Exercises on Utility

The utility of an oblivious mechanism, mapping query outcomes $\mathcal{Y}$ to reported values $\mathcal{Z}$, under a gain function g with set of guesses $\mathcal{W}$, is given by:

$$\mathcal{U} = \sum_{z \in \mathcal{Z}} \max_{w \in \mathcal{W}} \sum_{y \in \mathcal{Y}} p(y)p(z|y)g(w,y) \tag{1}$$

Note that an oblivious mechanism can be seen as a channel (in the sense of Quantitative Information Flow) from $\mathcal{Y}$ to $\mathcal{Z}$, and the formula above is exactly the posterior g-vulnerability of this channel, taking as prior the probability distribution $p(y)$ over the query outcomes.

For the "identity" gain function, having $\mathcal{W} = \mathcal{Y}$ as the set of guesses, and defined as $g(w,y) = 1$ if $w = y$ and 0 otherwise, the above formula can be simplified. We notice that $p(y)p(z|y)g(w,y) = p(w)p(z|w)$ if $y = w$ and 0 otherwise, so the formula becomes:

$$\mathcal{U} = \sum_{z \in \mathcal{Z}} \max_{y \in \mathcal{Y}} p(y)p(z|y) \tag{2}$$

(which is the posterior Bayes-vulnerability of the channel).

1. Compute the utility of the geometric mechanism for a counting query, with privacy degree ϵ, on the uniform prior distribution, with the gain function defined as the identity relation

Let n be the number of users, the result of a counting query is between 0 and n, hence $\mathcal{Y} = \{0, \ldots, n\}$.

The uniform distribution could be considered either as the distribution of the query outcomes (i.e. $p(y)$), or as the distribution of the users' value (in which case $p(y)$ becomes a binomial distribution). For simplicity we consider here $p(y)$ to be uniform, that is $p(y) = \frac{1}{n+1}$.

The geometric mechanism can output any integer (i.e. $\mathcal{Z} = \mathbb{Z}$), and for a counting query (i.e. for $\Delta_f = 1$) it is given by

$$p(z|y) = c\alpha^{|z-y|} \quad \text{where } \alpha = e^{-\epsilon}, c = \frac{1-\alpha}{1+\alpha}$$

Hence the utility under the identity gain function, given by (**??**), becomes:

$$\mathcal{U} = \frac{c}{n+1} \sum_{z \in \mathbb{Z}} \max_{y \in \mathcal{Y}} \alpha^{|z-y|} \tag{3}$$

Now let's consider the quantity $\max_{y\in\mathcal{Y}} \alpha^{|z-y|}$ for different values of z. Since $\alpha < 1$, the maximum is given by the y that minimizes $|z-y|$, i.e. the y that is closer to z. So for $0 \leq z \leq n$ we pick $y = z$, for $z \leq 0$ we pick $y = 0$ and for $z \geq n$ we pick $y = n$. Hence we have:

$$\max_{y\in\mathcal{Y}} \alpha^{|z-y|} = \begin{cases} \alpha^{|z-z|} = 1 & \text{if } 0 \leq z \leq n \\ \alpha^{|z-0|} & \text{if } z \leq 0 \\ \alpha^{|z-n|} & \text{if } z \geq n \end{cases}$$

So we can expand the sum of (**??**):

$$\begin{aligned} \sum_{z\in\mathbb{Z}} \max_{y\in\mathcal{Y}} \alpha^{|z-y|} &= \sum_{z=-\infty}^{0} \alpha^{|z|} + \sum_{z=1}^{n-1} 1 + \sum_{z=n}^{\infty} \alpha^{|z-n|} \\ &= n - 1 + 2\sum_{d=0}^{\infty} \alpha^d \\ &= n - 1 + 2\frac{1}{1-\alpha} \qquad \text{(geometric series for } \alpha < 1) \end{aligned}$$

so from (**??**) we finally get:

$$\begin{aligned} \mathcal{U} &= \frac{c}{n+1}(n + \frac{2}{1-\alpha} - 1) \\ &= \frac{c}{n+1}(n + \frac{2}{1-\alpha} - \frac{1-\alpha}{1-\alpha}) \\ &= \frac{c}{n+1}(n + \frac{1}{c}) \\ &= \frac{c\,n+1}{n+1} \end{aligned}$$

It's worth taking a look at this quantity as a function of c. The utility under the identity gain function is simply the probability to correctly guess the query outcome (that is, the posterior Bayes-vulnerability). Recall that

$$c = \frac{1-e^{-\epsilon}}{1+e^{-\epsilon}}$$

Since $c \leq 1$ we have $\mathcal{U} \leq 1$, as expected (it's a probability).

Consider the one extreme case: perfect privacy. The geometric mechanism is well-defined for $\epsilon > 0$, but as $\epsilon \to 0$ the noise increases and $p(z|y), p(z|y'), y \neq y'$ become closer to each other. At the limit we have $c = 0$ hence $\mathcal{U} = \frac{1}{n+1}$. Intuitively, the output of the mechanism is useless, we still have $\frac{1}{n+1}$ probability of guessing the correct value.

At the other extreme case (no privacy at all), when $\epsilon \to \infty$ then $p(\cdot|y)$ becomes a point distribution, giving $p(y|y) = 1$ and $p(z|y) = 0$ for $z \neq y$. At the limit c becomes 1, hence $\mathcal{U} = 1$: as expected we can now guess the correct value with probability 1.

