Bayesian Decision Analysis: a crash course

Cecina Babich Morrow

22 October 2024

The setting

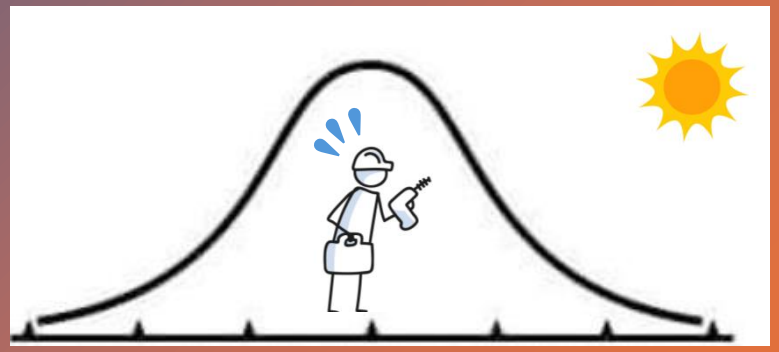

The year is 2041…

- Good news: Compass has received more funding, this time to run a construction company
- Bad news: the world is heating up, as are our statistics labourers

The setting

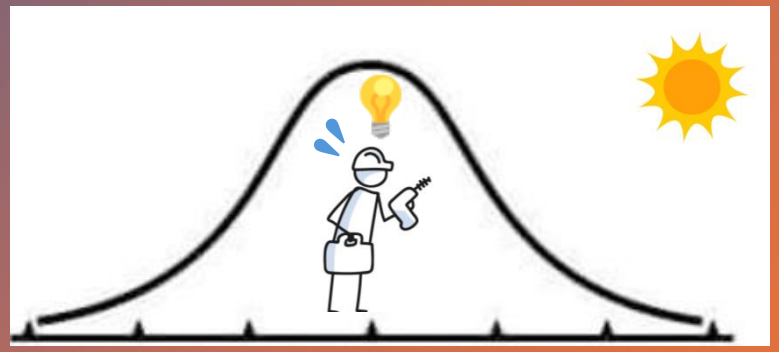

What should we do to adapt to heat stress?

Bayesian Decision Analysis!

