

Quant Trading Guide

Callum McDougall
email cal.s.mcdougall@gmail.com

November 2020

Contents

0	Introduction	2
1	What is quant trading?	3
1.1	Common terminology	3
1.2	What does a quant trader do?	3
1.3	Quant Trading vs Quant Research	4
1.4	Common misconceptions	4
1.4.1	Finance culture is terrible	4
1.4.2	The hours in finance are really bad	5
1.4.3	You need to be a maths genius	5
1.4.4	You need to be a really good coder	6
1.4.5	You need lots of finance experience	6
1.4.6	You need an amazing CV	7
1.4.7	If you make a small mistake in trading, you could lose your job	7
1.4.8	Working in finance is morally wrong	8
1.5	Good things about quant trading	8
1.6	Bad things about quant trading	9
1.7	What to do if you think you might be interested	10
2	Quant Trading Firms & Internships	11
2.1	What is an internship like?	11
2.2	Internship structure	12
2.3	Quant trading firms	13
3	How to prepare for interviews	15
3.1	Solving problems	15
3.2	Types of interview problems	16
3.2.1	Logic problems	16
3.2.2	Maths problems	17
3.2.3	Strategy game problems	18
3.2.4	Fermi problems	19
3.3	Market Making	20
3.4	Adding extra dimensionality	21
3.5	Fit questions	22
3.6	Finance Knowledge	23
4	Problems & Solutions	24
4.1	Logic Problems	25
4.2	Maths Problems	34
4.3	Strategy Game Problems	47
4.4	Fermi Problems	57

0 Introduction

Hello to whoever is reading this! Over the past few months, I've spent a frankly ridiculous amount of time preparing for quant trading interviews, and once they were all over I was left with a giant folder of interview notes that I had no idea what to do with, which is why I decided to write this guide.

Well, sort of. One other reason is that I think there are a ton of misconceptions about careers in quant trading (and quant finance more generally). The industry definitely isn't suitable for everyone, but I think there are way more people who should seriously consider it than who actually do.

In this document, I'll mainly be focusing on quant trading rather than research. I've made a few points about quant research in [section 1.3](#)), and if you think this sounds interesting then I'd definitely recommend finding other resources specific to research, because the advice here will be a lot less directly relevant.

It's worth noting here that I'm definitely not an expert on all this stuff. I wrote this guide mainly from the perspective of *"What would I like to have known when I started applying to these places?"*. Whatever position you're in when you read this, I hope you can get some use out of it!

- [Section 1](#) gives a brief overview of what a career in quant trading might look like. It focuses on several key misconceptions that many people have about quant trading.
- [Section 2](#) outlines what an internship in quant trading looks like, and lists some of the firms that offer internship roles (mainly in London and the rest of Europe).
- [Section 3](#) describes some of the key skills & competencies that are looked for during trading interviews.
- [Section 4](#) has a large selection of interview-style problems sorted by category, as well as their solutions.

This doesn't have to be read in order, so feel free to jump to whichever sections interest you. My only request is that you not jump straight to the problems at the end, because just practicing them without thinking about why these questions are asked and what kind of skills they are designed to test, is a pretty short-sighted way of preparing for interviews (for more information on this, see [Section 3](#)).

Happy reading!

1 What is quant trading?

1.1 Common terminology

I'm not going to provide a massive glossary here; there are some great resources online (like Investopedia or the website of basically any company listed in [Section 2](#)). However, I think there are a few terms that it's important to define clearly. If you know what all of these mean, feel free to skip this section.

- **Market**

This is a place where buyers and sellers can meet to trade things. People are constantly quoting bid and ask prices. The **bid price** is the highest amount someone is willing to buy at; for instance if someone has bid £10 it means they are willing to buy for any price \leq £10. The **ask price** is the price someone is willing to sell at, so an ask of £20 means they are willing to sell for any price \leq £20. These values are constantly changing as people submit new orders into the market. The ask will always be higher than the bid; this difference is called the **bid-ask spread**. A trade happens when someone **crosses the spread**, i.e. the inequalities of a buyer and seller cross.

For instance, suppose people are trading shares in company XYZ. If the highest bid for a share is £99 and the lowest offer is £101, and someone makes a bid of £102, they will be able to buy shares for this price (because the person who was offering £101 will be happy to sell shares to them for £102).

- **Liquidity**

Liquid markets are markets you can trade on easily. Typically it means there is a lot of order flow (i.e. lots of people are trading frequently) and a narrow bid-ask spread, so you can normally buy and sell at close to the current trading price.

For instance, the market for shares in a large US company like Apple will be typically be very liquid, but it might become less liquid during a time of market panic (e.g. like caused by Covid), because people will be less sure of what the true value of shares are, and they'll be less willing to trade. There is a cost to having a bid/offer open, which is related to the concept of *asymmetric information* - if you're very unsure about the true value of something, it's much more likely someone will come along with a much better idea of what the value is, and take advantage of the prices that you're offering.

- **Market making**

This is the practice of quoting a bid and ask for a particular good or service. It means the market maker guarantees to take the other side of a trade at certain prices. The advantage for the rest of the market is more liquidity, because the bid-ask spread is usually a lot narrower in the presence of a market-maker than it would be otherwise. The advantage for the market maker is profiting from the bid-ask spread - on average, they'll be buying the stock for slightly less than they sell it for.

- **Prop trading**

This is short for proprietary trading; it just means that the firm trades with its own money, rather than taking clients' money. Lots of people interested in quant finance confuse prop trading firms with **hedge funds**, which are often very quantitative (e.g. Renaissance Technologies, DE Shaw), but which do take on external money. Almost all the places I list in section 2 are prop trading firms.

1.2 What does a quant trader do?

Unfortunately, the industry is kind of opaque, so it can be pretty difficult to get a good idea of what it is people actually do. It can also be quite hard to give specifics, because there are many different areas you might specialise in quant trading, and 2 people with the same job titles can be doing very different things. I'm also not an expert on this area, because I've only done internships, and the actual job of quant trading is a bit different from what you'll do in an internship (more on this in section 2), so please take this section with a pinch of salt. However, I can give a basic outline of what a career in quant trading might look like.

Broadly speaking, quant trading is about making trading decisions in real time. You are required to think fast, react to events as they happen, and generally try to develop a picture of what is going on in the financial markets. Things that correlate with skill at quant trading include:

- strong quantitative skills (more on this later)
- being able to make decisions under uncertainty, based on intuition / quick judgements
- having a head for probabilities, risks, expected values, etc
- enjoying strategy games like poker, chess, MTG

One common confusion people have goes something like this: “*Why do people have jobs doing quant trading, when there are algorithms who trade way faster than they can?*”. The simple answer to this question is that you can’t automate everything. Some markets aren’t suitable for algorithms to trade in without human interaction, maybe because they are very illiquid or there isn’t enough data to create models / train algs. Also, even in markets where algorithms are trading, you need people to supervise the algorithm in case it goes off the rails in response to some significant market event, or for some technical reason, or just to step in if the algorithm builds up too much risk in one area. In fact, it’s often the times when algorithms start behaving weirdly that traders can add the most value.

1.3 Quant Trading vs Quant Research

Roles in quant finance broadly split into 2 categories: quant trading and quant research. I’ve already given an outline of quant trading, and the majority of this document is aimed at people who are interested in that, but in this section I give an extremely brief summary of quant research and how it differs from trading.

While quant trading focuses on making decisions in real time, quant research is much more focused on making models. A lot of the job is spent researching trading strategies, thinking about ideas, reading papers, backtesting and number-crunching, dealing with large messy data sets, etc. Interviewers will test you a lot more on stats and programming. Things that correlate with skill at quant research include:

- skill at programming and data science / statistical learning / machine learning
- enjoying thinking for a long time about hard problems
- enjoying self-directed research, sometimes with no clear answers
- creativity (in coming up with hypotheses to test, or developing trading strategies)

Unless explicitly stated, anything that follows in this document should be assumed to refer to quant trading, not quant research. As mentioned earlier, if this sounds interesting to you then I recommend you look for other resources on quant research - you may find the blog posts [here](#) and [here](#) helpful.

1.4 Common misconceptions

In this section, I’m going to list a few of the common misconceptions that people have about a career in quant finance, or common misguided reasons not to pursue this career. Not all these points are totally wrong (some of them do have merit) but I’ve chosen them because I think they are generally overblown, and can be quite misleading.

1.4.1 Finance culture is terrible

Lots of finance culture is pretty terrible, but quant trading is actually a subset with a really nice culture on the whole. Most places feel a lot more like tech firms than typical finance firms. Jeans and t-shirts will be the norm, not suits. The specifics vary between companies, but you definitely don’t get Wolf of Wall Street-type shenanigans going on. As a general rule, if you are doing a STEM degree at a good university and you like the people you’re working with, then you’ll probably like the people you meet in quant trading, because it’s essentially the same group.

As a small anecdote to illustrate the informal office environment, I was doing a virtual internship at a trading firm last year. On group videochats, we could hear some noises from the trading floor, and most traders had programmed their computers with sound effects that go off when a trade is made or a stock breaks a price level. These sound effects ranged from the normal to the bizarre, including:

- Bells and alarms
- Cash registers
- Machine guns and blasters
- Animal noises (cats, horses and cows)
- Mario sound effects
- Darude Sandstorm
- Borat (“I like!”, “WaWaWeeWa!”)
- John Bercow (“ORDER!”)
- Obi-Wan Kenobi (“May the force be with you”)

After hearing all this, it was hard for me to keep thinking of trading firms as extremely formal, corporate environments!

1.4.2 The hours in finance are really bad

Finance is known for absolutely insane hours, i.e. 14+ per day for the first few years in investment banking. Quant trading is *nothing like this!* However, the hours are probably worse than an average job, and it’s completely legitimate to be put off by this if good work hours are something that matters a lot to you. Mornings are usually around 7-8am, most people leave at 5-6pm. This varies between firms, as well as your role, e.g. which global markets you trade. The hours can also stretch longer during certain exceptional circumstances (e.g. chaotic market conditions, like those caused by covid, can lead to longer days). However, there’s generally a pretty good work-life balance, and you aren’t expected to think too much about work stuff when you go home. Having to work weekends is exceedingly rare.

1.4.3 You need to be a maths genius

This is maybe the most common concern people have, and also probably the most valid, because the industry is very competitive, and you do need maths skills. However, these aren’t the same skills you’d need e.g. for maths olympiads / exams / STEM-based degree subjects. Being good at this stuff does correlate with quant trading, but the correlation is far from perfect, so not being amazing at olympiads or top of your cohort in university shouldn’t stop you from applying.

Mental maths can be pretty important, but isn’t the be-all and end-all. Some firms (e.g. Optiver and Flow Traders) have internship applicants take numerical tests, which can be tough (especially for people who haven’t done similar things before), but they are mainly used as an initial filter stage, since trading firms get a lot of applicants. If you want to practice this kind of mental maths, [this website](#) provides a good resource. I would recommend the Optiver-style test - the questions are harder than they are in the real Optiver test, and the real Optiver test is about as hard as numerical tests get, so if you can do okay in this then you should be able to cope well with most mental maths you might come across. As a point of reference, if you can get 50+ marks after practice for a while, then that’s really good (and even if you can’t get this high, it might not be a deal breaker).

The more general maths skills at later interview stages are nowhere near university level (for more information, see section 2). If you have a good handle on probabilities, expected value, risk, etc, this will get you a long way. The [Jane Street pdf](#) on quant trading interviews provides a very good summary of these concepts. Some brainteasers might require knowledge of statistical distributions to solve efficiently: the most common are uniform, normal, binomial, and geometric.

As for amazing mathematical achievement like IMO participation or medals, don’t worry - these are absolutely not necessary. They’re great to have, and some trading firms do like publicising how many IMO medallists they’ve employed, but if they insisted on only hiring people who qualified to IMO, then trading floors would be a ghost town.

