ONLINE APPENDIX

Cognitive Model Decomposition of the BART: Assessment and Application

Don van Ravenzwaaij, Gilles Dutilh, and Eric-Jan Wagenmakers
University of Amsterdam

Correspondence concerning this article should be addressed to:
Don van Ravenzwaaij
University of Amsterdam, Department of Psychology
Roetersstraat 15
1018 WB Amsterdam, The Netherlands
Ph: (+31) 20–525–8870
E–mail should be sent to d.vanravenzwaaij@uva.nl.

Parameter Recovery Simulations

In this online appendix, we examine the 4–parameter (van Ravenzwaaij, Dutilh, & Wagenmakers, 2010), based on the models by Wallsten, Pleskac, and Lejuez (2005). We also discuss two simplifications: the 3–update, and the 3–stationary versions of the BART model (van Ravenzwaaij et al., 2010).

For each model, we simulated data for a grid of parameter values, fit the model to the simulated data and compared the resulting parameter estimates (specifically, the posterior mean) with the original values that were used to generate the data. For all simulations, parameters were recovered with a Bayesian implementation of the model. For each of the model fits in the next section, we used a single chain, consisting of 2000 iterations with a burn–in of 1000 samples. The simulations were conducted with a range of starting values for the MCMC chains. The results were qualitatively similar, unless reported otherwise.

4–Parameter Model

The 4–parameter model corresponds to “Model 3” from Wallsten et al. (2005, Table 2). The authors advocate this model as the best–fitting and most parsimonious model of the subset they investigated. To examine the ability of the model to recover its parameters, we generated data for a range of α’s and μ’s, shown on the x–axis of each panel in Figure 1. In addition, we used fixed generative values $\gamma^+ = 1.4$, $\beta = .6$, and $p^{burst} = .15$, as simulations with the 2–parameter model (van Ravenzwaaij et al., 2010) showed that these parameter values lead to the best parameter recovery.

We used the following priors: $\gamma^+ \sim U(0, 10)$, $\beta \sim U(0, 10)$, $\alpha \sim U(0, \mu)$, and $\mu \sim U(\alpha, 40)$, where U indicates the uniform distribution. We conducted 1000 simulations of a single synthetic participant completing 300 BART trials. With only 90 trials, parameter recovery was very poor and is not reported here. Recovery of all parameters is shown in Figure 1.