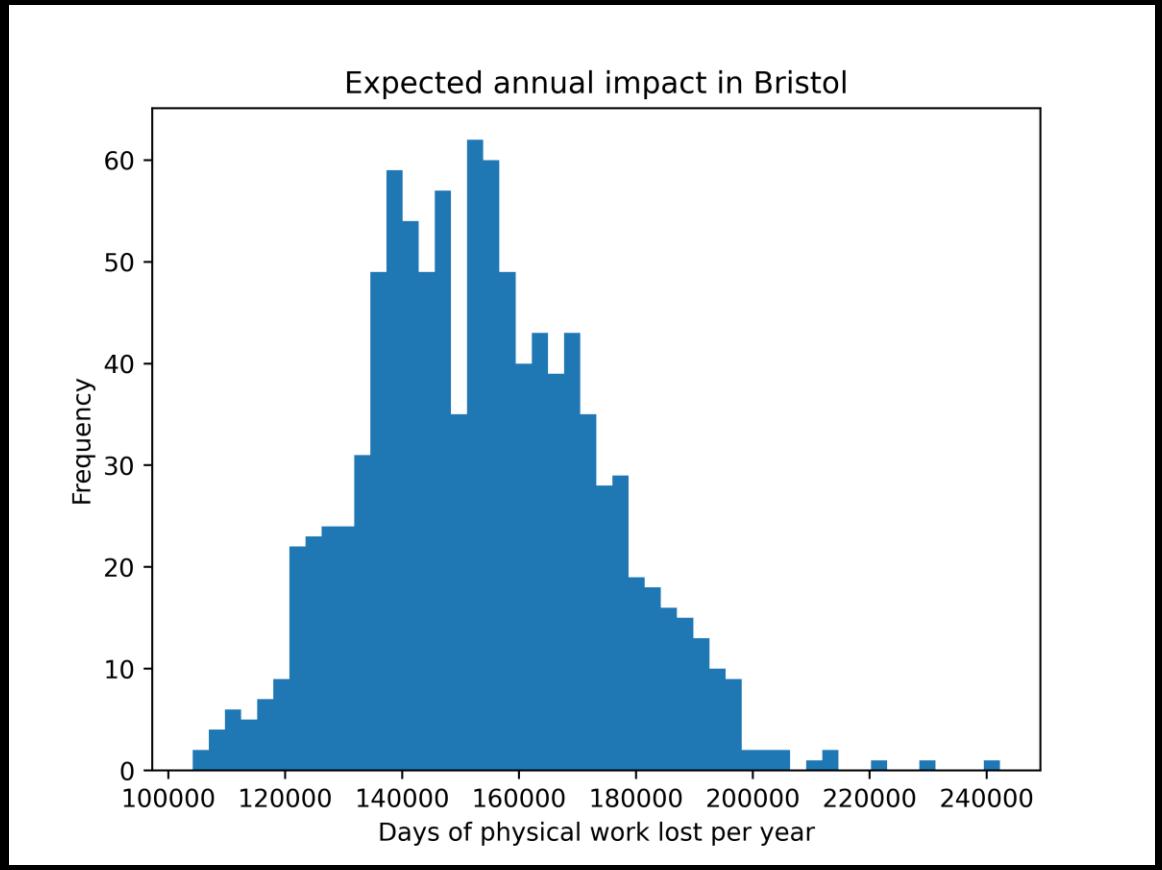

What could happen?

Space of possible states of nature: Θ

- Days of work lost per year

Need the distribution p(θ|x), for observations x.

We could be losing anywhere from 104,000 to 242,200 days of work annually.

What could we do?

Option	Cost per person	Added cost per day	Savings per day
Do nothing	£0	£0	£0
Modify working hours	£100	£20	£50
Buy cooling equipment	£500	£2	£80

Space of possible decisions: $\mathcal{D}$

1. Do nothing
2. Modify working hours
3. Buy cooling equipment

$\mathcal{D}$

Our decisions d:

1. Do nothing
2. Modify working hours
3. Buy cooling equipment

Θ

The state of nature θ:

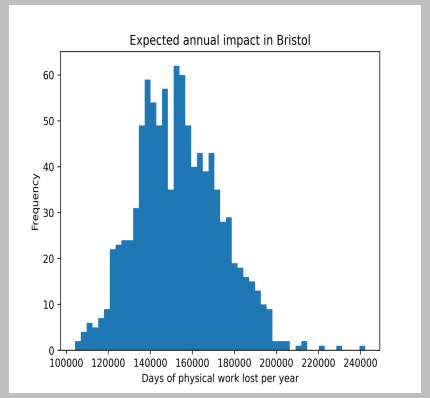

$L(\theta, d)$

Loss function

$\mathcal{L}$

Loss: How much will it cost to make decision *d*, if the number of days of work lost is θ?