Of course, it can't hurt to get more mathematically literate, and if this is something you're interested in then it's definitely worth your time. The areas I would prioritise (in roughly decreasing order of priority) are:

- **Bayesian statistics**

This comes up a lot in brainteasers in some form. Also, the core ideas of updating beliefs based on evidence are pretty important in trading. To this end, Bayesian statistics can give you some really valuable tools (e.g. odds ratios, which provide a natural and intuitive way to understand how much you should adjust probabilities based on the quality of evidence). I'm a bit biased here, because I personally think Bayesian statistics is awesome.¹

- **Statistics & data science**

Linear regression is used *a lot* in finance, so this can definitely be worth getting familiar with. At the more advanced level, some knowledge of some ML and data science can make your application stand out, although it's unlikely to be directly useful for interviews. This is a totally different story for quant research, when this stuff would be much higher-priority.

- **Markov chains**

Sometimes useful for brainteasers in interviews, but unlike stuff like uniform distributions, you'll never need Markov chains to solve problems.

- **Game theory**

A bit more niche, but a very interesting branch of study, that intersects with other fields like maths, economics, social science and computer science.

1.4.4 You need to be a really good coder

Coding is always a good thing to learn, but the vast majority of quant trading interviews won't require you to code (whether the job description explicitly mentions programming is a good litmus test). I'd still definitely recommend learning to program if you currently don't do much, because many traders spend a lot of their time coding, and so it helps to know how much you enjoy it. If you want to get started, I'd get familiar with Python and Jupyter Notebooks, since these tools are used very frequently in quant trading. As for libraries, important ones include Numpy, Pandas, Scipy, and Matplotlib. A cool programming project can stand out on your CV too (more on this in [Section 1.4.6](#)).

All of this is very different for quant research, where a lot more emphasis is put on programming literacy and stats knowledge.

1.4.5 You need lots of finance experience

Almost every firm will tell you this during the application process: finance experience is nice to have, but not essential. Most places will teach you the finance you need to know on the job. Some firms might like to see an *interest* or *passion* for finance, which might take the form of:

- Being involved in a finance society at your uni (e.g. CUFIS or CUATS in Cambridge)
- Reading news sources like the FT, Bloomberg, or Matt Levine's Money Stuff
- Trading your own account (or on a platform like QuantConnect or Quantopian²)
- Data-science projects related to finance
- Past internships

But while any of these will boost your application, none of them are a dealbreaker if you don't have them on your CV. It's unlikely to be a great use of your time to try these things just for the sake of impressing interviewers, so so them if you think they genuinely sound interesting.

¹(because it is!)

²pre-November 2020. [RIP Quantopian](#)

1.4.6 You need an amazing CV

First point to make: if you're studying a quantitative degree at a good university, that already puts you ahead of lots of applicants. Lots of people feel intimidated when they write a CV, because they don't think they have enough stuff to put on there. However it's worth remembering, part of what trading firms are looking for on a CV is just an interesting person who isn't defined by the degree they're studying, so if there's anything that you're passionate about, you can include it! Here are a few good things to consider putting on your CV, if you can:

- Olympiads (e.g. maths/informatics/physics)
- Finance-related stuff (see [Section 1.4.5](#))
- Strategy games (the classic one is poker, but also bridge, chess, diplomacy, MTG, etc)
- Sports (especially team / competitive sports)
- Society leadership (especially any societies related to the points above)
- Coding projects (especially any projects related to the points above)
- Tutoring (good communication and mentoring skills are super valuable)
- Any random but remotely impressive stuff³

Second point to make: CVs don't matter that much at later interview stages. Mostly they are useful to get you through the door, and if you have some relevant experience/internships then you might be asked about them (this varies between companies), but in general trading firms care much more about how you perform in the interview than your actual CV. For more on this, see [Section 3](#).

Third point to make - referrals can be another good way to get your foot in the door. If you know anyone who works there / did an internship, this can often suffice to getting a first round interview. Finding talented people is hard, so companies often like getting referrals as a way of making the finding process easier for them!

1.4.7 If you make a small mistake in trading, you could lose your job

There can be reasons to be concerned about job security (see [Section 1.6](#)), but I don't think this particular worry is very well-founded. Everyone makes mistakes, but one thing quant trading firms really value (in fact one of the things they specifically look for in interviews) is the ability to notice, admit to, and learn from your mistakes. If you make a mistake while trading and lose some money, but take full responsibility for it and make sure not to make the same mistake again, it's very unlikely you'd be fired or demoted for this.

There is a limit to this - obviously if you lose all the firm's money then it will be a different matter! However, most firms put measures in place to make sure that one single person can't screw up badly enough to really hurt the firm in a significant way. Junior traders will be doing simulated trading, or else overseen by senior traders, for quite a while when they first join.

All this said, some firms will treat mistakes more harshly than others. Since quant trading firms are frustratingly opaque, it's a good idea to speak to as many people as you can, and try and get an idea of what the culture is like in the firm. If you get the impression that mistakes aren't really tolerated, that can be a good sign that you should find a better place to work!

³I personally hold a Guinness World Record for the [lamest possible thing you can imagine](#). I put this down on my CV, and was asked about it by no less than 4 interviewers during my interview process for summer 2021 internships!

1.4.8 Working in finance is morally wrong

The truth is this is a pretty complicated point to address. There are definitely finance companies who have a net negative impact on the world, but quant trading definitely isn't in the category of worst offenders. A lot of what quant traders do is provide liquidity to markets, i.e. make it slightly easier for other market participants to buy and sell things on exchanges. This is a useful function (since these counterparties can often be e.g. pension funds), but on the margin it's hard to argue that quant trading firms are doing an absolutely necessary societal good.

However, going into quant trading to earn money and donate it to charity can still do a huge amount of good. This is a big part of what personally motivates me, and I think it's a really important factor to consider. As a quick [Fermi Estimate](#): it costs approximately \$4000 to save a life according to [GiveWell's calculator](#), and so giving 50% of a \$200k per annum salary = 25 lives saved per year.

As a caveat to this, if you're going into quant trading just to earn money (either for yourself or for charity), you probably won't be able to stay motivated. Quant trading is a pretty intense job, and to work here you need to enjoy what you do as well. Also, if you are the kind of person who feels they need to do something with a direct positive impact, then quant trading might not be for you. However, if you're interested in the work, and feel like you would be motivated by the idea of donating a lot to help make the world a better place, then this is great! For more on this, check out [80,000 Hours' quant trading resources](#).

1.5 Good things about quant trading

I've already mentioned some personal factors that tend to correlate with being a good quant trader. In this section, I'm going to list a few more general positive aspects of a career in the industry.

- **Intellectually stimulating**

Quant trading tends to draw on a wide variety of disciplines: maths, stats, computer science, data science, finance, game theory, etc. The work is consistently challenging and stimulating, and the general culture of excellence and high competitiveness means you are always developing and learning new skills.

- **Work with smart, talented people**

Thanks in part to the above point, quant trading tends to attract really smart people who are all passionate about what they do. The vibe is quite similar to a Cambridge maths cohort. More generally, if you're doing a STEM subject from a top university and you like the people you're studying with, there's a good chance you'll like the people you meet in quant trading.

- **Tight feedback loops**

If you have an idea, it's very easy to test it and see how well it is doing, because the objective function of "does this make money?" is pretty clear. It's easy to see if you're making progress, and you'll spend very little time wondering about whether what you're doing is actually useful.

- **Autonomy**

In most trading firms, you have the ability to work on a wide variety of different tasks/projects. If there's something you find interesting, and you can prove that it's a valuable thing to work on, then you'll be allowed to pursue it and given the resources to do so.

- **Flat structure**

There is generally little to no hierarchy and bureaucracy in these firms. Everyone is listened to, regardless of how long they've worked there. Age and experience counts for a lot, but most people will have the same job title. Primarily, you are judged for your character and how good your ideas are.

- **High pay**

I don't want to emphasise this point too much because it would probably attract the wrong people. If you only go into quant trading for the high salaries then you'll almost certainly burn out. That said, I couldn't write a document about quant trading and *not* mention salaries! They are almost offensively large, even at the intern level, and grow extremely fast over time as you stay at the firm. There are also usually perks, like free food, in-office gyms and fitness resources.

1.6 Bad things about quant trading

In the interests of balance, I thought it would be worth including this section. I think quant trading is a great career that way more people should consider than actually do, so obviously I'm biased. That being said, it's still not for everyone. There are some misguided reasons to choose not to go into quant trading (see [Section 1.4](#)), but there also some genuine reasons why you might not be a good fit for this career. Here are some of them:

- **Long hours / high intensity**

As stated above, the hours are nowhere near as bad as investment banks, but they are still a little longer than the average 9-5 job. 10 hours per day is typical. The day-to-day responsibilities of a trader are also very intense, and the job can be mentally & emotionally draining.

- **Not great job security**

This varies from firm to firm. Some firms actually have pretty good job security (and as a result, have a strong incentive to provide good education and training to employees). However, there are definitely some places who take the option of "hire lots of people; only keep the ones who do well." This, combined with the tendency to burn out (or just choose to transition to less intense jobs), helps to explain why the average age for employees in quant trading firms can be as low as late 20s.

- **Non-compete agreements**

Firms are often competitive about all their IP, and that includes employees. Many places will have you sign a non-compete, i.e. an agreement not to work for a competitor for a period of time after you leave. Some of these are infamously strict (e.g. on the order of 2 years or more). You will generally be paid a small fraction of your salary during this time.

- **Poor exit opportunities**

Quant trading doesn't give you as many obvious exit opportunities as other finance careers (like investment banking). The skill of "making a ton of fast decisions in real time" doesn't really transfer quite as naturally to other industries as the softer skills of team management and client interaction, so you might not develop much career capital. This heavily depends on your role, e.g. if you do a lot of coding, you might be able to move into data science / software somewhere else. One important point - in most firms, you'll have a decent bit of autonomy over your role, so you can gravitate more towards coding if you want to.

- **Flat structure**

I think most people would see this as a positive, but some people do enjoy the ability to have some kind of career progression. The longer you stay at trading firms, the more responsibility and knowledge you will accumulate, and (in all likelihood) the more money you'll make, but if you enjoy having a title bump every few years, then you might not like how this structure works.

- **Competitive Industry**

There are some industries with more jobs than people to take them; quant trading is not one of them! Most firms only offer internships to a small group of people (15-20 is typical, sometimes less, and I don't know any who take more than 40). As a result, the interviews can be extremely competitive, and you may have to apply to many different firms for a chance of securing a space. For anyone applying to quant trading roles, it's worth having some kind of backup plan, so you aren't putting all your eggs in one basket.

Another point - even if you don't think you will be able to get the offer and/or quant trading isn't your first choice, it can still be worth applying, because getting internship offers has great information value. If you can get an offer, the odds are pretty good that you're the kind of person who will enjoy the work.

1.7 What to do if you think you might be interested

If all this stuff sounds cool to you (and the bad stuff isn't too offputting!) then I'd recommend trying to learn more about the industry. One of the best ways is to go to events at your university. Most finance careers events are clustered in the first term (e.g. Jane Street estimathons), although some places do still run events later in the year. These events are a fun way to get to know professionals in the industry, as well as ask questions.

However, nothing is a substitute for getting a feel for yourself of what the industry is like, which is why applying for internships is basically the best thing you can do. In later sections I'll talk a *lot* more about the application process and what interviews are like. For now, I'll just give the key points: application process is low-effort (just a CV and cover letter needed, sometimes not even a cover letter), interviews can actually be kind of fun, and firms put a lot of emphasis on education and making a good internship experience. You'll get to interact with traders a lot, probably get a view into lots of different areas in the firm, and the projects interns are given tend to be really worthwhile and interesting.

Even the interview process itself can be super high value - you'll get the chance to ask the interviewer questions, and if you get an offer then you'll usually have more opportunities to talk with people in the firm about your main uncertainties. There will probably be a strong bias because they're sales pitching you the company, but you can still usually learn a lot of useful things.

One final point to make - some firms (like Jane Street and Optiver) run Spring Week events or Insight Days, usually aimed at first or second years. These can be great ways to get on the radar of these companies, and can often be a stepping stone to securing an internship. To find out about all this stuff, the best thing is to look on the company websites or consult your university careers service.

