Learning in Dynamic Environments:

A Tentative Implementation of the Volatile Kalman Filter in STAN

Margherita Calderan
margherita.calderan@phd.unipd.it

https://github.com/Mar-Cald/VKF_STAN

Article | Published: 02 March 2015

Anxious individuals have difficulty learning the causal statistics of aversive environments

Michael Browning, Timothy E Behrens, Gerhard Jocham, Jill X O'Reilly & Sonia J Bishop

Nature Neuroscience **18**, 590–596 (2015) | Cite this article

Impaired adaptation of learning to contingency volatility in internalizing psychopathology

Christopher Gagne, Ondrej Zika, Peter Dayan, Sonia J Bishop

RESEARCH ARTICLE

With an eye on uncertainty: Modelling pupillary responses to environmental volatility

Peter Vincent*, Thomas Parr, David Benrimoh, Karl J Friston

Article | Open access | Published: 08 July 2023

Blocking D2/D3 dopamine receptors in male participants increases volatility of beliefs when learning to trust others

Nace Mikus, Christoph Eisenegger, Christoph Mathys, Luke Clark, Ulrich Müller, Trevor W. Robbins, Claus Lamm & Michael Naef

Nature Communications **14**, Article number: 4049 (2023) | Cite this article

Article | Published: 31 July 2017

Adults with autism overestimate the volatility of the sensory environment

Rebecca P Lawson, Christoph Mathys & Geraint Rees

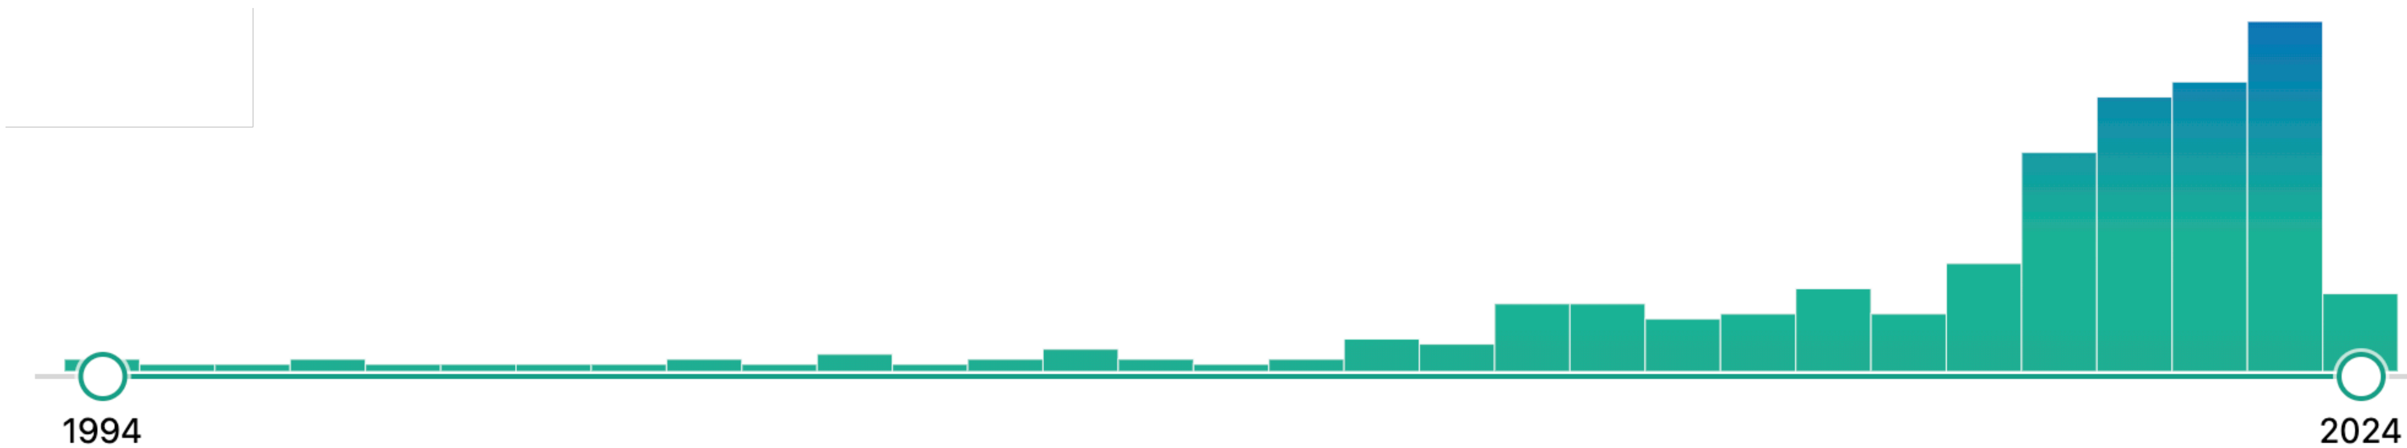

Overview

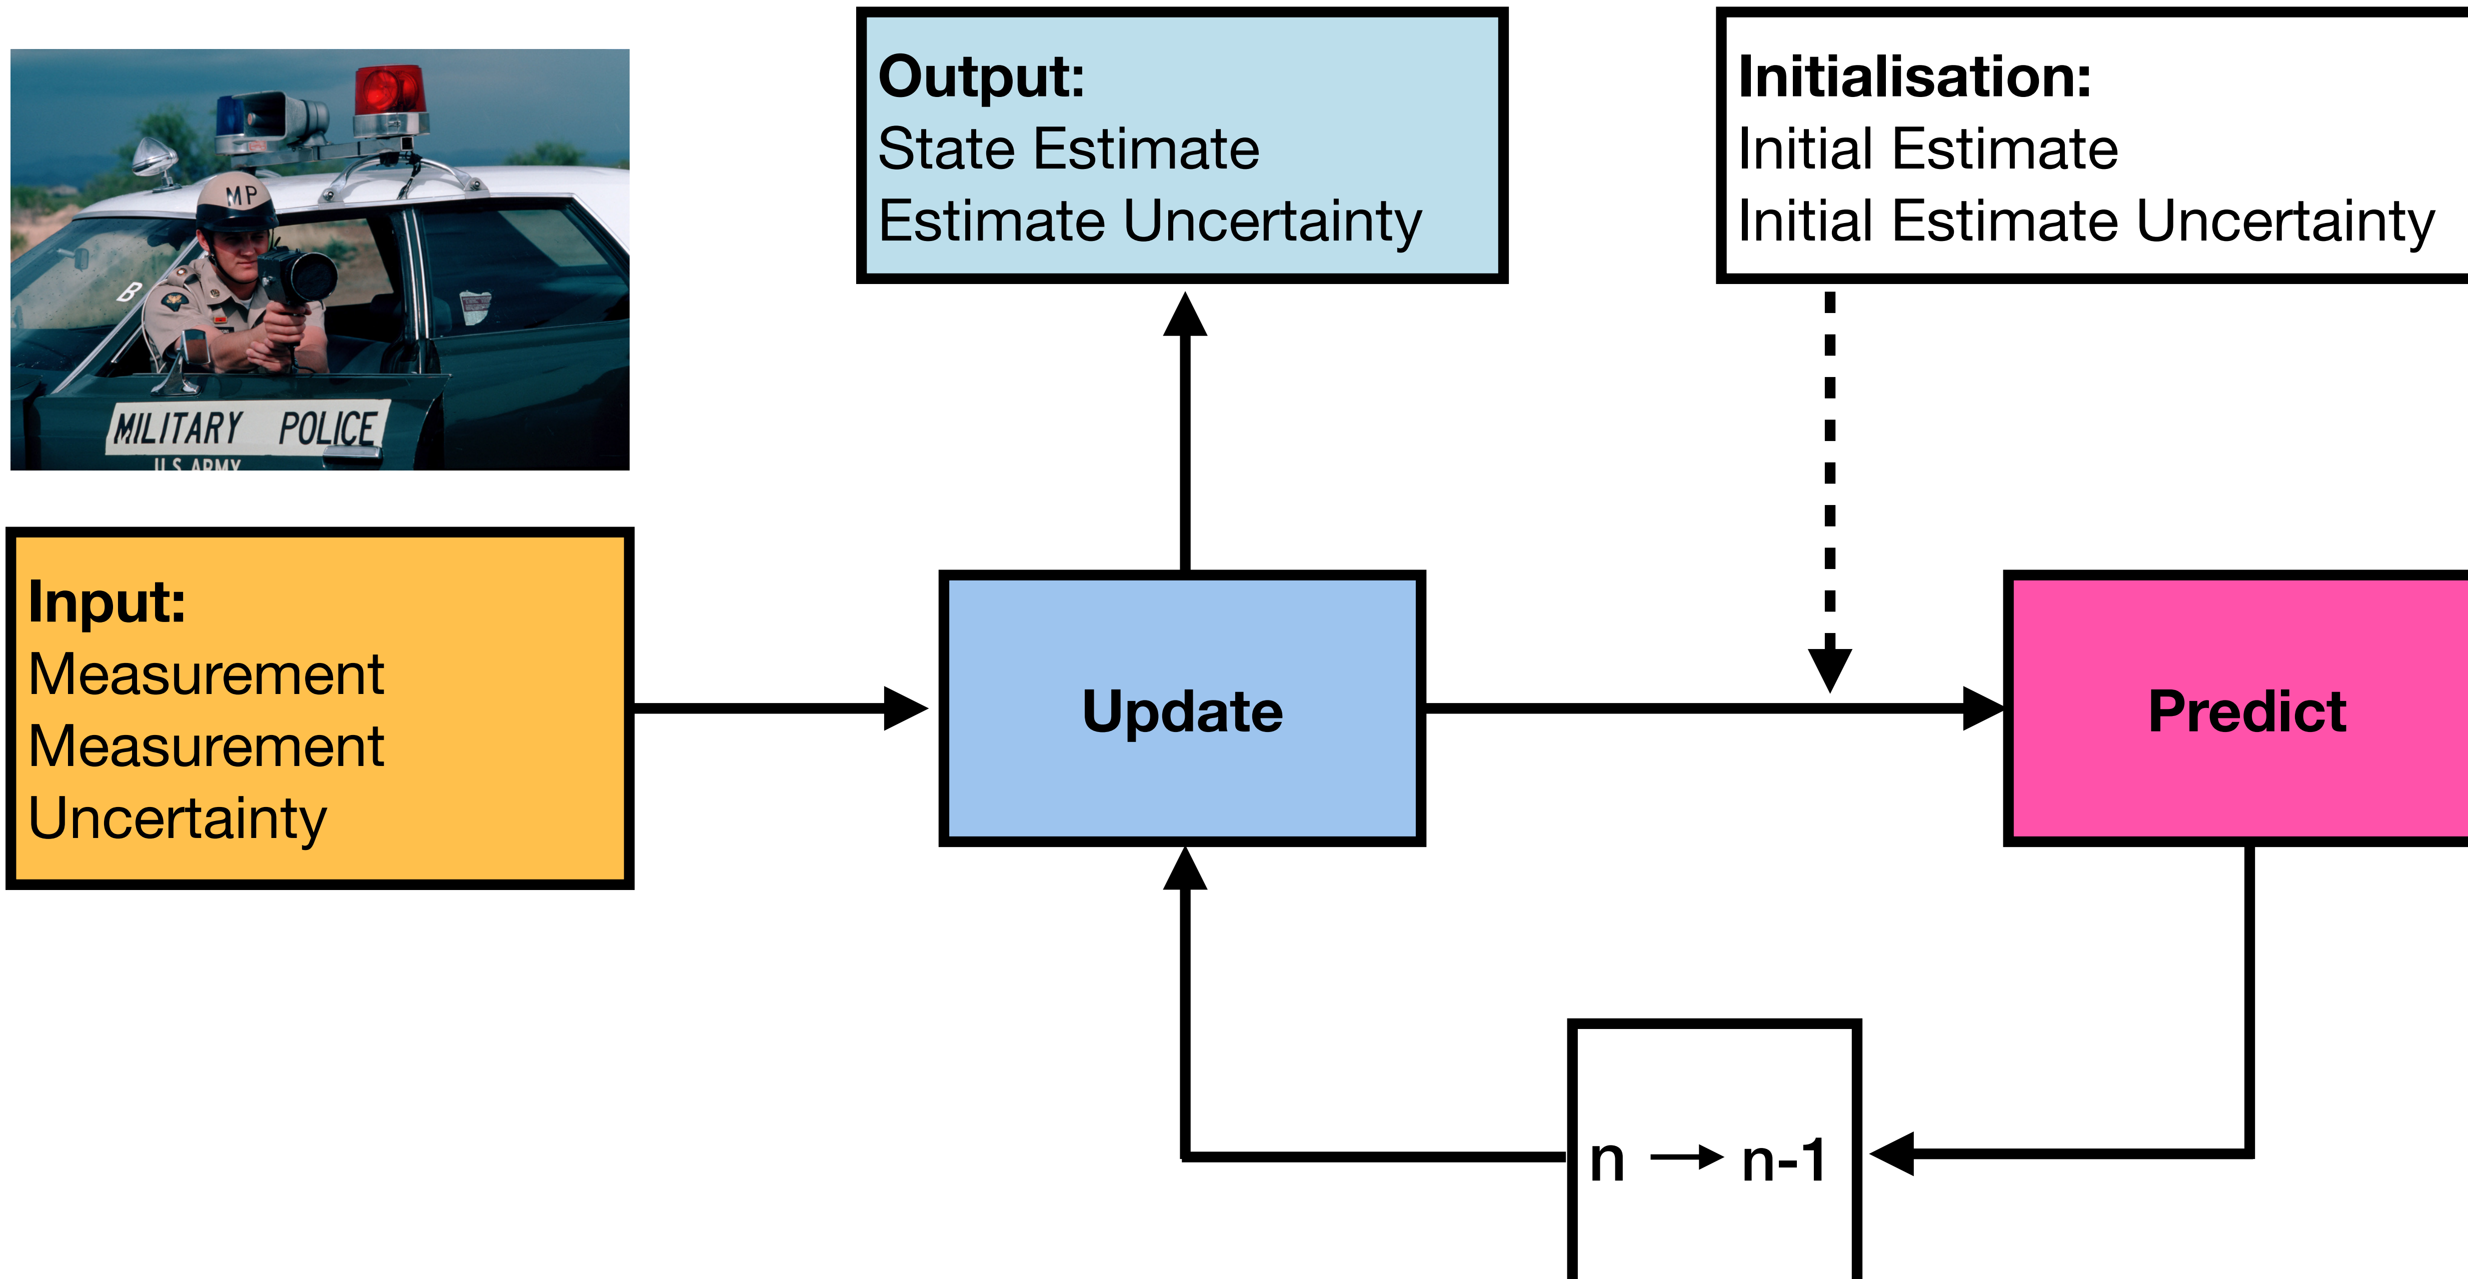
MILITARY POLICE
Output:
State Estimate
Estimate Uncertainty
Initialisation:
Initial Estimate
Initial Estimate Uncertainty
Input:
Measurement
Measurement
Uncertainty
Update
Predict
n → n-1

