Supplement to Nieuwenstein et al's "On making the right choice: A meta-analysis and large-scale replication study of the unconscious thought advantage"

Richard D. Morey

This document outlines the computation of the Bayes factor in Nieuwenstein et al's "On making the right choice: A meta-analysis and large-scale replication study of the unconscious thought advantage."

The model

Let y_a and y_d be the number of correct responses in the "attend" and "distract" conditions, respectively, out of N_a and N_d total trials. We can define a full model by assuming that y_a and y_d are binomially distributed with probabilities p_a and p_d, respectively:

$$y_a \sim \text{Binomial}(p_a, N_a)$$
$$y_d \sim \text{Binomial}(p_d, N_d)$$

The parameters p_a and p_d themselves arise from a standard probit model:

$$\Phi^{-1}(p_a) = \mu - \delta$$
$$\Phi^{-1}(p_d) = \mu + \delta$$

where Φ^{-1} is the inverse of the cumulative distribution function of the normal distribution, as shown in the figure below. This is a simple probit model. Probit models are a common way of building a model for a variable on the space of real numbers, and then mapping those variables into the (0,1) probability space needed for models of probability.

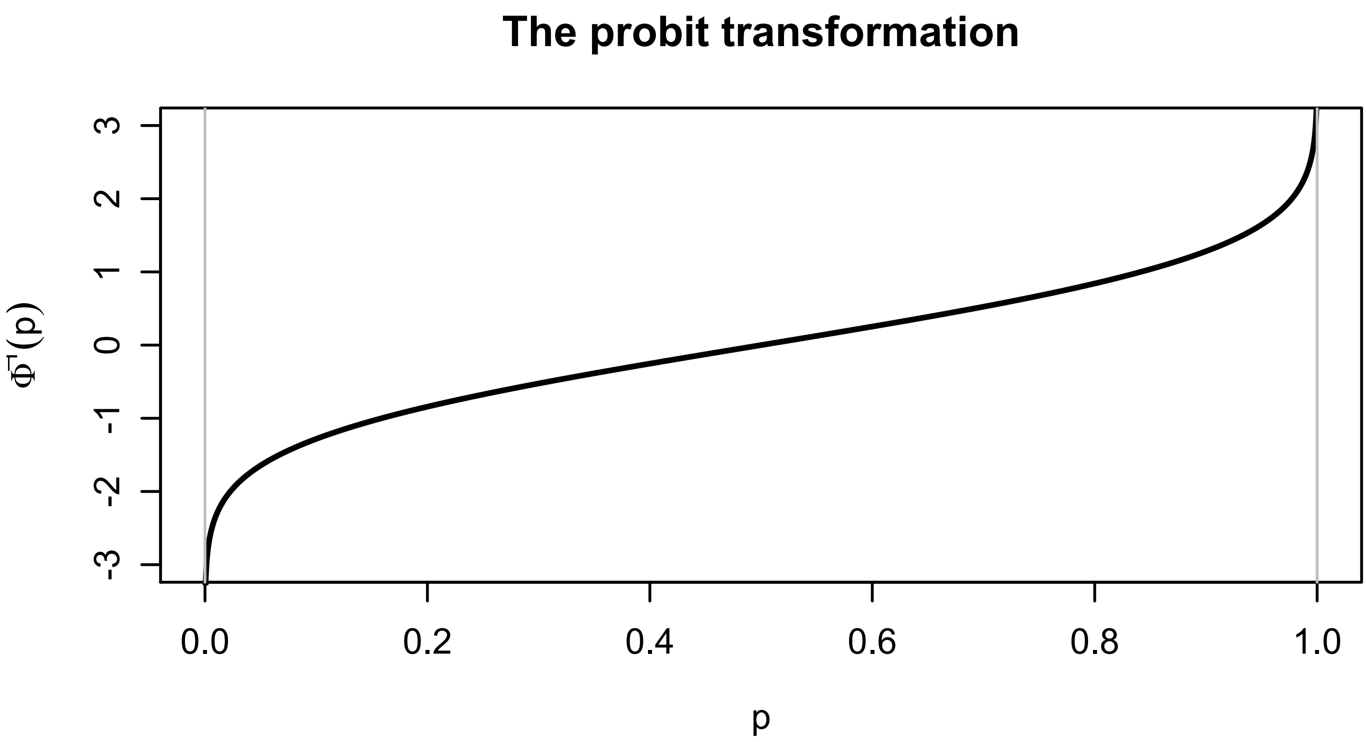

We wish to compare two hypotheses; the null hypothesis H_0 in which there is no advantage of distraction, and H_1 in which there is. Under the null hypothesis, $\delta = 0$ (that is, the conditions yield identical average performance). Under the alternative, $\delta > 0$. It remains to decide on prior distributions for μ under both hypotheses, and δ under the alternative hypothesis.

The approach we take here is to use prior studies by the proponents of the Unconscious Thought Effect (UTA) to build up an expectation for the effect size δ and average performance μ.

Explore the data

We selected six comparable studies on which to base the priors, shown in the table below. Also shown in the last row are the data from the large scale replication experiment.

	y_d	N_d	y_a	N_a	$\hat{p}_d - \hat{p}_a$	$(\hat{p}_d + \hat{p}_a)/2$
Dijksterhuis (2004, Exp 2)	16	27	16	34	0.122	0.532
Dijksterhuis, et al, (2006, Exp 1)	10	18	5	22	0.328	0.391
Nordgren et al. (2011, Exp 1)	8	12	7	18	0.278	0.528
Nordgren et al. (2011, Exp 2)	8	13	7	20	0.265	0.483
Strick et al. (2010, Exp 1)	32	49	8	47	0.483	0.412
Strick et al. (2010, Exp 2)	19	31	11	31	0.258	0.484
Large Scale Replication	112	203	109	196	-0.004	0.554