2 Quant Trading Firms & Internships

2.1 What is an internship like?

Quant trading internships tend to be pretty homogeneous. The details might vary, but the overall structure is usually the same. In the next section, I've sketched out a hypothetical 10-week⁴ internship in quant trading that I believe is representative of the average firm. However, before getting into that, there are a few more general points about quant trading internships that are worth mentioning.

Compared to most industries, quant trading firms put a lot of effort into their internships. They function as an extended interview process, and as a way for the firm to sell themselves to you. For this reason, the internships are kept fun and intellectually stimulating - you won't be doing photocopies or making coffee! Even if you aren't sure that you want to go into finance, you can still get a lot of value out of the internship. That said, the majority of the value to the firm doesn't come from what you do during the internship; it comes from the possibility of you becoming a full-time hire.

The internships are usually very well paid (often close to or even above 6 figures pro rata⁵). Salaries can be a good proxy for how high-end a firm is, although sometimes they can be inconsistent - for instance, one of the firms I list in section 2.3 actually doubled their intern salary in a single year recently. Salaries also aren't the most important way to choose between firms, because if you end up working in quant finance they'll be a very small fraction of the total amount of money you'll make.

As a last general point about quant finance interviews, it's worth emphasising that you will have the opportunity to ask lots of questions in many different situations - take advantage of this opportunity! In my previous internship, we were told that the firm had never had an intern who, in their eyes, asked too many questions - although they consistently get people who don't ask enough. There is a natural aversion to asking questions, if you don't want to seem dumb/in need of assistance, but this is the wrong attitude to have. Nobody is expecting you to have all the knowledge you need right away, but they are expecting you to recognise your blind spots and earnestly try to fill them.

⁴This seems to be the modal length. All internships I know of are between 6 and 12 weeks.

⁵Pro rata refers to the adjusted annual salary. For instance, a pro rata salary of £80k for a 10 week internship would work out to about £15k in total.

2.2 Internship structure

- **Education**

The first 2-3 weeks of the program usually consists of education, through internal classes or talks. A lot of this will be finance-based, covering concepts like markets, stocks, options, futures, indices, ETFs, market-making, arbitrage, hedging, etc. Firms will probably spend more time on the specific financial products they specialise in (see section 2.1). There will also be classes on programming - topics might include languages like Python, Excel, VBA, SQL, or Linux.

- **Project work**

As mentioned, this is the main way interns directly add value to the firm. You will be given a project (often individual, sometimes as a group) that you work on from the end of the education period to the end of the internship. These projects are usually taken from the list of things traders/researchers would be doing themselves if they had the time (although they are also optimised for being self-contained, fun, and appropriate for people without masses of prior finance experience). To give a flavour, some possible projects might be:

- Investigating the price movements of a specific financial product, with the goal of developing a trading strategy
- Improving the systems used to monitor existing strategies
- Building dashboards/GUIs that provide traders with useful real-time information

These projects are often chosen to be a good match with your specific skills/background/interests. You will be paired with a particular trader/desk for the duration of the project. They will be the person you discuss the project with and who provides you with direction and help when you need it. Also, being sat next to this trader on their desk means you have an opportunity to ask a lot of questions and get an insight into the trading process. At the end of the internship, you will probably get to give a presentation on your project.

- **Mock trading**

Almost all trading firms have some form of mock trading. This is probably the part of the internship that varies between firms the most, although the general purpose is to teach you to be a good trader: understand your biases, get better at making decisions under uncertainty, have a good intuition for risk, don't panic and freeze up, etc. The precise trading methods are proprietary to firms, but without giving names, I know of three different approaches that have been taken by some of the firms I list later in this section:

- Games with dice or cards that involve asymmetric information and difficult strategies
- Open-outcry trading system based on a theoretical stock or option
- In-house trading simulation software, designed to come as close as possible to real trading

For mock trading which requires no prior knowledge (e.g. as described in the first point), this will run for the entire internship. If some background finance knowledge is required, the mock trading might only start later in the internship.

- **Other stuff**

Most firms will organise social events throughout the internship. I can't really speak much to this point, because my trading internship was fully virtual thanks to Coronavirus. However, we still had a few social events organised. Often these will be themed around strategy games (commonly poker, but others as well) or estimation & Fermi problems (see section 2.2). There are also likely to be some social events with nothing to do with trading - some that I've heard include wine-tasting, cupcake decoration, sports-themed events, restaurant trips, theatre trips, and escape rooms.

2.3 Quant trading firms

This section gives a very brief overview of the most well-known quant trading firms that offer internships. I've mainly focused on quant trading roles based in Europe, and even then I haven't included every firm that offers internships. A few significant places that are missing from this section are Tibra Capital, DRW, Da Vinci Derivatives, Hudson River Trading, and more research-oriented firms like Jump Trading and GSA, as well as hedge funds that offer quantitative roles like DE Shaw and Two Sigma. There are also loads more options in the US, especially if visas aren't an issue for you. A fuller database of internships can be found [here](#). However, I think this list should provide a good starting point.

Another point before getting started with this list - it can be a really good idea to spam apply to loads of places and see what happens. It's a really competitive process, so doing this will maximise your chances of success, as well as giving you lots of practice for the ones you might care more about!

Jane Street has offices in London, New York and Hong Kong. They are known for having a very maths-focused, academic culture. They seem to have the vibe of "people who might have done a PhD, but didn't". They are also infamous for their extremely difficult interview process, which involves market-making on some tricky games/problems, usually involving dice, cards or coins. Their internship is 10 weeks, and takes place in their London offices, although usually involves some time in one of their other global offices.

Susquehanna International Group (SIG) are based in Ireland. They're one of the largest and oldest quant trading firms (quite a few of the others were started by employees who left SIG). They specialise in options, although like virtually all trading firms, they trade a wide variety of products. Probably the thing they're best known for is poker - they teach it during their internship, as a proxy for the skills a trader needs (thinking quantitatively, decision-making under uncertainty, etc). Their internship is 10 weeks long, and takes place in their Ireland offices.

Optiver are based in Amsterdam. Their interview process starts with a numerical test, which is generally regarded as the hardest of the first-round numerical tests offered by trading firms. Later rounds focus on the dynamics of trading and market making, Fermi estimates, and some options theory⁶. Some of the things they are known for include an extremely competitive graduate trader training program (with a very low conversion rate of grad to full time), and very advanced in-house trading simulation software that they use to train interns and graduates. Their internship takes place in their Amsterdam offices.

Flow Traders are based in Amsterdam. They specialise in Delta One products⁷, particularly [ETFs](#). The application process starts with a numerical test, and an aptitude test. The later rounds focus a lot on ETFs and the core business structure, which is something they expect you to read up on, as well as standard brainteasers and quantitative skills. Their internship 6 weeks long, and takes place in their Amsterdam offices.

IMC are based in Amsterdam. They run frequent coding challenges (especially during first term of the academic year), which can be a good way to get exposure to the company. The application process begins with a series of games/challenges to test your cognitive abilities. Later rounds focus on competency questions, brainteasers, market makers, and some light options theory. Their internship is 12 weeks long, and takes place in their Amsterdam offices.

Five Rings are based in the US, but they have offices in London. They were founded by an ex-trader at Jane Street. They are smaller and younger than most of the other companies on this list, although they are growing very fast. Their interview process is pretty competitive; the first round involves a series of rapid-fire questions, and the later rounds are usually structured around a single unusually difficult brainteaser that you will work through with the interviewer. Their internship takes place in their London offices.

⁶For more on options theory, see [Section 3.6](#)

⁷Delta One means their prices move in 1-1 correspondence with the underlying security. This is in contrast to derivatives like options, which have a non-linear relationship with the underlying.

Citadel Securities are a market-making firm. It is one of the two businesses operated by Citadel LLC (the other being the hedge fund Citadel). If you apply for Citadel Securities, your application is often also considered for the hedge fund. Interviews with Citadel can be unpredictable, since they seem to vary a lot based on the interviewer. Standard brainteasers are common. Citadel interviews will usually involve coding problems; they use the CoderPad platform. You are also more likely to be asked questions on finance topics, e.g. related to algorithmic trading strategies or infrastructure. Their internship is 10 weeks long, and takes place in their London offices.

Maven are based in London. They are the youngest firm on this list (founded in 2011). Their initial screening involves a numerical test and a series of one-way video interview questions. Their final round is held with several applicants at once, and starts with a trading game that all the applicants play together. Later one-on-one interview rounds are more focused on competency questions and brainteasers. Their internship is 9 weeks long, and takes place in their London offices.

3 How to prepare for interviews

3.1 Solving problems

I spent a good chunk of [Section 1](#) listing all the things that *aren't* required for getting through quant trading interviews. You might be wondering - what actually *is* required?

One of the most common things you'll have to do in interviews is solve brainteasers. These can vary in type, from finding precise solutions to logic puzzles to developing imperfect but effective strategies for playing complicated games. These problems are asked in interviews because they act as a proxy for a bunch of useful skills that traders need. Some of the most important are:

- **Problem-solving & creativity**

Being able to think outside the box and come up with new and imaginative solutions is a very important skill. It's probably slightly more important for research than for trading, but it's still a top skill that interviewers will be looking for.

- **Making quick judgements / coping under pressure**

This is super important for traders, for obvious reasons. Often, this will take the form of an explicit time limit. Some firms (e.g. Optiver and Maven) organise group interviews where you play market making games with other interns as a direct way of testing how you hold up under stressful market conditions. Timed numerical tests are another form of testing this ability.

In general, most firms will prefer slow but correct answers to fast and wrong answers, but sometimes you will be expected to come up with a quick-and-dirty solution even if it's not the absolute best one. For this reason, it's a good idea to try and be responsive to your interviewer's prompts, and see if you can get an idea of what they are trying to test for.

- **Good calibration**

Calibration is a measure of the accuracy of your own probability estimate about what you believe to be true. A well-calibrated person is neither overconfident nor underconfident; things that he assigns a probability of X to will (on average) happen X% of the time. To put it another way, calibration is the answer to the question "*how well do you know what you know?*".

One classic way firms test this is by asking you to give confidence intervals for Fermi problems (more on these in [Section 3.2.4](#), where I will also make a few more points about over/underconfidence and calibration). Additionally, if they ask questions like "*how sure are you that this is the right answer?*", or "*what do you think the probability is that the true value is larger than X?*", then this is what they're testing for.

- **Understanding of concepts like expected value and risk**

Being able to think in these terms is a trader's bread and butter. One common style of interview question involves the interviewer describing a game to you, and asking you to develop a strategy that maximises your expected return while keeping your risks minimal - and a good first step in problems like these is to estimate the EV and variance (more on these in [Section 3.2.3](#)).

- **Mental maths**

Different firms emphasise this to different degrees. Generally it's more important in initial screening stages, but still matters a fair amount in later interviews. It is strongly linked to the ability to think fast, deal with uncertainty, and apply quantitative reasoning to come up with a quick approximate solution to a problem.

3.2 Types of interview problems

Although all the skills I listed in the previous section are important, some firms will prioritise particular skills more highly than others, and so they will ask questions that are particularly aimed at testing those skills. What follows is a decomposition of the types of questions you are likely to be asked into 4 main categories. Very broadly, they are on a spectrum from what you might call objectivity to subjectivity, or from “answer matters” to “method matters”. On one end of the spectrum are logic problems: these usually have a single objectively correct answer that you are trying to find. On the other end are Fermi problems, which are much more about analysing your decision making and thought processes than coming up with a perfect answer. At the start of each subsection, I list some example questions. Many more can be found in [Section 4](#), along with solutions.

3.2.1 Logic problems

A plane has 100 seats, and each of 100 people have been assigned a unique seat. The first person to board is a drunkard, who sits in a randomly chosen seat. Every subsequent person sits in their own seat, unless someone is already sitting there, in which case they also sit in a random seat. You are last to board. What is the probability you get to sit in the seat you were assigned?

