In-Class Exercise: Power and Sample Size Calculation

Psychology 311

Spring, 2013

KEY

1. You believe that a particular training program will improve the physical fitness (as measured by oxygen uptake) of a group of military recruits. The measured oxygen uptake of the recruit population has been established to a high degree of accuracy over a number of years at a value μ_0. Your plan is to administer the new training program to a randomly selected group of 75 recruits, then perform a 1-sided t-test with $\alpha = 0.01$ of the hypothesis that $H_0: \ \mu \leq \mu_0$, which is the opposite of what you actually believe.

   (a) If the standardized effect size $E_s = (\mu - \mu_0)/\sigma$ is really 0.30, what is the power of your design to reject the null hypothesis? Use the `Power.T1` function defined in the *Code.txt* file.

   *Answer.*We begin by loading back in the power calculation functions from class. We can input *any* values of μ,μ_0 and σ that yield an E_s of 0.30.

```
> source("Code.txt")

> n <- 75
> alpha <- 0.01
> tails <- 1
> mu <- 0.30
> mu0 <- 0
> sigma <- 1
> Power.T1(mu,mu0,sigma,n,alpha,tails)

[1] 0.5886836
```

   (b) Assuming that you are "stuck" with an n of 75, about how large would the standardized effect size have to be before power is at least 0.90? Use a graphical approach together with the `Power.T1` function.

   Answer. Let's draw a plot of power versus E_s. Note that, if we keep μ_0 and σ constant at 0 and 1, respectively, varying μ is the same as varying E_s.

```
> curve(Power.T1(x,0,1,75,0.01,1),0.30,1.00,col='red',
+   xlab='Standardized Effect Size (E_s)',ylab='Power')
> abline(h=0.90)
```

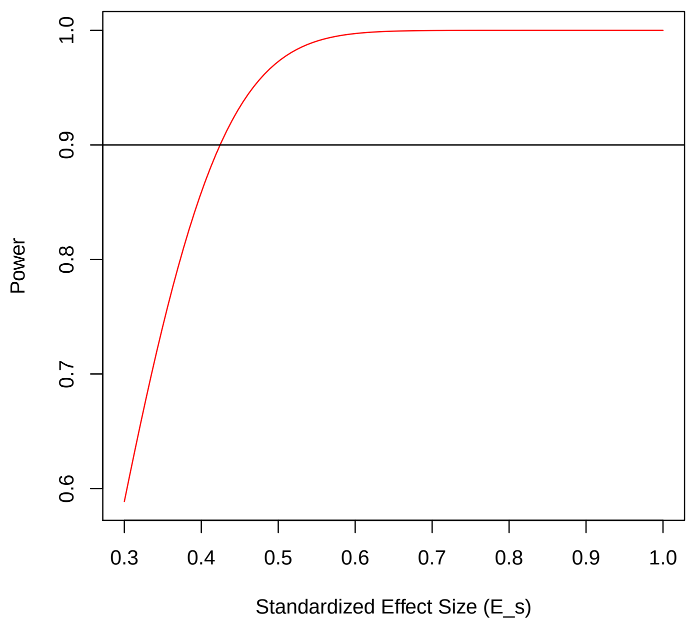

Our first plot shows that power crosses 0.90 roughly around an E_s of 0.40, but we need to home in on the plot.

```
> curve(Power.T1(x,0,1,75,0.01,1),0.35,0.45,col='red',
+   xlab='Standardized Effect Size (E_s)',ylab='Power')
> abline(h=0.90)
> abline(v=0.42)
```

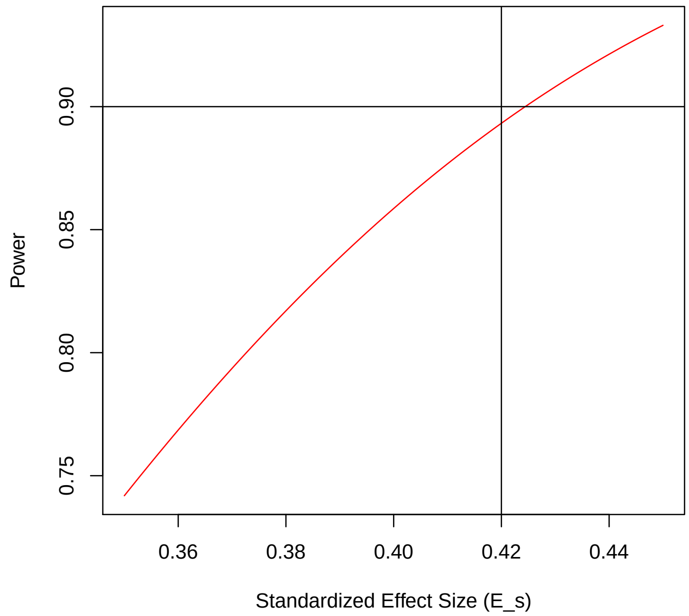

One more try should nail it down to 3 decimal place accuracy.

```
> curve(Power.T1(x,0,1,75,0.01,1),0.42,0.43,col='red',
+   xlab='Standardized Effect Size (E_s)',ylab='Power')
> abline(h=0.90)
```

