

Main Part 1 (Scala, 6 Marks)

Important

- Make sure the files you submit can be processed by just calling `scala <<filename.scala>>` on the command line.¹ Use the template files provided and do not make any changes to arguments of functions or to any types. You are free to implement any auxiliary function you might need.
- **Do not leave any test cases running in your code because this might slow down your program!** Comment out test cases before submission, otherwise you might hit a time-out.
- Do not use any mutable data structures in your submissions! They are not needed. This means you cannot create new `Arrays` or `ListBuffers`, for example.
- Do not use `return` in your code! It has a different meaning in Scala than in Java. It changes the meaning of your program, and you should never use it.
- Do not use `var`! This declares a mutable variable. Only use `val`!
- Do not use any parallel collections! No `.par` therefore! Our testing and marking infrastructure is not set up for it.

Also note that the running time of each part will be restricted to a maximum of 30 seconds on my laptop.

Disclaimer

It should be understood that the work you submit represents your **own** effort! You have implemented the code entirely on your own. You have not copied from anyone else. Do not be tempted to ask Github Copilot for help or do any other shenanigans like this! An exception is the Scala code I showed during the lectures or uploaded to KEATS, which you can freely use.

Reference Implementation

Like the C++ assignments, the Scala assignments will work like this: you push your files to GitHub and receive (after sometimes a long delay) some automated feedback. In the end we take a snapshot of the submitted files and apply an automated marking script to them.

¹All major OSes, including Windows, have a command line. So there is no good reason to not download Scala, install it and run it on your own computer. Just do it!


In addition, the Scala coursework comes with a reference implementation in form of jar-files. This allows you to run any test cases on your own computer. For example you can call Scala on the command line with the option `-cp drumb.jar` and then query any function from the template file. Say you want to find out what the function `get_january_data` produces: for this you just need to prefix them with the object name `M1` and call them with some arguments:

```
$ scala -cp drumb.jar
```

```
scala> M1.get_january_data("FB", 2014)
val res2: List[String] = List(2014-01-02,54.709999,....)
```

Hints

For Main Part 1: useful string functions: `.startsWith(...)` for checking whether a string has a given prefix, `_ ++ _` for concatenating two strings; useful option functions: `.flatten` flattens a list of options such that it filters way all `None`'s, `Try(...).getOrElse ...` runs some code that might raise an exception—if yes, then a default value can be given; useful list functions: `.head` for obtaining the first element in a non-empty list, `.length` for the length of a list; `.filter(...)` for filtering out elements in a list; `.getLines.toList` for obtaining a list of lines from a file; `.split(",").toList` for splitting strings according to a comma.

 **Note!** Fortunately Scala supports operator overloading. But make sure you understand the difference between `100 / 3` and `100.0 / 3`!


Main Part 1 (6 Marks, file `drumb.scala`)

A purely fictional character named Mr T. Drumb inherited in 1978 approximately 200 Million Dollar from his father. Mr Drumb prides himself to be a brilliant business man because nowadays it is estimated he is 3 Billion Dollar worth (one is not sure, of course, because Mr Drumb refuses to make his tax records public).

Since the question about Mr Drumb's business acumen remains open, let's do a quick back-of-the-envelope calculation in Scala whether his claim has any merit. Let's suppose we are given \$100 in 1978 and we follow a really dumb investment strategy, namely:

- We blindly choose a portfolio of stocks, say some Blue-Chip stocks or some Real Estate stocks.
- If some of the stocks in our portfolio are traded in January of a year, we invest our money in equal amounts in each of these stocks. For example if we have \$100 and there are four stocks that are traded in our portfolio, we buy \$25 worth of stocks from each. (Be careful to also test cases where you trade with 3 stocks.)
- Next year in January, we look at how our stocks did, liquidate everything, and re-invest our (hopefully) increased money in again the stocks from our portfolio (there might be more stocks available, if companies from our portfolio got listed in that year, or less if some companies went bust or were de-listed).
- We do this for 41 years until January 2019 and check what would have become out of our \$100.

Until Yahoo was bought by Altaba a few years ago, historical stock market data for such back-of-the-envelope calculations was freely available online. Unfortunately nowadays this kind of data is more difficult to obtain, unless you are prepared to pay extortionate prices or be severely rate-limited. Therefore this part comes with a number of files containing CSV-lists with the historical stock prices for the companies in our portfolios. Use these files for the following tasks.

 **Note:** Do not hardcode the path to the CSV-files. The testing framework will assume that these files are in the same directory as the `drumb.scala` file.