There are 12 balls, 11 of the same weight and 1 of a different weight (it might be lighter or heavier). You have a set of balance scales, which you can use to compare the weight of any two groups of balls. What is the maximum number of uses required to determine both which ball is different, and whether it is lighter or heavier?

These are on the most “objective” end of the spectrum. They have a single correct answer that you are trying to find. They are probably the kinds of things that instinctively come to mind when you hear the word “brainteaser”. Some things the interviewer may be looking for if they give you one of these include:

- Creativity and out-of-the-box thinking
- Clearly communicating your ideas and thought processes
- Responding to hints and direction
- Not getting stuck to your first good idea, unwilling to try another
- Being able to persevere with a hard problem

It’s worth emphasising that firms aren’t just looking for candidates who can do logic problems really well. You could silently blast every problem out of your way with ease and walk away from the interview without really providing any useful signals as to whether you’d make a good trader. On the other hand, if you find the questions difficult and take a long time to get the answer, but make a continuous effort to communicate your thoughts and take on board the interviewer’s hints, you’ll probably get a lot more credit.

One other note - I’m also including problems in this section that have some mathematical content, so there is a bit of overlap between this section and the next one. For more information on where I have demarcated them, see the next section.

3.2.2 Maths problems

What is the expected number of times you'll have to roll a fair 6-sided dice, until you get two 6s in a row?

You flip a fair coin 100 times. What (approximately) is the probability that you get more than 55 heads?

Questions in this section test very similar competencies to logic problems, but there are some significant differences:

- **Stronger quantitative skills are needed**

Most questions in the previous section require logical reasoning and creativity to solve, and several have virtually no mathematical content. Questions in this section will rely much more directly on your quantitative abilities. Sometimes you will need mental maths, but most of the time you will be given pen and paper. Calculators are rarely allowed.

- **The actual content of the problem is more important**

The problems will most often be based around concepts that are directly important to traders, like probabilities and expected values. These are the problems where extra mathematical knowledge (e.g. knowing some of the more important random distributions and their properties) can help a lot.

- **You can't always get a perfect answer**

Although there is usually an objectively correct answer, it may be impossible to calculate without using a calculator or computer program, and you may have to use a quick & dirty method that gets an approximate but not perfect answer.

Mathematicians tend to be perfectionists who don't like making rough approximations. If you've got to interview stage then it should be obvious from your CV that you're smart, so they'll be looking more for things like this.⁸

As a consequence, most of the points listed in the previous section are still relevant, with the following additions:

- Strong maths abilities
- Familiarity with concepts like probability and expected values
- Not being a perfectionist / a willingness to approximate
- Good intuitions (e.g. sanity-checking an answer)

⁸As a meta point here - it can be a good idea to think about your CV, and what is missing from it. If you were an interviewer reading your CV and trying to determine if this candidate would make a good trader, what would be your biggest uncertainties? If you aren't sure, try giving it to friends who have some experience in quant trading, and see what they think.

3.2.3 Strategy game problems

You roll an 8-sided dice, and keep track of the cumulative sum. You can stop and take the sum at any time, but if the sum is ever greater than 8 then the game ends, and you get a payout of zero. What is the best strategy, and what is the expected value?

You roll two 4-sided dice, and your payout is their product. You can pay £2 to re-roll either of the dice, as many times as you like. What is the best strategy, and what is the optimum value? What happens when the cost of re-rolling gets very large, or very small?

There are 10 cards, face down, numbered $\{1, 2, \dots, 10\}$. If you pay £1, you can point to any 2 cards and I will tell you which card is smaller and what the value of that smaller card is. At any point you can finish the game by choosing a card, and you win the value of that card. What is the expected value of this game?

This is the smallest of the four sections, because in a sense it's a slightly niche subset of the previous section. However, I think it has some interesting dynamics that make it different, so it's worth discussing separately.

Strategy game problems almost always center around the same concepts as maths problems in the previous section. The most common settings involve dice, cards, coins, betting games, etc. The interviewer will describe a game you can play, that gives you some variable payout at the end (which depends both on your own strategy and some inherent randomness). The defining feature that makes them different from the previous section is that you have to develop a strategy to play the game, then try to figure out what the expected value of the game is with this strategy. This leads to 2 sources of uncertainty in your answer: whether you have the right "best strategy", and what the value of the game is when you use that strategy. Even when the best strategy is obvious, sometimes the fact that other strategies are available is important.

Often, these problems are much more nuanced than anything in the previous two sections. For example, you need a very good handle on concepts like expected value and risk. If you were buying the right to play game, you would want to buy at below expected value, and you might care a lot more about downside risk (i.e. small probability of a very small payout). On the other hand, if you were selling the right, you would want to sell at above the expected value, and you would have to consider things like whether the person you are selling to has a better strategy than the one you've found or whether he can use a lower-expectancy strategy that has a small probability of very high payout, which would but you at risk of a huge loss. For more on these kinds of considerations, see [Section 3.3](#) on market making.

All of the points in the previous section are important, as well as (to a slightly lesser extent) the points in the section before. Other important skills include:

- Ability to think strategically
- Strong intuitive understanding of concepts like expected value and risk
- Using imperfect methods to get accurate approximations

It is also worth noting that sometimes you might be asked games with no clearly defined (or computationally tractable) expected payout value. These games are much more subjective, with no clear answer. In these cases, thinking strategically and having good instincts might be more important than maths skills.

3.2.4 Fermi problems

How many slices of pizza are eaten in America each year?

How many bus stops are there in London?

If you stacked £1 coins on top of each other, how many would it take to reach the moon?

Fermi problems, also known as back-of-the-envelope calculations, involve the estimation of one or more highly uncertain quantities. They were named after Enrico Fermi, who used this technique to estimate the strength of the atomic bomb that detonated at the Trinity test. The now-accepted value is 21 kilotons of TNT; the result he came to was 10 kilotons. This might seem like a very large error, it was actually very impressive, since Fermi problems often contain uncertainties of at least one order of magnitude⁹, and you may have to take the product of several of these to get an answer.

Many people mistakenly think Fermi problems are about how much trivia you know. It's true that you can get lucky if a Fermi question happens to be around a topic you have background knowledge on, but it should be made clear, **THIS IS NOT WHAT INTERVIEWERS ARE TESTING FOR!** Generally, traders will know what is going on in the markets, but this won't always be the case. Sometimes, things go wrong or start behaving weirdly, and you have a duty as a liquidity provider to keep quoting a bid and ask. In these cases, it's super important to be able to keep your shit together, and try to reason qualitatively towards a solution when you don't have all the answers.

Another form a Fermi problem might take is that of a confidence interval. This means a range of numbers, which you assign a certain level of confidence of containing the true value. For instance, if you are asked for a 90% confidence interval, and you give an answer of "10 at 100", it means you are 90% sure the true value is greater than 10, but less than 100. The precise definition of an X% confidence interval is one which, if you gave multiple independent intervals, would contain the true value X% of the time in the long-run. If your X% intervals are actually correct X% of the time, this is known as being *well-calibrated*. If your intervals are correct more or less often than the target probability, this means you are either over or under confident. There are a few good ways to train calibration, such as [this web app](#) (although unhelpfully, most of the trivia questions are US-centric). However, to give some anecdotal advice about calibration, I've noticed two things:

- Significantly more people are overconfident than underconfident. For instance, studies have shown that when you ask people for 90% confidence intervals, the average *hit rate* (i.e. percentage of confidence intervals that contain the true value) is more like 50%.
- Both overconfidence and underconfidence are bad, but at the extreme, overconfidence is probably worse. Better to miss a good opportunity than lose all your money!¹⁰

For this reason, my advice to the average person would be to adjust in the direction of slightly less confidence in your own answer, i.e. giving wider intervals than you might instinctively give (particularly if you're asked for high-confidence intervals like 90%). This probably isn't good advice for everyone, but my guess is that more people will find it helpful than harmful!

If an interviewer asks you Fermi problems, some skills they might be looking for are:

- Good calibration
- Good intuitions for uncertainties, and how to combine them
- Ability to be systematic, and break down a problem into parts

⁹The technical term for a factor of 10. For instance, if you think a quantity is between 100 and 1000, then your uncertainty is one order of magnitude.

¹⁰For instance, the final round interview at Jane Street takes the structure of an extended betting game where your P&L carries between interviews, and you are given the number one imperative of never going bankrupt.

3.3 Market Making

Market making is a really important topic to be comfortable with in interviews. For a basic overview of how it works, here's the [Jane Street pdf](#) (again!). This is probably the interview skill that is most closely linked to what a career in trading will actually be like, so it's super important to have a good handle on it. That said, not all firms will test it directly. Jane Street do a lot of it in their final round (generally they will ask you to market-make around basic uncertain quantities like a dice roll, or in more difficult cases, answers from section 3.2.1, or the right to play a game like in section 3.2.2), Optiver often ask you to make markets and trade around the answers to Fermi problems, and as previously mentioned, Optiver and Maven both make group trading games part of their interview process.

Because this can take so many forms, it's hard to list specific points that interviewers test for, but here are a few general things that it can be worth considering. Keep in mind they vary between firms (e.g. Jane Street style market making questions care a lot about concepts like hedging and risk profiles, whereas Optiver care more about things like trading dynamics and adjustment of markets).

- **Market adjustment**

Say you estimate the fair value of an uncertain quantity as 50 at 60, and the interviewer chooses to buy at 60¹¹. You should then choose to adjust your expected value upwards, because the interviewer's decision is evidence¹² that the true value is greater. If you adjust your expected value, you should adjust your market accordingly.

It's worth noting that the interviewer's decision may not actually constitute evidence to adjust your market, e.g. if they've chosen to buy/sell randomly - it can be a good idea to explicitly ask about the decision theory that they're using. Even if they aren't choosing randomly, the amount you should adjust varies based on how informed the interviewer is - for more on this, see the point on asymmetric information.

- **Hedging**

Say you have bought a contract on the value of a 6-sided dice (i.e. you will roll a dice, and win an amount of money equal to the number you get). Now the interviewer asks if you want to trade a contract with a value of 1 if the dice is at least 4, and 0 if the dice is 3 or less. The expected value of this contract is 0.5 (because it is worth 0 or 1 with equal probability), but you should prefer to sell it at this price than buy it¹³ because this will partially offset your risks from the first contract. If you are asked to make a market, you might centre it around a value below 0.5. This is the essence of using one market to try to hedge your position.

- **Keeping track of positions**

Some firms will let you record things on paper, some will ask you to remember the trades in your head. Either way, they are testing for your ability to keep track of several things at once and split your focus efficiently. The most important thing is usually that you keep track of your current position, i.e. you don't confuse buying and selling.

- **Understanding of asymmetric information**

Question - suppose an interviewer asks you if you want to make a market on the roll of a 6-sided dice, 2 wide. He knows the value of the dice, and can choose whether to trade with you or not. Should you make a market? Answer - absolutely not! This is a classic case of asymmetric information; the interviewer will only trade with you in the cases where you make a loss. Other cases can be more nuanced, e.g. the interviewer has partial knowledge that you don't have, but you still have an advantage because your role as a market maker allows you to charge a bid-ask spread.

¹¹If you trade on a market given by somebody else, you have to buy on the ask (i.e. the higher price), or sell on the bid (the lower price). This is known as **crossing the spread**. Make sure not to mix this up; it's a classic failure mode for people in interviews, even people with lots of experience!

¹²More specifically, it is *Bayesian evidence* - this is another situation when familiarity with Bayesian statistics provides a really useful framing to certain problems.

¹³Assuming you are risk-averse. This is usually the case, but again it can be worth asking the interviewer if they intend you to have risk-averse preferences - if nothing else, the interviewer will probably be impressed that you're making these kinds of considerations.

If this all seems very intimidating, don't worry! Most people are completely new to market making when they apply for quant trading roles, and firms won't expect you to have a super deep understanding of all these topics. I have included them here just for an idea of the kinds of concepts you might have to think about during later interview stages. If you have read all these points and feel like you understand the concepts, then that already puts you ahead of a good many candidates!

