

WITH DON VAN RAVENZWAALJ

There is a large psychological literature studying individual differences in the judgments people make, and the preferences they have, in tasks involving risk. One controlled laboratory task developed for these purposes is the Balloon Analogue Risk Task (BART: Lejuez et al., 2002). Every trial in this task starts by showing a balloon representing a small monetary value, as in Figure 16.1. The subject can then either transfer the money to a virtual bank account, or choose to pump, which adds a small amount of air to the balloon, and increases its value. There is some probability, however, that pumping the balloon will cause it to burst, causing all the money to be lost. A trial finishes when either the subject has transferred the money, or the balloon has burst.



Potential earnings: \$0.10

Balloon number: 2 of 30

Number of pumps: 2

Total winnings: \$0.15

Pump up the balloon

Collect \$\$\$

Fig. 16.1 The Balloon Analogue Risk Task (BART).

In the original version of the BART, the probability of the balloon bursting increased with every pump, but we consider a simplified version in which that probability is constant, and the expected gain of every decision to pump is zero. In the standard behavioral analysis of the BART, the risk propensity for a subject is measured as the average number of pumps for those balloons that did not burst. It is also possible to measure risk propensity using cognitive models of people's decisions on the BART (Rolison, Hanoch, & Wood, 2012; van Ravenzwaaij, Dutilh, & Wagenmakers, 2011; Wallsten, Pleskac, & Lejuez, 2005).

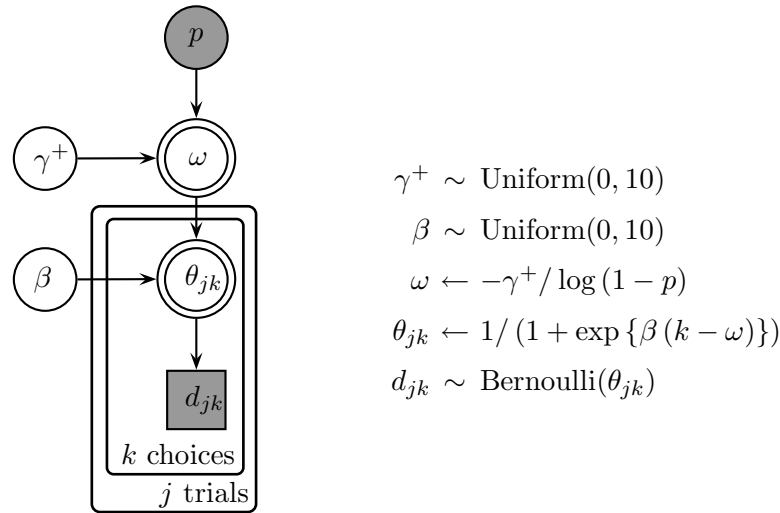


Fig. 16.2 Graphical model for the two-parameter BART model.

16.1 The BART model

We focus on a simple model using just two parameters, used by van Ravenzwaaij et al. (2011). One parameter, γ^+ , controls risk taking and the other, β , controls behavioral consistency. It is assumed the subject knows the constant probability p that the balloon will burst any time it is pumped. The number of pumps the subject considers optimal, ω , depends on this probability, and on the propensity for risk taking, such that $\omega = -\gamma^+ / \log(1 - p)$. Larger values of γ^+ lead to larger numbers of pumps being considered optimal, and so to greater risk seeking.

The probability that a subject chooses to pump on the k th opportunity within the j th trial depends on the number of pumps considered optimal, and on the behavioral consistency of the subject. These two factors are combined using the logistic function $\theta_{jk} = 1 / (1 + \exp\{\beta(k - \omega)\})$. High values of β correspond to less variable responding. When $\beta = 0$, $\theta_{jk} = 0.5$, and both pumping and cashing in choices are always equally likely. As β becomes large, the choice becomes completely determined by whether or not k exceeds the number of pumps the subject considers optimal. Finally, the observed decision made on the k th choice within the j th trial simply follows the modeled choice, so that $d_{jk} \sim \text{Bernoulli}(\theta_{jk})$.

A graphical model that implements this account of decision-making on the BART is shown in Figure 16.2. Note that the probability θ_{jk} depends only on the choice being made within a trial (i.e., k) and not the trial itself (i.e., j), and so the model is more general than it needs to be.