Measurement Noise

Kalman Gain $k_n = w_{n-1} / (w_{n-1} + noise)$

Predicted speed $m_n = m_{n-1} + k_{n-1} (o_n - m_{n-1})$

Variance $w_n = (1 - k_{n-1})\, w_{n-1}$

High Gain

variance in estimate >

variance in measurement

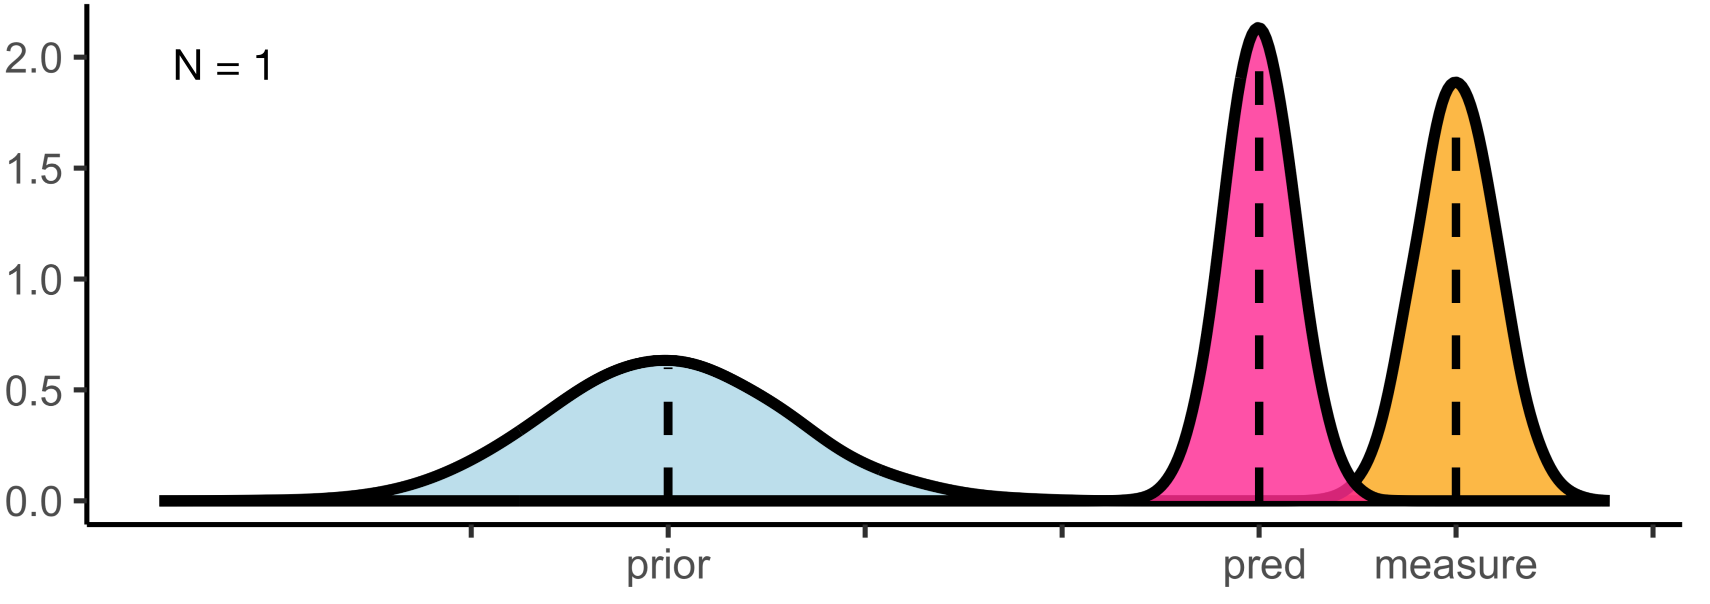

Measurement Noise

Kalman Gain $k_n = w_{n-1} / (w_{n-1} + noise)$

Predicted speed $m_n = m_{n-1} + k_{n-1} (o_n - m_{n-1})$

Variance $w_n = (1 - k_{n-1}) w_{n-1}$

High Gain

variance in estimate >

variance in measurement

Low Gain

variance in estimate <

variance in measuremen

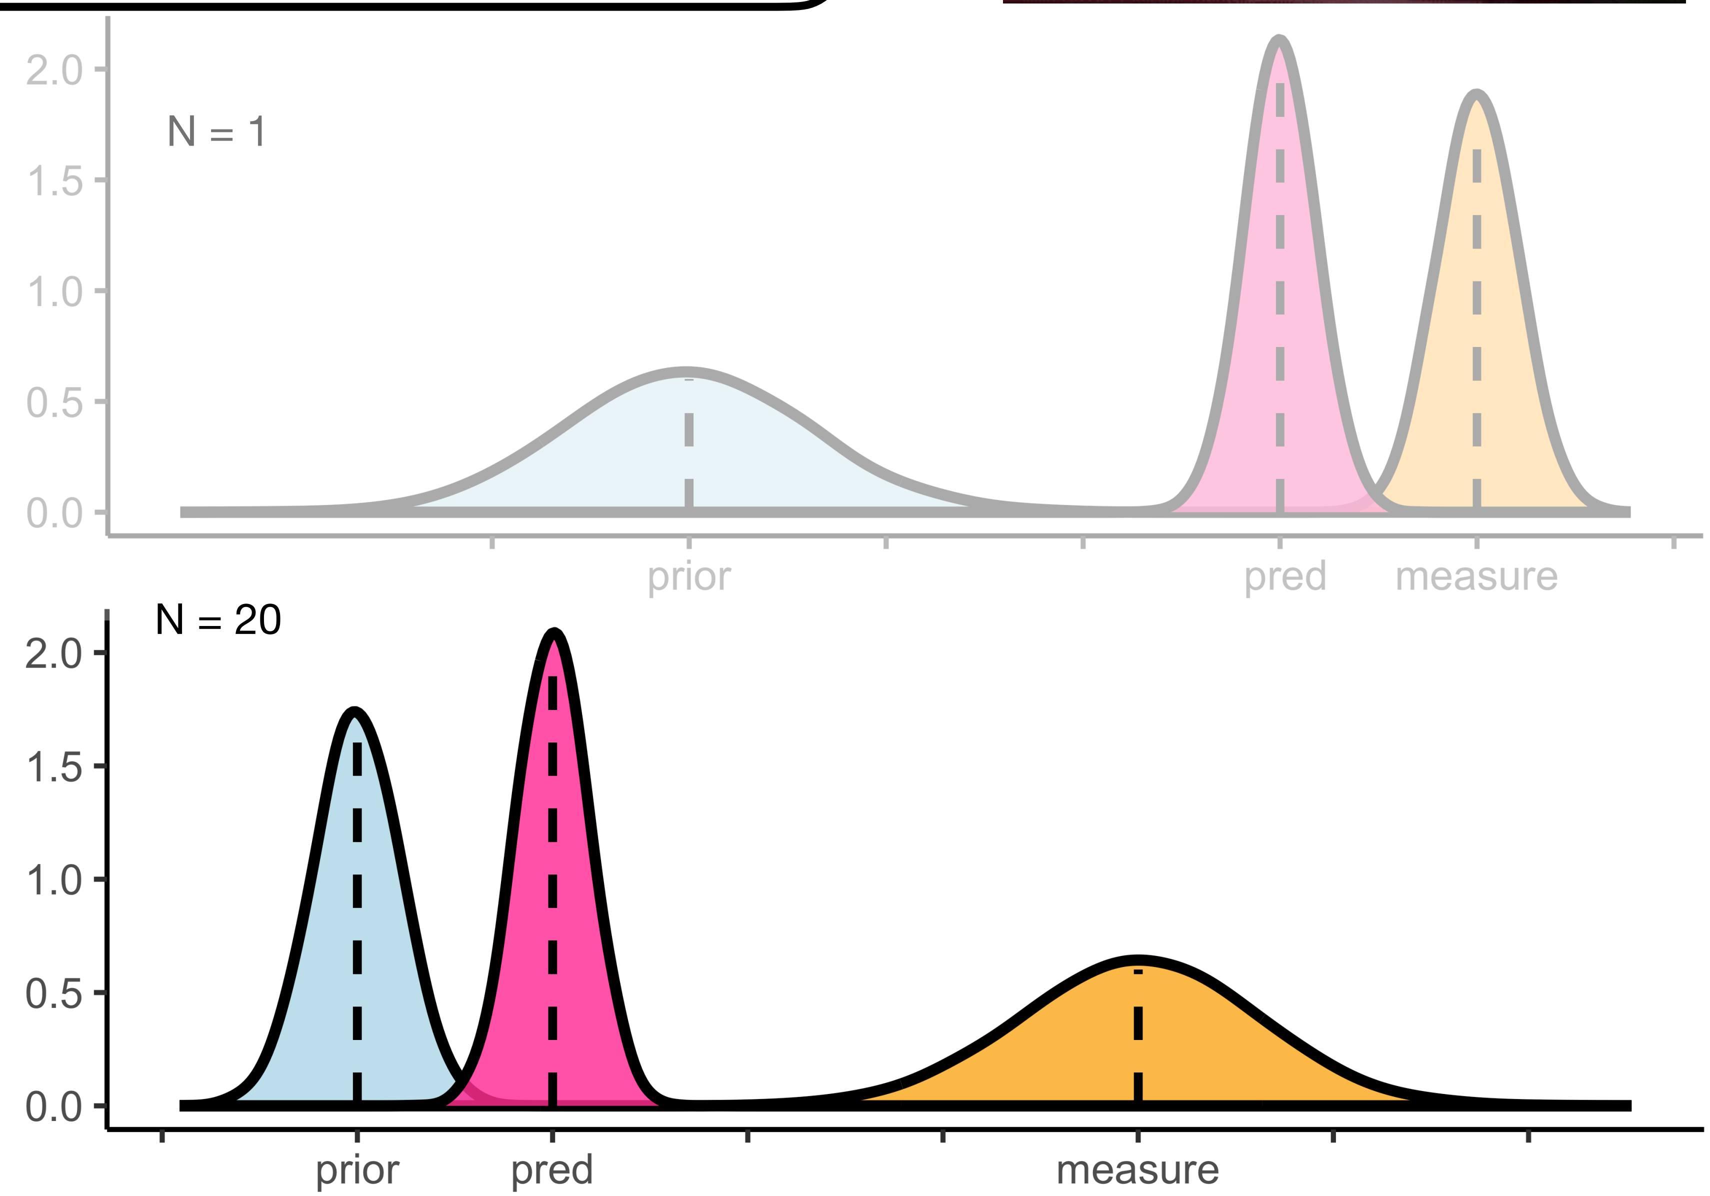

Process Noise

$$k_n = (w_{n-1} + z) / (w_{n-1} + z + \text{noise})$$

$$m_n = m_{n-1} + k_{n-1} (o_n - m_{n-1})$$

$$w_n = (1 - k_{n-1}) (w_{n-1} + z)$$

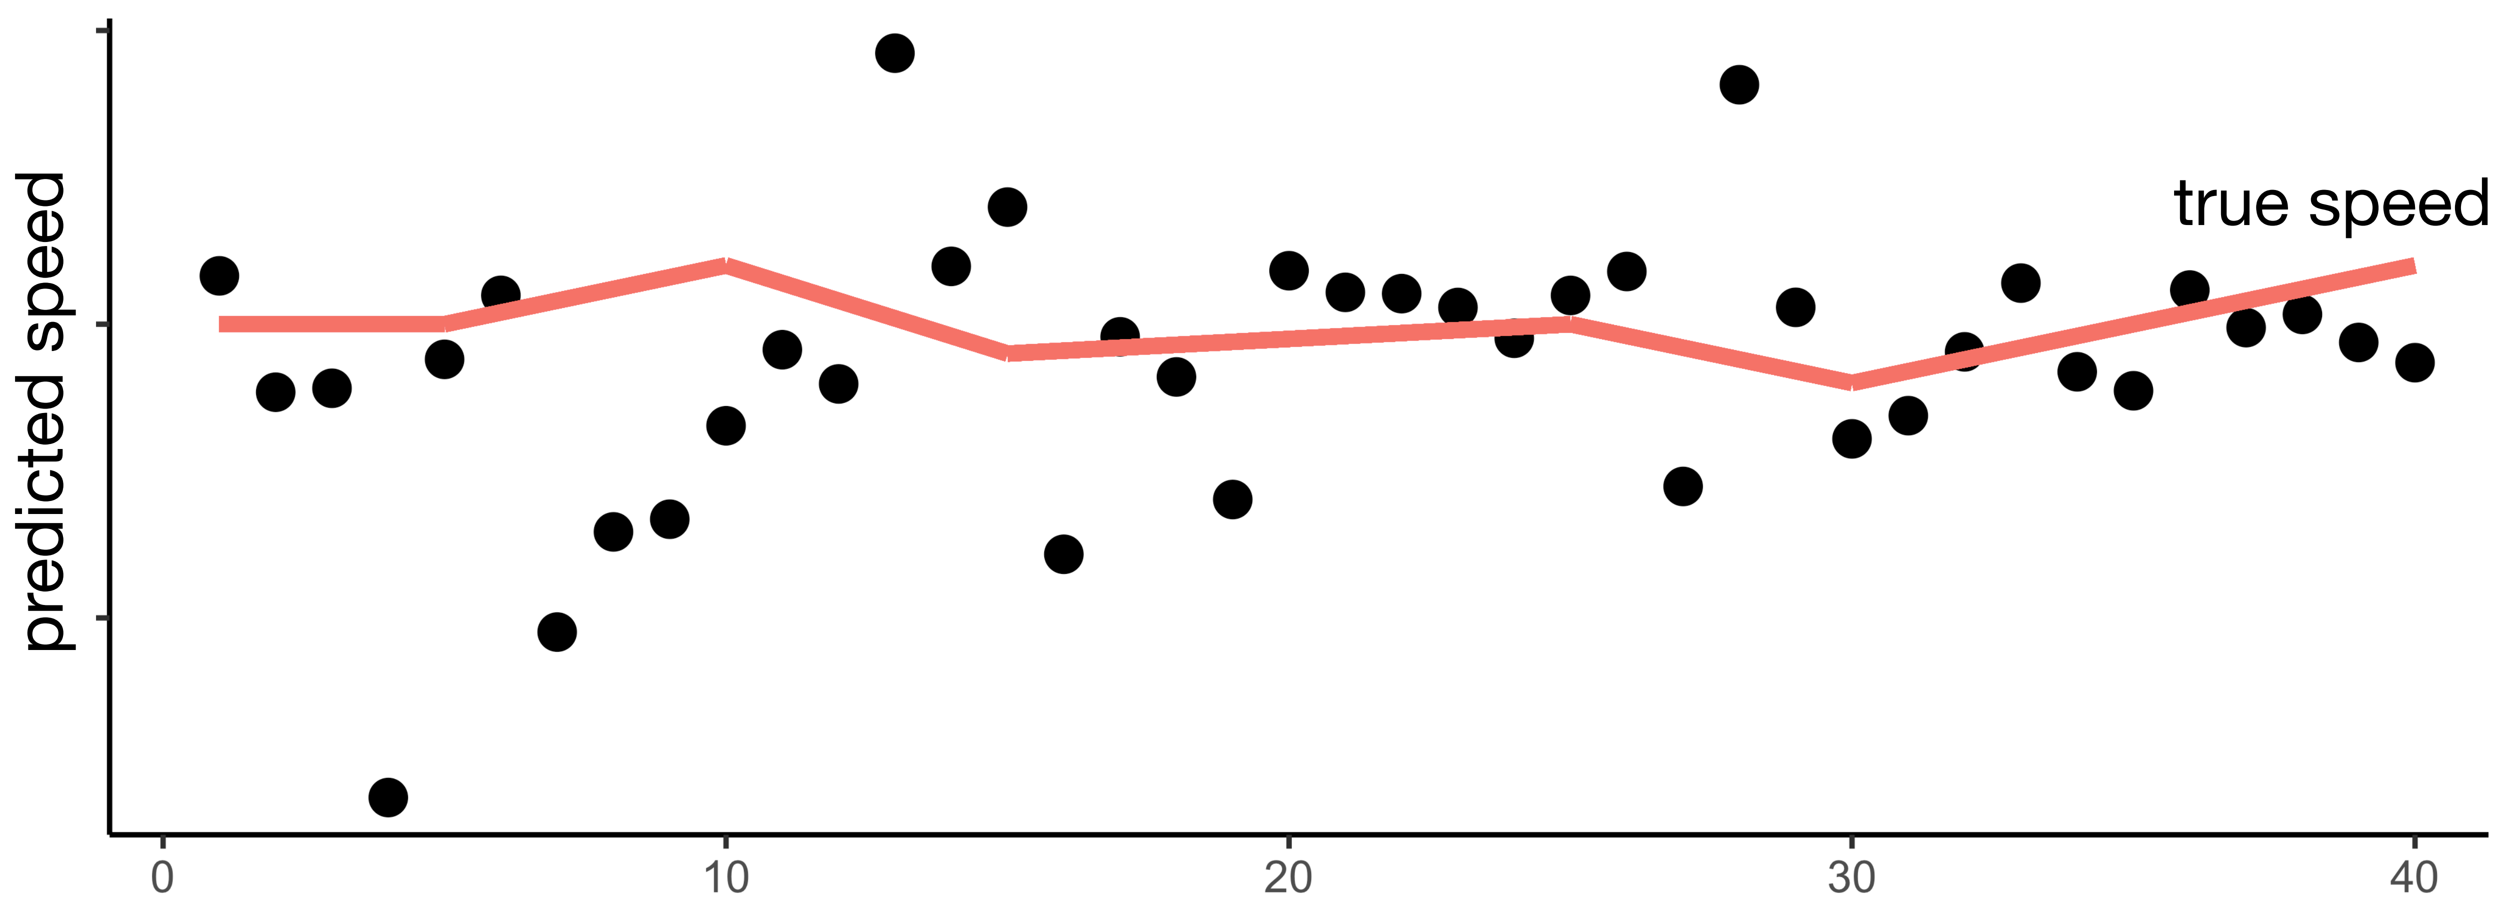

Kalman Filter

$$k_n = (w_{n-1} + z) / (w_{n-1} + z + \text{noise})$$

$$m_n = m_{n-1} + k_{n-1} (o_n - m_{n-1})$$

$$w_n = (1 - k_{n-1}) (w_{n-1} + z)$$

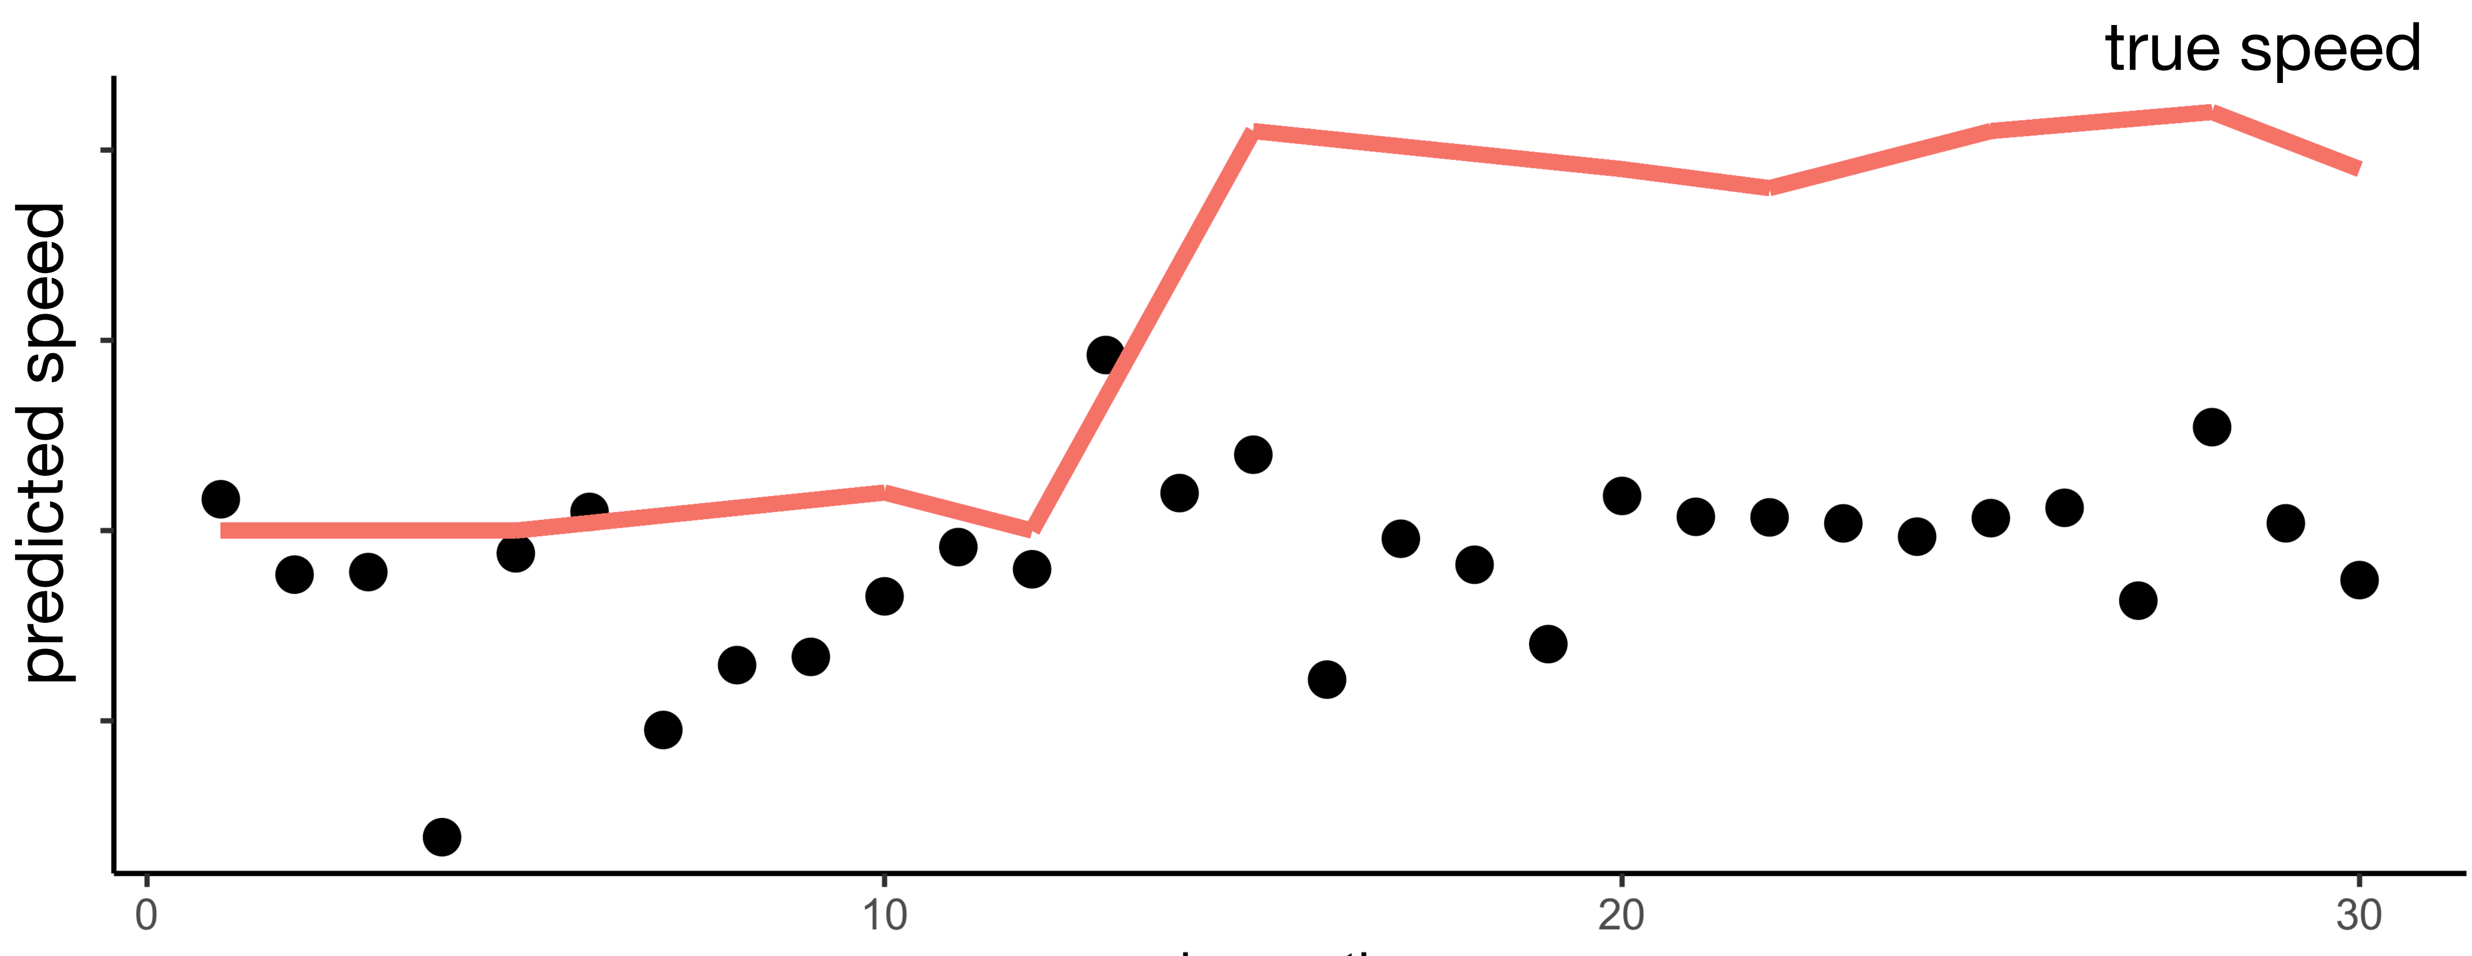

Volatile Kalman Filter (VKF)

Piray, P., & Daw, N. D. (2020)

$$k_n = (w_{n-1} + z_{n-1}) \,/\, (w_{n-1} + z_{n-1} + \text{noise})$$

$$m_n = m_{n-1} + k_{n-1}\,(o_n - m_{n-1})$$

$$w_n = (1 - k_{n-1})\,(w_{n-1} + z_{n-1})$$

It is not noise, it's a change in the environment!