3.4 Adding extra dimensionality

Practicing a lot of problems is definitely effective for interview preparation, but the benefits are limited. In particular, just going through all the problems in this document / any other resource is a pretty bad practice, because often you can get too caught up in the problem, and forget about its higher purpose. As I've already mentioned, firms don't want to see you brush aside problems easily. They are meant to be proxies for the skills a trader requires, e.g. those that I've listed in previous sections, and the more problems you practice while neglecting those direct skills, the less effective problems become as an assessment tool. In this subsection, I'm going to focus on ways you can add dimensionality to problems, i.e. find different and interesting ways of tackling them that provide much better training of the skills that interviewers are actually looking for.

- **Time conditions**

Try solving the problems under extreme time pressure. Maybe you set yourself 3 minutes to solve a problem, or even set yourself one minute to try and get as accurate an approximation as possible. As well as being good practice of working under pressure, this also tests your calibration and judgement, because to set yourself an appropriate time limit, you need to be able to guess how long a problem will take you.¹⁴

- **Confidence intervals**

Obviously you can give confidence intervals for Fermi problems, but they can also be used in the other 2 cases. For logic puzzles, ask yourself "what percentage chance do I put on my answer being correct?". For maths/strategy problems, try to give confidence intervals of 50%, 75% and 90% for the true answer, as well as a point estimate. Track these intervals over time, and see if they are correct the right proportion of time or if you are systematically over/underconfident.

- **Communication**

Being concise communicating well is extremely important in interviews. A few ways to work on this this even if you have nobody else to practice with are:

- Practice narrating your thoughts about a problem out loud
- Practice writing out a condensed solution afterwards, summarising the key ideas that were involved

- **Practice interviews**

Having a friend to interview with is great, because they can simulate a real interview, explicitly testing skills like how well you respond to hints. It's also quite hard to simulate the stressful, unpredictable environment in an interview by asking yourself questions and setting yourself time limits. Make sure your friend is willing to give you brutal feedback though!

As well as mock interviews, try practicing with real interviews! Even if there's one firm in particular you want to get an offer from most, it's a good idea to apply for lots of different places, because the application process is normally quite low-effort.¹⁵

¹⁴The *planning fallacy* is when you underestimate how long a task will take you, often despite historical evidence implying it will take longer. Try not to fall into the classic trap of setting very short time periods, then consistently failing to get an answer in that time. If you do this, you aren't getting much benefit from timing yourself.

¹⁵Although this largely depends on how draining you find interviews.

3.5 Fit questions

These are questions about your personality, experiences, interests and passions. In other finance roles (like investment banking) you will be asked a huge number of these. I remember when I was preparing for Spring Weeks I had to prepare answers to all the generic ones (like “*what is your biggest weakness*”, “*when did you overcome a significant challenge*”, “*describe a time you have worked as a team*”). It’s not a bad idea to look at some of these questions and think about how you’d answer them, because they do sometimes come up, but you don’t need to overdo preparation in this area! The vast majority of fit questions in quant trading interviews simplify to a rephrasing of one of three:

- *Tell me about yourself / walk me through your CV*
- *Why are you interested in a career in finance / trading?*
- *Why are you interested in this firm specifically?*

There is obviously no correct answer to any of these questions, and it’s important to answer them honestly rather than just inventing stuff you think the interviewer wants to hear. Lots of people just want to go into quant trading for the money, so it’s important you convince the interviewer that you are interested in this career for other reasons.¹⁶

However, I’ll just say a little about the 3rd question. I think it’s almost a bit unfair that trading firms ask this, because they can often be very opaque by design and it’s difficult to learn enough about them to come to an informed decision about which one you are best suited for. There are some decent ways to go about answering this question, I’ve listed some below (in roughly decreasing order of preference).

- **Go to company events**

This is the number one best thing you can do, because you get an opportunity to learn more about the culture of the firm and how they approach trading. These events can range from extremely informal social events (e.g. Jane Street game nights) to programming/trading competitions you need to apply for (e.g. IMC or Citadel both run coding challenges). Participating in competitions, or applying for Spring Weeks / Insight Days, can actually accomplish 3 things at once:

- learn more about the company
- get on their radar (and potentially fast-tracked to interviews)
- get a better idea of whether trading interests you

- **Hear from other people**

If you know people in these companies who really enjoy their work, or people who have interned there and really enjoyed it, mention that. If not, you can try reaching out to people. For instance, if there was someone at your university careers fair but you missed an opportunity to speak with them, you can try connecting with them on LinkedIn. Generally very few people in quant finance cold message asking for advice (e.g. compared to other areas of finance like consulting or investment banking). If you’re nice and polite, your odds of getting a response are pretty good.

- **Go on their website**

Most firms will have some information about their culture and work environment on their website. The more personalised and less generic you can make your answer, the better. For instance, some firms will have videos where traders speak about what they like about the company and their experiences, referencing this directly is a lot better than just quoting the firm’s main values, because it shows you’ve actually done your research. Try to make sure you express genuine interest and ask good questions - their main concern is likely to be spam applicants who know nothing about quant trading.

Another approach is to express an interest in the products that they trade, but this has a few risks. For one, it’s dangerous to do this if you don’t actually know much about the products. For another, there’s a possibility that you’ve mis-characterised the firm in some way (e.g. thinking they specialise in one particular product when they actually do a wide range), and this can leave a bad impression.

¹⁶If you actually are just going into quant trading for the money, then you’re shit outta luck.

3.6 Finance Knowledge

As mentioned in section 1, most firms won't require you to have knowledge in lots of specific areas. Traders are recruited from all sorts of different backgrounds, and requiring finance knowledge in the interviews would unfairly discriminate against many candidates. Most firms seem to adopt the philosophy of *"It's much easier to teach smart people finance than to teach finance people to be smart!"*. However, some firms might specifically ask that you read up on certain topics in interviews, and even for the ones that don't, a basic understanding of some important areas in finance can be helpful. Reading up on the following can also give you a good idea of whether a career in quant finance might interest you (although bear in mind, depending on a trader's exact job description, some of these may be less relevant than others).

- **Basic Finance & Markets Concepts**

A good starting point would be knowing what the following words mean: stock, bond, interest rate, index, ETF, future, option, hedging, speculation, exchange, over-the-counter, orderbook, alpha, beta, risk, volatility

- **Options Theory**

A good starting point would be knowing what the following words mean: call, put, parity, intrinsic value, time value, Greeks, Delta, Theta, Vega, implied volatility

[Investopedia](#) is a good start for a lot of these (as well as just searching for them on Google). The newsletters mentioned earlier in [Section 1.4.5](#) may also help, because you will be able to see some of this language used in context.

4 Problems & Solutions

I've split problems in this section into 4 different subsections, in correspondence to how I've characterised them in Section 3. All problems will have solutions available, as clickable links next to the questions, and most will also have one or more hints available.

Within each section, I've also classed the questions in terms of difficulty (either easy, medium or hard). Very broadly, I would say that easy questions are ones you might be expected to solve in 5 minutes with little to no assistance, whereas hard questions might take up an entire 30-40 minute interview, and involve lots of discussion and hints from the interviewer.

It's worth mentioning here that I've mainly chosen problems that I've been asked, or else that I've heard about from other people. Most problems here aren't widely known or generally available online / in books (although there are a few exceptions). For this reason, I would recommend that you supplement these problems with other resources. A few good ones include:

- **Fifty Challenging Problems In Probability (book)**

The questions in this book are very wide-ranging, some very easy and some incredibly hard. The distribution of difficulty doesn't seem to be very logical, and this can make going through the book an odd experience, but nevertheless there is still a lot of value in some of the problems. Any questions related to geometry are not as relevant to quant trading interviews, and can be skipped.

- **A Practical Guide to Quantitative Finance Interviews (book)**

The problems in this book are usefully categorised. For quant trading interviews, I would most strongly recommend Chapter 2 (Brain Teasers) and Chapter 4 (Probability Theory). There is also a good introduction on general interview skills.

- **A Collection Of Dice Problems**

This is a great resource for dice problems on the harder difficulty end. Chapters 2.1 and 2.2 are both good, 2.3 can be skipped, and 2.4 is *extremely* good. In particular, I would recommend this document if you are preparing to interview at Jane Street. These questions are also perfect candidates for some of the "extra dimensionality" that I talked about in [Section 3.4](#), i.e. trying to get approximate answers in time conditions, or giving confidence intervals. This document also contains solutions, although some of them are slightly ridiculous and impractical, involving Markov chains that take up an entire page! In these cases, it is better practice to try and come up with an approximate solution, than thinking about how you would compute an exact one. Working out which problems can actually be solved with pen and paper is good practice in itself.

- **Glassdoor or WallStreetOasis**

Both good websites where you can find a large supply of relevant interview problems, as well as peoples' reviews of companies and their interview processes.

- **LeetCode**

This is a great source of coding practice. It's less relevant for interviewing with the firms in Section 2 (with the exception of Citadel Securities), but improving coding skills is always a good idea!

- **Fermi Problems**

This pdf has a really long list of Fermi problems, which are great practice to work through. To add variety to your practice, you can try giving both point estimates and confidence intervals.

4.1 Logic Problems

EASY

1. Light Bulbs

There are 100 light bulbs, and 100 people, each numbered from 1 to 100 respectively. The light bulbs all start as unlit. The people each walk past the light bulbs, and toggle the bulbs with numbers divisible by their number. For instance: the first person will toggle every bulb, the second person will toggle bulbs 2, 4, 6, etc. Which bulbs will be on at the end of this process?

[Hint](#)
[Solution](#)

2. One Heavier Ball

You are given 9 balls, 8 of the same weight and 1 which is heavier. You have a set of balance scales, which you can use to compare the weight of any two groups of balls. What is the maximum number of uses required to determine both which ball is heavier?

[Hint](#)
[Solution](#)

3. 100 Prisoners and a Light Bulb

There are 100 prisoners in a cell. Each day, one of them is chosen at random to go into an interrogation room containing a single light bulb. Initially the light is switched off, and can be toggled by each prisoner. At any point, a prisoner in the interrogation room may make the claim that every single other prisoner has been in the room at some point. If the claim is wrong, all the prisoners are executed. If the claim is correct, they all go free. The prisoners can discuss strategy beforehand. How can they guarantee that they'll go free?

[Hint](#)
[Solution](#)

4. The Lion And The Sheep

There are N lions in a field, and 1 sheep. If a lion eats a sheep, they then become the sheep. A lion will prefer eating the sheep to going hungry, but it will prefer going hungry to being eaten. What will happen?

[Hint](#)
[Solution](#)

5. Fish in a Bucket

There are some fish in a bucket. 3 fishermen come up to the bucket in turn. The first one tries to divide them into groups of 3, but finds that there is one left over, so he takes a third plus the left-over fish. The second fisherman tries to divide the remaining fish into 3, finds there is one left over, so takes it plus a third. The third fisherman does exactly the same, taking a third plus the one left over. What is the minimum number of fish that could have been in the bucket at the start?

[Hint](#)
[Solution](#)

MEDIUM

6. Ants on a Line

There are 50 ants on a 1m line. The 25 left-most ants are moving right, and the 25 right-most ants are moving left. When 2 ants collide, they will both reverse direction. How many collisions will there have been in total once all ants have fallen off the end of the line?

[Hint](#)
[Solution](#)

7. Hungry Truck Driver

A truck driver has been tasked with transporting 3000 apples between towns A and B, which are 1000 miles apart. However, if the truck driver is driving while there are apples in his truck, he will eat them at a rate of one per mile. The truck has a maximum capacity of 1000 apples at any one time. The driver may deposit apples at any point between A and B. How can he get the maximum possible number of apples from A to B?

[Hint](#)
[Solution](#)

8. Single bullet

You are in a field with 100 criminals. You have a gun with a single bullet in it, so you can't stop all the criminals escaping at once. Each criminal is assumed to be purely selfish and rational; they prefer escaping to being in captivity, but they prefer being in captivity to dying. Is there something you can say that will mean no criminal tries to escape?