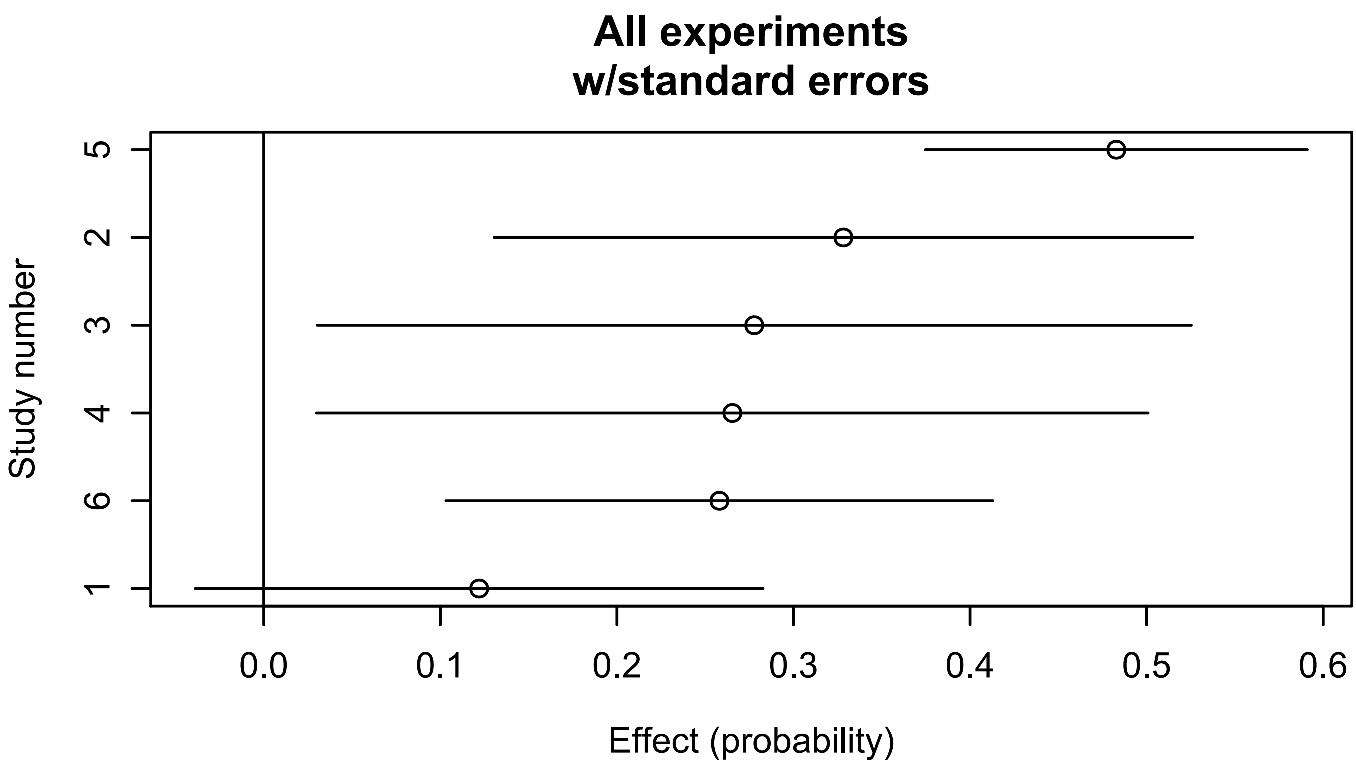

Prior settings for one-sided test

In order to build a prior distribution for the average performance μ and effect d, we use the six studies in the table above. We assume that these studies came from a population that of studies, each of which has its own unique average performance μ_i and effect d_i. These two populations have a normal and truncated normal distribution, respectively:

$$\mu_i \sim \text{Normal}(\mu_\mu, \sigma^2_\mu)$$
$$\delta_i \sim \text{Normal}_+(\mu_\delta, \sigma^2_\delta)$$

where Normal$_+$ is a normal distribution truncated to the positive real numbers. This instantiates the UTA assumption that the effect is positive. Of interest is to estimate the parameters μ_μ, σ^2_μ, μ_δ, and σ^2_δ, as these will be used as the prior for the large-scale replication experiment. The following JAGS code (Plummer, 2003) was used to estimate the parameters:

```
model{
  # Binomial model on observations
  for(m in 1:M){
    yDistract[m] ~ dbin(pDistract[m],NDistract[m])
    yAttend[m] ~ dbin(pAttend[m], NAttend[m])

    #probit model lining probabilities to parameters
    probit(pAttend[m]) <- mu[m] - d[m]
    probit(pDistract[m]) <- mu[m] + d[m]

    # Priors from population
```

```
      mu[m] ~ dnorm(muMu, precisionMu)
      d[m] ~ dnorm(muD, precisionD) T(0,)
    }

    muMu ~ dnorm(0,1)
    precisionMu ~ dgamma(1,1)
    muD ~ dnorm(0,1)
    precisionD ~ dgamma(1,1)

  }
```

The result of the analysis above will be samples from the posterior distributions of the parameters μ_μ, σ^2_μ, μ_δ, and σ^2_δ. In order to build expectations for the parameters of experiments similar to the six previous studies, we used the posterior means (that is, the means of the samples obtained above) as estimates of the population parameters on which to build the priors. (It is also possible to treat μ_μ, σ^2_μ, μ_δ, and σ^2_δ as themselves uncertain, but this would penalize the UTA model more, due to resulting diffuse predictions for the data. Since our goal is only to find a reasonable expectation on the basis of the previous studies, we use the posterior means.) The posterios means of the population parameters are found in the table below.