The script `BART_1.txt` implements the graphical model in WinBUGS:




Fig. 16.3 The number of pump decisions (left panel), posterior distribution for risk propensity γ^+ (middle panel), and posterior distribution for behavioral consistency β (right panel) for subject George.

```
# BART Model of Risky Decision-Making
model{
  # Optimal Number of Pumps
  omega <- -gplus/log(1-p)
  # Choice Data
  for (j in 1:ntrials){
    for (k in 1:options[j]){
      theta[j,k] <- 1-(1/(1+max(-15,min(15,exp(beta*(k-omega))))))
      d[j,k] ~ dbern(theta[j,k])
    }
  }
  # Priors
  gplus ~ dunif(0,10)
  beta ~ dunif(0,10)
}
```

The code `BART_1.m` or `BART_1.R` applies the model to data from a single subject, known as George, provided in the file `GeorgeSober.txt`. Figure 16.3 shows the results that are produced. The left panel shows the empirical distribution of the number of pump decisions across all trials. The middle panel shows the posterior distribution for γ^+ , George's propensity for risk. The right panel shows the posterior distribution for β , George's behavioral consistency.

Exercises

Exercise 16.1.1 Apply the model to data from a different subject, Bill, provided in the file `BillSober.txt`. Compare the estimated parameters for George and Bill. Who has the greater propensity for risk?

Exercise 16.1.2 What happens if two pumps are added to each trial for George's data? Make this change to the `npumps` variable in Matlab or R, and examine the new results. Which of the two parameters changed the most?

Box 16.1

Merit in application

“The merits of any statistical method are determined by the results it gives when applied to specific problems.” (Jaynes, 1976, p. 178)

Exercise 16.1.3 Modify George’s data in a different way to affect the behavioral consistency parameter.

16.2 A hierarchical extension of the BART model

Alcohol abuse can stimulate risk taking behavior. For example, alcohol abuse has been found to increase risk taking during driving (e.g., Burian, Liguori, & Robinson, 2002) and to increase participation in unsafe sex (e.g., McEwan, McCallum, Bhopal, & Madhok, 1992; but see Leigh & Stall, 1993). Examining performance on the BART, Lejuez et al. (2002) found that people with potentially problematic drinking habits also took more risks blowing up the balloons.

To examine the effects of alcohol on risk taking behavior more systematically, van Ravenzwaaij et al. (2011) carried out a within-subject manipulation, in which each subject completed the BART while sober, tipsy, and drunk. For a man weighing 70 kg, the blood alcohol concentration level for “drunk” required the consumption of 180 ml of vodka. We analyze here a subset of the data using a hierarchical extension of the individual two-parameter model.

The hierarchical extension is straightforward and requires only group-level distributions for the risk taking parameter γ^+ and behavioral consistency parameter β . We assume these parameters are drawn from Gaussian distributions, but are constrained to be positive. The graphical model now contains an extra plate that corresponds to the different levels or conditions of intoxication, and is shown in Figure 16.4.

The script `BART_2.txt` implements the graphical model in WinBUGS:

```
# Hierarchical BART Model of Risky Decision-Making
model{
  # Choice Data
  for (i in 1:nconds){
    gplus[i] ~ dnorm(mug, lambdag)I(0,)
    beta[i] ~ dnorm(mub, lambdab)I(0,)
    omega[i] <- -gplus[i]/log(1-p)
    for (j in 1:ntrials){
      for (k in 1:options[i,j]){
        theta[i,j,k] <- 1-(1/(1+max(-15,min(15,exp(beta[i]*(k-omega[i])))))
        d[i,j,k] ~ dbern(theta[i,j,k])
      }
    }
  }
}
```