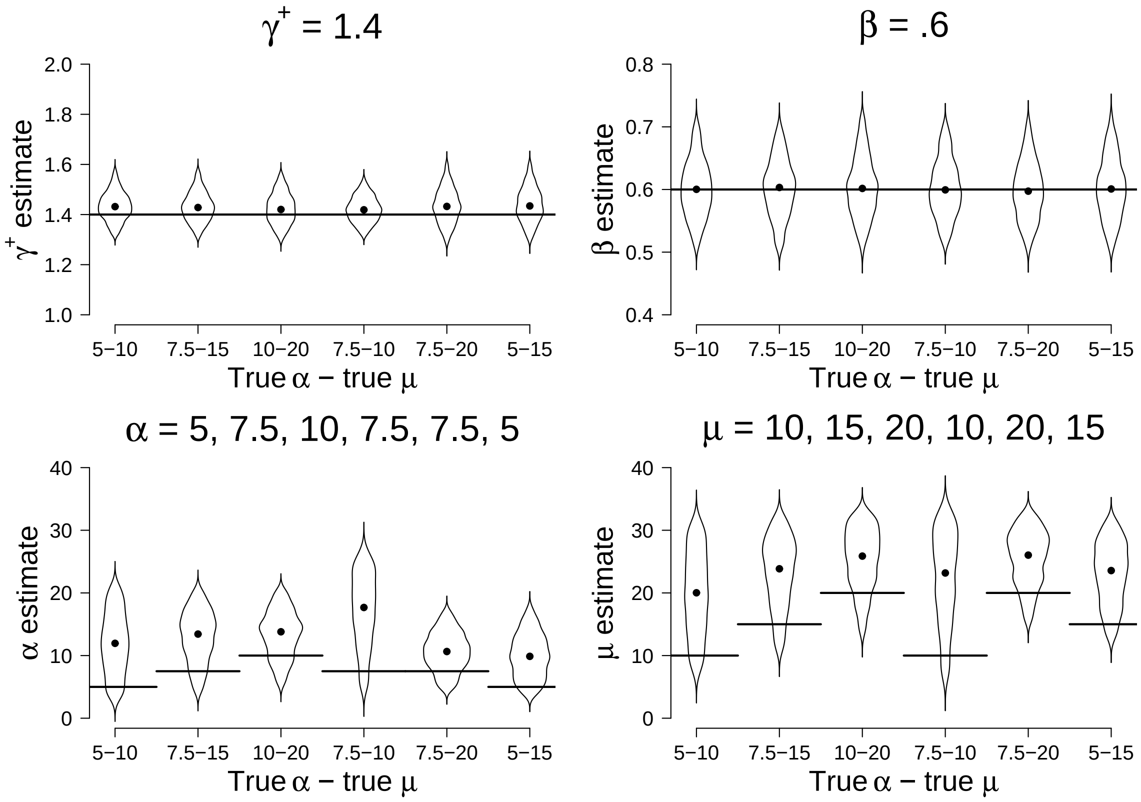

Figure 1. The 4–parameter BART model recovers parameters γ^+ and β, but fails to recover parameters α and μ (results based on a 300–trial BART). The dots represent the median of 1000 posterior means. The violins around the dots are density estimates for the distribution of the 1000 posterior means, with the extreme 5% truncated. The horizontal lines represent the true parameter values.

For each of the 1000 simulations, we used the posterior mean as a point estimate for each parameter. In Figure 1, the dots represent the median of the 1000 point estimates, and the violins that surround the dots represent density estimates for the distribution of the 1000 posterior means, with the extreme 5% truncated. The horizontal lines represent the true parameter values that are also indicated above each panel. The figure shows good parameter recovery for γ^+ and β, with only a slight overestimation of γ^+. The α and μ parameters are systematically overestimated. The overestimation of α increases when the true value of μ gets smaller (in the bottom left panel, compare the fourth, second, and fifth violin from the left or compare the leftmost and rightmost violins). The overestimation of

μ increases when the true value of α gets larger (in the bottom right panel, compare the first and the fourth violin from the left). Both phenomena suggest that parameter recovery suffers when the true value of α is close to the true value of μ.

Table 1 presents the correlation between the different parameters for the posterior means. The table shows a negative correlation between γ^+ and β, and a substantial positive correlation between α and μ. It seems the model is capable of estimating the ratio between α and μ, yet has problems finding the specific size of the parameters.

Table 1: *Parameter correlations in the 4–parameter model with 300 trials per simulation.*

α	5	7.5	10	7.5	7.5	5
μ	10	15	20	10	20	15
γ^+ vs. β	-0.66	-0.66	-0.65	-0.71	-0.59	-0.64
γ^+ vs. α	0.10	0.10	-0.03	0.11	-0.08	-0.01
γ^+ vs. μ	0.21	0.26	0.18	0.17	0.30	0.27
β vs. α	0.01	-0.05	0.01	-0.04	0.02	0.01
β vs. μ	0.01	-0.04	0.02	-0.04	0.02	0.00
α vs. μ	0.97	0.95	0.93	0.99	0.83	0.89

Thus, the 4–parameter model experiences severe problems trying to recover the generating values for α and μ. These problems may stem from the fact that the model gathers most information about α and μ from the first couple of trials. After the first trials, the impact of parameters α and μ on p_k^{belief} is dwarfed by the data (see van Ravenzwaaij et al., 2010, Eqn. 2). Therefore, the model might be better capable of estimating the α and μ parameters in a hierarchical multiple–subject design, where there are more "first trials". To examine this possibility, we generated data for a range of different numbers of participants and trials. In these simulations, all participants had identical parameter values (a best–case scenario); interest centered on the group mean for the parameters. For each participant, the true values were: $\gamma^+ = 1.4$, $\beta = .6$, $\alpha = 25$, and $\mu = 30$. Different values yielded qualitatively similar results.

