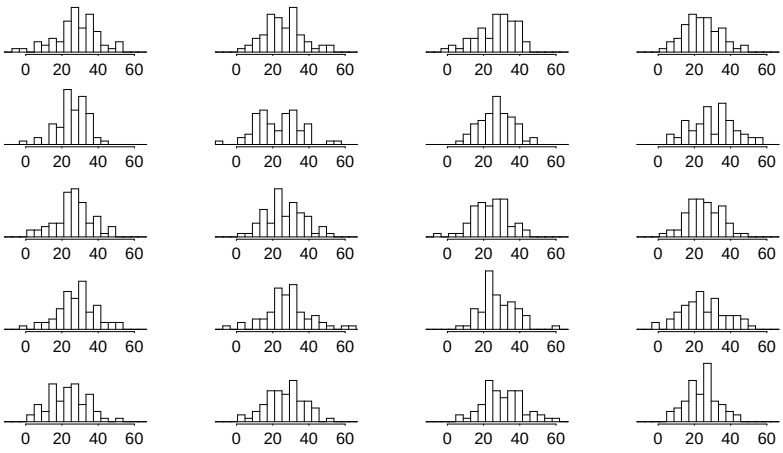

Figure 8.4 *Twenty replications, y^{rep}, of the speed-of-light data from the predictive distribution under the normal model; compare to observed data, y, in Figure 8.3. Each histogram displays the result of drawing 66 independent values y_i^{rep} from a common normal distribution with mean and standard deviation (μ, σ) estimated from the data.*

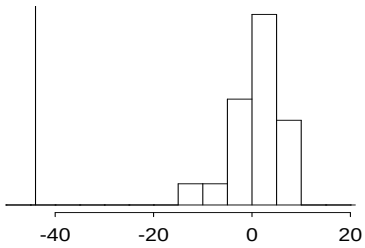

Figure 8.5 *Smallest observation of Newcomb's speed-of-light data (the vertical line at the left of the graph), compared to the smallest observations from each of 20 posterior predictive simulated datasets displayed in Figure 8.4.*

R code

```
Test <- function (y){
  min (y)
}
test.rep <- rep (NA, n.sims)
for (s in 1:n.sims){
  test.rep[s] <- Test (y.rep[s,])
}
```

We then plot a histogram of the minima of the replicated datasets, with a vertical line indicating the minimum of the observed data:

R code

```
hist (test.rep, xlim=range (Test(y), test.rep))
lines (rep (Test(y), 2), c(0,n))
```

Figure 8.5 shows the result: the smallest observations in each of the hypothetical replications are all much larger than Newcomb's smallest observation, which is indicated by a vertical line on the graph. The normal model clearly does not capture the variation that Newcomb observed. A revised model might use an asymmetric contaminated normal distribution or a symmetric long-tailed distribution in place of the normal measurement model.