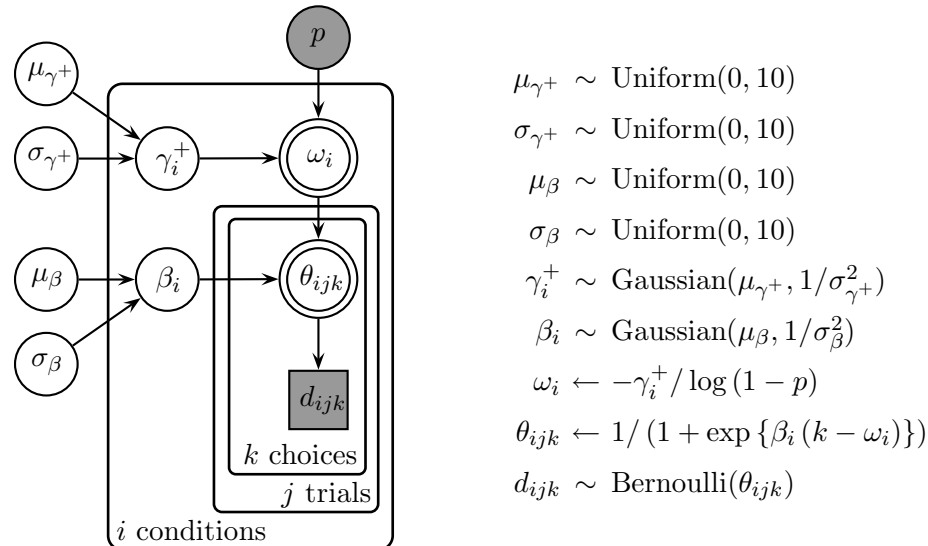


Fig. 16.4 Graphical model for the hierarchical two-parameter BART model.

```

# Priors
mug ~ dunif(0,10)
sigmag ~ dunif(0,10)
mub ~ dunif(0,10)
sigmab ~ dunif(0,10)
lambdag <- 1/pow(sigmatag,2)
lambdab <- 1/pow(sigmab,2)
}

```

The code `BART_2.m` or `BART_2.R` applies the model to George's data under different levels of intoxication, provided in the files `GeorgeSober.txt`, `GeorgeTipsy.txt`, and `GeorgeDrunk.txt`. It produces an analysis like that in Figure 16.5, showing the empirical distribution of the number of pumps, and posterior distributions for the two parameters, for all three intoxication conditions.

Exercises

Exercise 16.2.1 Apply the model to the data from the other subject, Bill. Does alcohol have the same effect on Bill as it did on George?

Exercise 16.2.2 Apply the non-hierarchical model in Figure 16.2 to each of the six data files independently. Compare the results for the two parameters to those obtained from the hierarchical model, and explain any differences.

Exercise 16.2.3 The hierarchical model in Figure 16.4 provides a structured relationship between the drinking conditions, but is still applied independently to each subject. Many of the applications of hierarchical modeling considered in our case studies, however, involve structured relationships between subjects, to capture individual differences. Develop a graphical model that extends Fig-

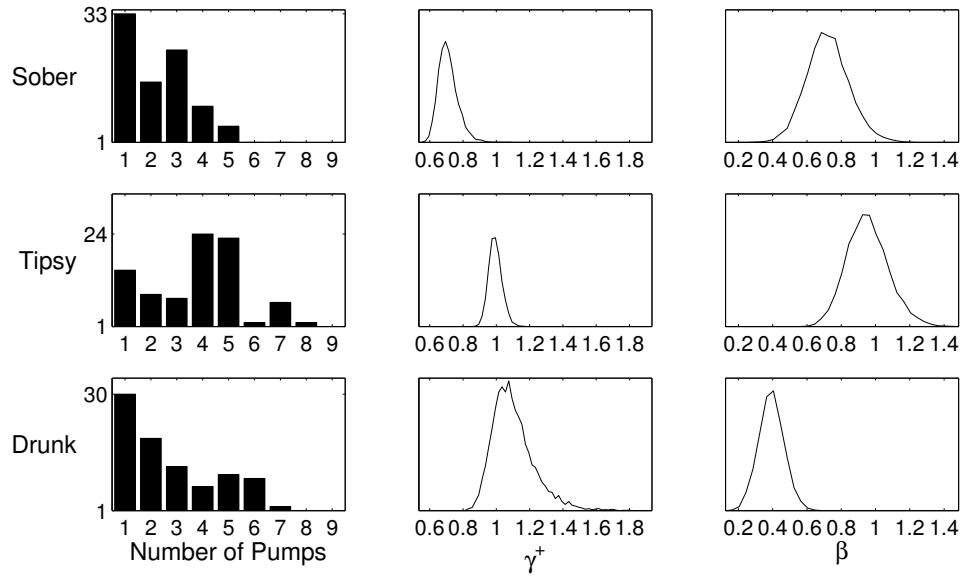


Fig. 16.5 The number of pump decisions (left panels), posterior distribution for risk propensity γ^+ (middle panels), and posterior distribution for behavioral consistency β (right panels) for subject George, in the sober (top row), tipsy (middle row), and drunk (bottom row) conditions.

ure 16.4 to incorporate hierarchical structure both for drinking conditions and subjects. How could interactions between these two factors be modeled?