For our hierarchical model fit, we used the following priors: $\gamma_i^+ \sim N(\gamma_\mu^+, \gamma_\sigma^+)$, $\beta_i \sim N(\beta_\mu, \beta_\sigma)$, $\alpha_i \sim N(\alpha_\mu, \alpha_\sigma)$, $\mu_i \sim N(\mu_\mu, \mu_\sigma)$, $\gamma_\mu^+ \sim U(0, 10)$, $\beta_\mu \sim U(0, 10)$, $\alpha_\mu \sim U(0, \mu_\mu)$, $\mu_\mu \sim U(\alpha_\mu, 40)$, $\gamma_\sigma^+ \sim U(0, 10)$, $\beta_\sigma \sim U(0, 10)$, $\alpha_\sigma \sim U(0, 10)$, and $\mu_\sigma \sim U(0, 10)$.

We used the following initial values for participant specific and mean parameters: $\gamma_i^+ = \gamma_\mu^+ = 1.2$, $\beta_i = \beta_\mu = .5$, $\alpha_i = \alpha_\mu = 23$, and $\mu_i = \mu_\mu = 28$. Initial values for standard deviation parameters were: $\gamma_\sigma^+ = \beta_\sigma = \alpha_\sigma = \mu_\sigma = .1$.

The results are presented in Table 2, which shows that recovery of parameters α and μ does *not* improve when more participants or more trials are added; in most cases, the estimates are closer to the initial values than to the true values. We ran this simulation with different sets of initial values, but always found estimates to be closer to the initial values than to the true values. In contrast, each simulation showed recovery of γ_μ^+ and β_μ to be accurate up to two decimals.

In sum, the 4–parameter model is capable of recovering parameters γ^+ and β, but it fails to recover α and μ, even in a hierarchical design with many participants and trials. To

Table 2: *Parameter estimates in the hierarchical 4–parameter model with various numbers of participants and trials per simulation. True parameter values were $\alpha_\mu = 25$, $\mu_\mu = 30$, $\alpha_\sigma = 0$, $\mu_\sigma = 0$. Initial value of $\alpha_\mu = 23$, initial value of $\mu_\mu = 28$. True values of all standard deviation parameters were 0.*

Participants	18	50	200	500	1000	50	50	50	50
Trials	300	300	300	300	300	500	1000	2000	5000
α_μ	24.05	23.77	23.57	23.62	22.99	24.05	24.02	23.92	24.14
α_σ	0.23	0.07	0.12	0.18	0.07	0.14	0.05	0.14	0.22
μ_μ	28.85	28.29	28.38	28.40	27.63	28.83	28.87	28.69	29.10
μ_σ	0.20	0.04	0.06	0.05	0.05	0.15	0.11	0.26	0.11

investigate whether a simplification would improve parameter recovery, we now turn to the 3–update model.

3–Update Model

The first 3–parameter model, which we will call the 3–update model, simplifies Eqn. 2 from van Ravenzwaaij et al. (2010) by assuming that $\alpha = y \times \mu$, where y is a constant (e.g., set equal to the actual bursting probability p^{burst}). The ratio of α to μ is then fixed, and only the rate with which participants learn from the data needs to be estimated. The parameters to be estimated for the 3–update model are μ, γ^+, and β.

Parameter recovery performance was studied by generating data for a range of μ's. Analogous to the 4–parameter model simulation, $\gamma^+ = 1.4$, $\beta = .6$, and $p^{burst} = .15$. We fixed y to p^{burst}, so that $\alpha = .15\mu$.