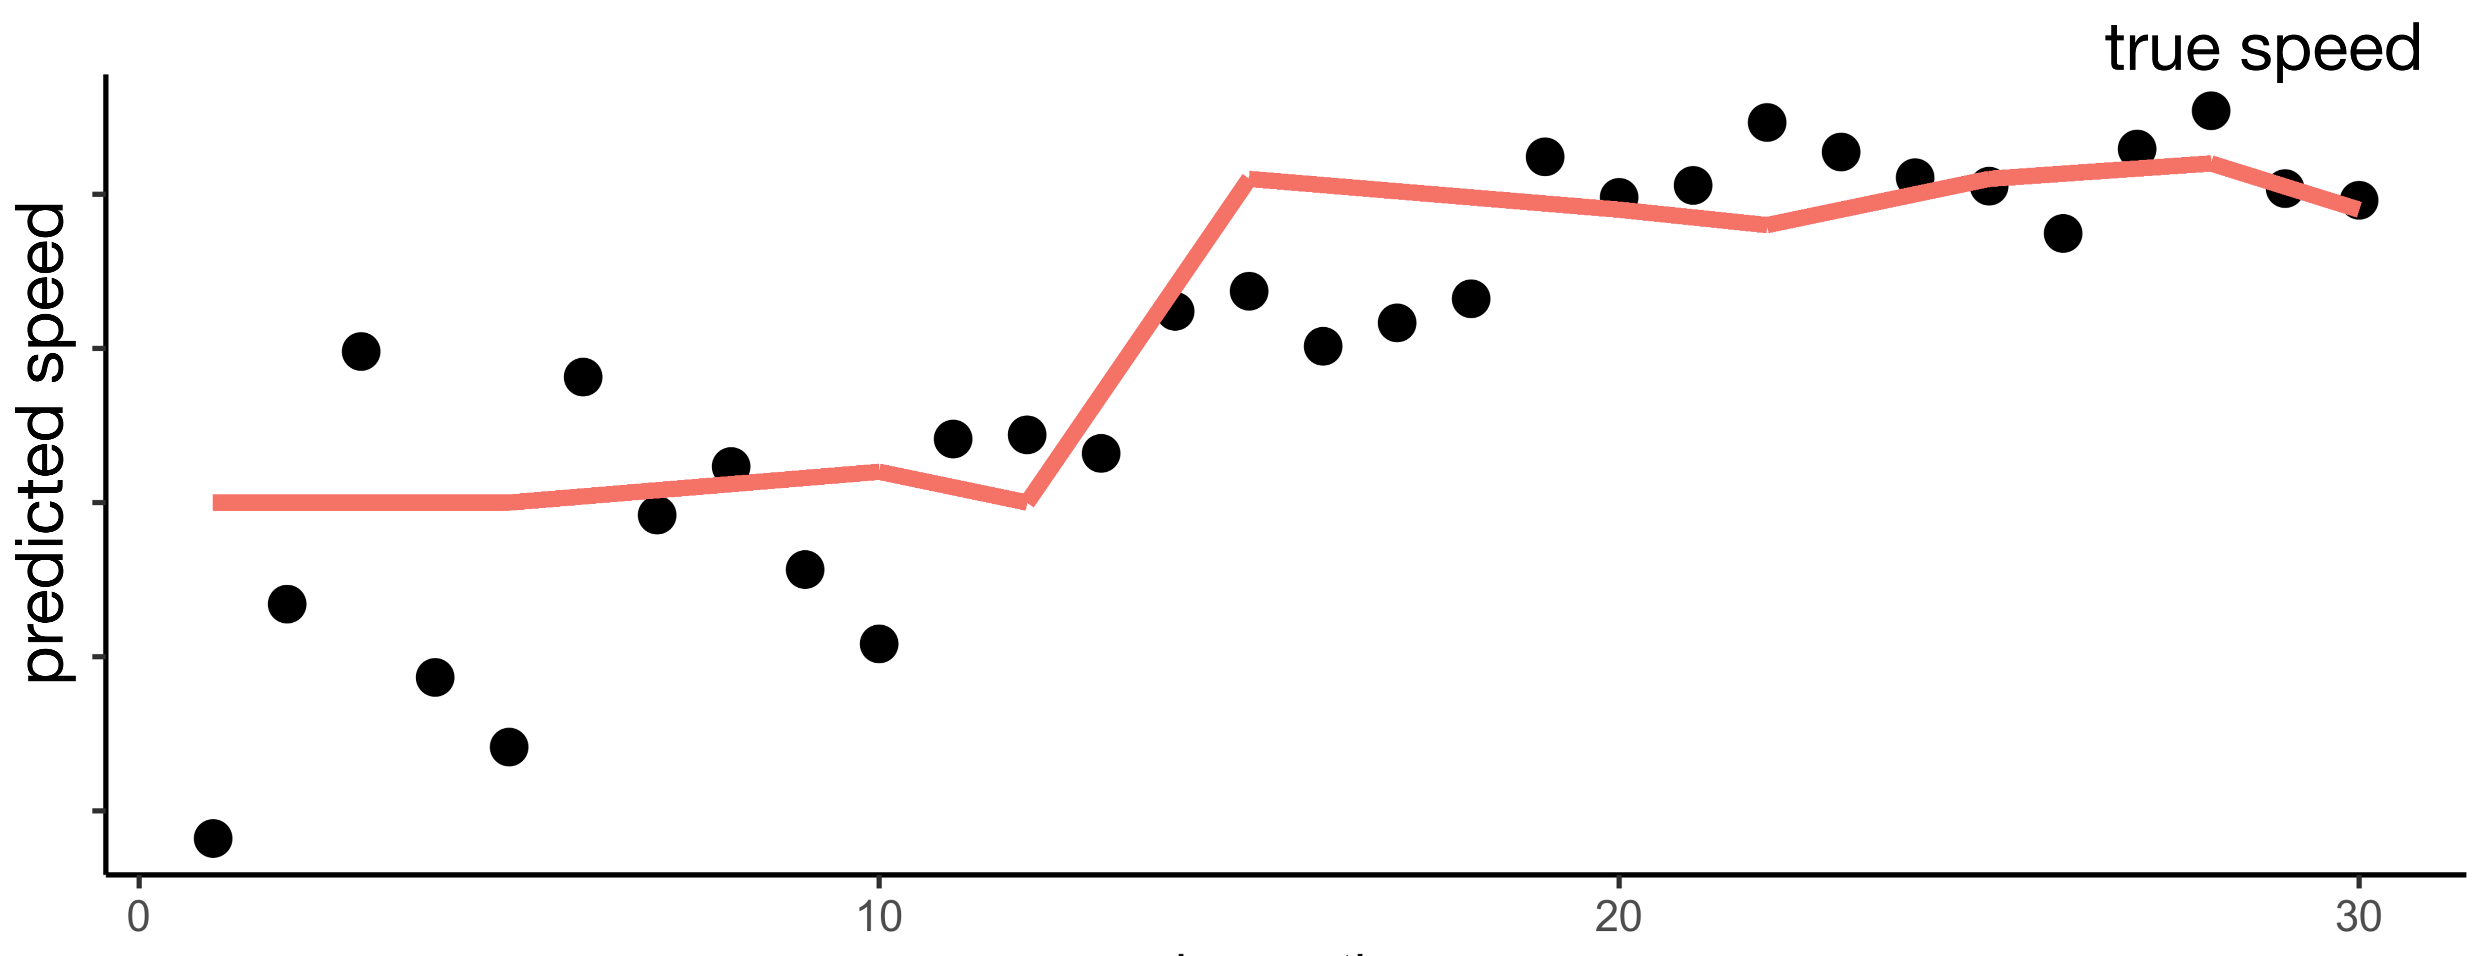

An Example

Go

NoGo

Probability of Go

VKF as Perceptual Model

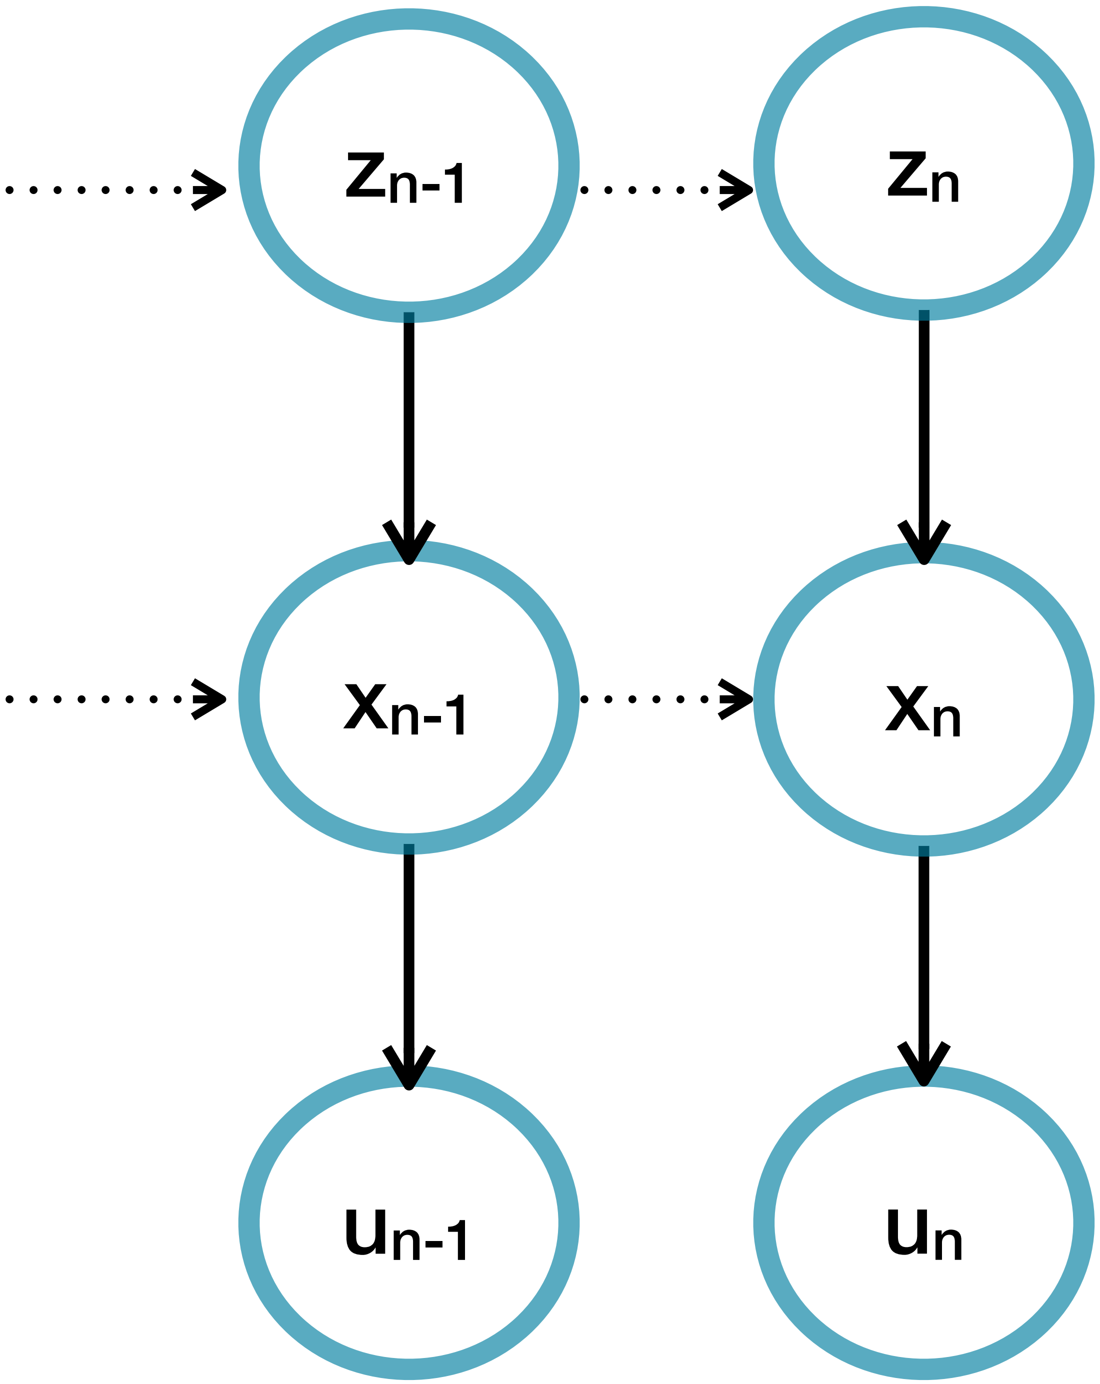

Volatility

$z_n = N(z_n \mid z_{n-1}, \sigma^2_v)$

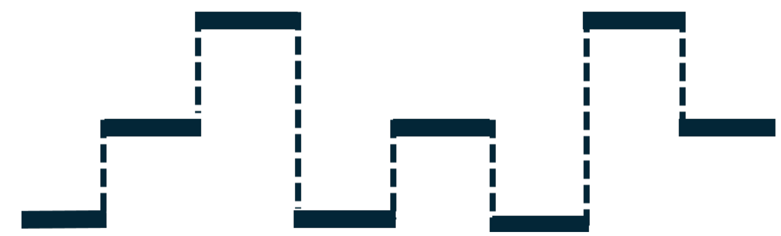

Mean

$x_n = N(x_n \mid x_{n-1}, z^{-1}_n)$

Go(1) vs NoGo (0)

Model

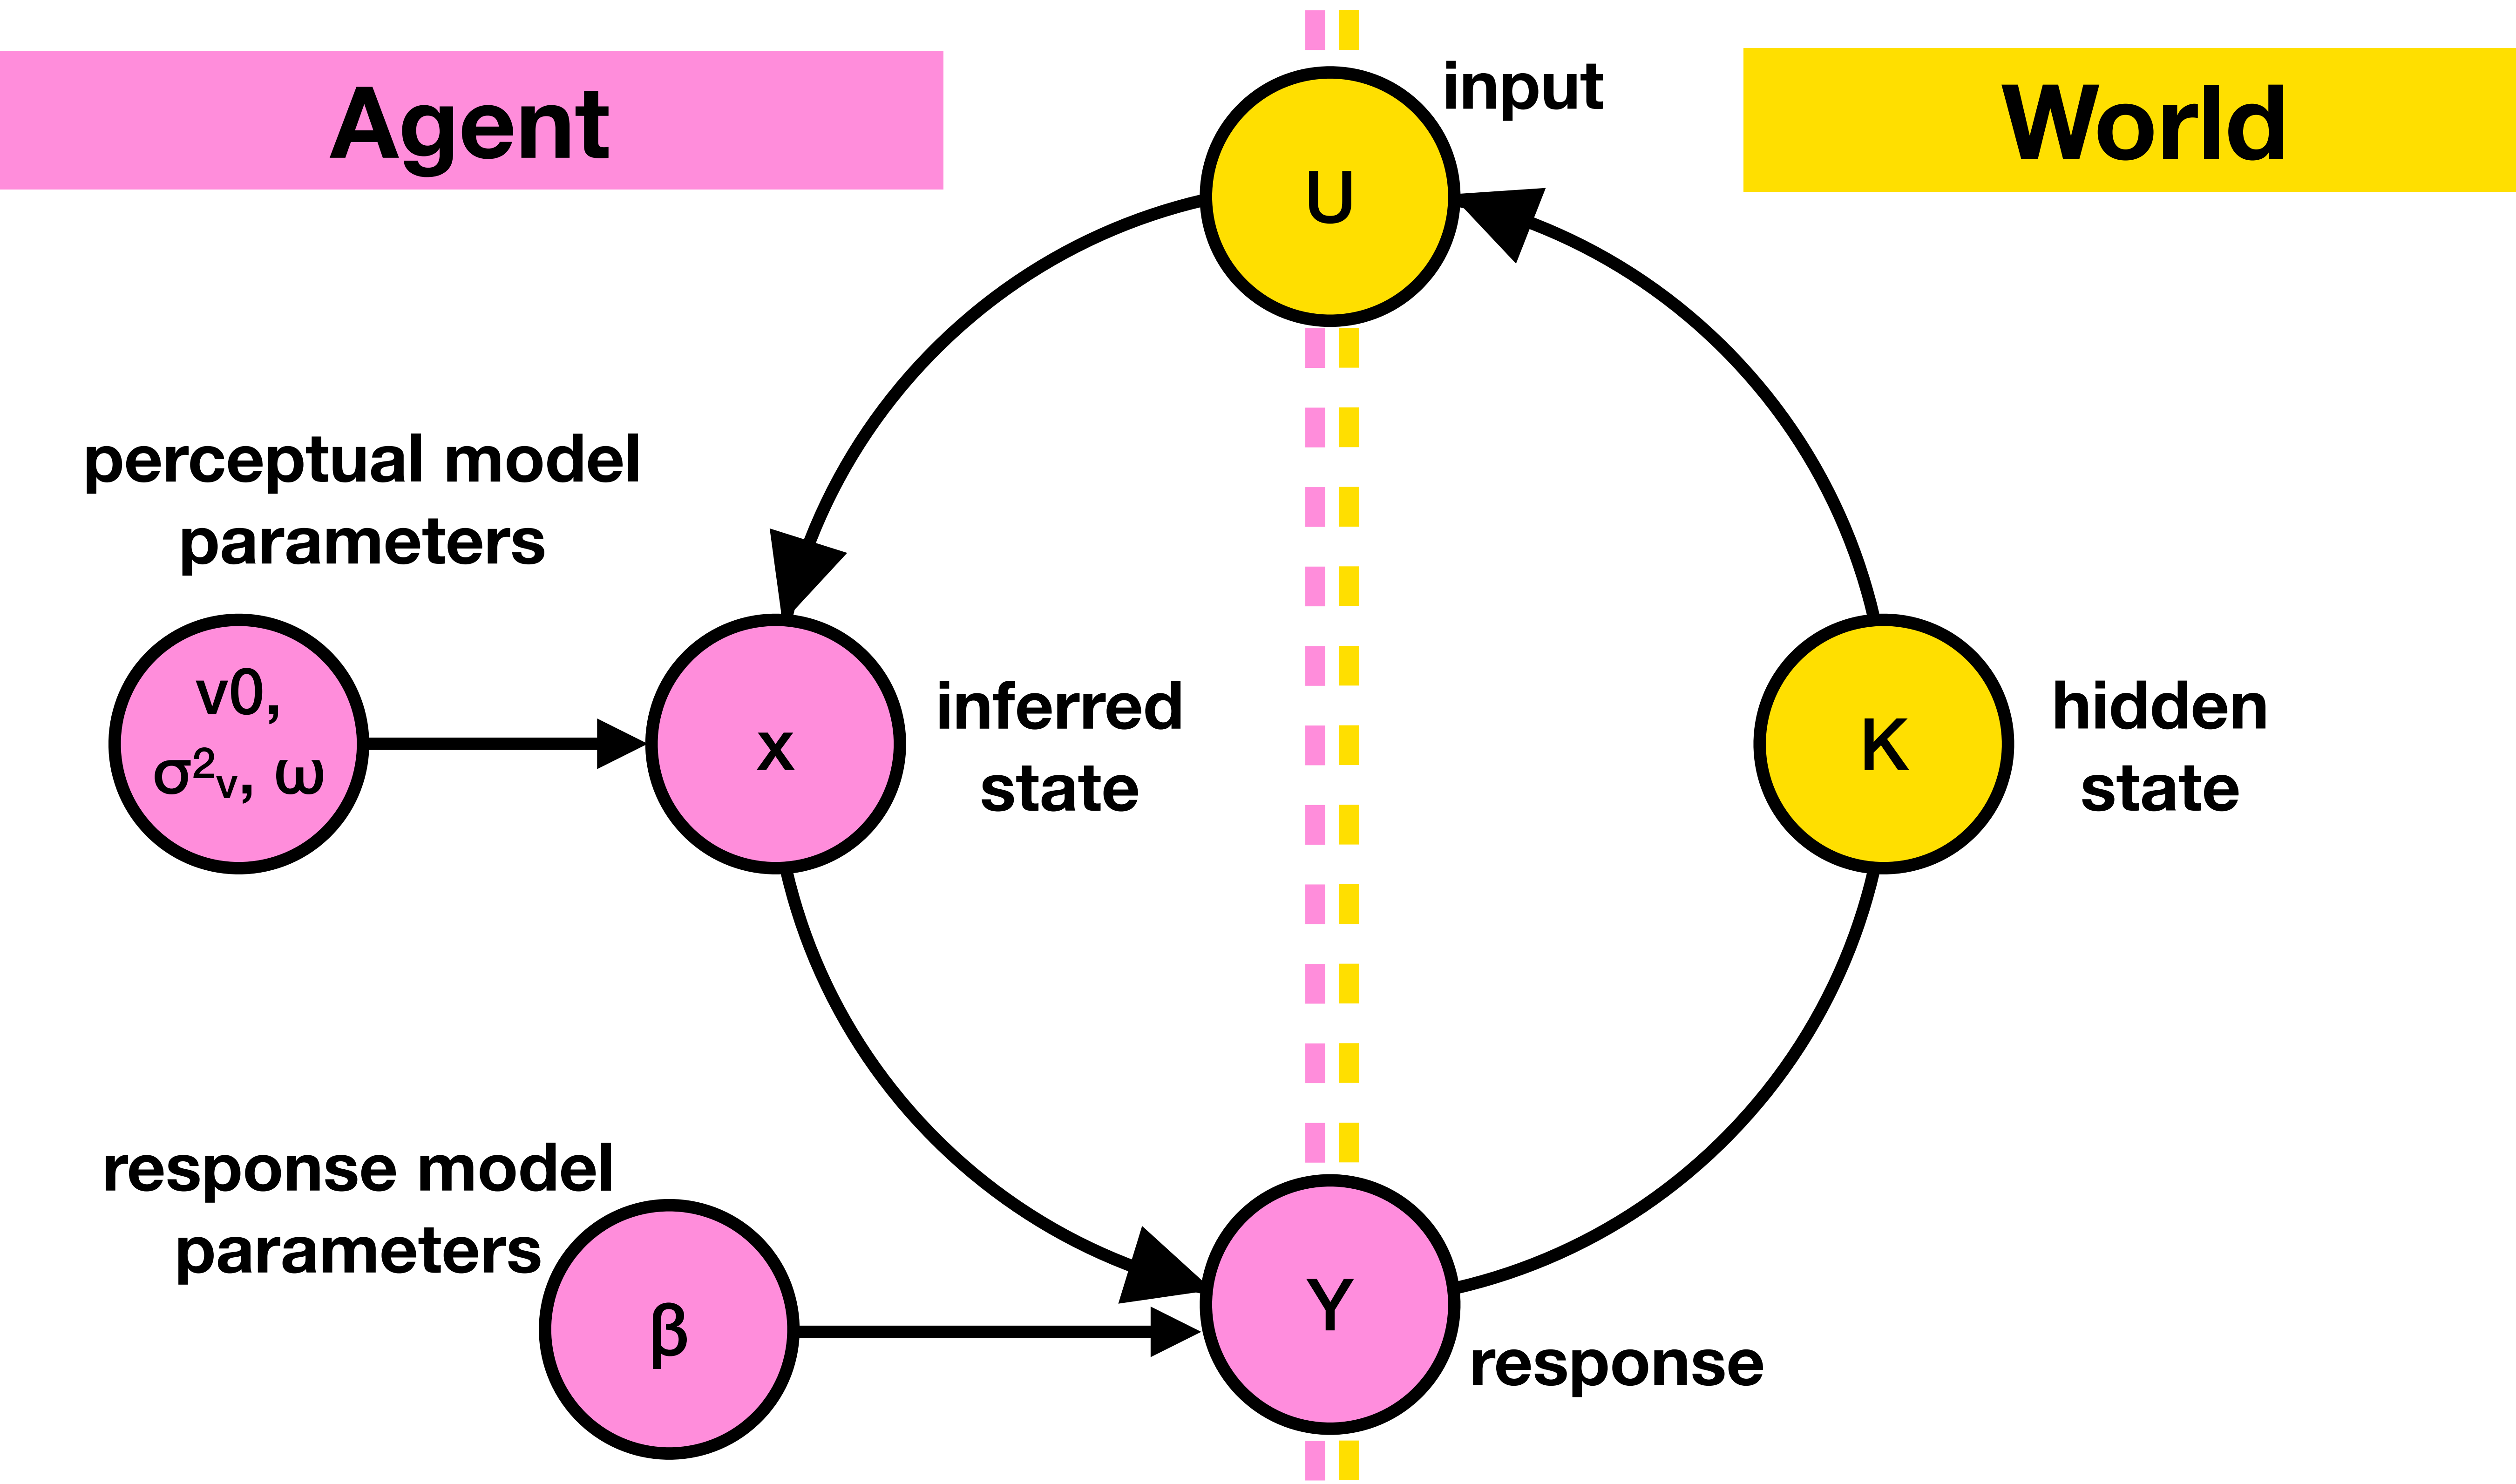

Frässle, S., et al. (2021); Mathys, C. D., et al., (2014)

