda_t^{(i)} = \frac{\alpha_0^{(i)} + k\, data_{1:t}^{(i)}}{\sum(\alpha_0^{(i)} + k\, data_{1:t})}$$

$= [1\ 1\ 8]; \quad data_{1:5} = [4\ 1\ 0]; \quad k = 0.5;$

$$\lambda_6^{(blue)} = \frac{[10\ 1.5\ 1]}{12.5}; \ \lambda_6^{(yellow)} = \frac{[3\ 8.5\ 1]}{12.5}; \quad \lambda_6^{(green)} = \frac{[3\ 1.5\ 8]}{12.5} = [0.24\ 0.12\ 0.64]$$

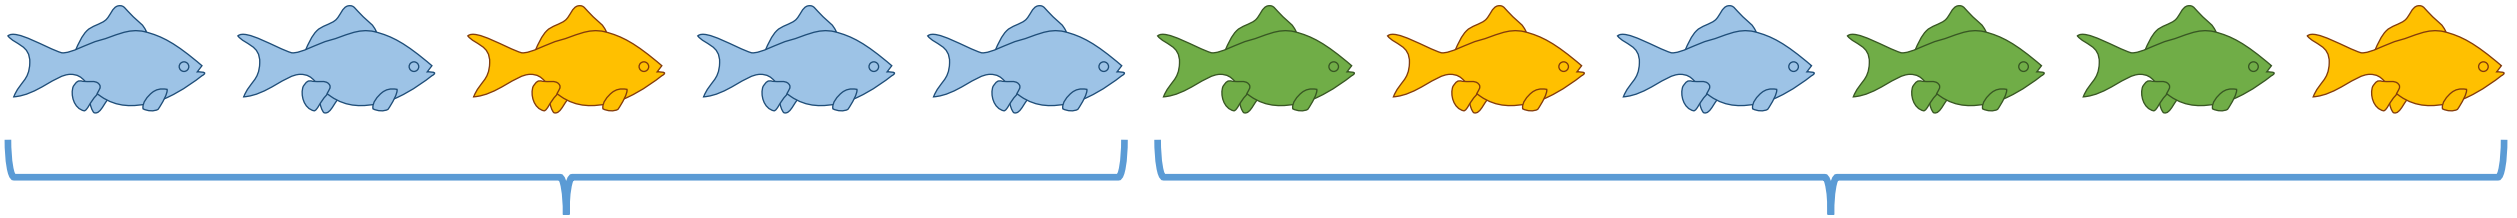

Volatility: 80%-10%-10%

Volatility: 50%-25%-25%

Modelling principles in action: hierarchies

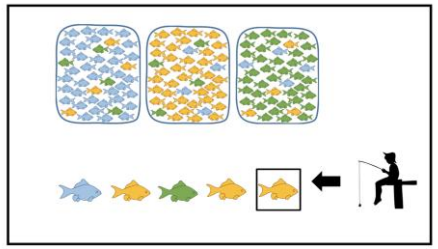

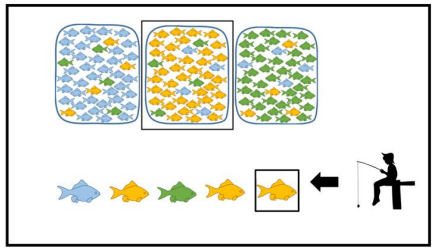

$\lambda_t = f(\alpha_0, k, Data_{1:t})$

Dirichlet update:

$$\lambda_t^{(i)} = \frac{\alpha_0^{(i)} + k\, data_{1:t}^{(i)}}{\sum(\alpha_0^{(i)} + k\, data_{1:t})}$$

e.g.:

$$\alpha_0^{(blue\ pond)} = [8\ 1\ 1]; \qquad data_{1:10} = [5\ 3\ 3]; \qquad k = 0.5;$$

$$\lambda_6^{(blue\ pond)} = \frac{[10.5\ 2.5\ 1.5]}{15} = [0.7\ 0.167\ 0.1]$$

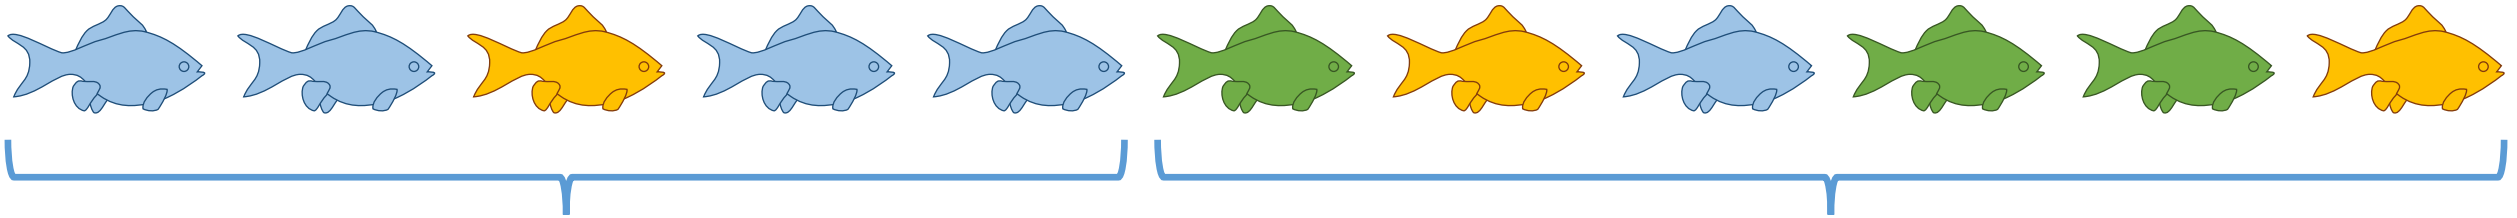

Volatility: 80%-10%-10%

Volatility: 50%-25%-25%

Summary

- Biological agents continuously collect information from the environment to form and update their own beliefs.
- Estimated distributions of probabilities (i.e., beliefs) can be updated using Bayes' theorem, assuming:
    - optimal behavior, relative to the objectives and the information available.
    - (usually) normal distributions.
- Subject-specific model parameters are estimated tuning the model to match the specific choice selections of each experimental subject.
    - Estimated precision in prior beliefs and in the likelihood of events affect the pace of update in non-linear ways.
    - Hierarchies can be used to determine multiple level of decision-making processes and estimations

Thank you!

Funding:

R01MH133625
R21MH129898

National Institute on Drug Abuse
Advancing Addiction Science

National Institute of Mental Health