[Hint](#)
[Solution](#)

9. Prisoner Hats

There are N prisoners standing in a line, facing forward, each wearing either a black hat or a white hat. Each can only see the prisoners in front of him. The prisoners are allowed to call out (in any order) a guess as to the colour of their own hat. When every prisoner has guessed, the ones that were correct get to go free. Other than calling out colours, no communication is allowed, although they may discuss strategy beforehand. What is the method that maximises the expected number of prisoners surviving? What if there were 3 hat colours: black, white and grey?

[Hint](#)
[Solution](#)

10. Coin Flipping

Adam and Ben each have a fair coin. Adam throws it 101 times, and Ben throws it 100 times. What is the probability that Adam gets strictly more heads than Ben?

[Hint](#)
[Solution](#)

11. Plane Boarding Problem

A plane has 100 seats, and each of 100 people have been assigned a unique seat. The first person to board is a drunkard, who sits in a randomly chosen seat. Every subsequent person sits in their own seat, unless someone is already sitting there, in which case they also sit in a random seat. You are last to board. What is the probability you get to sit in the seat you were assigned?

[Hint](#)
[Solution](#)

12. **Poison bottles #1**

You have 2 mice, and 9 bottles, exactly one of which is poisoned. Each mouse can taste any combination of bottles at once, and if it is poisoned it will die after exactly one minute. Once a mouse is dead, they can't taste any more bottles. How many minutes do you need to determine exactly which bottle is poisoned?

Note - this question has a harder follow-up [here](#).

[Hint](#)
[Solution](#)

13. **The Two-Egg Problem**

You have 2 identical eggs, and can drop each of them from any floor of a 100-storey building. Once you break an egg, you can't drop it again. What is the minimum number of drops required to find the lowest floor on which an egg will break (or, if it won't break on any floor, find this out)?

[Hint](#)
[Solution](#)

HARD

14. **Couples Handshake Problem**

You are with your spouse at a party, along with 6 other couples. Some people shake hands at this party, although nobody shakes hands with themselves or their spouse. You observe that everyone (excluding yourself, but including your spouse) shook hands with a different number of people. How many people did your spouse shake hands with?

[Hint](#)
[Solution](#)

15. **One Odd Ball**

There are 12 balls, 11 of the same weight and 1 of a different weight (it could be lighter or heavier). You have a set of balance scales, which you can use to compare the weight of any two groups of balls. What is the maximum number of uses required to determine which ball is different, *and* whether it is lighter or heavier?

[Hint](#)
[Solution](#)

16. **Poison Bottles #2**

This is a continuation of [Poison Bottles #1](#) from the medium difficulty section. Suppose you had M mice and T minutes. What is the largest number of bottles for which you could be guaranteed to find the poisoned one?

This is a very hard problem, so there are multiple hints, which should be used chronologically.

[Hint #1](#)
[Hint #2](#)
[Solution](#)

Hints

1. Try seeing what happens to the first 10 light bulbs. Do you notice any pattern?

[Question](#)
[Solution](#)

2. Label the balls 1, 2, ..., 9. Suppose you weigh {1, 2, 3} against {4, 5, 6}. What are the possible results of this weighing, and what can you deduce from these results?

[Question](#)
[Solution](#)

3. Try designating one person to be the “guesser”. He only guesses when he is sure everyone else has gone into the room. How can he work out when this has happened?

[Question](#)
[Solution](#)

4. Think about what will happen with 1 lion, then 2, then 3. Can you find a pattern?

[Question](#)
[Solution](#)

5. Let the number of fish at the start be N . Given that the first fisherman was able to divide the fish up in the way he was, what can you say about N ? Can you express this mathematically?

[Question](#)
[Solution](#)

6. Consider what actually happens when 2 ants collide. Can you model the collisions with a simpler situation, which is functionally the same?

[Question](#)
[Solution](#)

7. Try reframing the problem as multiple drivers making a single journey. There are 3000 apples, so at least 3 drivers must set out at the start, and they are consuming 3 apples per mile. How soon can you “kick a driver out”?

[Question](#)
[Solution](#)

8. First try a simple case: 2 criminals (call them A and B). The only problematic case is when A and B both try to escape at once. What would happen if A knew that if they both tried to escape, you’d always shoot him, and not B?

[Question](#)
[Solution](#)

9. For the 2-colour case, the first prisoner to say a colour will always have a 50% probability of survival. Is there any way this prisoner could say a colour that transmits encoded information to the other prisoners, allowing them to all guess correctly?

[Question](#)
[Solution](#)

10. Say the number of heads and tails that Adam gets are H_A and T_A respectively, and Ben gets H_B heads and T_B tails. If Adam gets more heads than Ben, what can you say about $H_A + T_B$?

[Question](#)
[Solution](#)

11. Before you board, someone else will sit in either your seat or the drunkard's seat. How likely are these two possible outcomes, and what do they each imply for where you will sit?

[Question](#)
[Solution](#)

12. If both mice drink from the same bottle, then if they both die you know which bottle is poisoned. What is the problem with having both mice drinking from two of the same bottles?

[Question](#)
[Solution](#)

13. One simple method would be to drop the first egg at floor 10, 20, ..., 100 (or until it breaks). This would tell you the floor is in a range of ten, and you could then test each of these ten floors with the second egg, starting from the lowest one. The problem with this method is that it's wasteful. If the answer is between 1 and 10, it will take at most 10 drops to find, but if the answer is between 91 and 100, it can take as many as 19. Can you find a better system for dropping the first egg, that "smooths out" this asymmetry?

[Question](#)
[Solution](#)

14. There are 13 people other than yourself. Each person can shake hands with at most 12 other people, so the other 13 must have shaken hands with 0, 1, ..., 12 people respectively. What can you say about the people who shook hands with 0 & 12 others?

[Question](#)
[Solution](#)

15. Consider the earlier related question on identifying a ball of a different weight. Try using a similar argument to lower-bound the number of measurements it can take. Can you then find a way to attain this bound?

[Question](#)
[Solution](#)

16. **HINT #1**

Think about the easy problem (Poison Bottles #1), as well as very basic cases (e.g. if you only had one mouse, or one bottle). Can you spot a pattern in the results you get?

HINT #2

When do each of the M mice die, and what information does this give you? Can you represent this information in a number of base $(T+1)$?

[Question](#)
[Solution](#)

Solutions

1. If a number is not a perfect square, then you can pair up all of its factors. For instance, the factors of 20 are $\{1, 20\}$, $\{2, 10\}$, $\{4, 5\}$. This means it has an even number of factors, so the corresponding bulb will be toggled an even number of times, i.e. it will be switched off at the end. However, perfect squares have an odd number of factors, because when you pair up all their factors you have the square root left over. So the bulbs left switched on at the end will be precisely those labelled with square numbers.

[Question](#)

2. Label the balls $\{1, 2, \dots, 9\}$. Suppose you weigh $\{1, 2, 3\}$ vs $\{4, 5, 6\}$. If one of the groups is heavier, the odd ball is in that group. If both groups are the same, the odd ball is 7, 8 or 9. In either case, we know the odd ball is one of three. We can then weigh 2 balls in this group of 3 against each other, and by a similar argument, this tells us which one is heaviest.

[Question](#)

3. Some prisoner is selected to be the “guesser”; they will always turn off the light whenever they find it on. Each other prisoner is to turn the light on the very first time they find it off, otherwise they are to leave it in the state they found it. When the “guesser” turns the light off for the 100th time, they can safely assume that all prisoners have been in the room.

[Question](#)

4. If there is one lion, they will eat the sheep. If there are 2 lions, no lion will eat the sheep, because then that lion would itself be eaten. If there are 3 lions, then the sheep will be eaten, because the lion that eats the sheep will then not be eaten by the other 2 lions. By induction, we can see that the sheep will be eaten if and only if N is odd.

[Question](#)

5. Let the number of fish at the start be N . Given they can be divided into 3 parts with 1 left over, we know $N = 3X + 1$, for some integer X . The number of fish left after the first fisherman takes his share is $2X$. We know $2X$ has remainder 1 modulo 3, meaning X must have remainder 2 modulo 3, so $X = 3Y + 2$. This means $2X = 6Y + 4$, so the number of fish left over after the second fisherman takes his share is $4Y + 2$. Finally, we know this must also have remainder 1 modulo 3. The smallest value of Y such that this is true is $Y = 10$. We can work backwards, and check that this gives a valid answer $N = 25$.

[Question](#)

6. When two ants collide, we can model them as just passing straight through each other. Since the ants are indistinguishable, this is functionally the same as the previous situation. We can now see that each ant will pass through exactly 25 ants (the ones moving in opposite directions to themselves), so the total number of collisions experienced by all ants is $25 \cdot 50 = 1250$. We need to halve this, because this method will count each collision twice. So total number of collisions is 625.

[Question](#)

7. We can reframe the problem as multiple drivers making a single journey. There are 3000 apples, so at least 3 drivers must set out from A, and they are eating 3 apples per mile. After $1000/3 = 333$ miles,¹⁷ there are 2000 apples left, so we only need 2 drivers, i.e. we can “kick one out”. After another $1000/2 = 500$ miles, there are 1000 apples left, so we kick out another driver. Finally, we drive the remaining 167 miles. The number of apples left when we arrive at B is $1000 - 167 = 833$.

¹⁷For the sake of convenience, I’m rounding - obviously these aren’t exact numbers.

To check that this reframing of the problem works, we can examine what happens in the original single-driver setting. The driver sets out from A with 1000 apples, stops when he drives 333 miles, and deposit 667 apples. He then drives back to A, and repeats this process twice more. This means we have gone from having 3000 apples at A to having 2000 apples at a point 333 miles closer to B. He can now repeat this process by driving from A to the point where the apples were deposited, and in 2 trips get 500 miles closer to B (now with 1000 apples left). Finally, he can take these 1000 apples and drive directly to B.

Question

8. The trick is to make sure in every group there is one scapegoat, i.e. one person who knows he'll be the one to get shot. Suppose a group tried to escape, then the scapegoat wouldn't want to escape, so he would leave the group. The remaining group would then have a new scapegoat, who would also leave. Eventually everyone would leave the group, i.e. nobody would try to escape.

The easiest way of doing this is to assign a number between 1 and 100 to each prisoner, and telling them that if any group tries to escape, you will shoot the escapee with the highest number.

Question

9. For the 2-colour case, the last prisoner (who can see everyone else) should count the number of black hats he sees. If it is an odd number, he calls out black. If even, he calls out white. The second-to-last prisoner can now count the number of hats he sees. If the parity agrees with what the last prisoner called out, then he knows his hat is white. If not, his hat is black. Once he has called out the correct colour, the third-from-last prisoner can similarly deduce the colour of his hat, using the knowledge given to him by the last and second-to-last prisoners. This continues all the way to the front of the line. Thus, all prisoners can go free with certainty, except for the first one who has odds of 50%. We clearly can't do better than this, because the first prisoner to guess will always have odds of 50%.

For the 3-colour case, a similar idea works - we have to use modular arithmetic. First, we assign each colour to a unique value modulo 3 (i.e. 0, 1 or 2). There are several strategy variants that can work here, but maybe the simplest one is for the last prisoner to count the hats he sees, calculate the value of (number of black hats - number of white hats) modulo 3, then call out the colour corresponding to this number. Let this number be X. The second-to-last prisoner then makes the same calculation, and gets the number Y. We have 3 possibilities:

- $Y \equiv X - 1 \pmod{3} \implies$ 2nd-to-last prisoner's hat is black
- $Y \equiv X + 1 \pmod{3} \implies$ 2nd-to-last prisoner's hat is white
- $Y \equiv X \pmod{3} \implies$ 2nd-to-last prisoner's hat is grey

So he can deduce his hat colour, and guess correctly. The remaining prisoners can proceed in a similar way.