We used the following priors: $\gamma^+ \sim U(0, 10)$, $\beta \sim U(0, 10)$, and $\mu \sim U(0, 40)$. We conducted 1000 simulations of a single synthetic participant completing 90 trials. Recovery for all parameters is shown in Figure 2.

Figure 2 shows that for all true values of μ, γ^+ is overestimated, with higher values of μ leading to a larger bias. Although the 4–parameter model overestimated values for γ^+ slightly, with only 90 trials instead of 300, the problem is worse for the 3–parameter model. Parameter β is slightly overestimated for low values of μ, but is underestimated for higher values of μ. The model does not seem to be capable of picking up changes in μ (shown in the bottom left panel), making the inclusion of this parameter superfluous at best. Different starting values led to somewhat different results for μ, but the main result was the same: μ could not be recovered.

Table 3 presents the correlation between the different parameters for the posterior means. Analogous to the pattern found for the 4–parameter model, this table shows a clear negative correlation between γ^+ and β. Parameter μ does not correlate substantially with either of the other two parameters.

To examine whether recovery of the μ parameter could be improved by increasing the number of trials in the BART, we next ran 1000 simulations of a participant completing 300 trials, as we did for the 4–parameter model. Recovery of all parameters is shown in Figure 3. As can be seen from the smaller violins around γ^+ and β in Figure 3, parameter recovery

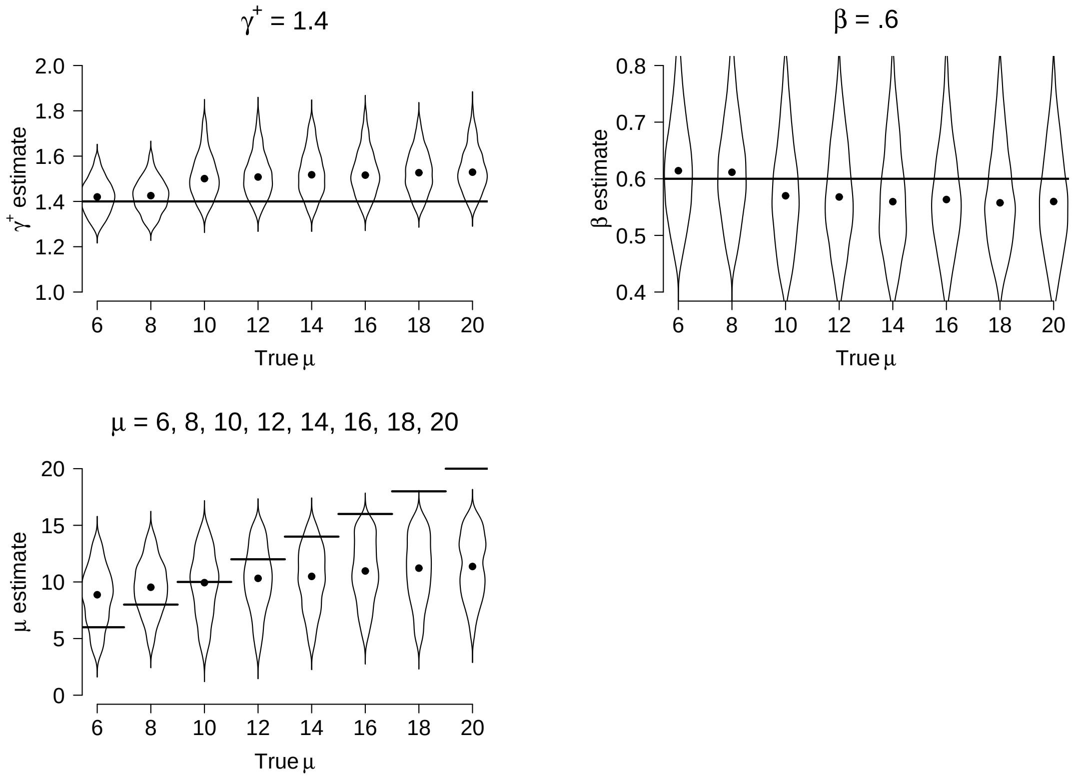

Figure 2. The 3–update BART model fails to recover parameters γ^+, β, and μ (results based on a 90–trial BART). The dots represent the median of 1000 posterior means. The violins around the dots are density estimates for the distribution of the 1000 posterior means, with the extreme 5% truncated. The horizontal lines represent the true parameter values.

