

Exercise 1

You and your team are doing professional data consulting on a wide spectrum of projects. Last week, you've been given a call from a company called We Really Care Inc (short WRC Inc.) This company is one of the oldest health insurance companies in the world and they've aggregated insane amount of data about their activity all across the globe.

After countless months of trying to collect all of the bits and pieces into uniform data schema, you've managed to minimize all the relevant logs into a single well-formed data set that consists of four distributed collections:

```
case class Broker(brokerId: Int, fullName: String, ...)
case class Client(clientId: Int, fullName: String, ...)
case class Contract(contractId: Int, brokerId: Int, clientId:
Int, planId: Int, ...)
case class Plan(planId: Int, yearIntroducedIn: Int, price:
Double, ...)

val brokers: RDD[Broker] = ...
val contracts: RDD[Contract] = ...
val clients: RDD[Client] = ...
val plans: RDD[Plan] = ...
```

The company employs brokers whose job is to seek and arrange contracts with clients. Contracts are tied to specific plans in the company's product line-up.

Question 1.1

In the past year WRC has been having hard time financially. Due to this, upper management decided to let some people go. To make an educated decision about whom to fire they need to analyze the performance of the current employees.

Part 1

Compute an RDD that maps broker full names to their total revenue (i.e. total income from all of their contracts). Try to compute this as efficiently as possible.

```
val brokerNameRevenue: RDD[(String, Double)] = ???
```

Part 2

Efficiently compute an RDD that outputs the names of the 20% lowest-performing employees.

```
val low20: RDD[String] = ???
```

Question 1.2

WRC is quite an old company and has a large list of life-long clients who managed to get a contract a long time ago, when everything was cheaper. Naturally, WRC doesn't benefit as much as it could from those clients.

To optimize its revenue, WRC decided to do a proactive outreach to such clients and try to upsell them a newer, fancier plan that is actually just the same thing but more expensive.

Part 1

Compute an RDD of client's full names whose latest contract is older than 10 years old.

```
val upsellClientNames: RDD[String] = ???
```

Part 2

For every upsell client, find the cheapest plan of 2017 that's still more expensive than the current one they have at the moment, if any.

```
val upsellSuggestion: RDD[(String, Plan)] = ???
```