(a)

$$\begin{aligned} Y &\sim Bi(20, \theta_G) \\ Z &\sim Bi(18, \theta_I) \end{aligned}$$

(b)

$$\begin{aligned} \hat{\theta}_G &= \frac{17}{20} \\ \hat{\theta}_I &= 1 \\ \bar{\theta}_G|Y &= \frac{9}{11} \\ \bar{\theta}_I|Y &= \frac{19}{20} \end{aligned}$$

(c) In S-plus,

```
> thetaG_rbeta(10000, 18, 4)
> thetaI_rbeta(10000, 19, 1)
> R_thetaG/thetaI
> summary(R)
  Min. 1st Qu. Median   Mean 3rd Qu. Max.
 0.451  0.8048 0.8663 0.8632  0.9247  1.4
> hist(R, 25)
> length((1:10000)[R<1])/10000
[1] 0.9325
> quantile(R, c(.025, .975))
     2.5%    97.5%
 0.6607913 1.056657
```

The historgram of R is shown in Figure 1. It's centered around 0.86, with the range of about 0.5 to 1.2. 93.25% of the simulated Rs are less than 1. The 95% credible interval of R is $(0.66, 1.06)$. From this We can conclude that the examiner is in fact better at detecting truth-tellers than liars on the basis of this data.

(d)

```
> pG_thetaG/(thetaG+1-thetaI)
> summary(pG)
   Min. 1st Qu. Median   Mean 3rd Qu. Max.
 0.6793  0.9208 0.9581 0.9453  0.9826    1
> hist(pG)
> quantile(pG, c(.025, .975))
     2.5%     97.5%
 0.8244609 0.9984045
```

The histogram of pG is in Figure 2. The Bayesian posterior mean is 0.9581. The 95% credible interval of pG is $(0.824, 0.998)$. We can thus conclude that he is very likely to be guilty.

(e) $\hat{P}_G = 1$. This is not a very good estimate given the data because it says with certainty that a suspect is guilty if the polygraph says he is. The posterior probability, though, shows that there is inevitably uncertainty. It's not a good idea to adopt MLE because we are dealing with real people, who will have to serve real sentence if wrongly convicted.