does benefit from the increase in trials. However, the estimation of γ^+ and β is still more biased than for the 4–parameter model. In other words, fixing the α parameter seems to do more harm than good. Also, the recovery of μ remains poor. Table 4 presents the correlation between the different parameters for the posterior means. Parameter correlations did not change substantially with the increase in the number of trials.

Analogous to the 4–parameter model, we ran a hierarchical version of the model to see whether that would improve parameter recovery. The results were disappointing; recovery of parameters γ^+ and β improved slightly, but recovery of parameter μ was still poor. Specifically, the parameter estimate of μ remained very close to the starting value.

In sum, the 3–update model is capable of recovering parameters γ^+ and β, though not as well as the 4–parameter model. The 3–update model proved incapable of recovering the μ parameter. In the next section, we examine the 3–stationary model.

Table 3: *Parameter correlations in the 3–update model with 90 trials per simulation.*

μ	6	8	10	12	14	16	18	20
γ^+ vs. β	-0.75	-0.75	-0.80	-0.80	-0.82	-0.79	-0.80	-0.78
γ^+ vs. μ	0.02	0.00	0.10	0.11	0.15	0.13	0.22	0.13
β vs. μ	-0.06	-0.06	0.03	-0.04	-0.05	0.02	-0.11	0.02

Table 4: *Parameter correlations in the 3–update model with 300 trials per simulation.*

μ	6	8	10	12	14	16	18	20
γ^+ vs. β	-0.78	-0.79	-0.78	-0.76	-0.77	-0.78	-0.78	-0.77
γ^+ vs. μ	0.02	0.06	0.10	0.11	0.15	0.15	0.13	0.20
β vs. μ	0.07	0.05	0.03	0.03	0.02	0.02	0.03	-0.02

3–Stationary Model

The second 3–parameter model, which we will call the 3–stationary model, assumes that DM's belief about the probability that pumping the balloon will make it burst is fixed over trials. In other words, DM does not learn. This means we can drop the subscript k from p^{belief}, and replace the entire right part of Eqn. 2 from van Ravenzwaaij et al. (2010) with a single value that needs to be estimated. The parameters to be estimated for the 3–stationary model are γ^+, β, and p^{belief}.

To examine parameter recovery performance we generated data for a range of p^{belief}'s. Analogous to previous simulations, $\gamma^+ = 1.4$, $\beta = .6$, and $p^{burst} = .15$.

We used the following priors: $\gamma^+ \sim U(0, 10)$, $\beta \sim U(0, 10)$, and $p^{belief} \sim U(.05, .45)$. We ran 1000 simulations of 1 participant completing 90 trials. Recovery of all parameters is shown in Figure 4.

Figure 4 shows that when γ^+ is overestimated, so is p_i^{belief}; when γ^+ is underestimated, so is p_i^{belief}. Estimation of β is satisfactory. Parameter estimation could not be improved by increasing the number of trials. The parameter correlations presented in Table 5 demonstrate why.

As can be seen from the table, for the lower true values of p^{belief}, the correlations between γ^+ and p^{belief} are close to 1. This is because both parameters have the same effect on the data: both a high propensity for risk taking (i.e., high γ^+) and a pronounced belief that the burst probability is low (i.e, low p^{belief}) lead to the same kind of behavior: more pumping (see van Ravenzwaaij et al., 2010, Eqn. 3). Therefore, if both parameters were heightened, the model would not be able to pick up a difference, as their effects would cancel out. This correlation becomes lower for higher values of p^{belief}, as the effect of γ^+ will be overwhelmed by the effect of p^{belief} as p^{belief} gets closer to 1.