	Average performance μ	**UTA Effect δ**
Prior Mean	-0.09	-0.78
Prior Standard deviation	0.63	0.81

Analysis of replication data

The figure below shows the chosen prior distributions on the probit space (top row) and on the probability space (bottom row). The gray points in the lower plots show the estimates from the six studies on which the prior is based. The priors are more spread out than the data on which they are based due to the large amount of uncertainty in the original studies.

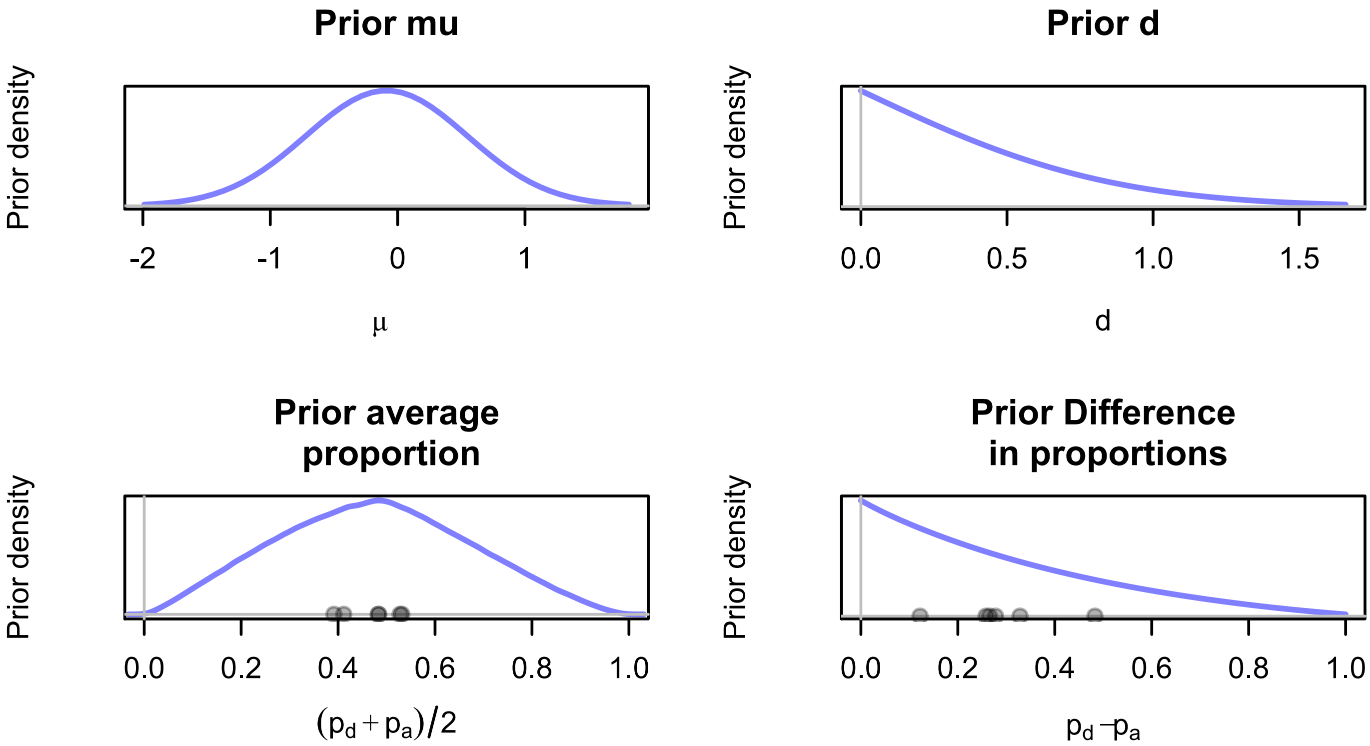

The figure below shows the effect of the data. The prior distribution (blue) is changed by the data into the posterior distribution (red line). The evidence in favor of the null hypothesis that $\delta = 0$, the Bayes factor, is the factor by which the posterior is higher than the prior at $\delta = 0$, which is in this case a factor of 7.8386. This indicates that given the prior distributions chosen, a rational observer considering H_0 against H_1 would have their odds swayed by a factor of 7.8386 in favor of H_0. This is not surprising, given that the difference in the proportions is very close to 0.