Question

10. Say the number of heads and tails that Adam gets are H_A and T_A respectively, and Ben gets H_B heads and T_B tails. We know $H_B + T_B = 100$, so Adam getting more heads than Ben is equivalent to the statement $H_A + T_B > 100$. However, since Ben's coin is fair, this is equal to the probability of $H_A + H_B > 100$. In other words, the probability of Adam getting more heads is equal to the probability that more than half of the total 201 coin flips are heads. Since it is impossible for exactly half to be heads, this means the answer is 0.5.

Question

11. Of the first 99 people to board before you, one of them must sit in either your seat or the drunkard's seat. Call this person X (note, X could be the drunkard). The drunkard chooses randomly, and whenever anyone else sits in a seat other than their own they choose a seat randomly, so X must have chosen a seat randomly. We know they either sat in your seat or the drunkard's seat, so these two must have equal probability.

If they picked the drunkard's seat, then everyone after them can sit in their own seat, including you. If they picked your seat, then everyone after them can sit in their own seat, except for you (since you have to sit in the drunkard's seat).

From this, we deduce there is a 50% probability you will be able to sit in your own seat. This is true for any number of passengers.

[Question](#)

12. Label the mice A and B, and the bottles 1-9. Have the first mouse test {1, 2, 3} and the second mouse test {1, 4, 5}. We have 3 cases:

- Both die, in which case we know bottle 1 is poisoned.
- One dies, in which case we have once mice left, and 2 candidates for the poisoned bottle. Have the living mouse test one of the 2 bottles, then we can tell which is poisoned .
- None die, in which case the candidates for poisoned bottle are {6, 7, 8, 9}. Have mouse A test {6, 7}, mouse B test {6, 8}. There are 4 possible outcomes for mouse A and B, and each outcome corresponds to a different one of these four bottles being poisoned.

So using this method, we can determine the poisoned bottle in 2 minutes, which is the best we can do. For a deeper dive into this problem, see [Question 16](#).

[Question](#)

13. Drop the first egg at floors 14, $14 + 13 = 27$, $14 + 13 + 12 = 39$, ..., $14 + 13 + \dots + 4 = 99$, then 100. If the first egg breaks on the Nth drop, then we have narrowed down the answer to a range $15 - N$ wide (unless it breaks at 100, or it doesn't break at all, in which case we know the answer is 99 or 100 respectively, so we are done). For example, if it breaks on the 2nd drop, we know the answer is in the range {15, 16, ..., 27}, i.e. 13 possible answers.

Once the first egg breaks, we can then test all floors in our range in increasing order. If we test all but one and it doesn't break, we know the answer is the highest floor in this range. Therefore, we can get the answer in N drops of the first egg, and $14 - N$ drops of the second, so 14 drops in total.

How do we know it's impossible to do better? Consider making 13 total egg drops. We label a drop 0 if the egg doesn't break, and 1 if it does. This means the information we get from 13 drops can be represented as a 13-digit binary number with at most two 1s. There are $\binom{13}{2} = 78$ binary numbers with two 1s, 13 with a single 1, and one with no 1s. This gives us 92 binary numbers in total, which is not enough to distinguish between 100 different floors.

[Question](#)

14. There are 13 people other than yourself. Each person can shake hands with at most 12 other people, so the other 13 must have shaken hands with 0, 1, ..., 12 people respectively. Let us label these people (0), (1), ..., (12).

(12) must have shaken hands with everyone except their spouse, and we know (12) didn't shake hands with (0), so (0) must be the spouse of (12). Now, (11) must have shaken hands with everyone except (0) and their spouse, and we know (11) didn't shake hands with (0) or (1) (because (1) only shook hands

with (12)), so we deduce (1) is the spouse of (11). We can continue in this way, pairing off (0)+(12), (1)+(11), ..., (5)+(7). The only person left unpaired is (6), so your spouse must be (6).

Note - from this information, we can actually work out precisely what handshakes took place at the party, including the fact that you also shook hands with 6 people. Can you see how?

Question

15. The answer is 3 measurements. The precise method takes a long time to write out, so I've decided to be lazy and link to it [here](#). If you got it, then congrats!

Question

16. The answer is $(T + 1)^M$. First, we will show this is an upper bound, then that it is attainable.

We can give each mouse a number between 0 and T, by the following method: if the mouse survives then we give it the number 0; if not then we give it the number of the minute in which it died. Putting all these numbers together for each mouse, we get a number in base-(T+1) with M digits. The number of M-digit base-(T+1) numbers is $(T + 1)^M$, so this is the maximum number of bottles we can distinguish between.

Now we want to show this is attainable. Label the mice as (1), (2), ..., (M). We want a strategy in which the data we get from the mice (as described in the previous paragraph) spells out the number of the bottle that is poisoned. In other words, we want mouse (K) to give us a number J precisely when the poison bottle has Kth digit J when written in base-(T+1). We can do this as follows:

“in the Jth time period, have mouse (K) drink from all the bottles with Kth digit equal to J”

If the Kth digit of the poisoned bottle is $J \geq 1$, the mouse will die in the Jth time period. If the mouse doesn't die, that means the Kth digit must be 0 (because the mouse has drunk from all bottles where the Kth digit isn't 0). So, we are done.

Question

4.2 Maths Problems

EASY

1. All Six Numbers

What is the expected number of rolls of a fair d6 (6-sided dice) until all 6 faces have appeared at least once?

[Hint](#)
[Solution](#)

2. Double Sixes

What is the expected number of rolls of a d6, until you see two 6s in a row?

[Hint](#)
[Solution](#)

3. Dice Stopping Game

You start with £0, and continually roll a d6. If you roll a 1, 2 or 3 then you add £1 to your pot. If you roll a 4 or 5 then you stop, and win your pot. If you roll a 6 then you stop, and win nothing. What are your expected winnings?

[Hint](#)
[Solution](#)

4. Bayesian Coins

There are 1000 coins: 999 are fair, and one has heads on both sides. You choose a coin at random and flip it 10 times; it comes up heads every time. What is the probability you chose the coin with two heads?

[Hint](#)
[Solution](#)

5. Coins Sequences

You are playing a game with the interviewer. Each of you is assigned a sequence of 3 coin flips (e.g. HTH), and you flip a fair coin until someone's sequence appears. If your sequence appears first then you win; if the interviewer's appears first then he wins. Initially, your sequence is HTT and the interviewer's sequence is HHT - what is the probability that you win? Now suppose the interviewer chooses the sequence HHH. What sequence should you choose to maximise the probability that you win?

[Hint](#)
[Solution](#)

6. 100 Flips

You flip a fair coin 100 times. What (approximately) is the probability that you get more than 55 heads?

[Hint](#)
[Solution](#)

7. Drawing to an Ace

You shuffle a pack of 52 cards, then flip over one card at a time. What is the expected number of cards you'll have to flip until you get an ace?

[Hint](#)
[Solution](#)

MEDIUM

8. Bidding For A Car

Suppose the true value of V a car is uniformly distributed between 0 and 1000. You can bid any amount for the car, and if you bid the true value or more then you pay your bid and get the car. You know a very good salesman, and are confident in your ability to sell the car for 50% more than its true value. What should you bid to maximise your expected profit? Now suppose $V \sim \text{Unif}[100, 1000]$, what should you bid? What if $V \sim \text{Unif}[100, 10000]$?

[Hint](#)
[Solution](#)

9. Substrings

What is the minimum positive integer N s.t. 2^N contains '2018' as a substring? Try to give your answer in the form of a confidence interval.

[Hint](#)
[Solution](#)

10. LCM of d10

You continually roll a fair d10. What is the expected number of rolls until the lowest common multiple of all numbers that have appeared is greater than 2000?

[Hint](#)
[Solution](#)

11. Maximum Of Polyhedral Dice

You roll each of the polyhedral dice $\{d4, d6, d8, d12, d20\}$ once. What is the expected value of the largest dice?

[Hint](#)
[Solution](#)

HARD

12. Consecutive Differences

Call a "consecutive difference" the absolute value of the difference between two consecutive rolls of a d6. For example, the sequence of rolls 143511 has the corresponding sequence of consecutive differences 31240. What is the expected number of rolls until all 6 consecutive differences have appeared?

[Hint](#)
[Solution](#)

13. **Increasing Dice Values**

You roll a 6-sided dice, and sum your rolls as long as you keep rolling strictly larger values. For instance, you might roll 1-3-4, then roll 2 and stop, meaning the number of rolls is 3, and total sum is 8. What is the expected number of rolls, and what is the expected sum? Now what about the limit for the case when you roll an N-sided dice, as N tends to infinity?

[Hint](#)
[Solution](#)

14. **Snakes and Ladders**

You are playing a game of snakes and ladders. You start at square 1, and each turn you roll a d6 and move the corresponding number of squares. When you reach at least square 25, you stop. There is also a snake which connects squares 5 and 20 (i.e. if you land on 20, you immediately move down to 5), and there is a ladder from 10 to 15 (so if you land on 10, you immediately move up to 15). What is the expected number of rolls until you finish the game?

[Hint #1](#)
[Hint #2](#)
[Solution](#)

15. **Sequences Problem #1**

A sequence is defined in the following way: $x_0 = 0$, and:

$$x_{n+1} = \begin{cases} x_n + A_n, & \text{if } A_n \mid x_n \\ x_n - A_n, & \text{otherwise} \end{cases}$$

where A_n is an integer chosen at random with uniform probability from $\{1, 2, \dots, 10\}$. What is the expected value of x_{100} ?

[Hint](#)
[Solution](#)

16. **Sequences Problem #2**

I start with the number 100, and continually flip a fair coin. If I get heads, I replace my number X with X+1. If I get tails, I replace it with 1/X. What is the expected value after I flip a coin 8 times?

Note - this is probably the hardest question in this section. Getting a perfect answer is pretty intractable, although there are methods that provide very good approximations. There are several hints available, which should be read chronologically along with the solution.

[Hint #1](#)
[Hint #2](#)
[Hint #3](#)
[Solution](#)

Hints

1. Your first roll will always be a “new” number. What is the expected number of rolls until you see a second new number? How about a third?

[Question](#)
[Solution](#)

2. Try splitting into 2 cases: where your last roll was a 6, and where your last roll wasn't a 6. See if you can find two equations, and solve them to get expected value in both cases.

[Question](#)
[Solution](#)

3. First, consider a simplified version of the problem where you also win the pot if you roll a 6. What is the value of this game? Now how does the possibility of losing the pot on 6 make a difference?

[Question](#)
[Solution](#)

4. There are 2 possibilities:

- (a) You picked the biased coin and it came up heads 10 times
- (b) You picked one of the fair coins and it came up heads 10 times

What is the probability of each? If you know one of them happened, what is the probability that it was (a) and not (b)?

[Question](#)
[Solution](#)

5. Both sequences start with a head. Let p be the probability of winning from this situation, and consider the next two flips - can you formulate an equation for p ?

[Question](#)
[Solution](#)

6. Try using the [central limit theorem](#) to approximate the number of heads as a normal distribution.

[Question](#)
[Solution](#)

7. Suppose the 4 Aces divide the pack into 5 groups, with A, B, C, D, E cards respectively (each group is separated by an Ace). Can you use symmetry to work out what the expected value of A is?

[Question](#)
[Solution](#)

8. Suppose you bid an amount X . What is the probability of winning the car? Given you've won it, what is the expected amount you'll re-sell it for? Can you use these to calculate your expected profit?

[Question](#)
[Solution](#)

9. A helpful simplifying assumption is that each string of 4 digits is independent, and has a probability of $(1/10)^4$ of being equal to '2018' (because each digit in the string has a $1/10$ chance of matching the corresponding digit in '2018'). This makes the computation much more tractable.

[Question](#)
[Solution](#)

10. What is the LCM of $\{1, 2, \dots, 10\}$? How many of the numbers from $\{1, 2, \dots, 10\}$ do we need to roll to make the LCM greater than 2000?

[Question](#)
[Solution](#)

11. Try to break the situation down into different cases, and examine them independently. For instance, what are the chances of the d20 being greater than 12? If you know it's greater than 12, what is the expected value of the maximum dice?