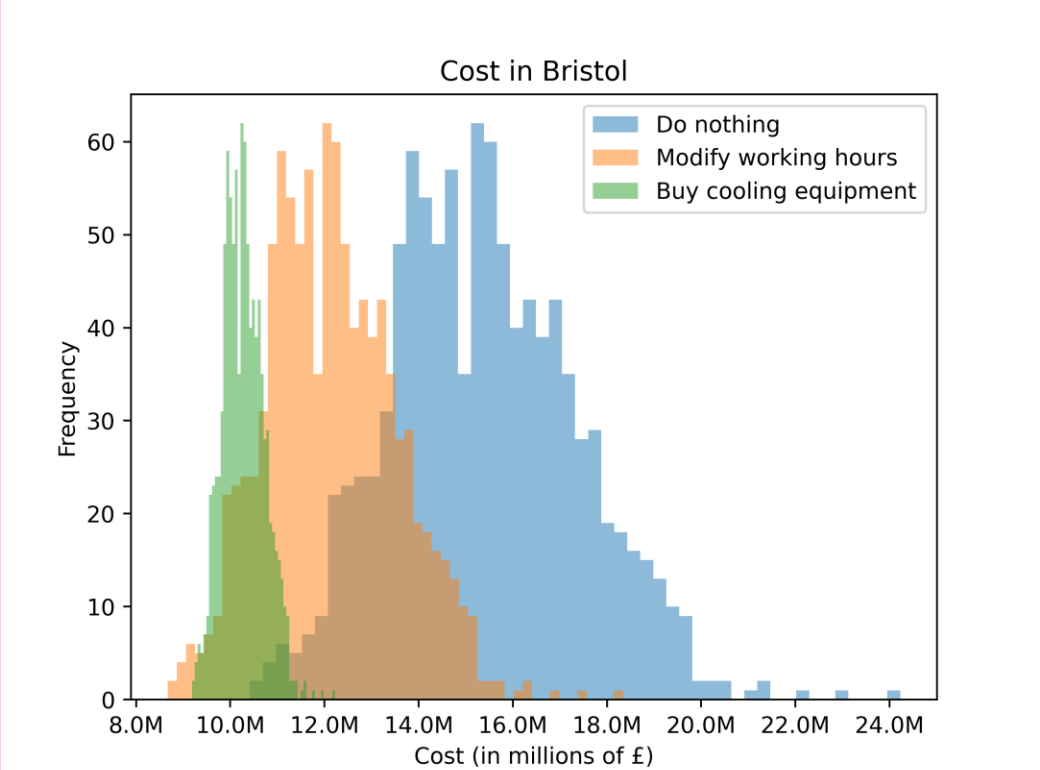

What might we lose?

Loss function to represent loss (financial or otherwise) of each decision, depending on the true state of nature

$\mathcal{D}$

Our decisions d:

1. Do nothing
2. Modify working hours
3. Buy cooling equipment

Θ

The state of nature θ:

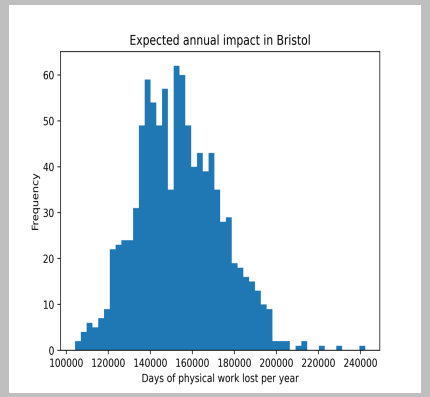

$\mathcal{L}$

Loss L(θ, d):

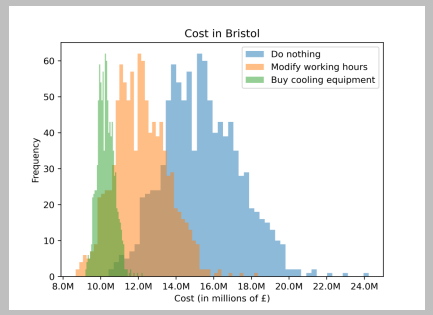

U(L(θ, d))

Utility function

Utility

Relative utility of decision *d* under state θ:

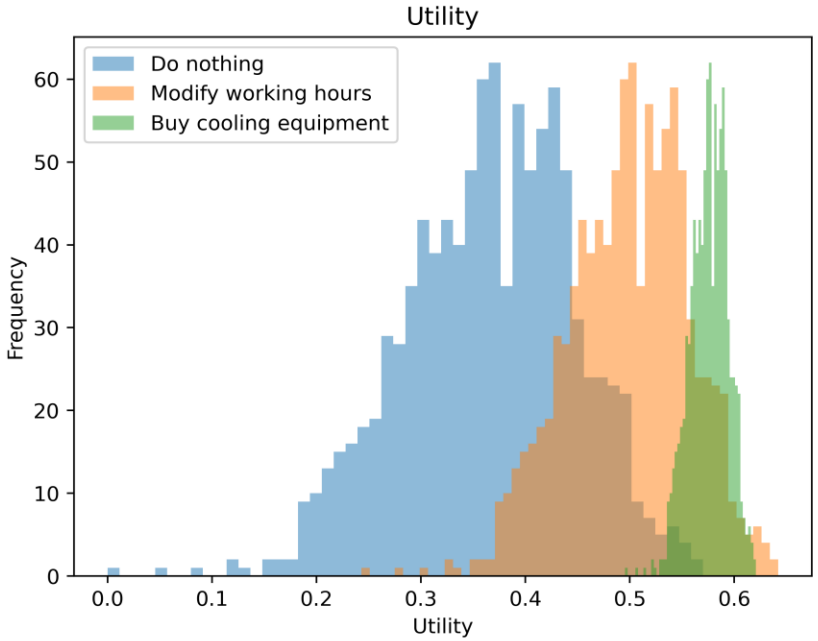

What should we do?

Pick the decision that maximises expected utility!