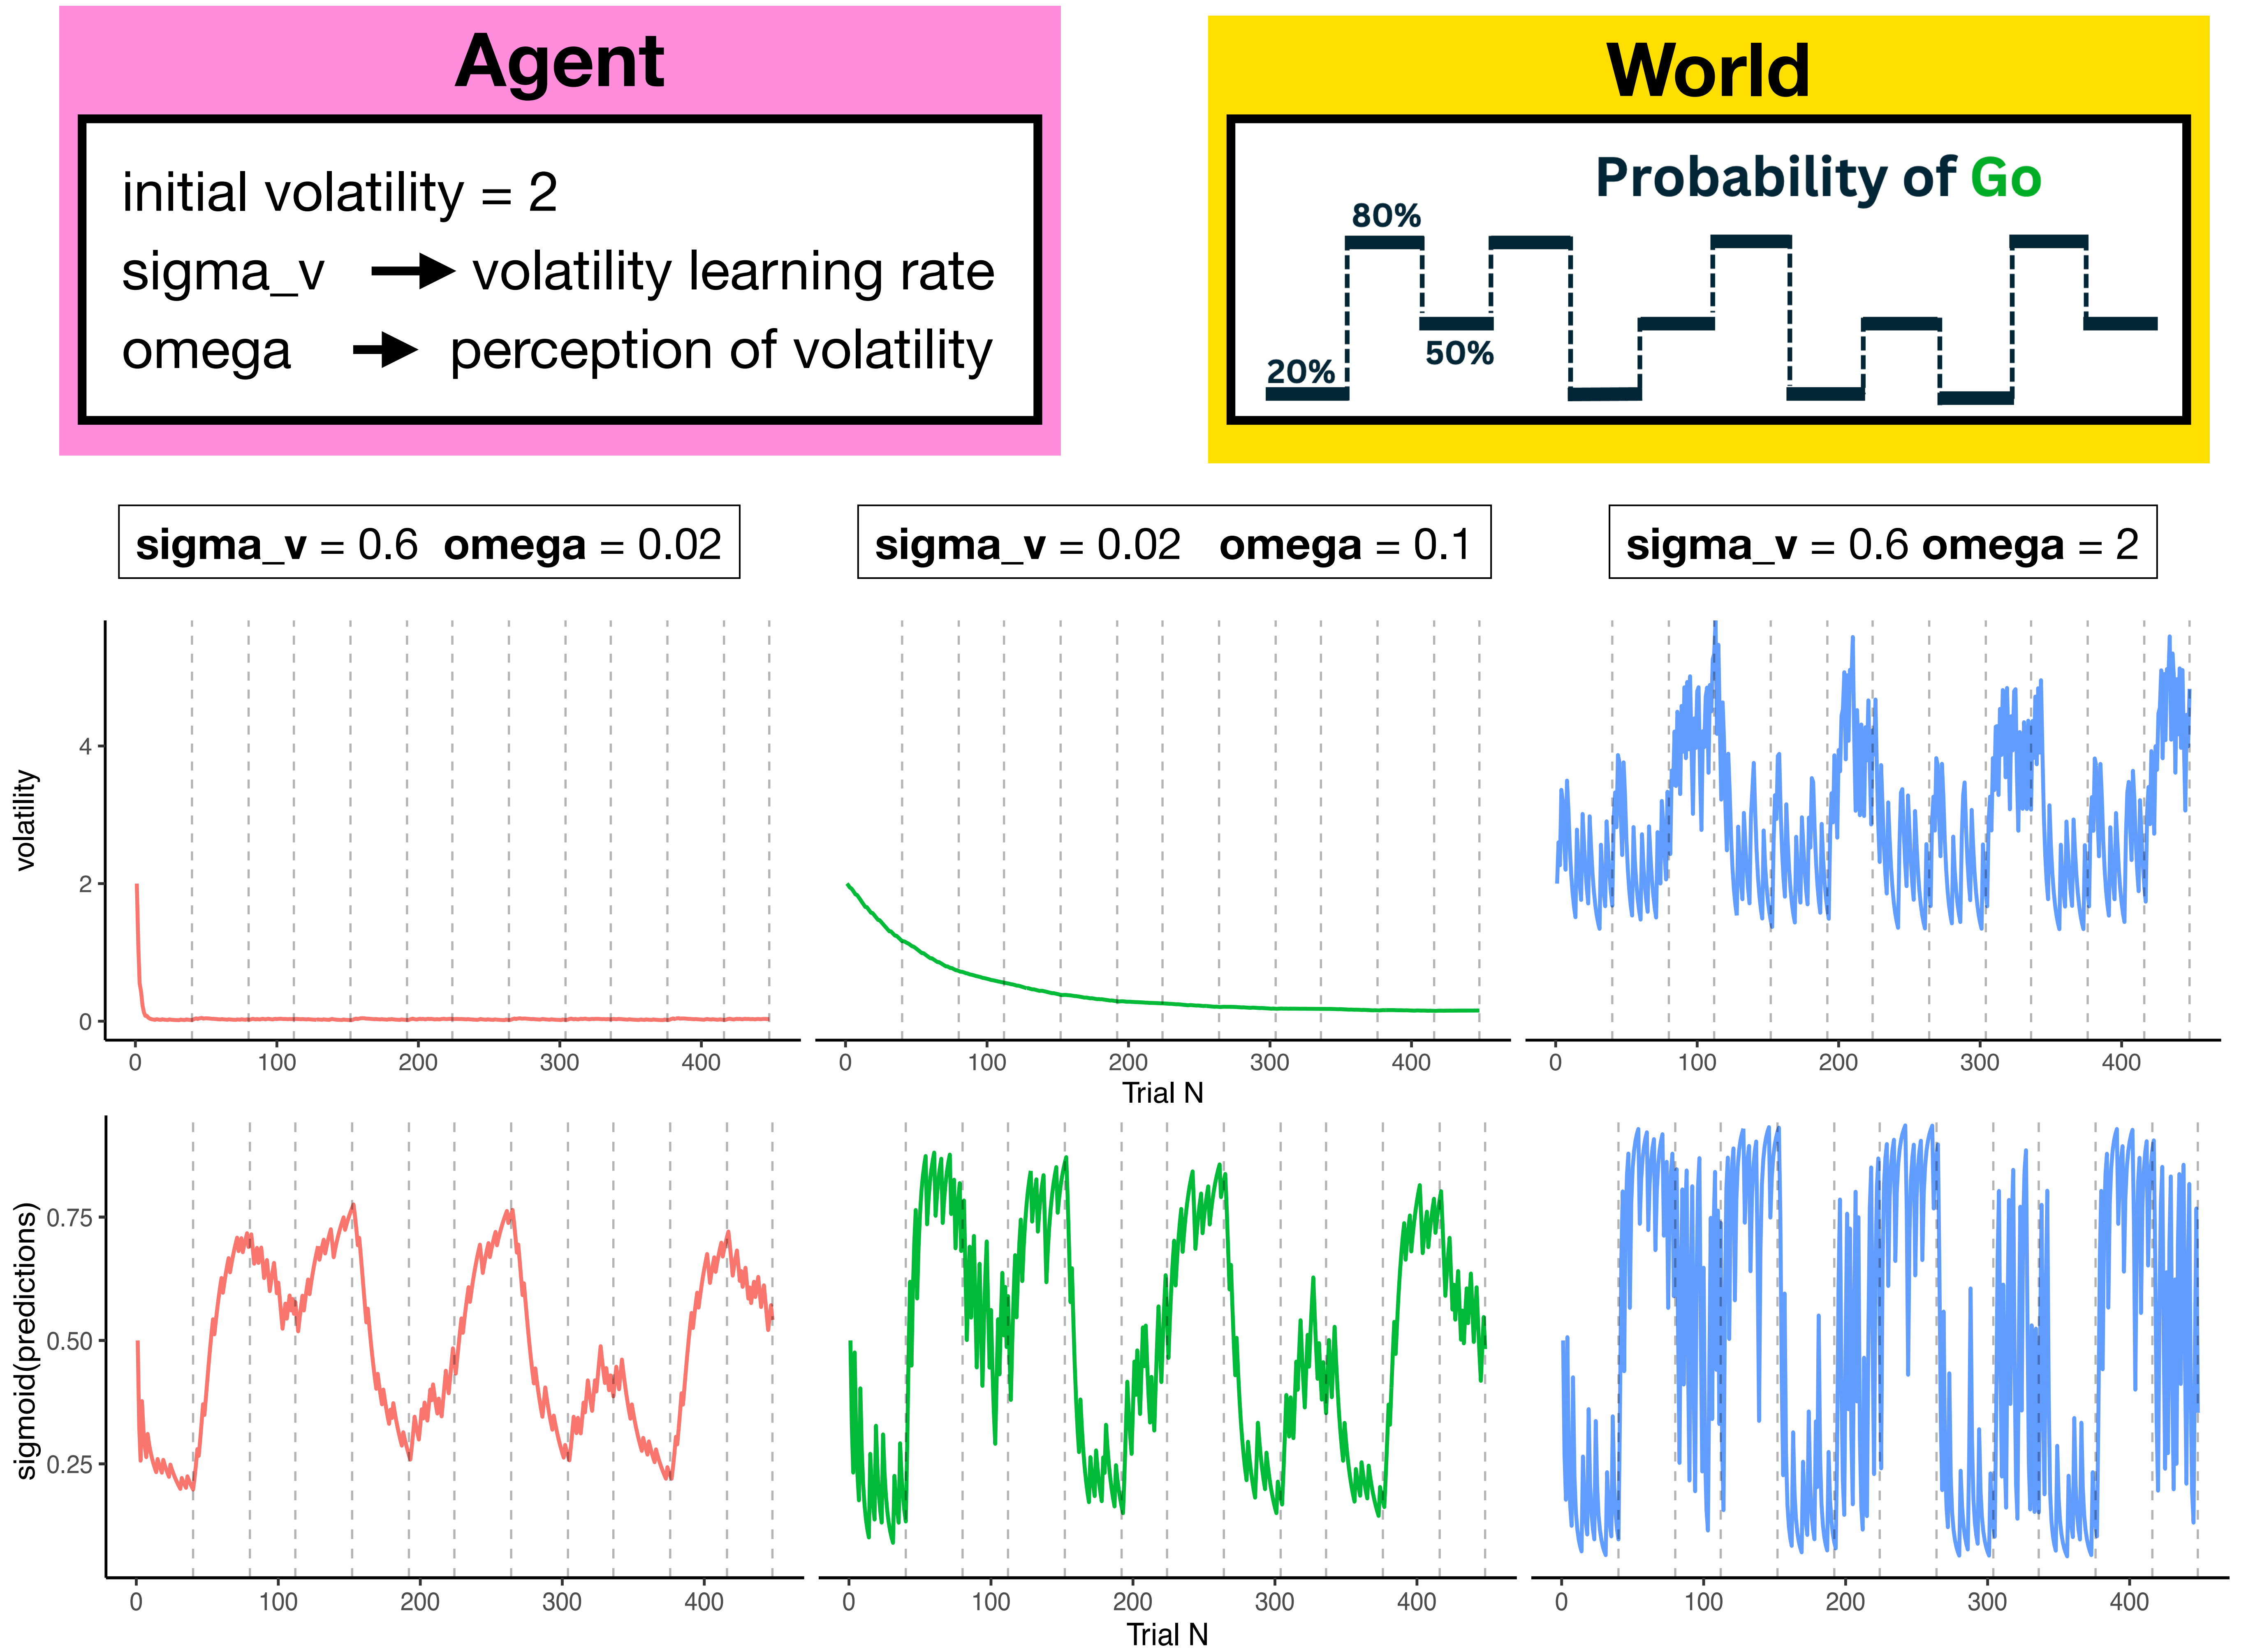

Agent
initial volatility = 2
sigma_v → volatility learning rate
omega → perception of volatility
World
Probability of Go
80%
50%
20%
sigma_v = 0.6 omega = 0.02
sigma_v = 0.02 omega = 0.1
sigma_v = 0.6 omega = 2
volatility
sigmoid(predictions)
Trial N
Trial N

Perceptual Model

omega → perception of volatility
v0 → initial volatility
sigma_v → volatility learning rate

Initial Values

w = **omega**

v = **v0**

m = 0 (i.e., .50)

```
for (t in 1:N) {
    o = GO[t];                // input
    mpre = m;                 // prediction
    wpre = w;                 // variance
    predictions[t] = m;
    volatility[t] = v;        // volatility
    delta_m = o - sigmoid(mpre);                 // prediction error (pe)
    k = (wpre + v) / (wpre + v + omega);         // Kalman Gain
    m = mpre + sqrt(wpre + v) * delta_m;         // prediction update
    w = (1 - k) * (wpre + v);                    // variance update
    wcov = (1 - k) * wpre;                       // covariance
    delta_v = (m-mpre)^2 + wpre + w - 2*wcov - v;    // volatility pe
    v     = v + sigma_v * delta_v;               // volatility update
}
```

Response model (reaction times, RT)

```
for (n in 1:N) {

  real T = RT[n] - ndt;  // decision time = RT - non-decision time

  real mu = intercept + predictions[n] * beta

  log_lik[n] = lognormal_lpdf( T | mu, sigma);

}
  target += sum(log_lik);
}
```