Analogous to the 4–parameter and the 3–update models, we ran a hierarchical version of the model to see whether that would improve parameter recovery. The results were

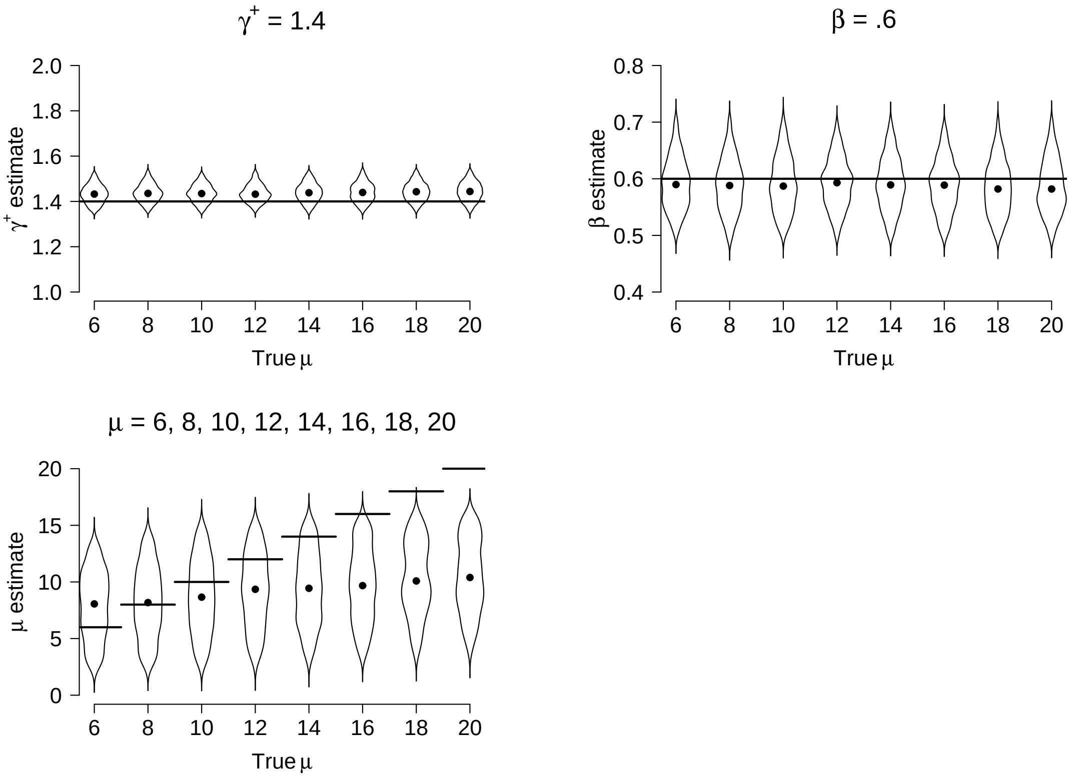

Figure 3. The 3–update BART model recovers parameters γ^+ and β, but fails to recover parameter μ (results based on a 300–trial BART). The dots represent the median of 1000 posterior means. The violins around the dots are density estimates for the distribution of the 1000 posterior means, with the extreme 5% truncated. The horizontal lines represent the true parameter values.

disappointing; recovery of parameter β was still okay, but recovery of parameters γ^+ and p^{belief} was still poor and these parameters remained heavily positively correlated.

In sum, the 3–stationary model can adequately recover β, but cannot differentiate between γ^+ and p^{belief}.

Conclusion

The simulation results show that parameter recovery of the 3– and 4–parameter models is suspect. Therefore, we chose to analyze the results from the empirical study (van Ravenzwaaij et al., 2010) with the 2–parameter model only.