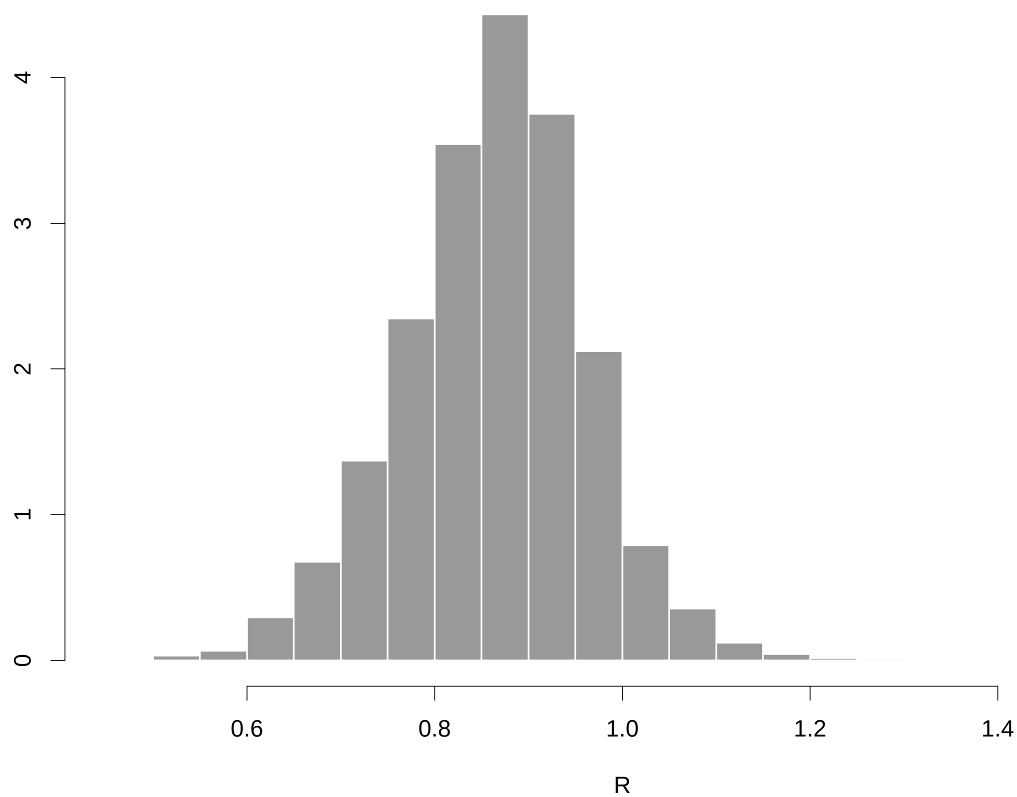

Figure 1: Histogram of R

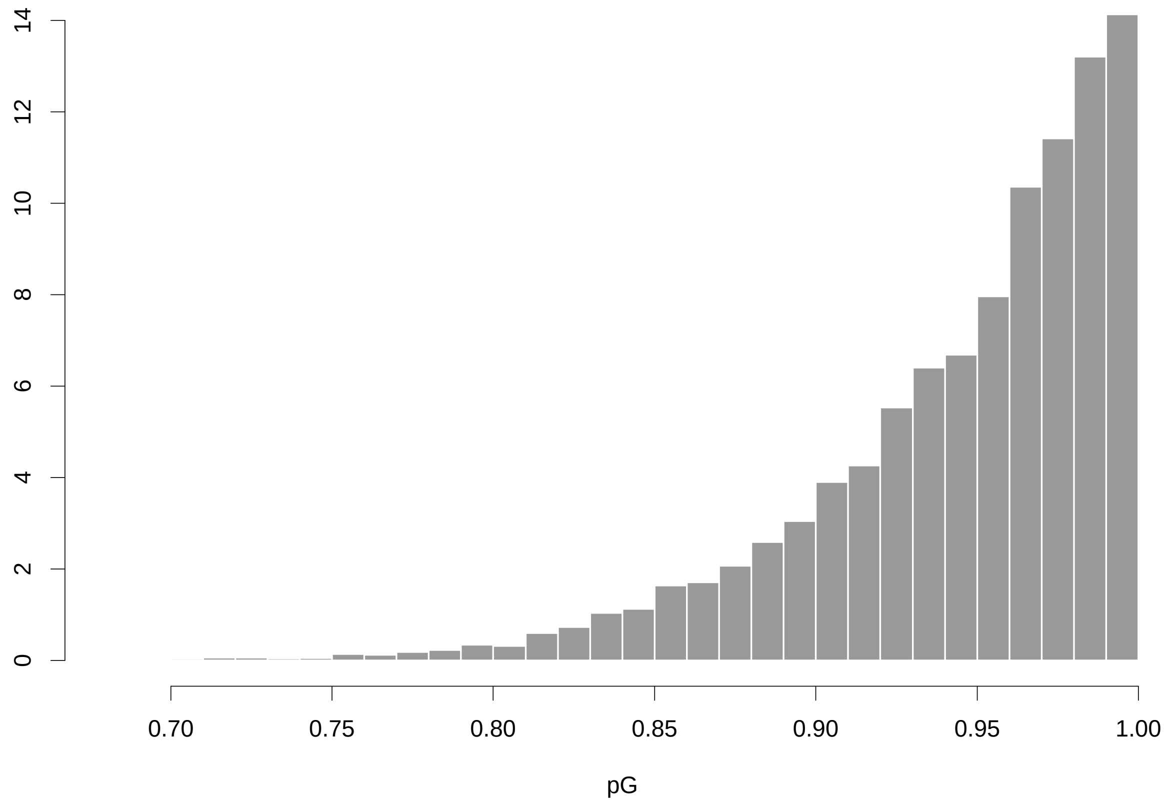

Figure 2: Histogram of pG