

# BEE 4750/5750

## Homework 4 Solutions

2026-10-14

### Problem 1 (30 points)

#### Problem 1.1 (8 points)

With  $Y = B_0(1 - R)$ ,  $B_0 = 50$  a constant,  $\mathbb{E}[R] = 0.6$  and  $\text{Var}(R) = 0.01$ :

$$\mathbb{E}[Y] = B_0(1 - \mathbb{E}[R]) = 50(1 - 0.6) = 20 \text{ mg/L}$$

$$\text{Var}(Y) = B_0^2 \text{Var}(R) = 50^2(0.01) = 25 (\text{mg/L})^2, \quad \text{so } \sigma_Y = 5 \text{ mg/L}.$$

For  $Z = Y^2$ , use the identity  $\text{Var}(Y) = \mathbb{E}[Y^2] - (\mathbb{E}[Y])^2$ , rearranged:

$$\mathbb{E}[Z] = \mathbb{E}[Y^2] = \text{Var}(Y) + (\mathbb{E}[Y])^2 = 25 + 400 = 425 (\text{mg/L})^2,$$

against  $(\mathbb{E}[Y])^2 = 400 (\text{mg/L})^2$ . They differ by exactly  $25 (\text{mg/L})^2$ .

This implies that running a model once at the mean input gives  $h(\mathbb{E}[X])$ , but the quantity you almost always want is the average outcome  $\mathbb{E}[h(X)]$ . These are equal only when  $h$  is linear. For any curved model the single run at the mean is biased.

#### Problem 1.2 (7 points)

The 95% confidence interval is  $0.08 \pm 1.96(0.006) = 0.08 \pm 0.0118$ , so  $[0.068, 0.092]$ .

The standard error scales as  $\sigma_Y/\sqrt{n}$ , so reducing it by a factor  $f$  costs  $f^2$  times as many samples. Here  $f = 0.006/0.002 = 3$ , so

$$n_{\text{new}} = 3^2 \times 2,000 = 18,000 \text{ samples}.$$

#### Note

Estimating  $\sigma_Y$  from  $\hat{p}$  instead, as  $\sqrt{\hat{p}(1-\hat{p})}$ , gives  $n = 0.08(0.92)/0.002^2 = 18,400$ , which agrees with 18,000 to the precision of the rounded standard error 0.006.

A finer spatial grid will not help, as it reduces *discretization* error. The 0.006 here is *Monte Carlo sampling* error, which comes from having drawn a finite number of random inputs.

#### Problem 1.3 (7 points)

The annual exceedance probability corresponding to a 50-year return period is  $p = 1/50 = 0.02$ , or 2% per year.

Assuming independent years, the probability of *no* exceedance is  $(1 - p)^{20}$ , so

$$\mathbb{P}(\text{at least one}) = 1 - (1 - 0.02)^{20} \approx 0.33.$$

The problem with the client's comment is that a return period is a long-run **average** frequency. Last year's flood carries no information about this year's due to the assumption about independence. The chance of overtopping next year, or any subsequent year, is still 2%.

### Problem 1.4 (8 points)

| Uncertainty                                    | Type             |
|------------------------------------------------|------------------|
| 1 Daily variation in influent flow             | <b>Aleatory</b>  |
| 2 Strength of the ice-albedo feedback          | <b>Epistemic</b> |
| 3 Rate of sediment recycling in a shallow lake | <b>Epistemic</b> |
| 4 Which of next year's storms is largest       | <b>Aleatory</b>  |

1. The flow genuinely varies day to day. Collecting more data could help us characterise that variability better but does not necessarily reduce or remove it.
2. The feedback has a definite strength, set by the physics of the climate system; what is uncertain is our knowledge of it.
3. For a given lake, recycling proceeds at a rate set by its sediments and chemistry. We are uncertain because it is hard to measure, not because of true randomness in the rate.
4. Which storm will turn out to be the largest is intrinsically random.

### Problem 2 (20 points)

#### Problem 2.1 (5 points)

```
loss_limit, steepness, midpoint = 200_000, 0.8, 3

function depth_damage(depth)
    if depth <= 0
        return 0
    end
    return loss_limit / (1 + exp(-steepness * (depth - midpoint)))
end

hazard = LogNormal(1.2, 0.3)
```




Figure 1: The depth-damage function and the hazard distribution over the same depths.

```
# Change in damage per metre of depth, measured over a 2 cm step around the midpoint
peak_slope = (depth_damage(midpoint + 0.01) - depth_damage(midpoint - 0.01)) / 0.02
```

The damage curve is steepest around its midpoint,  $h_0 = 3$  m, where the logistic curve is most nearly vertical — roughly the 2 m to 4 m range. At the midpoint the damage rises by about \$40000 per metre of depth. That band is also where most of the hazard distribution’s mass sits, so the structure is squarely in the region where a small change in depth produces a large change in loss. Away from that band, in shallow water or as the house nears a total loss, additional depth still adds damage, but more slowly.

## Problem 2.2 (8 points)

```
Random.seed!(4750)

function damage_sample(n)
    losses = zeros(n)
    for i in 1:n
        losses[i] = depth_damage(rand(hazard))
    end
    return losses
end

sample_sizes = [1_000, 2_000, 4_000, 8_000, 16_000, 32_000]
estimates = zeros(length(sample_sizes))
half_widths = zeros(length(sample_sizes))
for i in 1:length(sample_sizes)
    losses = damage_sample(sample_sizes[i])
    estimates[i] = mean(losses)
    half_widths[i] = 1.96 * std(losses) / sqrt(sample_sizes[i])
end
```

| $n$   | estimate | 95% half-width | change on doubling |
|-------|----------|----------------|--------------------|
| 1000  | \$114724 | ±\$2115        | —                  |
| 2000  | \$114684 | ±\$1515        | \$39               |
| 4000  | \$114931 | ±\$1063        | \$247              |
| 8000  | \$114658 | ±\$743         | \$273              |
| 16000 | \$115279 | ±\$529         | \$621              |
| 32000 | \$115541 | ±\$378         | \$263              |

At  $n = 32,000$  the expected annual damage is \$115541, with a 95% confidence interval of roughly  $\pm \$378$ . Seed 4750.

The quantity being estimated is the mean annual damage. Across the whole table the estimate never moves by more than a few hundred dollars on a six-figure loss — under 1% — and the reported half-width shrinks steadily as  $1/\sqrt{n}$ . By  $n = 8,000$  the estimate is stable well inside any precision this decision needs, and the extra doublings past that change nothing of consequence.

### Problem 2.3 (7 points)

```
Random.seed!(20)
big_sample = damage_sample(200_000)
mean_damage = mean(big_sample)
tail_damage = quantile(big_sample, 0.99)
```

The mean annual damage is \$115086, while the 99th percentile is \$189819 — about \$74733 more, or 1.6 times the mean, and close to the \$200,000 ceiling of the damage function.

Why is a percentile less precise? The estimate of the mean uses each of the  $n$  draws. In contrast, the 99th percentile is determined by the draws that land above or near it. Only 1% of the sample lies above it, so draws in that part of the distribution are sparse, and exactly where those few happen to fall moves the estimate.

Flood depth only captures the hazard, it cannot say where the properties are that would be exposed to flooding or how vulnerable those properties are to a certain level of flooding. The depth-damage function is the **vulnerability** component: it converts a physical hazard (a depth) into a consequence (a loss).