1. In each case below the seemingly equivalent floating point operations in R may result in irregular outputs. Discuss the reasons for the same.

   (a)

```
> x = 29/13
> y = 13*x
> z = y - 29
> z
```

```
> u = 29/1300
> v = 1300*u
> w = v - 29
> w
```

   (b)

```
> a = 2.6 + 0.2
> b = a + 0.2
> c = b + 0.2
> c
> d = 2.6 + 0.6
> d
```

```
> options(digits=17)
> y = 2.6 + 0.2
> z = y + 0.2
> w = z + 0.2
> w
> t = 2.6 + 0.6
> t
```

   (c)

```
> 20.55-19.2-1.53
```

```
> 20.55-1.53-19.2
```

   (d)

```
> .Machine$double.xmin
> .Machine$double.xmin/2
```

```
> .Machine$double.xmax
> .Machine$double.xmax*2
```

2.

```
> z= 0/0
> y = 0/0
> y == z
```

```
> x = 1/0
> w = 1/x
> w
```

```
> x = 1/0
> y = 1/0
> x == y
```

3.

```
> options(digits=20)
> .Machine$double.eps
> 1 + .Machine$double.eps
> 1 + .Machine$double.eps*2
> 1 + .Machine$double.eps/2
```

```
> options(digits=20)
> .Machine$double.neg.eps
> 1 - .Machine$double.neg.eps
> 1 - .Machine$double.neg.eps*2
> 1 - .Machine$double.neg.eps/2
```

4.

```
> options(digits=20)
> 1000 + .Machine$double.eps
```

```
> options(digits=20)
> library(pracma)
> a = eps(1000)
> 1000 + a
> 1000 + a*2
> 1000 + a/2
```

5. Using the R-code in file halfdiff.R in the shared dropbox folder show that the floating point line has gaps.
6. Using the R-code in file exp.R in the shared dropbox folder understand the error that occurs when we approximate e with $\left(1+\frac{1}{n}\right)^n$