2. Same exercise, but with the gain function defined as the converse of the distance.

The goal here is not to come up with an exact closed-form formula, but to see how the utility changes because of the gain function. Under this gain function we still have $\mathcal{W} = \mathcal{Y}$ but the gain is given by

$$g(w, y) = n - |y - w|$$

The closer our guess w is to the real answer y, the higher the gain.

From (**??**) we get

$$\begin{aligned}
\mathcal{U} &= \sum_{z \in \mathcal{Z}} \max_{w \in \mathcal{W}} \sum_{y \in \mathcal{Y}} p(y)p(z|y)(n - |y - w|) \\
&= n - \sum_{z \in \mathcal{Z}} \min_{w \in \mathcal{W}} \sum_{y \in \mathcal{Y}} p(y)p(z|y)|y - w| \\
&= n - \frac{c}{n+1} \sum_{z \in \mathcal{Z}} \min_{w \in \mathcal{W}} \sum_{y \in \mathcal{Y}} \alpha^{|y-z|}|y - w|
\end{aligned}$$

We need to investigate the value w that gives the minimum:

$$\sum_{y \in \mathcal{Y}} \alpha^{|y-z|}|y - w|$$

for each z. To minimize this quantity we need that the factor $|y - w|$ is as small as possible when $\alpha^{|y-z|}$ is big, which happens if w is as close as possible to z. Hence, similarly to the previous case, for $0 \leq z \leq n$ we pick $w = z$, for $z \leq 0$ we pick $w = 0$ and for $z \geq n$ we pick $w = n$.

Finally we can expand the formula of utility to:

$$\mathcal{U} = n - \frac{c}{n+1}\Big(\sum_{z=-\infty}^{-1} \sum_{y=0}^{n} \alpha^{y-z} y + \sum_{z=0}^{n} \sum_{y=0}^{n} \alpha^{|y-z|}|y - z| + \sum_{z=n+1}^{\infty} \sum_{y=0}^{n} \alpha^{z-y}(n - y)\Big)$$

3. Find a mechanism for the same counting query, with the same degree of privacy, but lower utility

Intuitively, the geometric mechanism is optimal because the noise is exactly as much as required by ϵ, not more. That is, the constraints of differential privacy for adjacent y's are satisfied with equality:

$$p(z|y) = e^{\epsilon} p(z|y + 1) \quad 0 \leq y < n, z \in \mathbb{Z}$$

To degrade its utility, we could, for instance, add more noise to a certain query outcome. For instance, for $y = 0$ we can use the same distribution that we

use for $y = 1$ (i.e. $p(z|0) = p(z|1)$). This modified version of the geometric mechanism is given by:

$$p(z|y) = c\alpha^{|z-\max\{1,y\}|}$$

That is, $p(z|0) = p(z|1) = c\alpha^{|z-1|}$ and $p(z|y) = c\alpha^{|z-y|}$ for $y > 1$.

Overall, the mechanism still satisfies differential privacy for the same ϵ, since the constraints for $p(z|1)$ and $(z|2)$ are still matched with equality. For any $\epsilon' < \epsilon$ we would have $p(z|1) > e^{\epsilon'} p(z|2)$.

On the other hand, utility is now lower, because we cannot distinguish 0 from 1: any value $z \leq 1$ should be mapped to either 0 or 1. Redoing the computation of the first exercise, we have:

$$\max_{y\in\mathcal{Y}} \alpha^{|z-\max\{1,y\}|} = \begin{cases} \alpha^{|z-z|} = 1 & \text{if } 1 \leq z \leq n \\ \alpha^{|z-1|} & \text{if } z \leq 1 \\ \alpha^{|z-n|} & \text{if } z \geq n \end{cases}$$

So the sum of (**??**) becomes:

$$\begin{aligned} \sum_{z\in\mathbb{Z}} \max_{y\in\mathcal{Y}} \alpha^{|z-y|} &= \sum_{z=-\infty}^{1} \alpha^{|z-1|} + \sum_{z=2}^{n-1} 1 + \sum_{z=n}^{\infty} \alpha^{|z-n|} \\ &= n - 2 + 2\sum_{d=0}^{\infty} \alpha^d \\ &= n - 2 + 2\frac{1}{1-\alpha} && \text{(geometric series for } \alpha < 1) \end{aligned}$$

and continuing similarly to the exercise 1, we get

$$\mathcal{U} = \frac{c(n-1)+1}{n+1}$$

Even under no privacy, when $\epsilon \to \infty$ and $c \to 1$, we have $\mathcal{U} = \frac{n}{n+1}$ (compared to $\mathcal{U} = 1$ for the original geometric mechanism), since two out of the $n+1$ elements are still completely indistinguishable!

4. We saw that post-processing cannot decrease privacy. Can it decrease the utility? Motivate your answer

Post-processing can create more confusion between the reported values. This does not decrease privacy (it can only become harder to infer the value of an individual) but it can decrease utility (it also becomes harder to infer the real outcome of the query).

A trivial example would be a constant post-processing function mapping every z to 0. The result of applying this post-processing is a non-interferent channel: it outputs 0 independently from the 0. This has perfect privacy but clearly no utility at all.