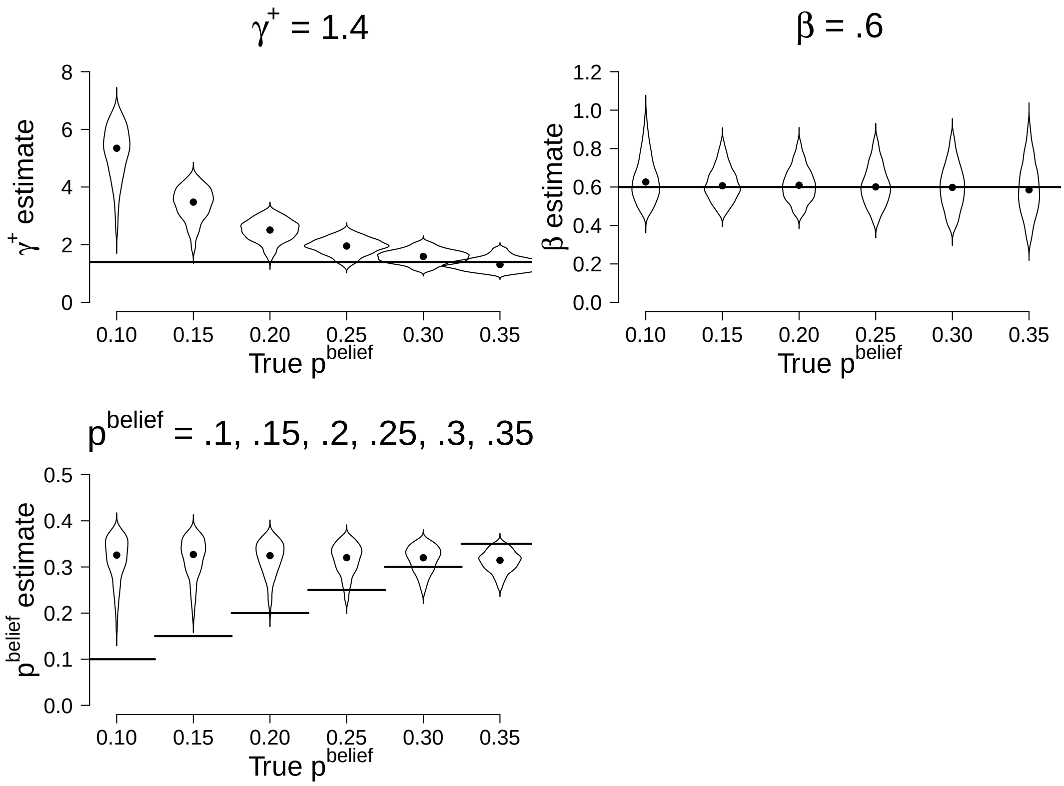

Figure 4. The 3–stationary BART model recovers parameter β, but fails to recover parameters γ^+ and p^{belief} (results based on a 90–trial BART). The dots represent the median of 1000 posterior means. The violins around the dots are density estimates for the distribution of the 1000 posterior means, with the extreme 5% truncated. The horizontal lines represent the true parameter values.

Table 5: *Parameter correlations in the 3–stationary model with 90 trials per simulation.*

p^{belief}	0.1	0.15	0.2	0.25	0.3	0.35
γ^+ vs. β	-0.24	-0.32	-0.39	-0.50	-0.58	-0.67
γ^+ vs. p^{belief}	0.95	0.93	0.89	0.80	0.67	0.46
β vs. p^{belief}	-0.03	-0.06	-0.03	-0.03	0.02	0.03

References

van Ravenzwaaij, D., Dutilh, G., & Wagenmakers, E.-J. (2010). Cognitive model decomposition of the BART: Assessment and application. Manuscript submitted for publication.

Wallsten, T. S., Pleskac, T. J., & Lejuez, C. W. (2005). Modeling behavior in a clinically diagnostic sequential risk–taking task. *Psychological Review*, *112*, 862–880.