Tasks

- (1) Write a function `get_january_data` that takes a stock symbol and a year as arguments. The function reads the corresponding CSV-file and returns the list of strings that start with the given year (each line in the CSV-list is of the form `someyear-01-someday, someprice`). [0.5 Marks]
- (2) Write a function `get_first_price` that takes again a stock symbol and a year as arguments. It should return the first January price for the stock symbol in the given year. For this it uses the list of strings generated by `get_january_data`. A problem is that normally a stock exchange is not open on 1st of January, but depending on the day of the week on a later day (maybe 3rd or 4th). The easiest way to solve this problem is to obtain the whole January data for a stock symbol and then select the earliest, or first, entry in this list. The stock price of this entry should be converted into a double. Such a price might not exist, in case the company does not exist in the given year. For example, if you query for Google in January of 1980, then clearly Google did not exist yet. Therefore you are asked to return a trade price with type `Option[Double]`...`None` will be the value for when no price exists; `Some` if there is a price. [1 Mark]
- (3) Write a function `get_prices` that takes a portfolio (a list of stock symbols), a years range and gets all the first trading prices for each year in the range. You should organise this as a list of lists of `Option[Double]`'s. The inner lists are for all stock symbols from the portfolio and the outer list for the years. For example for Google and Apple in years 2010 (first line), 2011 (second line) and 2012 (third line) you obtain:

```
List(List(Some(312.204773), Some(26.782711)),  
      List(Some(301.0466), Some(41.244694)),  
      List(Some(331.462585), Some(51.464207)))
```

[1 Mark]

- (4) Write a function that calculates the *change factor* (delta) for how a stock price has changed from one year to the next. This is only well-defined, if the corresponding company has been traded in both years. In this case you can calculate

$$\frac{price_{new} - price_{old}}{price_{old}}$$

If the change factor is defined, you should return it as `Some(change_factor)`; if not, you should return `None`. [1 Mark]

- (5) Write a function that calculates all change factors (deltas) for the prices we obtained in Task (2). For the running example of Google and Apple for the years 2010 to 2012 you should obtain 4 change factors:

```
List(List(Some(-0.03573991804411003), Some(0.539974575389325)),  
      List(Some(0.10103414222249969), Some(0.24777764141006836)))
```

That means Google did a bit badly in 2010, while Apple did very well. Both did OK in 2011. Make sure you handle the cases where a company is not listed in a year. In such cases the change factor should be None (recall Task (4)). [1 Mark]

- (6) Write a function that calculates the “yield”, or balance, for one year for our portfolio. This function takes the change factors, the starting balance and the year as arguments. If no company from our portfolio existed in that year, the balance is unchanged. Otherwise we invest in each existing company an equal amount of our balance. Using the change factors computed under Task (2), calculate the new balance. Say we had \$100 in 2010, we would have received in our running example involving Google and Apple:

$$\begin{aligned} \$50 * -0.03573991804411003 + \$50 * 0.539974575389325 \\ = \$25.21173286726075 \end{aligned}$$

as profit for that year, and our new balance for 2011 is \$125 when converted to a Long. Since `yearly_yield` should produce a Long, there are a number of ways how to round doubles. One way to do the calculation is to calculate the profit first as Double, and then round the result down to a Long (using `.toLong`) and add it to the balance (which is also a Long). [1 Mark]

- (7) Write a function that calculates the overall balance for a range of years where each year the yearly profit is compounded to the new balances and then re-invested into our portfolio. For this use the function and results generated under (6).

[0.5 Marks]

Test Data: File `drumb.scala` contains two portfolios collected from the S&P 500, one for blue-chip companies, including Facebook, Amazon and Baidu; and another for listed real-estate companies, whose names I have never heard of. Following the dumb investment strategy from 1978 until 2019 would have turned a starting balance of \$100 into roughly \$39,162 for real estate and a whopping \$462,199 for blue chips. Note when comparing these results with your own calculations: there might be some small rounding errors, which when compounded lead to moderately different values.

Moral: Reflecting on our assumptions, we are over-estimating our yield in many ways: first, who can know in 1978 about what will turn out to be a blue chip company. Also, since the portfolios are chosen from the current S&P 500,

they do not include the myriad of companies that went bust or were de-listed over the years. So where does this leave our fictional character Mr T. Drumb? Well, given his inheritance, a really dumb investment strategy would have done equally well, if not much better. Anyhow, one would assume that this guy is by now locked up in a prison and the key thrown away, but alas he is still around annoying commonsense people. What a pity.