Bayes decision under utility U:

$$d^* \in \arg\max_d \sum_{\theta \in \Theta} U(L(\theta, d)) p(\theta|x)$$
$$= \arg\max_d \bar{U}(d)$$

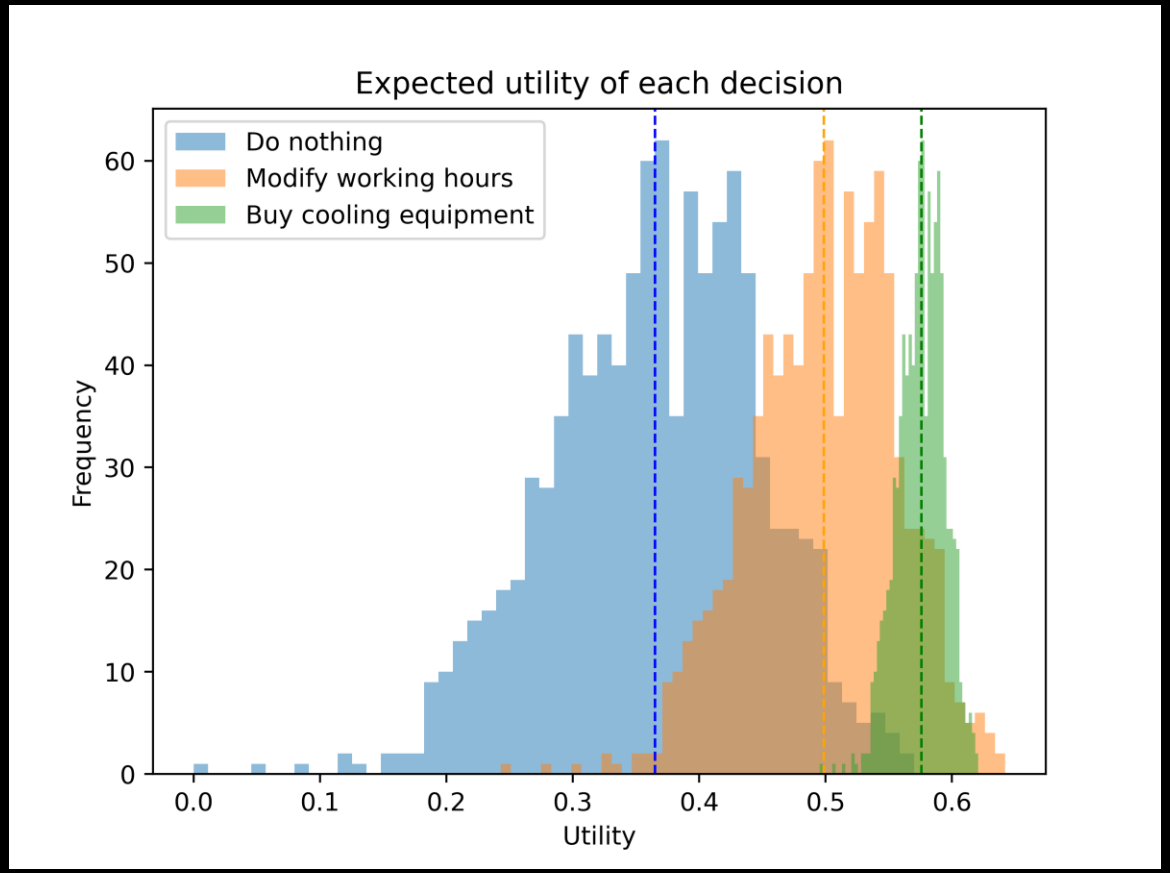

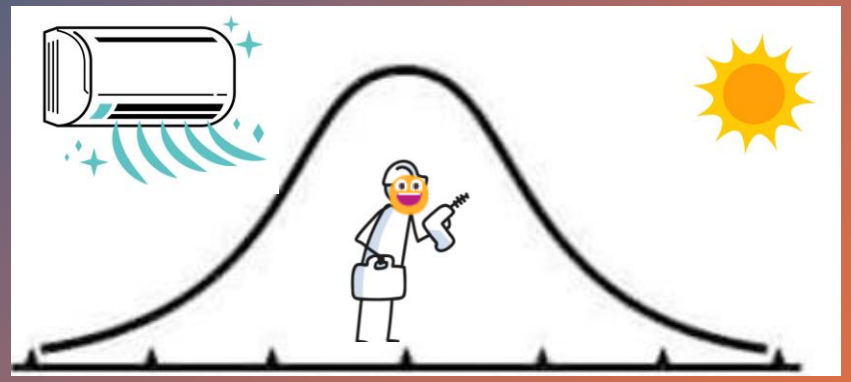

The decision

Buying cooling equipment maximises our expected utility!