Volatility and Predictions

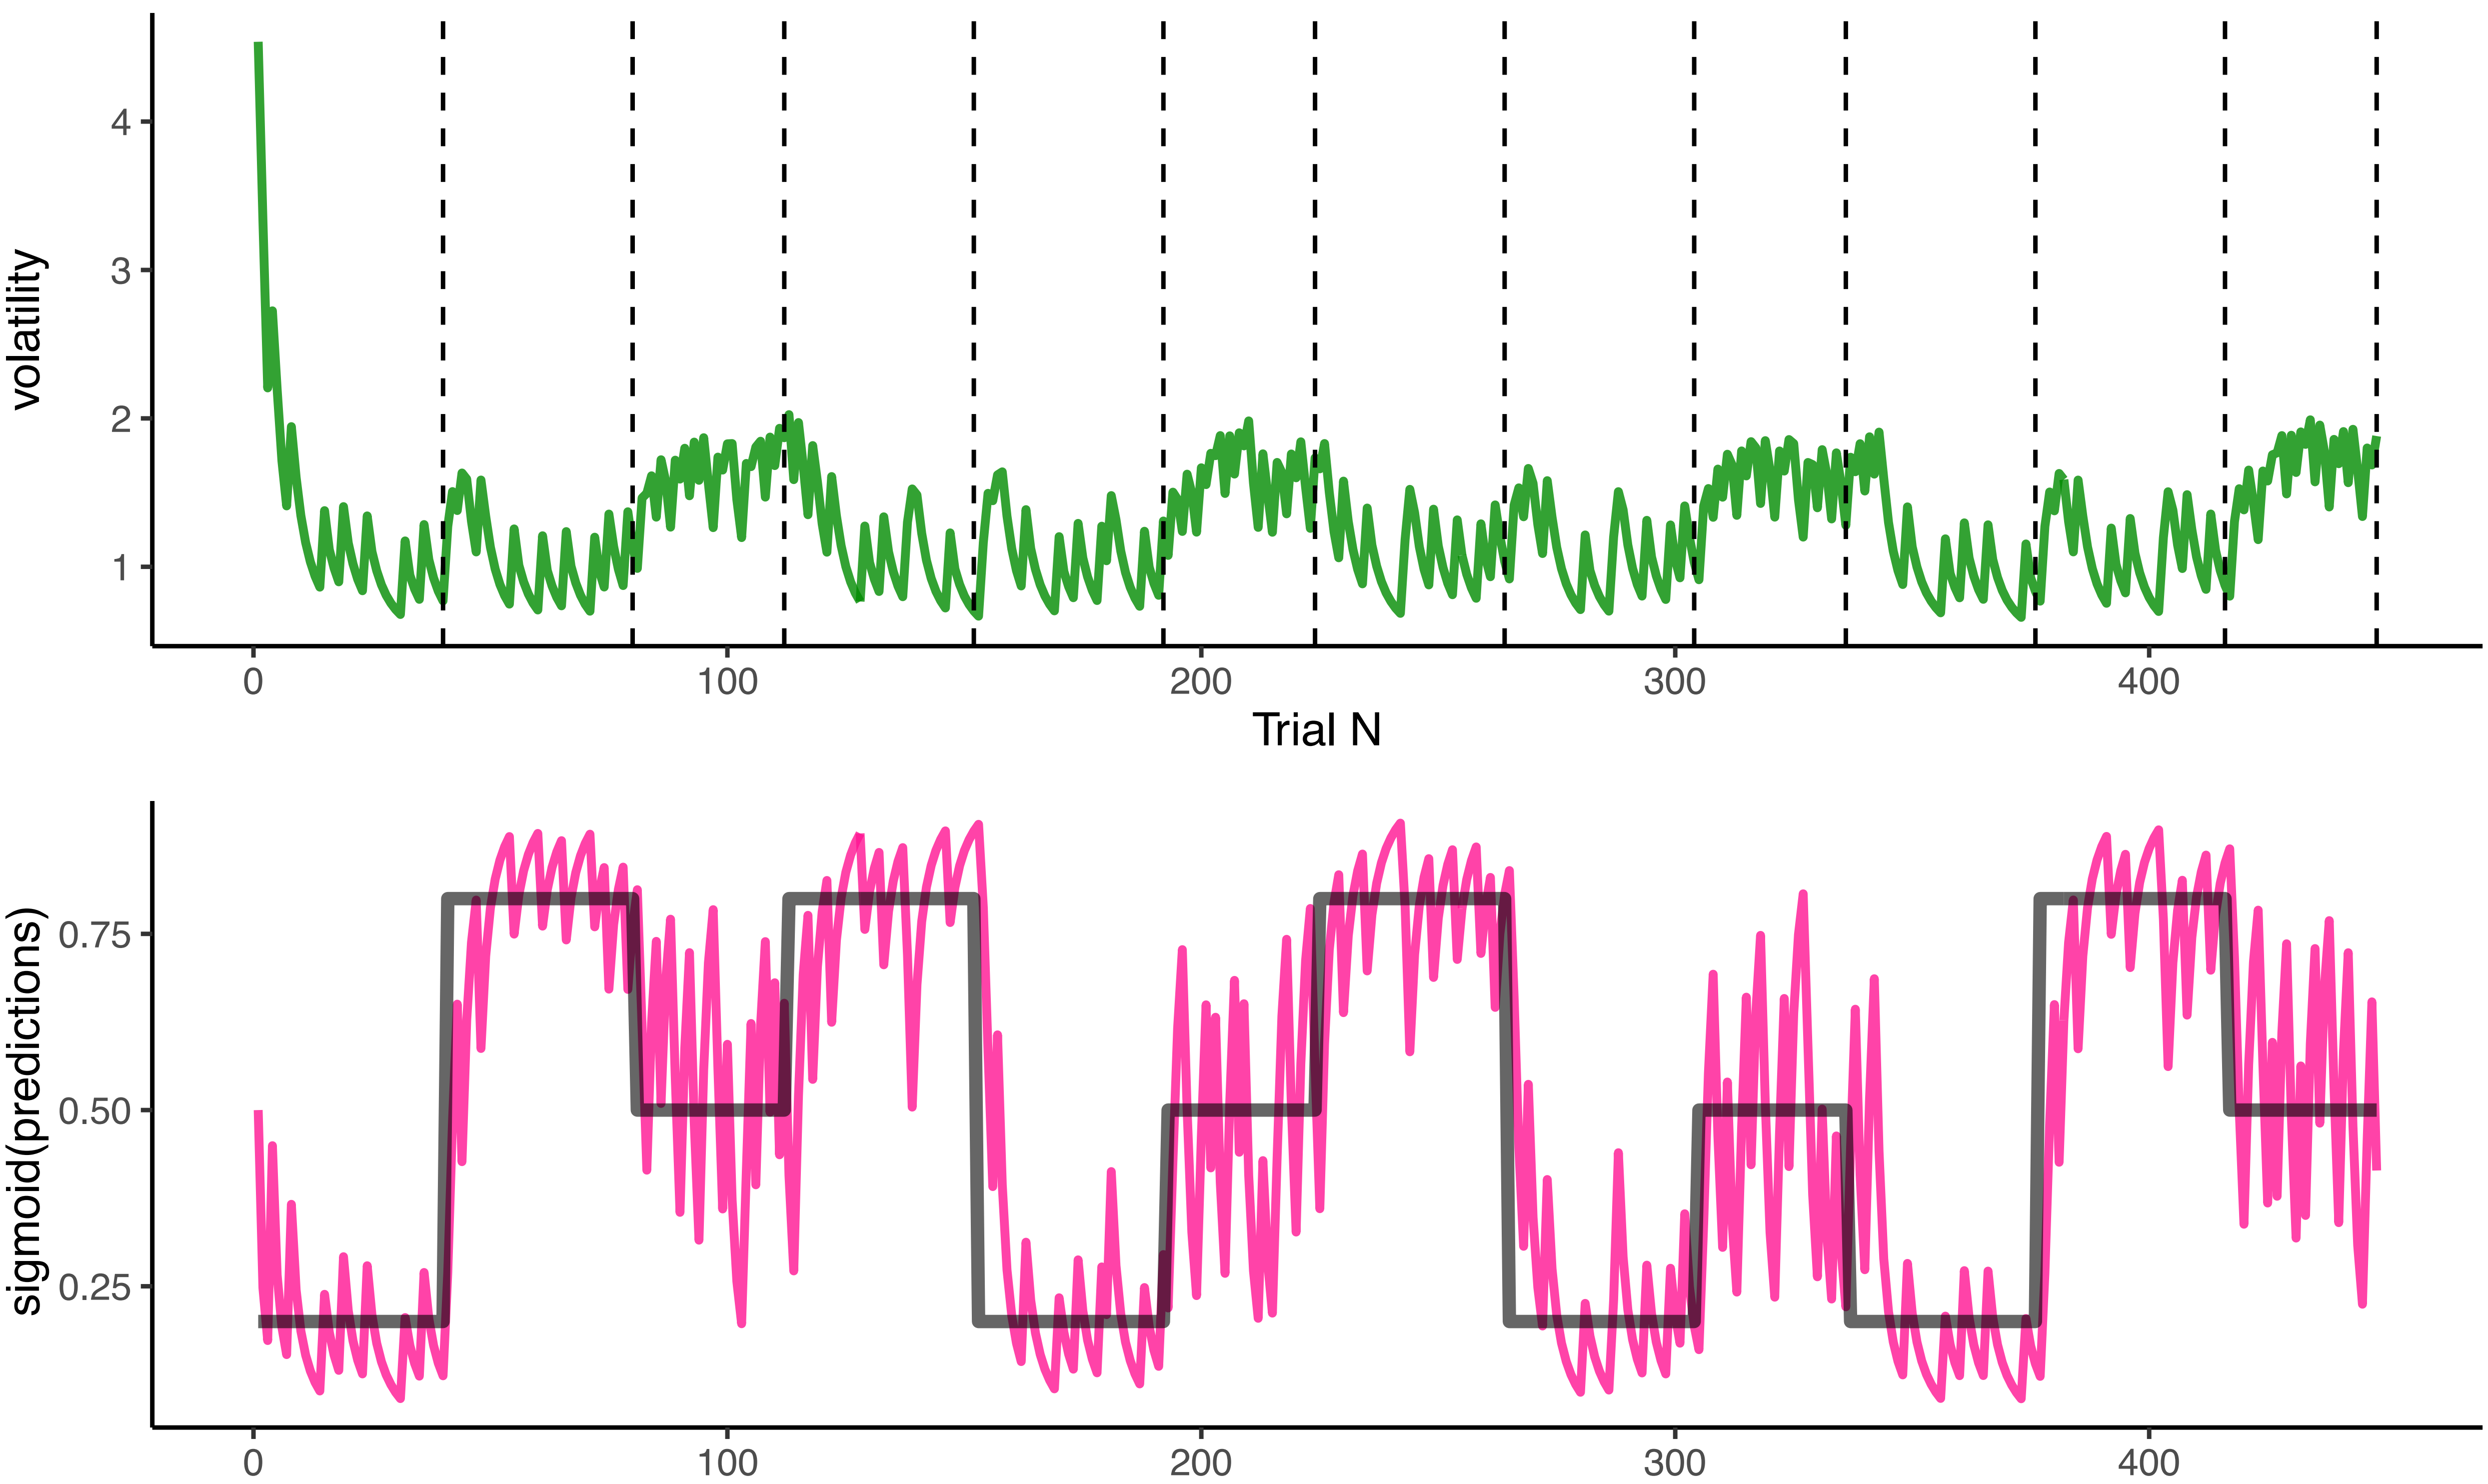

But…

omega → perception of volatility
v0 → initial volatility
sigma_v → volatility learning rate

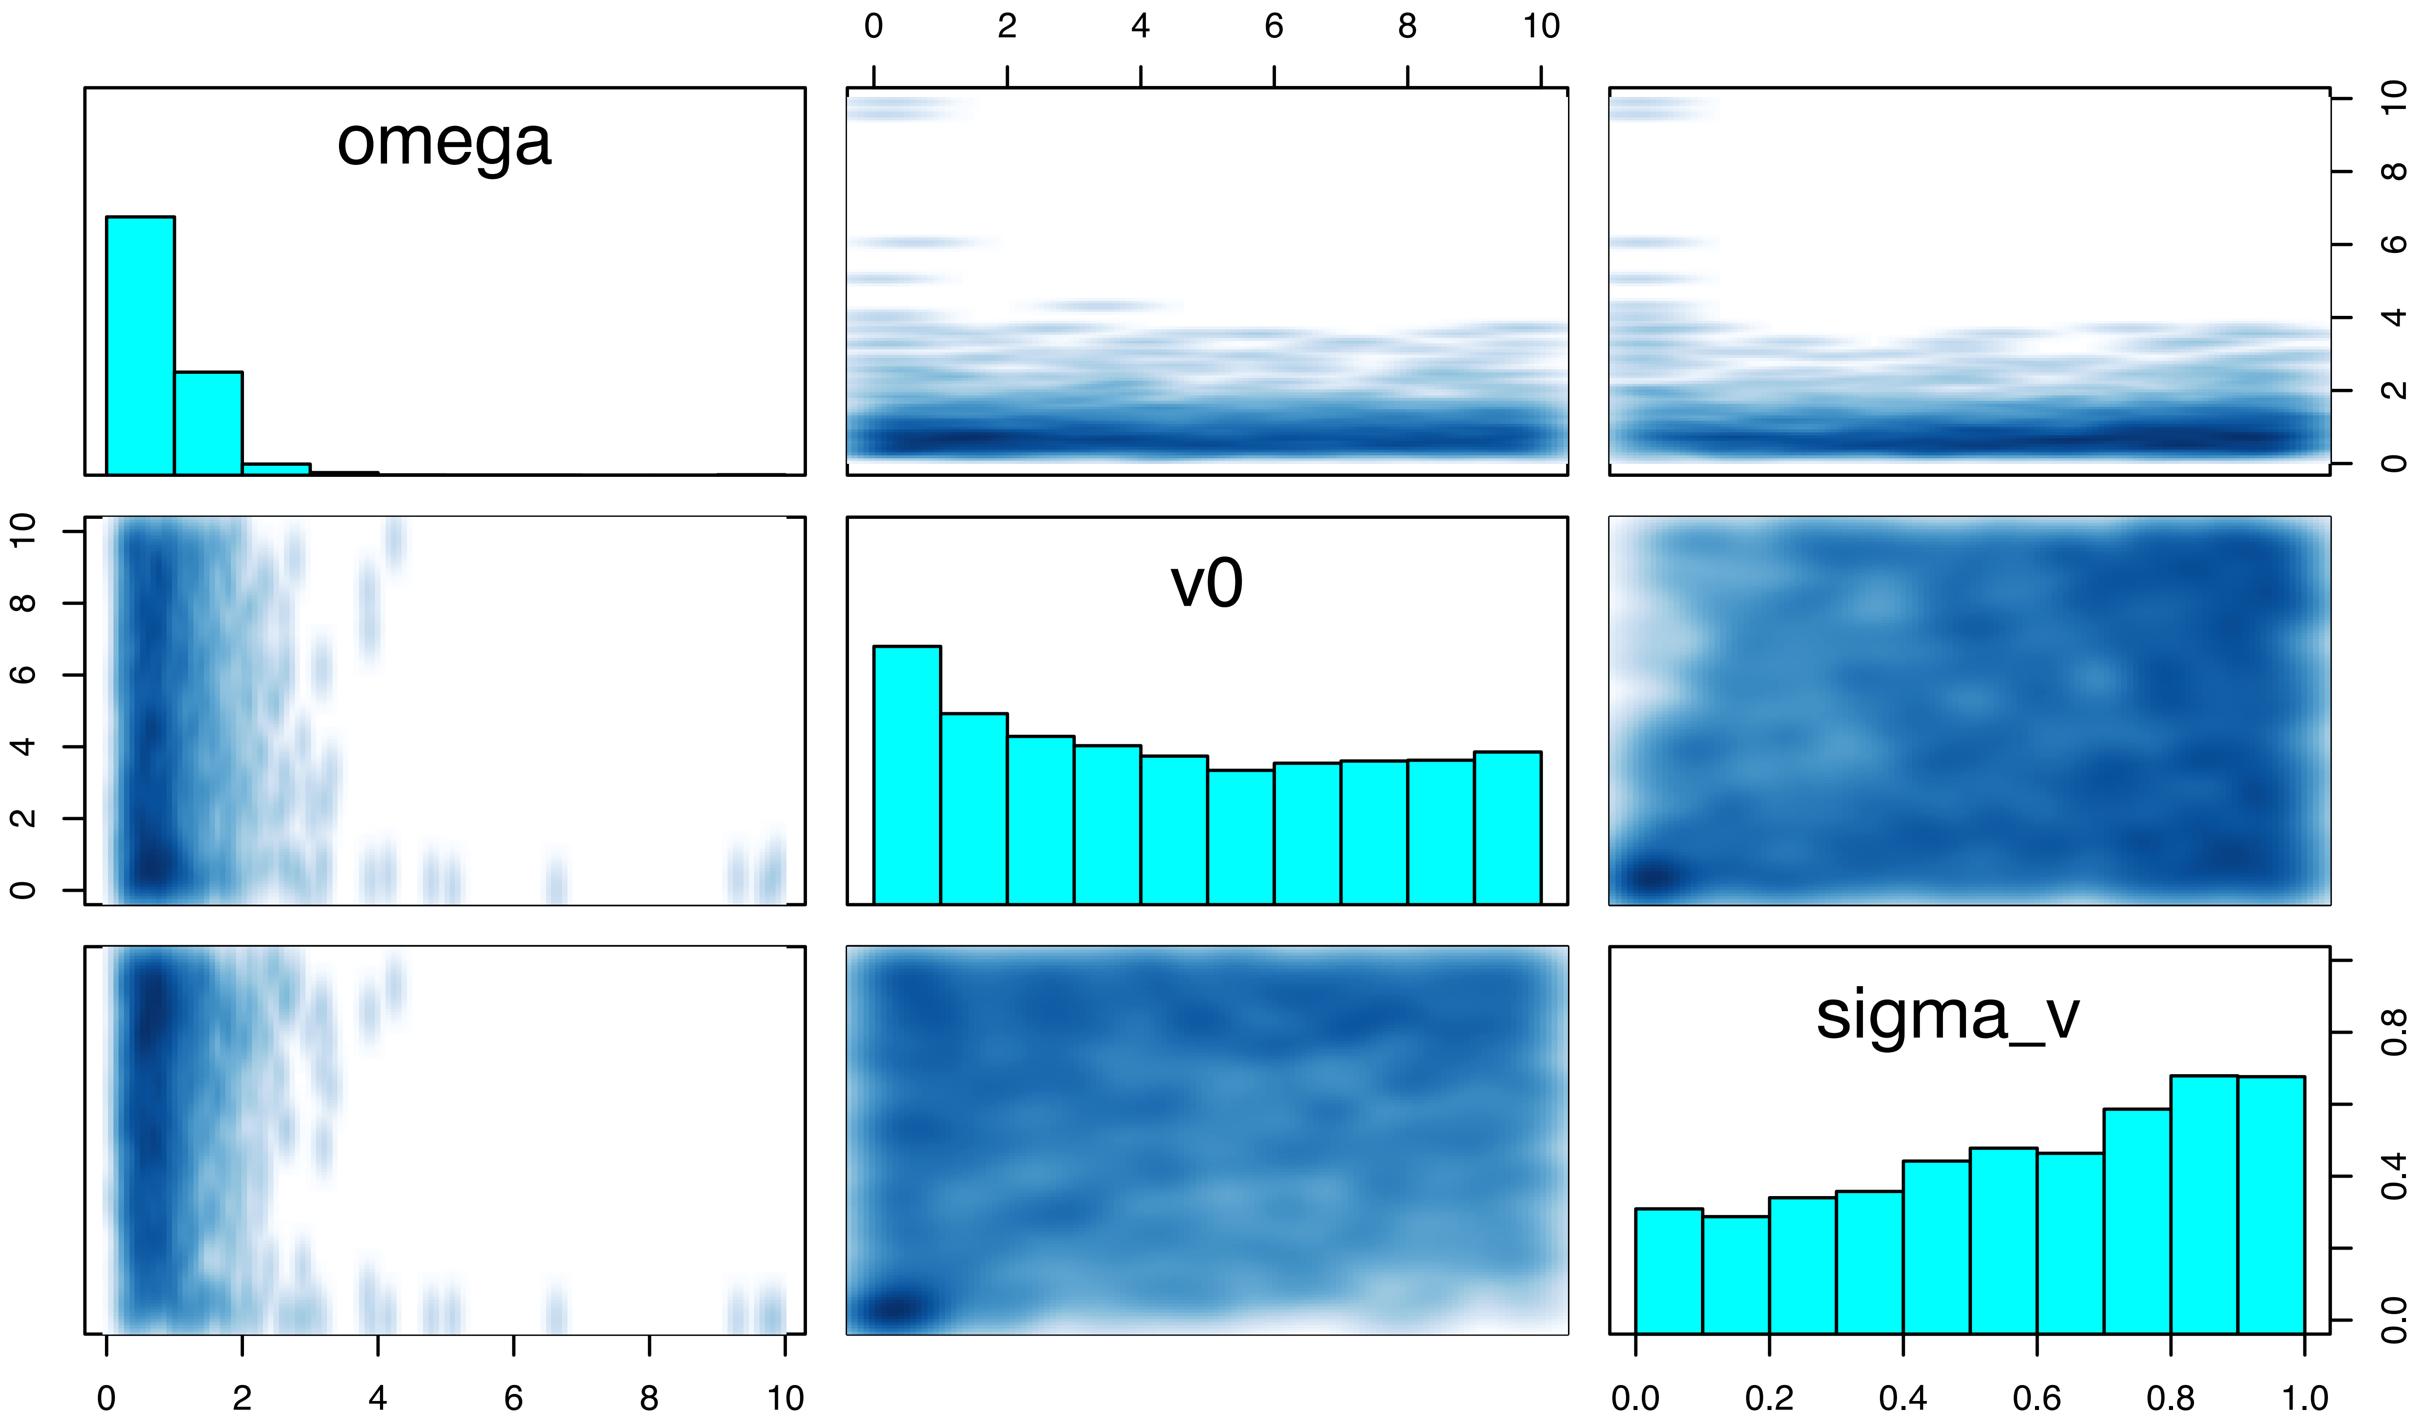

Simulation and Parameter Recovery

Examples with priors centred on the true values vs. not

Parameter Recovery

N_subj = 1, N_trials = 22400

Priors centred on true values

v0 → initial volatility
sigma_v volatility learning rate
omega → perception of volatility