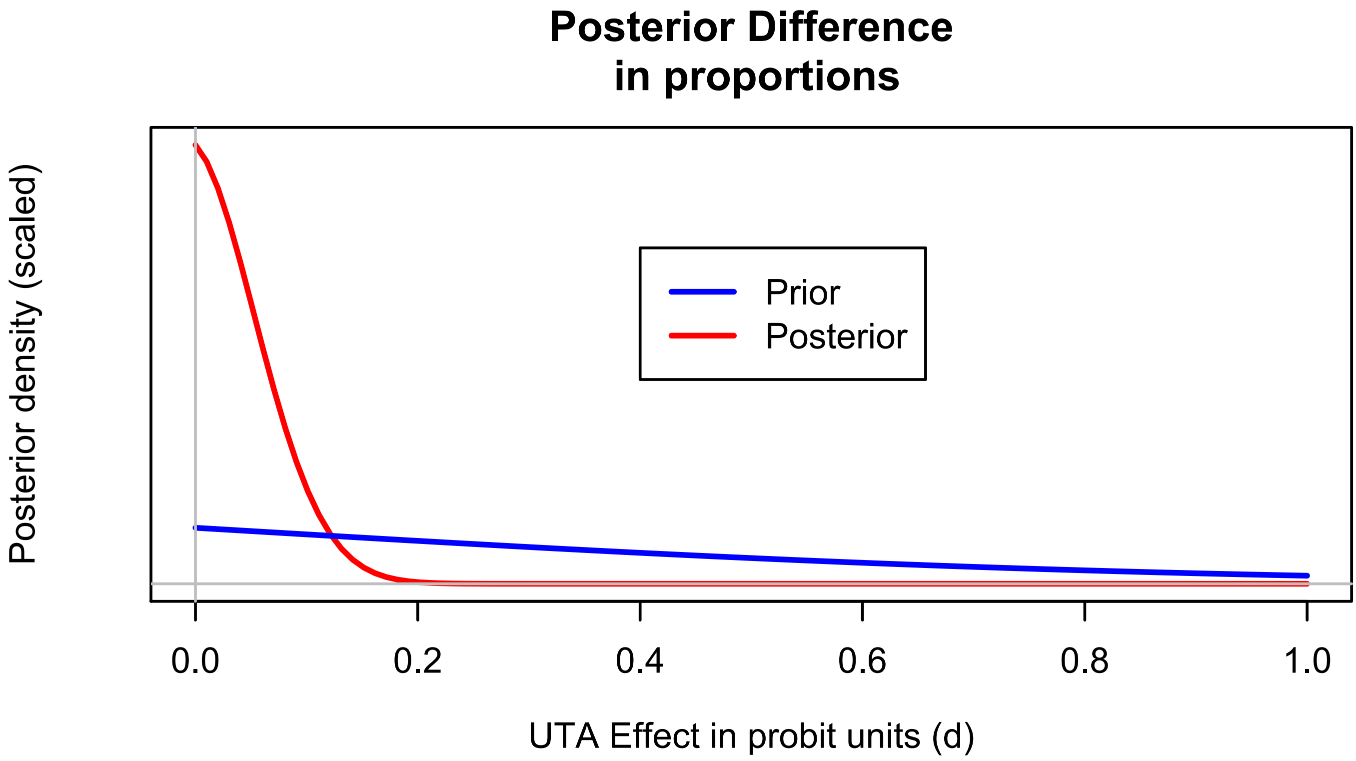

References

Plummer, Martyn (2003). JAGS: A Program for Analysis of Bayesian Graphical Models Using Gibbs Sampling. *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*.