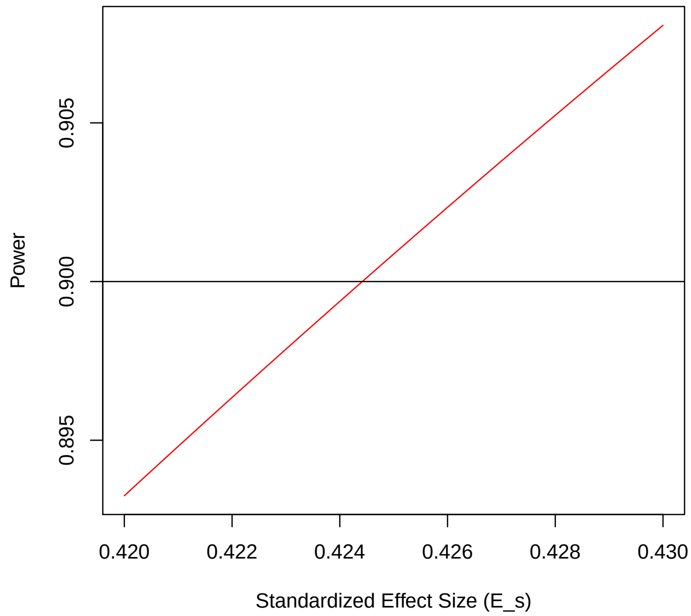

We can see that the value is between 0.424 and 0.425.

(c) Assume that $E_s = 0.30$ as before, but that now you have the ability to increase sample size. How large would sample size have to be to guarantee power of 0.90 to detect the experimental effect?

Answer. A similar graphical approach works just fine here. Now we vary n.

```
> curve(Power.T1(0.3,0,1,x,0.01,1),75,150,col='red',
+   xlab='Sample Size (n)',ylab='Power')
> abline(h=0.90)
```

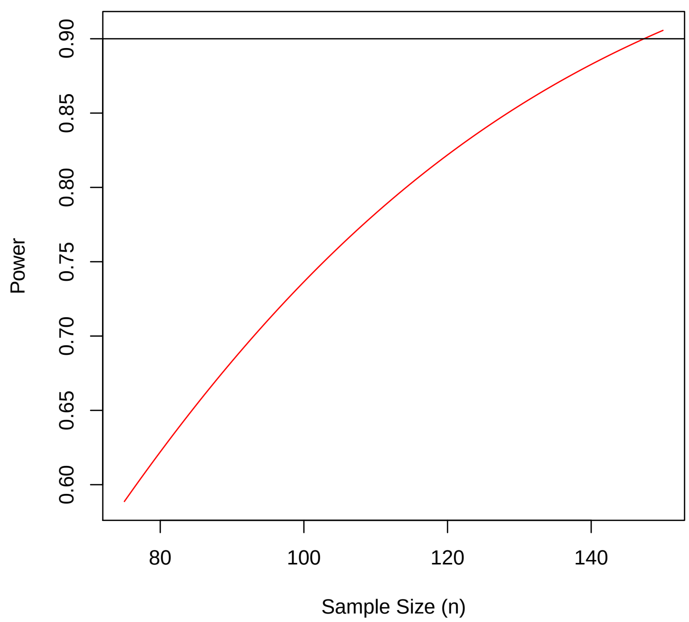


```
> curve(Power.T1(0.3,0,1,x,0.01,1),145,150,col='red',
+   xlab='Sample Size (n)',ylab='Power')
> abline(h=0.90,col='blue',lty=2)
> abline(v=145:150,col='lightgrey')
```

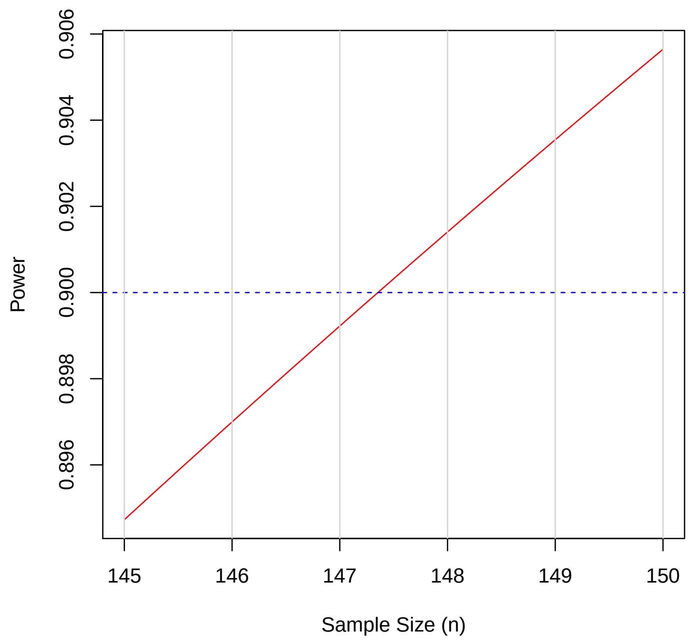

After homing in, we can see that power is just below 0.90 with $n = 147$, and just above 0.90 with $n = 148$.