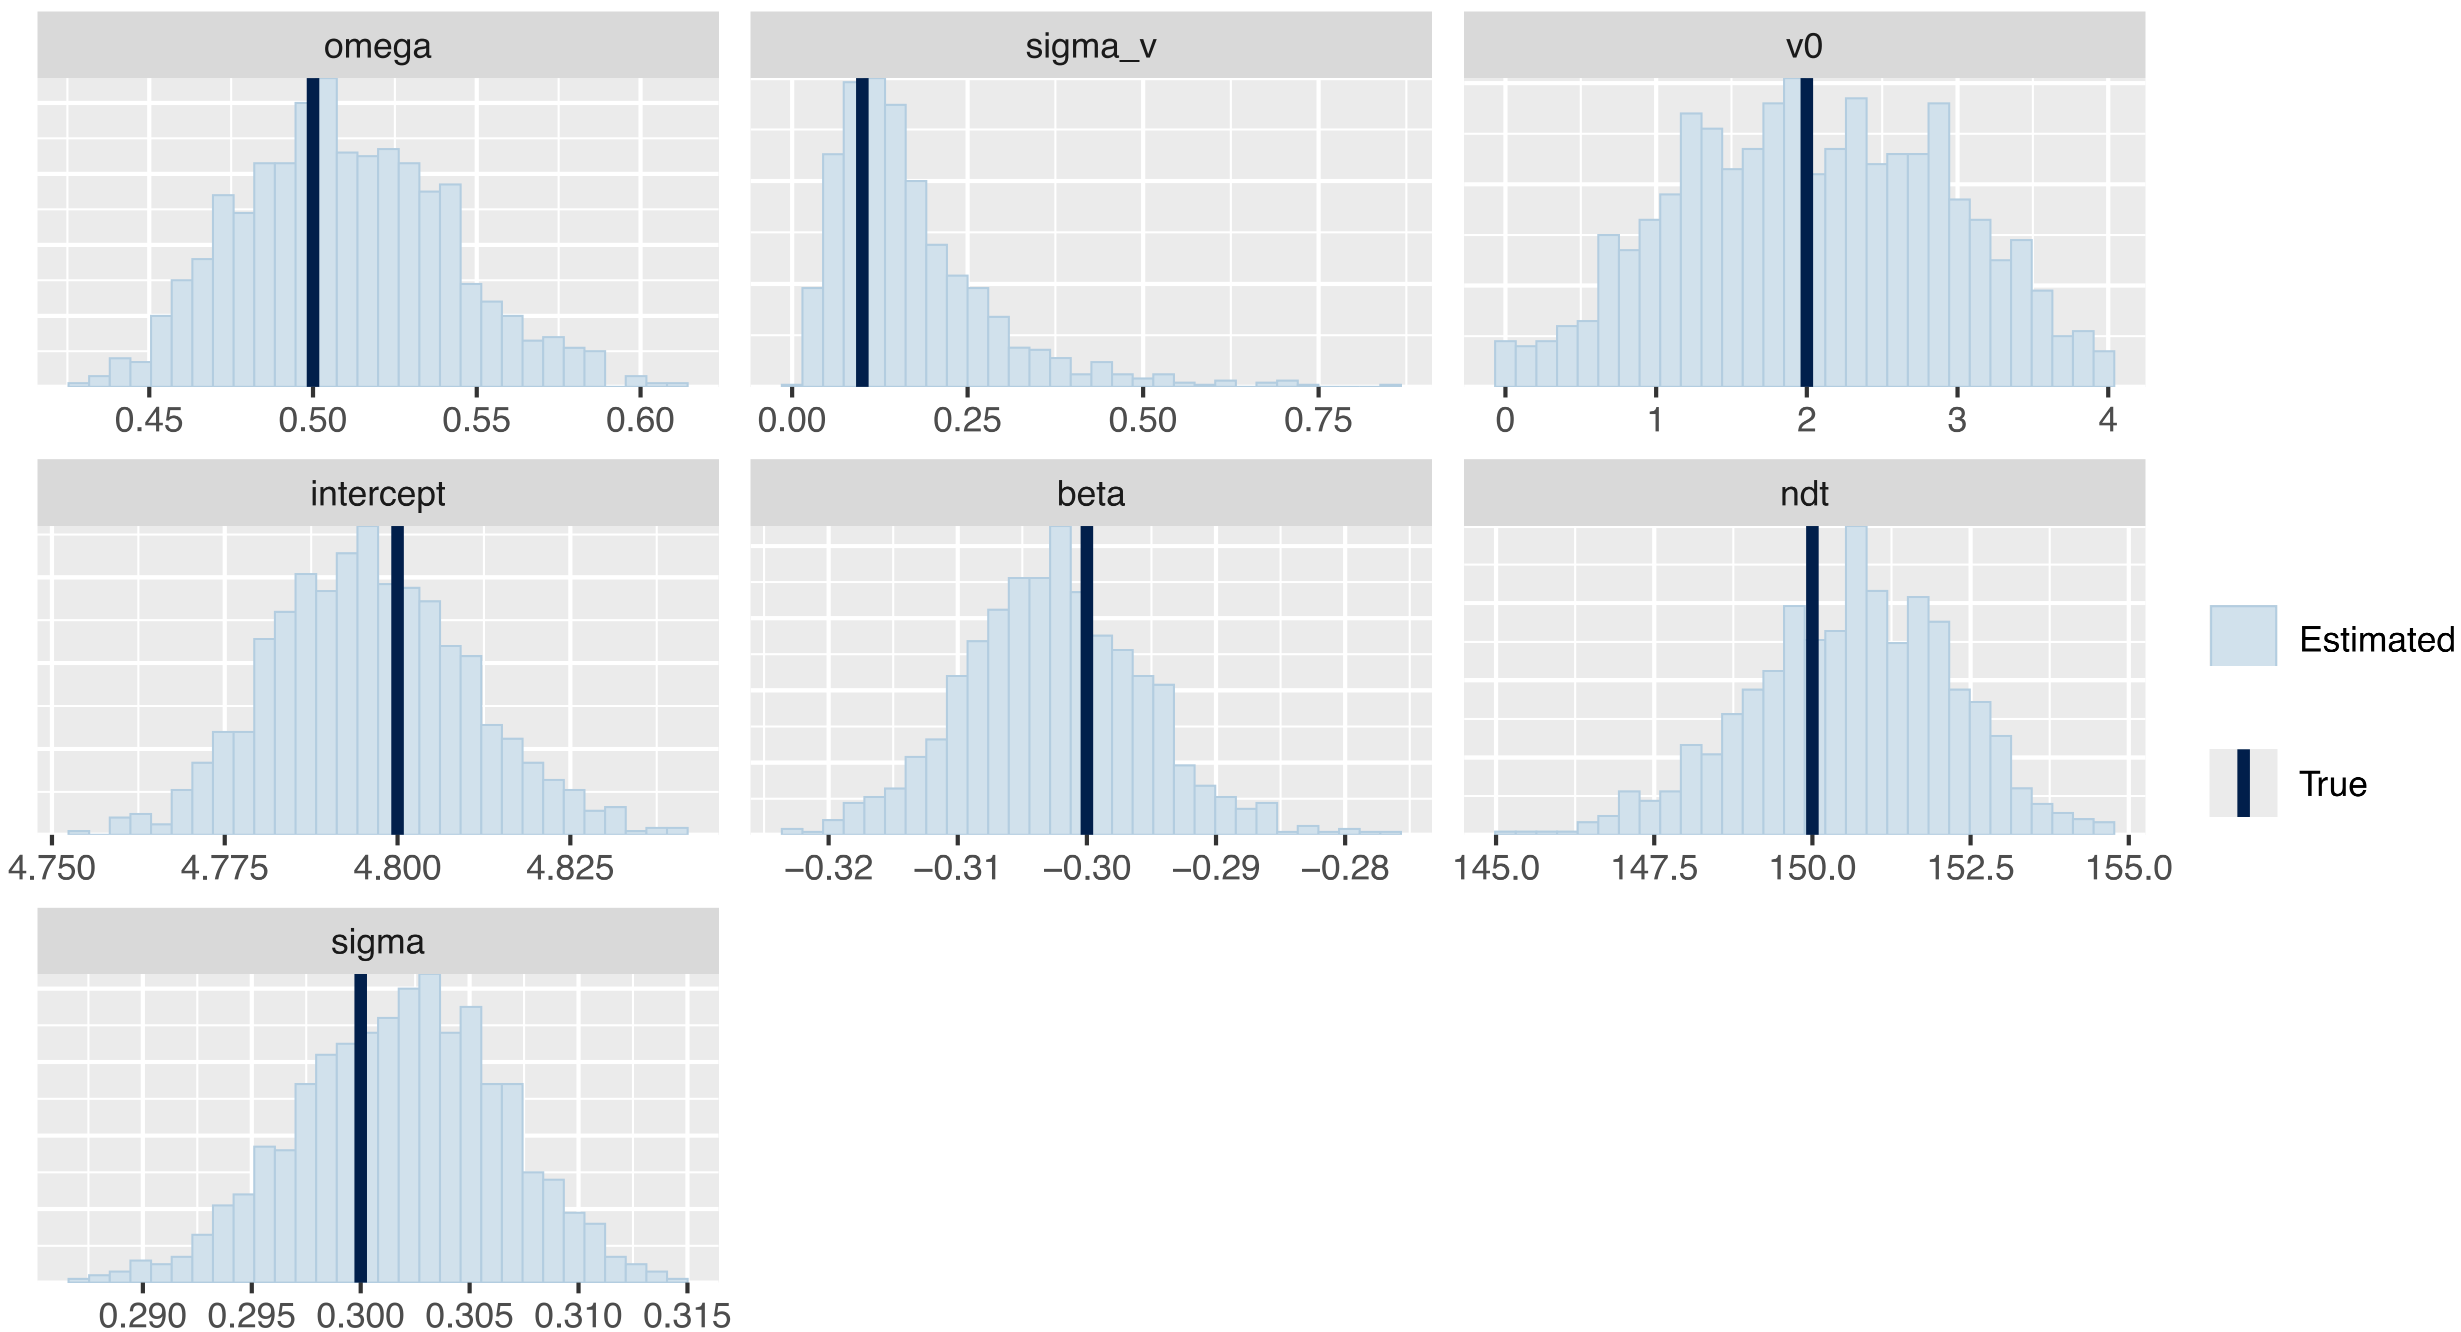

Parameter Recovery

N_subj = 1, N_trials = 22400

v0 → initial volatility
sigma_v volatility learning rate
omega → perception of volatility

Incorrect priors: omega ~ N(3,1); sigma_v ~ N(0.9,0.5); v0 ~ N(0.3,1)

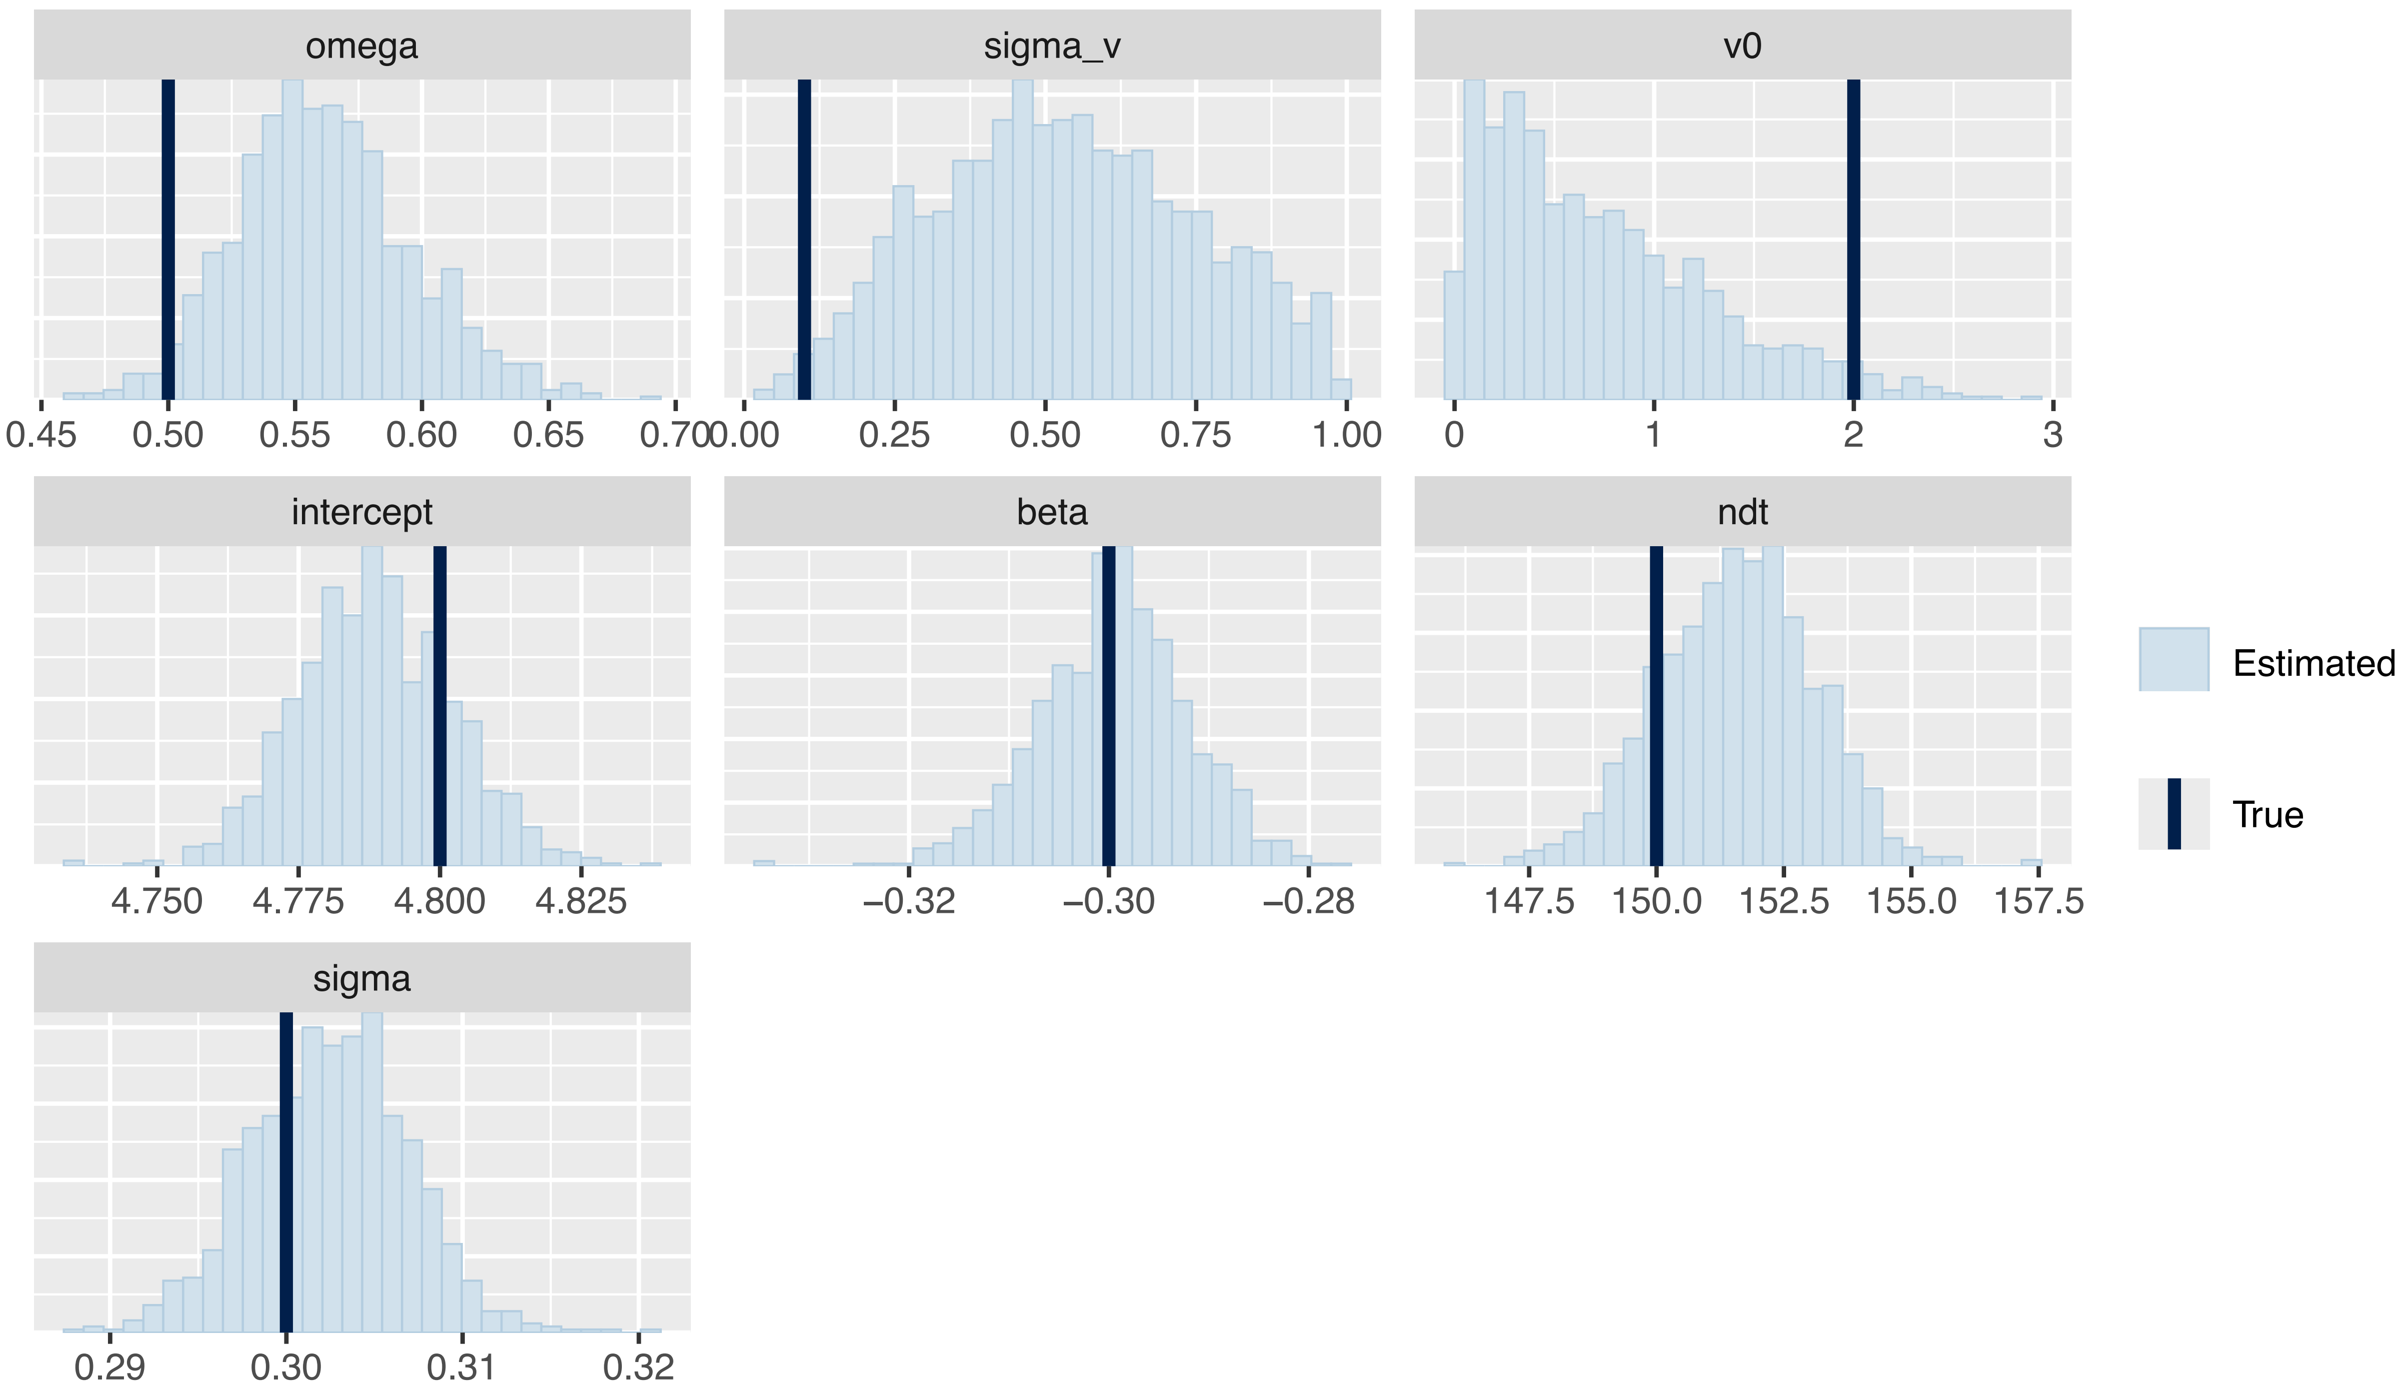

Perceptual Model

```
for (t in 1:N) {
    o = GO[t]; // input
    mpre = m;            // prediction
    wpre = w;            // variance
    predictions[t] = m;
    volatility[t] = v;        // volatility
    delta_m = o - sigmoid(mpre); // prediction error (pe)
    k = (wpre + v) / (wpre + v + omega); // Kalman Gain
    m = mpre + sqrt(wpre + v) * delta_m; // prediction update
    w = (1 - k) * (wpre + v);  // variance update
    wcov = (1 - k) * wpre; // covariance
    delta_v = (m-mpre)^2 + w + wpre - 2*wcov - v;  // volatility pe
    v         = v + sigma_v * delta_v;  // volatility update
}
```

Initial Values

w = **omega**

v0 = 2

m = 0 (i.e., .50)

Parameter Recovery

N_subj = 1, N_trials = 448

Priors centred on true values

v0 → 2
sigma_v volatility learning rate
omega → perception of volatility

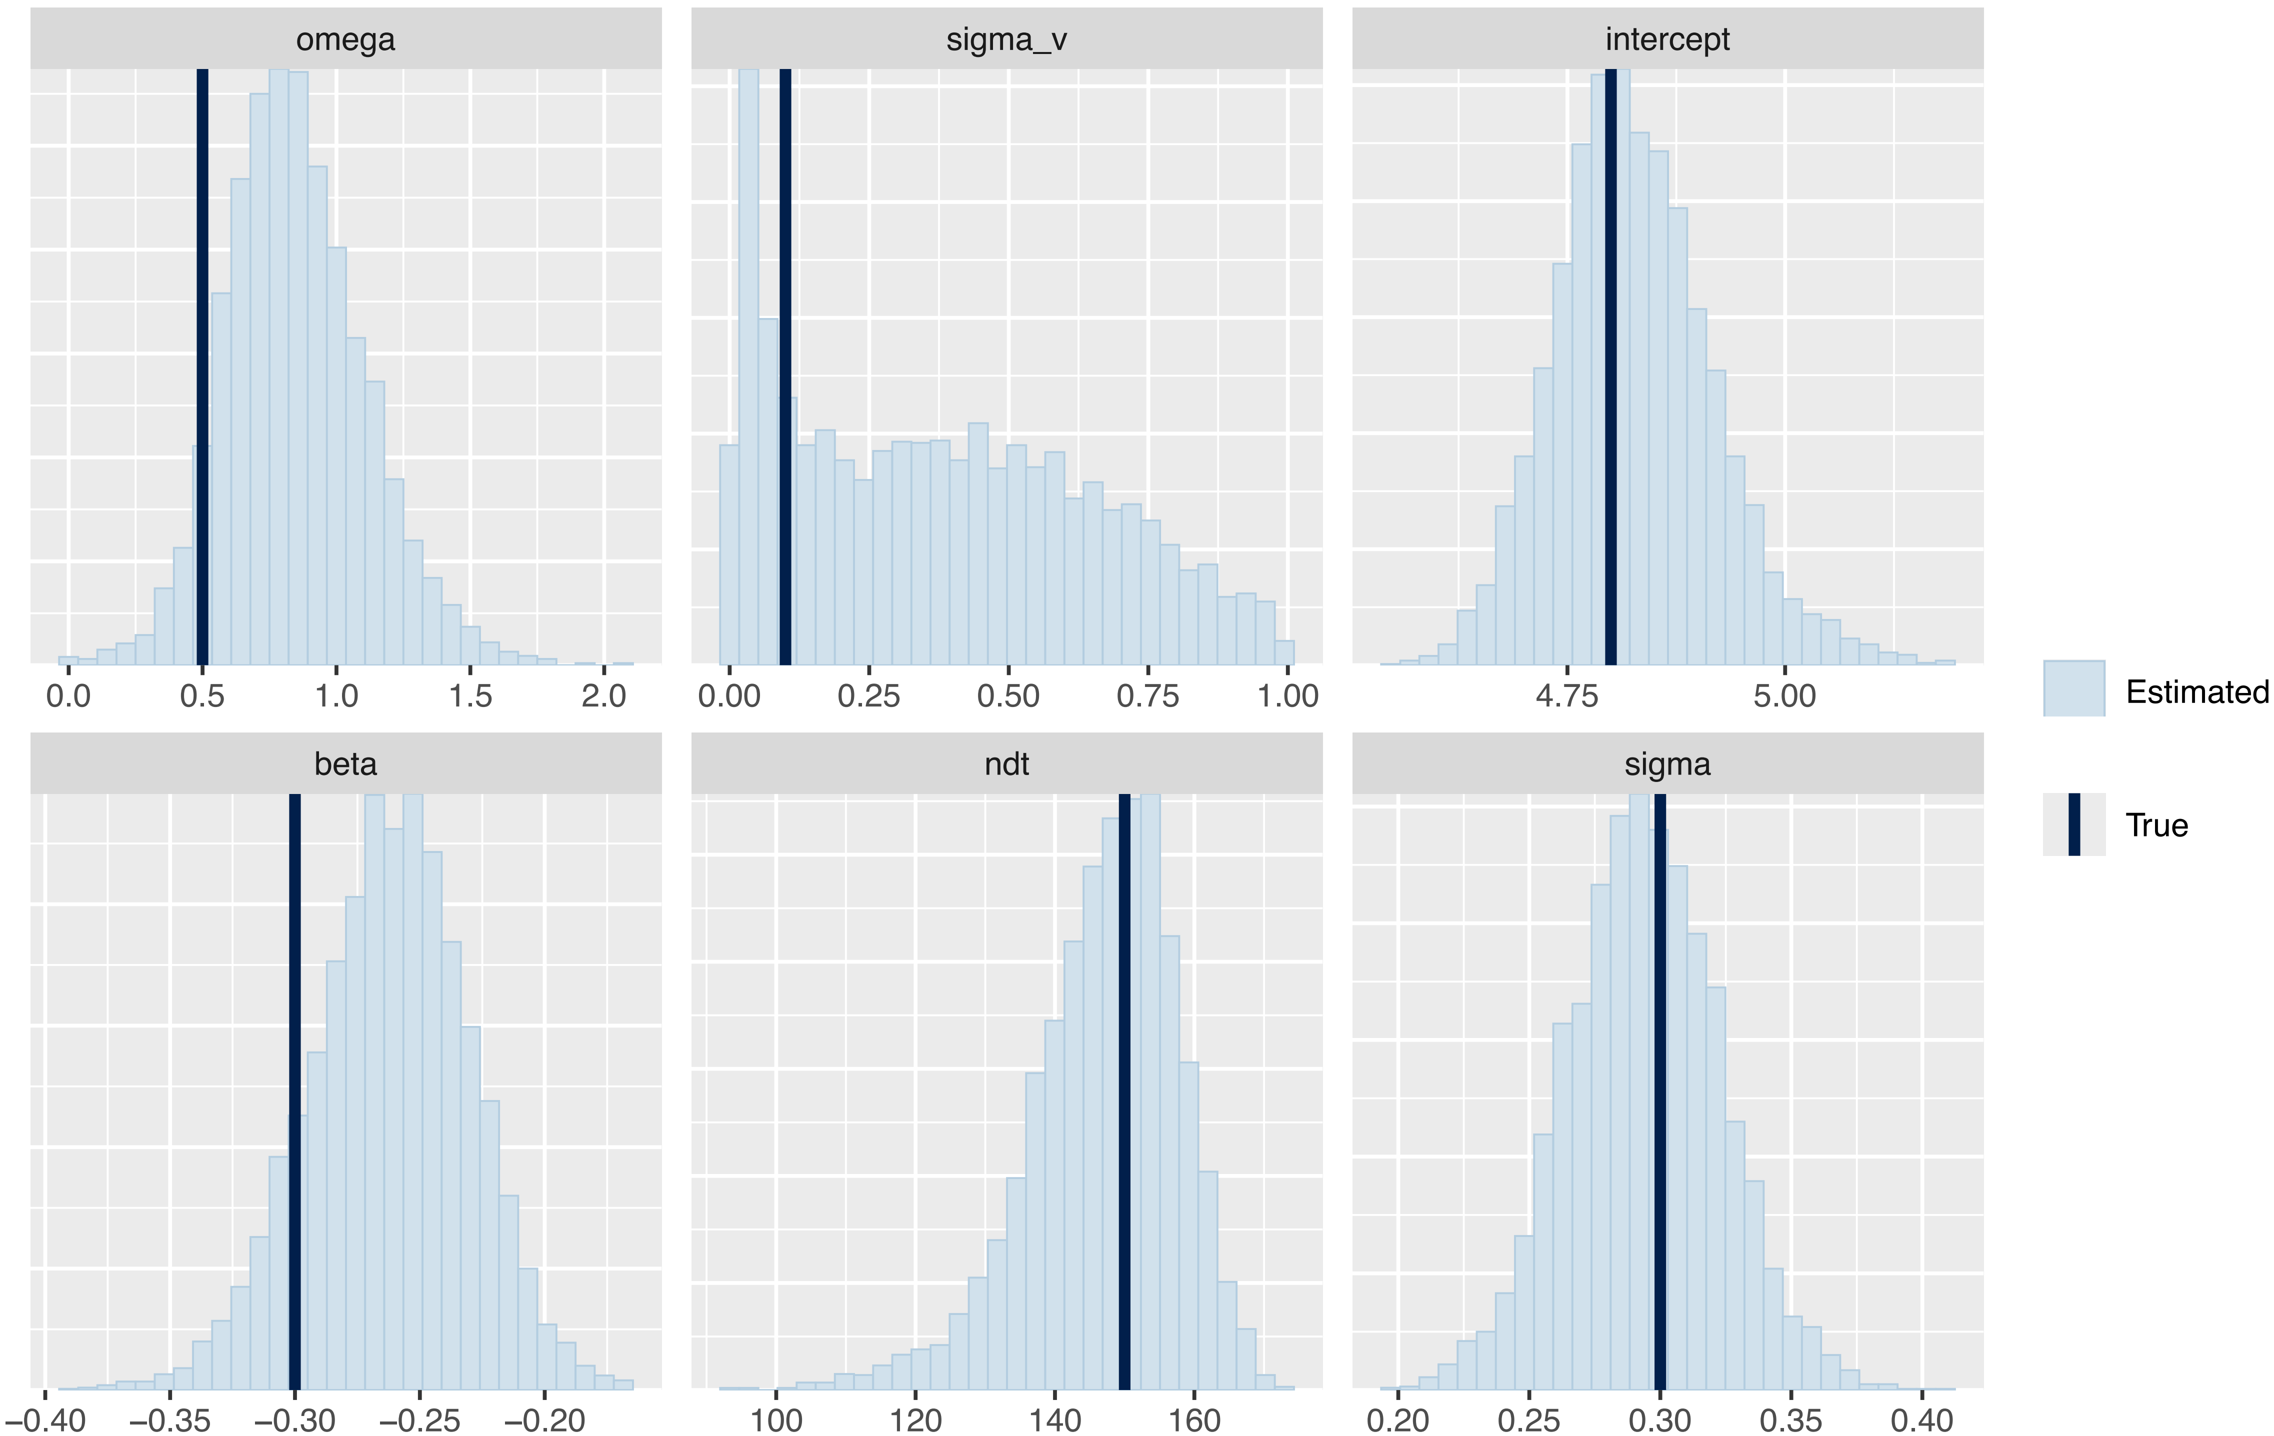

Parameter Recovery

N_subj = 1, N_trials = 448

v0 → 2
sigma_v volatility learning rate
omega → perception of volatility

Incorrect priors: omega ~ N(3,1); sigma_v ~ N(0.9,0.5)

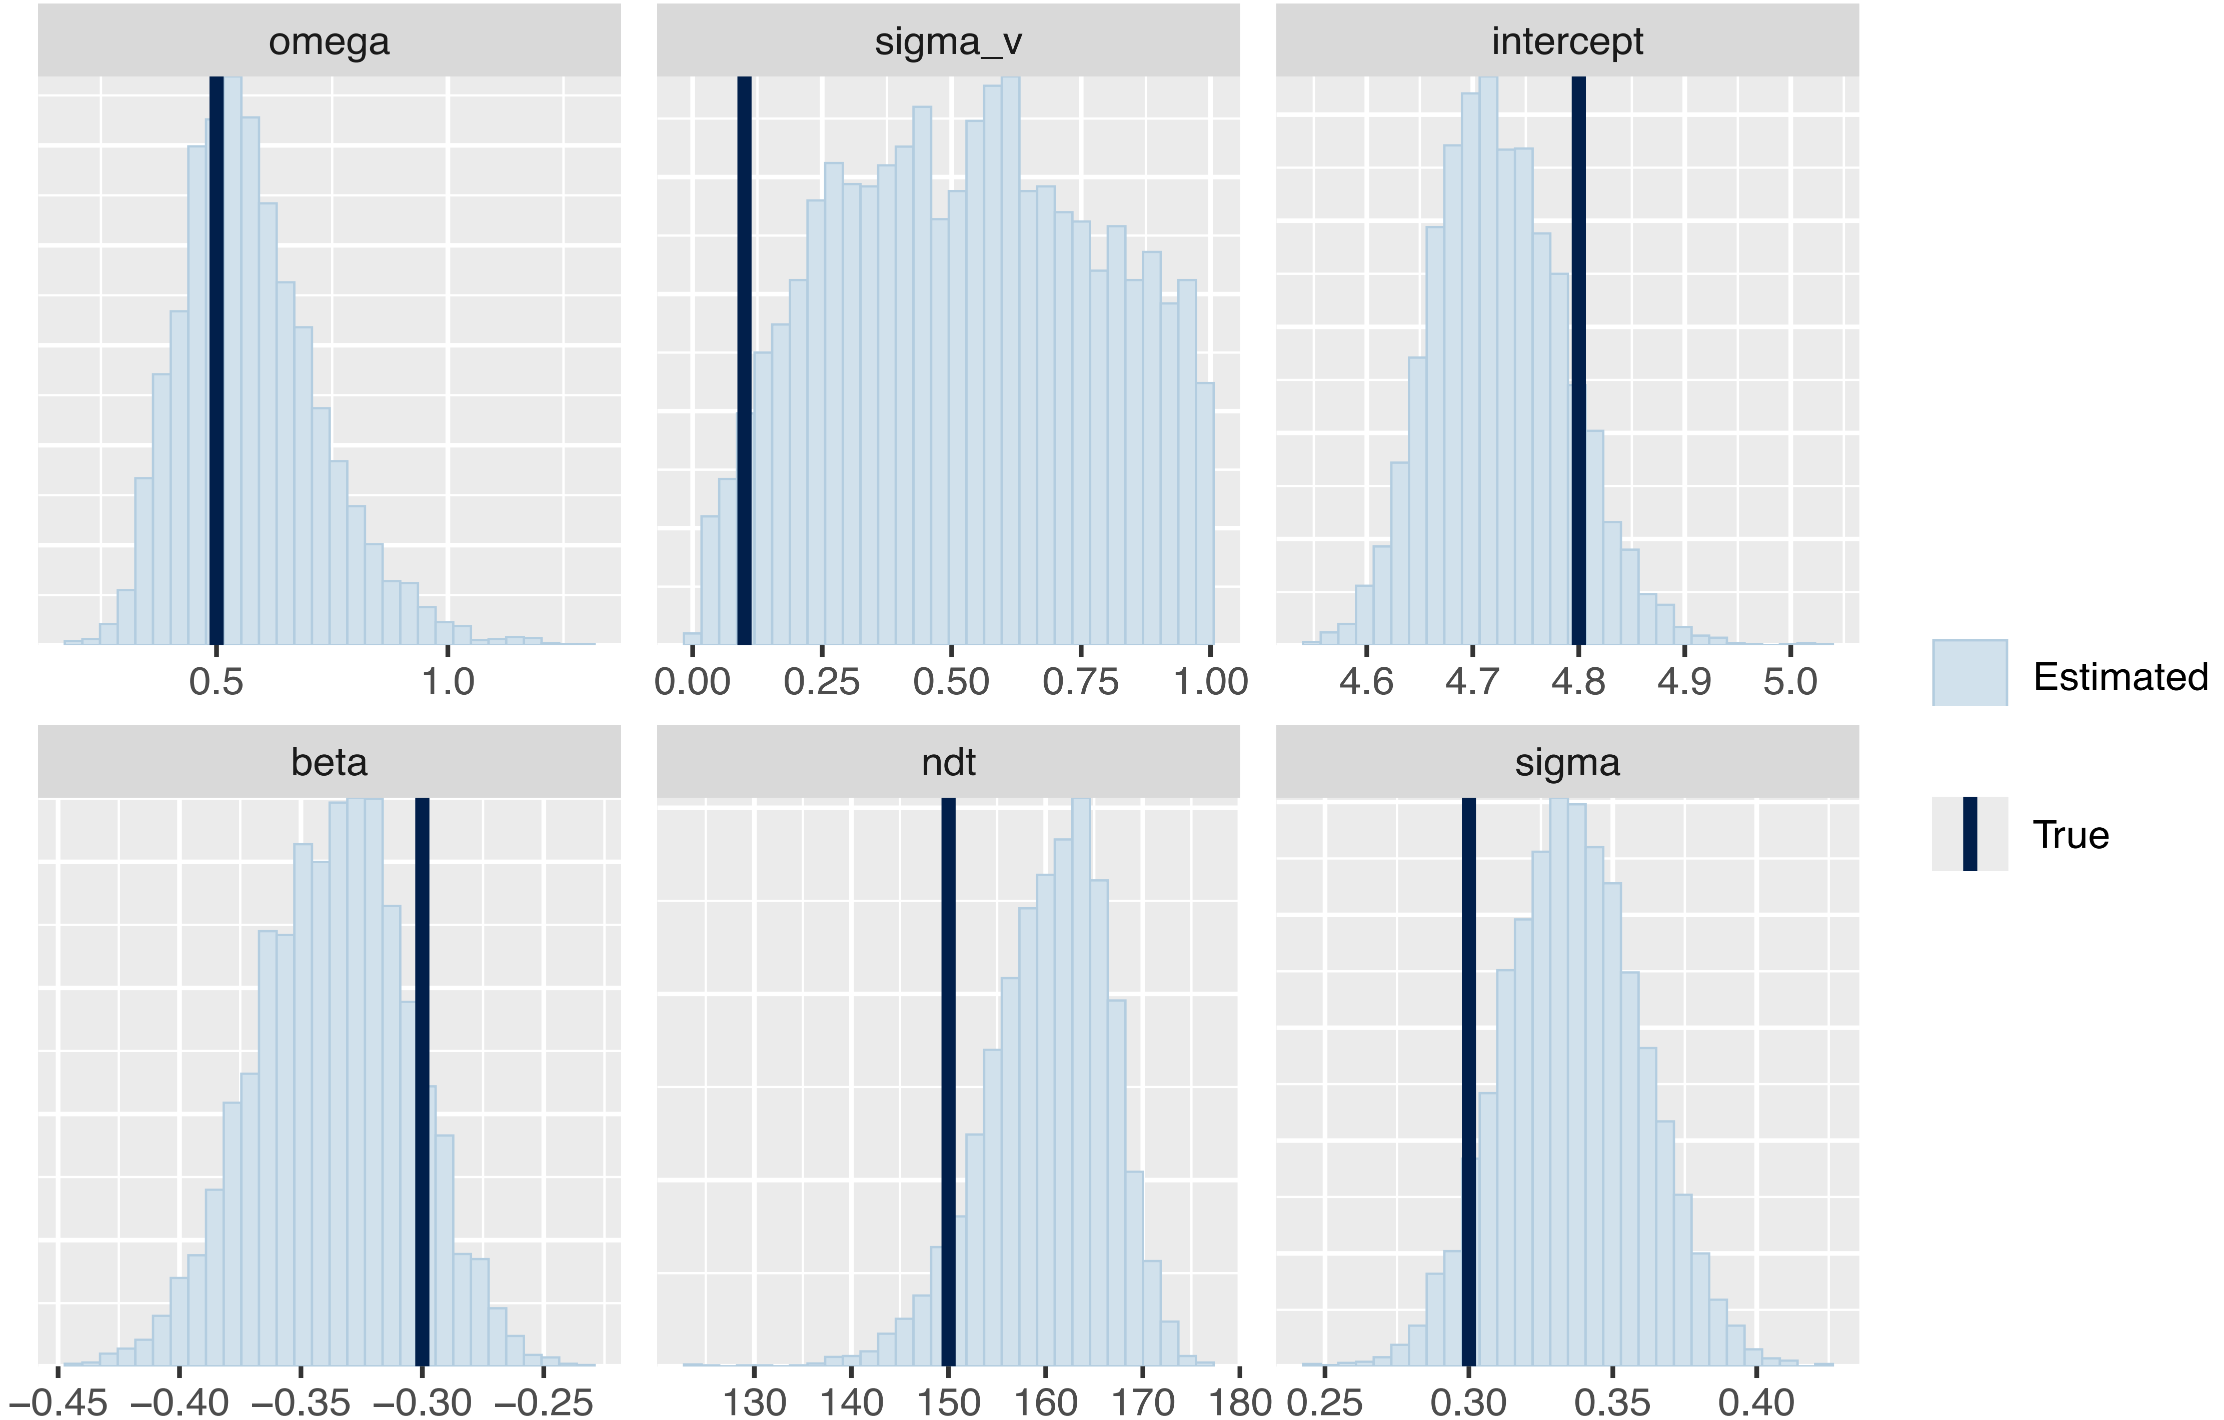

Parameter Recovery

N_subj = 1, N_trials = 4480

v0	→	2
sigma_v		volatility learning rate
omega	→	perception of volatility

Incorrect priors: omega ~ N(3,1); sigma_v ~ N(0.9,0.5)

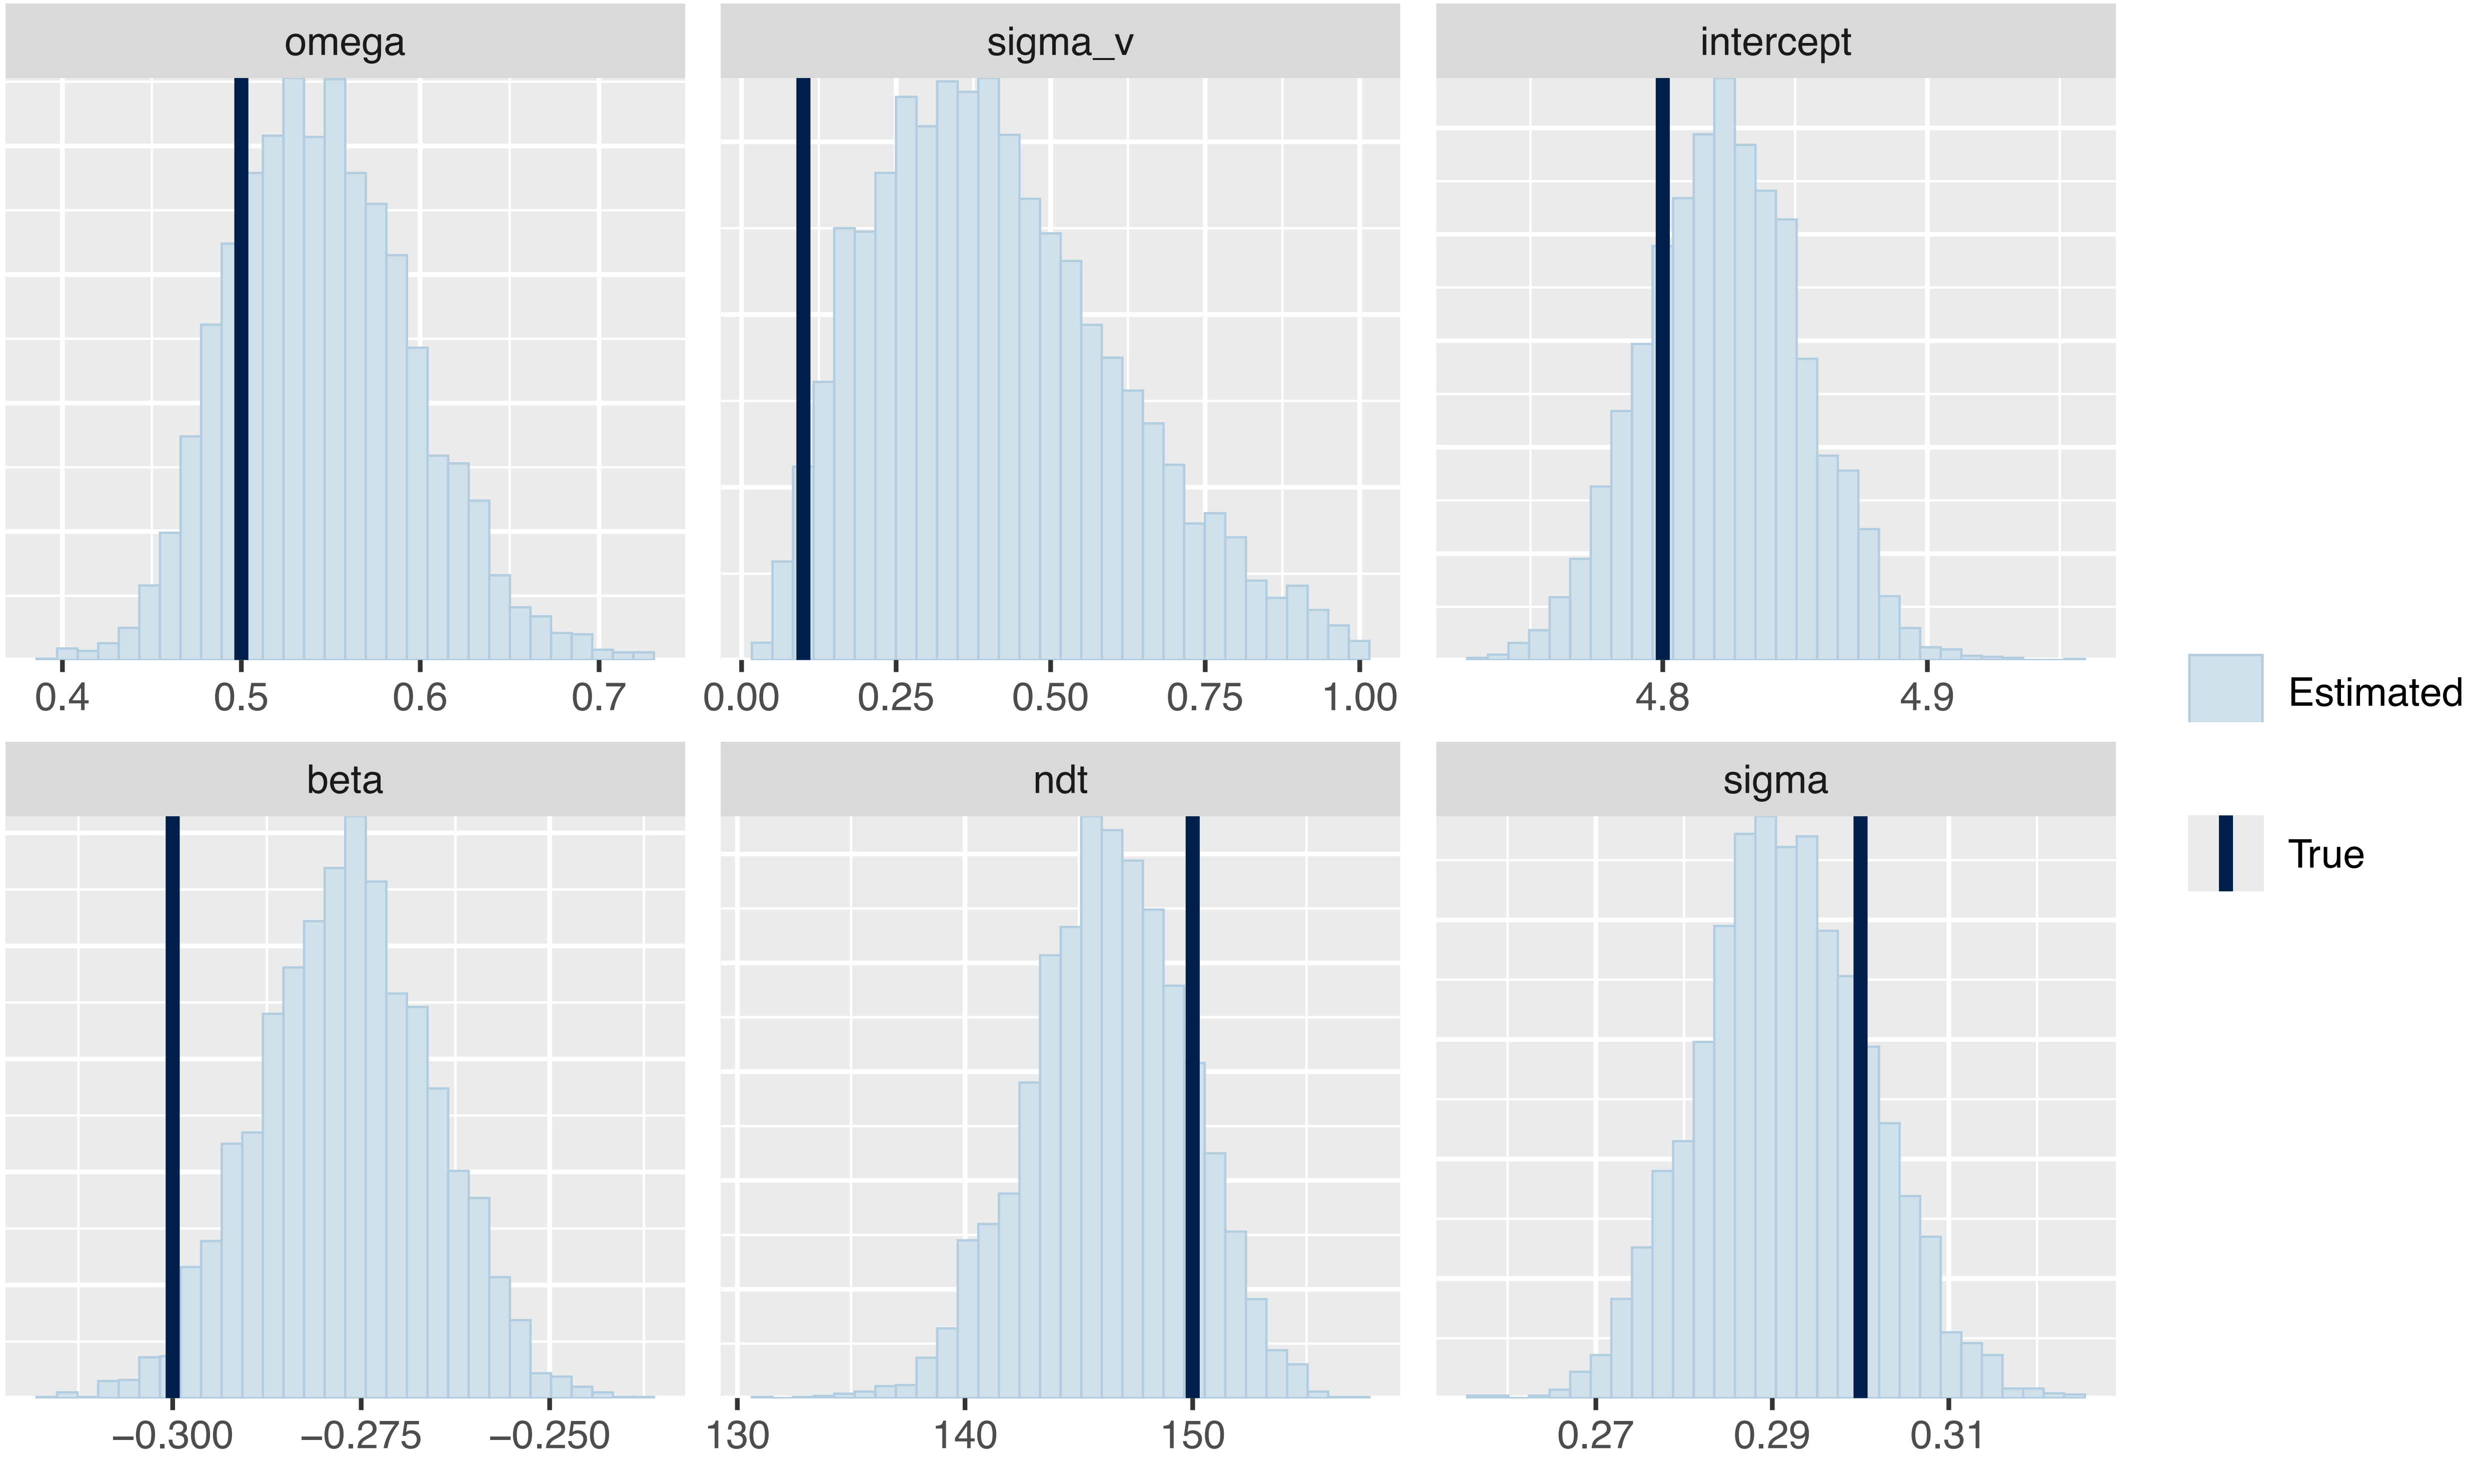

N_subj = 30; N_trials = 448; V0 = 4;

Priors on true values: **intercept** = 5.30; **ndt** = 150; **beta** = -0.3; **omega** = 0.5; **sigma_v** = 0.1

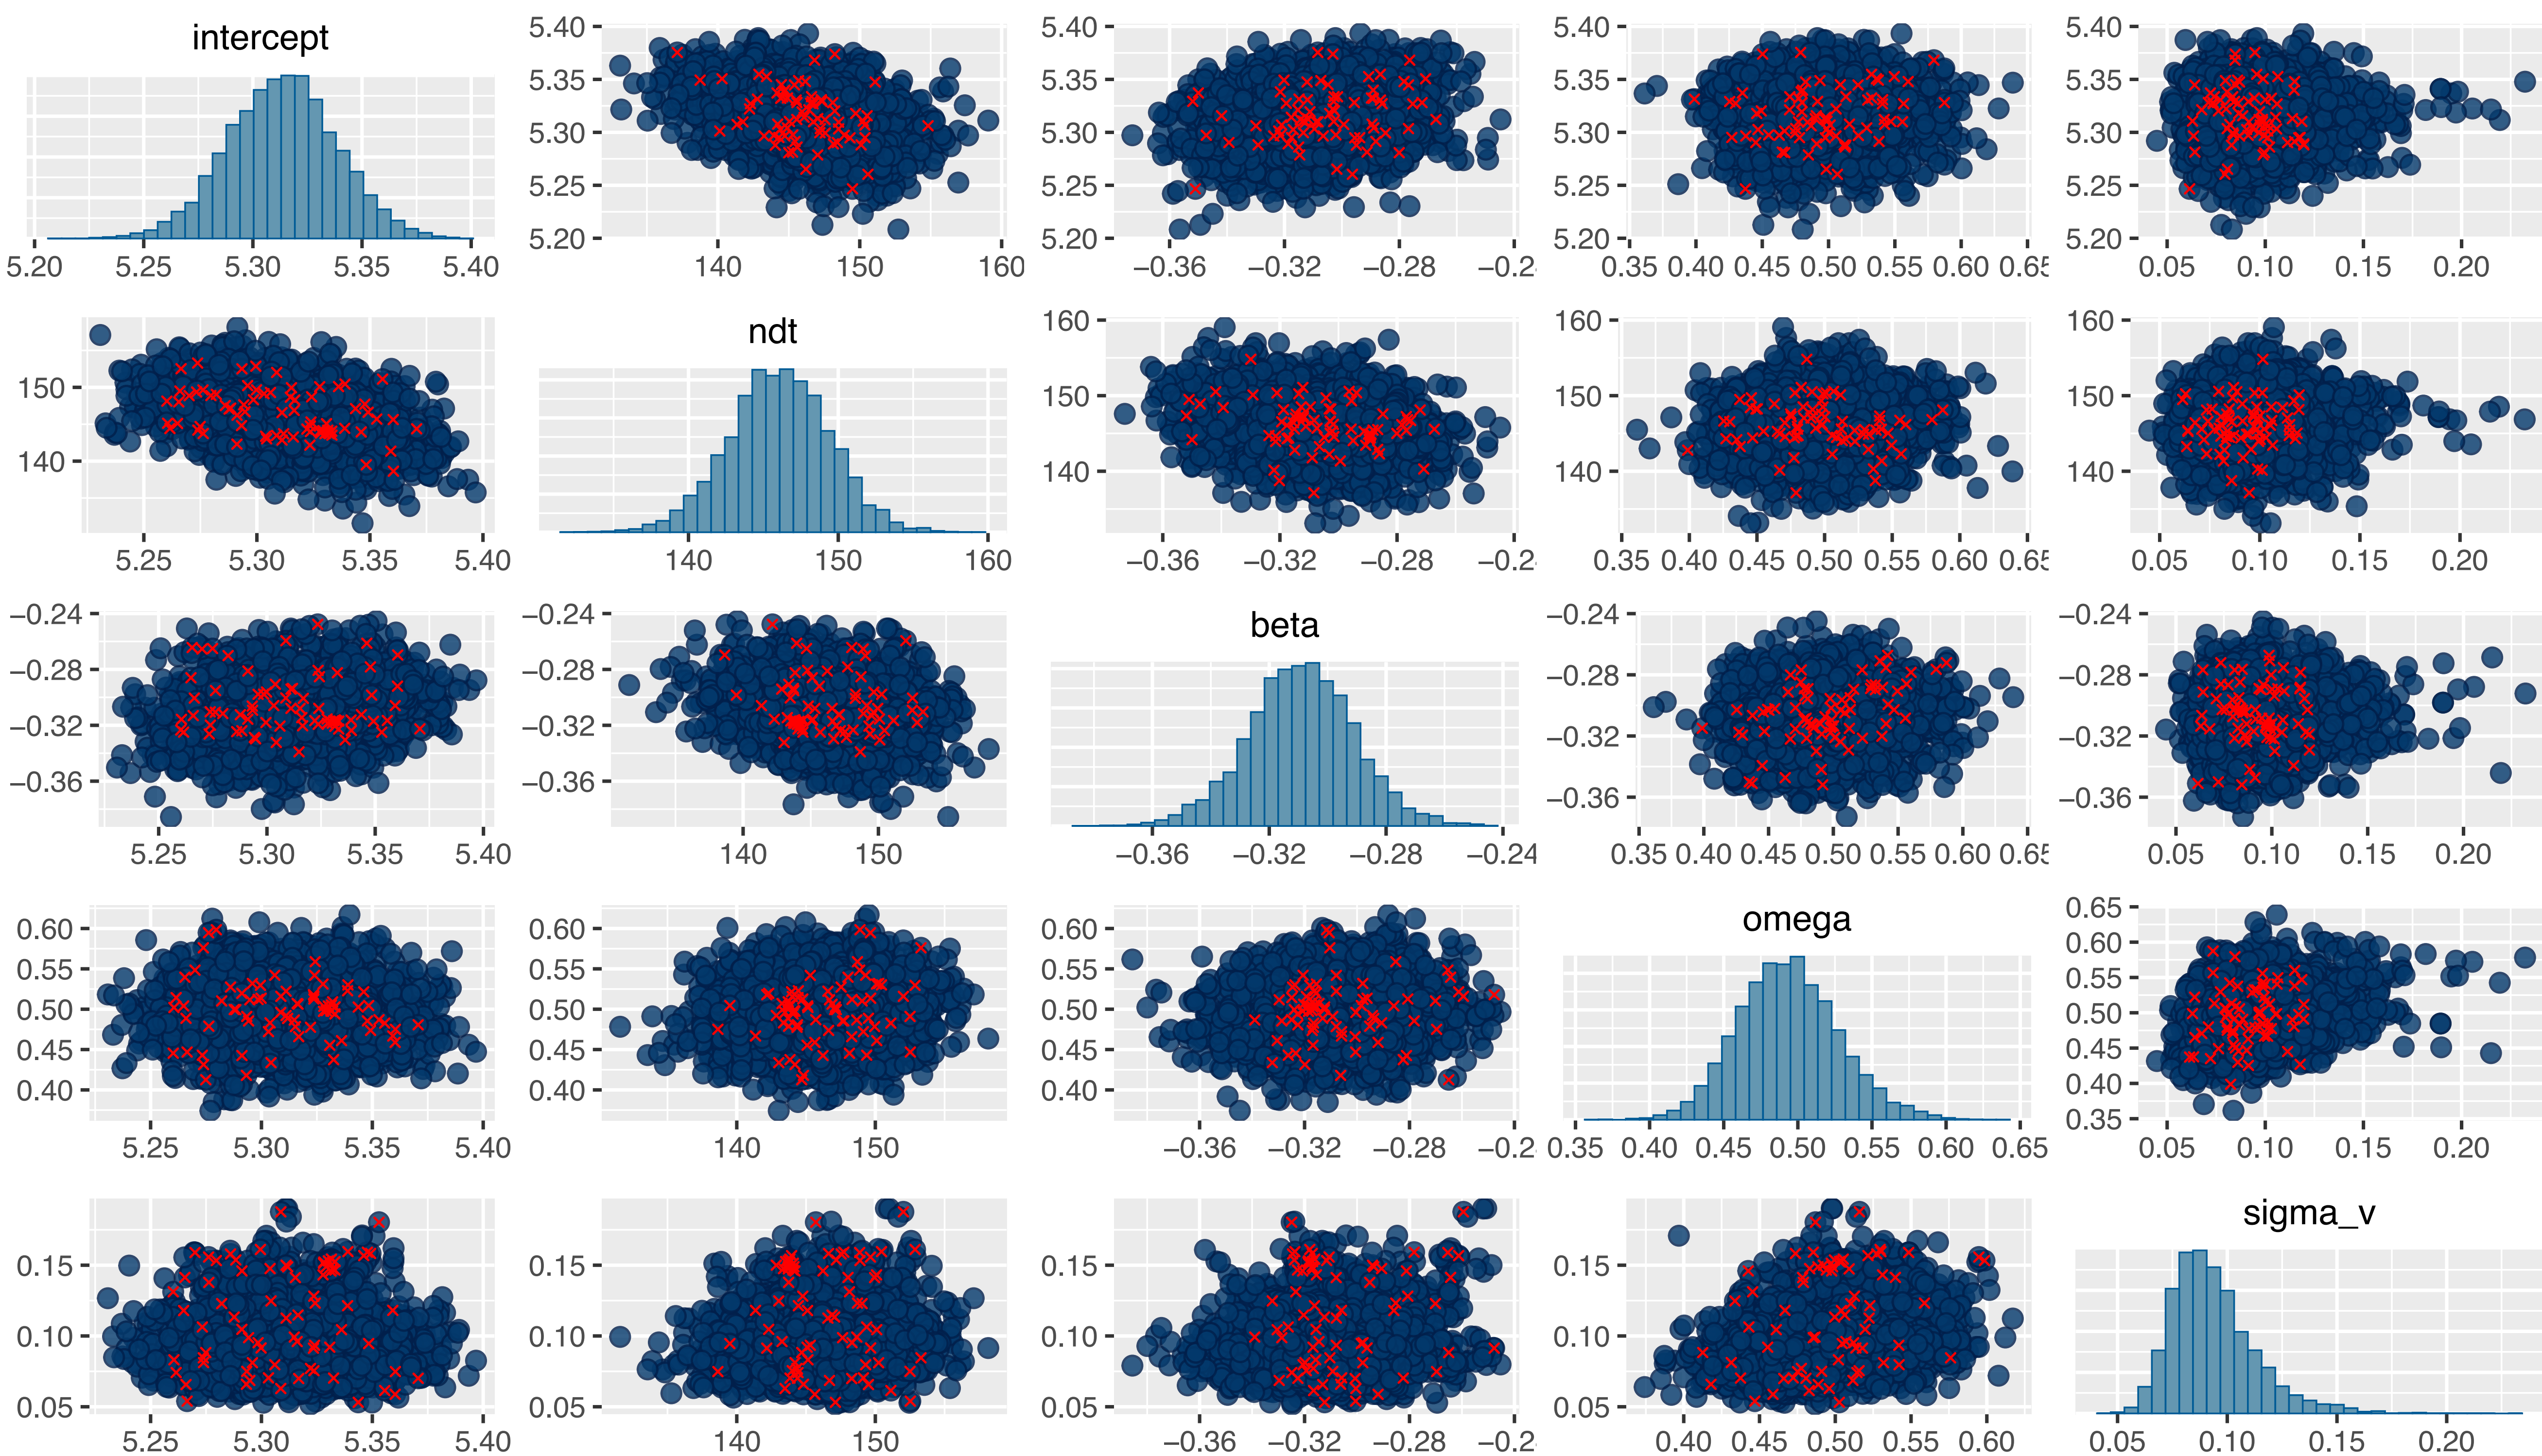

Take home message:

Cognitive modeling is cool but…

Test models before trust them!

Thanks!

Roberta Sellaro

Nicola Cellini

Antonino Visalli

References

- Becker, A., (2023). *Kalman Filter from the Ground to Up*, https://www.kalmanfilter.net/default.aspx
- Behrens, T. E., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. (2007). Learning the value of information in an uncertain world. *Nature neuroscience*, *10*(9), 1214–1221. https://doi.org/10.1038/nn1954
- Forsgren, M., Juslin, P., & van den Berg, R. (2023). Further perceptions of probability: In defence of associative models. *Psychological review*, *130*(5), 1383–1400. https://doi.org/10.1037/rev0000410
- Frässle, S., et al. (2021). TAPAS: An Open-Source Software Package for Translational Neuromodeling and Computational Psychiatry. Frontiers in Psychiatry, 12:680811. https://doi.org/10.3389/fpsyt.2021.680811
- Gallistel, C. R., Krishan, M., Liu, Y., Miller, R., & Latham, P. E. (2014). The perception of probability. *Psychological review*, *121*(1), 96–123. https://doi.org/10.1037/a0035232
- Kruschke, J. K. (2008). Bayesian approaches to associative learning: From passive to active learning. *Learning & behavior*, *36*(3), 210-226.
- Mathys, C. D., Lomakina, E. I., Daunizeau, J., Iglesias, S., Brodersen, K. H., Friston, K. J., & Stephan, K. E. (2014). Uncertainty in perception and the Hierarchical Gaussian Filter. *Frontiers in human neuroscience*, *8*, 825.
- Nassar, M. R., Wilson, R. C., Heasly, B., & Gold, J. I. (2010). An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *30*(37), 12366–12378. https://doi.org/10.1523/JNEUROSCI.0822-10.2010
- Nassar, M. R., Waltz, J. A., Albrecht, M. A., Gold, J. M., & Frank, M. J. (2021). All or nothing belief updating in patients with schizophrenia reduces precision and flexibility of beliefs. *Brain : a journal of neurology*, *144*(3), 1013–1029. https://doi.org/10.1093/brain/awaa453
- Nassar, M. R., Bruckner, R., Gold, J. I., Li, S. C., Heekeren, H. R., & Eppinger, B. (2016). Age differences in learning emerge from an insufficient representation of uncertainty in older adults. *Nature communications*, *7*(1), 11609.
- Piray, P., & Daw, N. D. (2020). A simple model for learning in volatile environments. *PLoS computational biology*, *16*(7), e1007963.
- Piray, P., & Daw, N. D. (2021). A model for learning based on the joint estimation of stochasticity and volatility. *Nature communications*, *12*(1), 6587.