[Question](#)
[Solution](#)

12. This problem is hard to get an accurate answer for. However, we can get a decent estimate for the true answer, using a similar method to [Question 1](#). As a first step, try writing down the probability of each consecutive difference, given two random independent consecutive rolls.

[Question](#)
[Solution](#)

13. For the expected number of rolls, try using the formula $\mathbb{E}[X] = \sum_{n=1}^{\infty} P(X \geq n)$ (for any random variable X which only takes positive integer values).

For the expected sum of rolls, try calculating the expected sum if the first number you roll is N , starting with $N = 6$ then working down.

[Question](#)
[Solution](#)

14. **HINT #1**

Start with the simple case of no snakes and no ladders. Then, add the snake (think about geometric random variables). Finally, add the ladder. If you are still finding the problem intractable, go to hint 2.

HINT #2

The expected value of a single dice roll is 3.5. What does that imply about the expected number of rolls before you reach 25? What does it imply about the probability that you will hit the number 20?

[Question](#)
[Solution](#)

15. Suppose x_n was just a randomly chosen integer. For any fixed value of A_n , what is the probability x_n is divisible by A_n ? Can you use this to calculate an expected value for $x_{n+1} - x_n$? Now can you use this to calculate an expected value for x_{100} ?

[Question](#)
[Solution](#)

16. **HINT #1**

Try drawing a probability tree listing all possible combinations up to the third coin flip. Do you notice anything about the types of numbers that come up? The next hint offers more guidance on what kind of pattern you should be looking for.

HINT #2

You can sort the numbers that come up into 3 “types”, where these types are significant because you can define the probability that you progress from one type to another. In Markov chain terminology, these are their *transition probabilities*. However, you don’t need to know about Markov chains to solve this problem; the important thing is to identify the three different types. If you aren’t sure what these types are, you can read the next hint to find out.

HINT #3

The three types are “close to 0”, “close to 1”, and “close to 100”. We can label these as (0), (1) and (100). The important difference between the first two types is that (0) (which includes numbers such as 0.01) can lead to (100), but (1) (which includes numbers such as 1.01) can never lead to (100). In Markov chain terminology, (1) is an *absorbing state*.

Try to count how many of each of the 2^N states at each level of the probability tree is in each of the three types. Can you spot a pattern?

[Question](#)
[Solution](#)

Solutions

1. Your first roll will always be a "new" number. For each subsequent number, there is a $5/6$ chance it will be "new", and $1/6$ chance it will be "old". This means the number of rolls until you get a new number is a geometric random variable (hyperlink) with parameter $\frac{5}{6}$. The expected number of rolls until a new number is $6/5$. Once you get a second number, a similar calculation shows the expected number of rolls until a third number is $6/4$. Proceeding in this way, we find that the expected number of rolls until every single number has been thrown is $(6/6) + (6/5) + \dots + (6/1) = 14.7$.

[Question](#)

2. Let X be the expected number of rolls before getting a double 6, given your last roll was a 6. Let Y be the expected number of rolls before getting a double 6, given your last roll wasn't a 6 (or you haven't rolled yet).

If the last number wasn't a 6 (i.e. state Y), then the next number is either a 6 (bringing us to state X), or it isn't a 6 (in which case we stay at state Y). This gives us the equation:

$$Y = \frac{1}{6} \cdot X + \frac{5}{6} \cdot Y + 1$$

If we are in state X , then we either roll a 6 and finish, or roll a non-6 and go back to state Y . This implies:

$$X = \frac{5}{6} \cdot Y + 1$$

Solving these equations, we find $X = 36$, $Y = 42$, so the answer is 42.

Quick sanity check on this answer¹⁸: the probability of getting two 6s in a row straight away is $1/36$, so the situation is close to a geometric random variable with probability $1/36$, meaning we expect a value of approximately 36. On this basis, our answer of 42 seems reasonable.

[Question](#)

3. First, consider a simplified version of the problem where you also collect your winnings if you roll a 6. The probability of rolling a 4, 5 or 6 is $1/2$, so the expected number of rolls before getting one of these numbers is 2. This means the expected number of times you roll a 1, 2 or 3 before stopping is 1. Now, consider the effect of losing your winnings if you roll a six. $2/3$ of the time you'll finish on a 4 or 5, and keep your winnings (expected value 1), but $1/3$ of the time you'll finish on a 6 and get nothing. Therefore, your expected winnings are:

$$\left(\frac{2}{3} \times 1\right) + \left(\frac{1}{3} \times 0\right) = \frac{2}{3}$$

[Question](#)

¹⁸Sanity checks are always great things to do in interviews - they are very quick ways of verifying your answer, plus they show the interviewer you aren't overconfident, you can stand back and take a big-picture perspective, and you have good intuitions.

4. This is a classic problem in Bayesian probability. The probability of choosing a fair coin is 999/1000, and the probability that the fair coin comes up heads 10 times is $2^{-10} = 1/1024$. The probability of choosing the biased coin is 1/1000, and if you choose the biased coin then it will always come up heads. The probability that you picked the coin with two heads is given by:

$$\frac{\frac{1}{1000} \times 1}{\frac{1}{1000} \times 1 + \frac{999}{1000} \times \frac{1}{1024}} = \frac{1}{1 + \frac{999}{1024}} \approx \frac{1}{2}$$

Note - another way you can solve this problem is by using *odds ratios*.

Question

5. Your sequence and the interviewer's sequence both start with H, so WLOG¹⁹ we can assume a single head has just been flipped. Let p be the probability you win from this state. We now condition on the subsequent flips.

- HH \implies you have definitely lost (because as soon as the next tail comes up, the interviewer wins)
- HT \implies we condition on the next flip:
 - HTT \implies you win
 - HTH \implies we are back where we started

This gives us the equation:

$$p = \frac{1}{2} \times 0 + \frac{1}{4} \times 1 + \frac{1}{4} \times p$$

which has solution $p = 1/3$.

You can never have more than probability 7/8 of winning (because the interviewer might get his sequence in the first 3 flips). If the interviewer chooses HHH, then you can attain this upper bound of 7/8 by choosing the sequence THH. The reason for this is as follows: if T appears in any of the first 3 flips, any sequence HHH that subsequently appears must be prefixed with a tail.

Question

6. The number of heads in 100 coin flips follows a Bin(100, 1/2) distribution. By the central limit theorem, we can approximate this as a normal distribution, with mean $\mu = 100 \times \frac{1}{2} = 50$, and variance $\sigma^2 = 100 \times \frac{1}{2} \times (1 - \frac{1}{2}) = 25$. The standard deviation is $\sqrt{25} = 5$. So getting > 55 heads is approximately a one standard deviation event, which happens with probability approximately 16%. This is reasonably close to the true value (about 13.5%).

Question

7. Suppose the 4 aces divide the pack into 5 groups, with A, B, C, D, E cards respectively. For any permutation of these numbers, e.g. (B, D A, E, C), the corresponding deck is equally likely to appear (because we can pair up decks from the groups corresponding to the two different permutations, or to use the technical term, form a *bijection*). It follows that each of the five numbers must have the same probability distribution.²⁰ However, we know they sum to 48 (because there are 48 non-Aces in the pack), so each of them must have expected value $48/5=9.6$. Finally, we add in the Ace itself, and get the answer 10.6.

Question

¹⁹Without Loss Of Generality

²⁰For more on the principle of symmetry, see *Fifty Challenging Problems In Probability*.

8. Suppose you bid an amount X . The probability of the car being worth less than X is $X/1000$, and conditional on the value being in this range, the expected value is the midpoint, i.e. $X/2$. This means the expected amount you can sell it for is $1.5 * X/2 = 0.75X$. This is less than the amount X that you originally paid. We deduce that betting any positive amount gives you negative expectancy, so the best decision is to bet nothing. This is an example of adverse selection - you are only winning the car in the situations where the car holds very little value.

Now consider the second case. Suppose you bid X , then the probability of winning the bid is $(X-100)/900$. If you know the car is worth at most X , then it is uniform between 100 and X , so the expected re-sale value is $1.5 \times (X + 100)/2$. Your expected profit is:

$$\begin{aligned} E &= \frac{X - 100}{900} \times \left(1.5 \times \frac{X + 100}{2} - X \right) \\ &= \frac{X - 100}{900} \times \left(75 - \frac{1}{4}X \right) \\ &= \frac{1}{3600}(X - 100)(300 - X) \end{aligned}$$

This is a quadratic in X , which has roots at 100 and 300, so is maximised at $X = 200$.

In the third case, the denominator of the expression above is different (9900 rather than 900), but everything else is unchanged, so we get the same result of $X = 200$.

[Question](#)

9. Using the simplifying assumption in the hint, we note that if the powers of two were written out back to back in one long line, the expected number of digits we would have to write until we get the string '2018' is 10^4 (by property of geometric distributions). Now, we have to calculate how many powers of 2 we need to write before we reach 10^4 digits. Again, we make use of $2^{10} \approx 1000$, which implies 2^N has approximately $0.3N$ digits. So if we write out the powers of 2 from 2^1 to 2^N , the number of digits is approximately:

$$\sum_{r=1}^N 0.3r = 0.3 \times 0.5N(N + 1) \approx 0.15N^2$$

Finally, we solve $0.15N^2 = 10^4$, and get $N = 258$. This is quite close to the true value, which turns out to be 204.

Although this might seem like a very rough and unreliable method, it is actually more reliable than it seems. The variance²¹ of N is actually less than the variance that the corresponding geometric random variable would have, because smaller powers of 2 have fewer digits (meaning the answer is unlikely to be much smaller than 258), and larger powers of 2 have more digits (meaning the answer is unlikely to be much larger than 258).

[Question](#)

²¹Technically, saying "variance" here is wrong, since N isn't actually random, but it is from our perspective.

10. $\text{LCM}\{1, 2, \dots, 10\}$ is $2^3 \cdot 3^2 \cdot 5 \cdot 7 = 2520$. This is larger than 2000, but if we take out any single factor, it is smaller. This means, to get $\text{LCM} > 2000$, we need all these factors to appear in our rolls, i.e. we need to roll $\{7, 8, 9\}$ and either a 5 or a 10. It is hard to work out the exact expected number of rolls this will take, but we can approximate it by pretending we don't need to roll 5 or 10. This shouldn't change the expected value much, because the probability of getting at least one of $\{5, 10\}$ before getting all three of $\{7, 8, 9\}$ is quite high. This simplification gives us an estimate of $(10/1) + (10/2) + (10/3) = 18.3$ (using the same method as in [Question 1](#)). This is close to the true value, which is approximately 18.8, but it is a slight underestimate as we would expect.

[Question](#)

11. There are several ways to approach this problem. One possible solution involves splitting into 3 cases:
- (a) **d20 > 12**
This happens with probability $8/20 = 2/5$, and in this case the d20 is the maximum, so the expected value of the maximum is $(13+20)/2=16.5$.
 - (b) **max{d20, d12} ≤ 8**
This happens with probability $(8/20) \times (8/12) = 4/15$. In this case, we functionally have a d4, d6, and three d8s, so anywhere in the range 6-7 seems a reasonable guess as to the expected value of the maximum - let's say 6.5.
 - (c) **d20 ≤ 12, and max{d20, d12} > 8**
This happens with probability $1 - 2/5 - 4/15 = 1/3$, and in this case we know the maximum is between 9 and 12, so the midpoint 10.5 seems like a reasonable guess.

Combining these 3 estimates and weighting them by their probability, we get an expected value of:

$$\frac{2}{5} \times 16.5 + \frac{4}{15} \times 6.5 + \frac{1}{3} \times 10.5 \approx 11.83$$

which is very close to the true value of approximately 11.91.

To refine this estimate, you would want to take the case with highest uncertainty and decompose it. Here, that would probably be case (b). One sensible idea is to decompose it into the cases where at least one dice is greater than 6, and no dice is greater than 6.

If you wanted a faster method, you could combine (b) and (c) into a single case. This would be more computationally efficient, but your answer might be less accurate.

[Question](#)

12. There are 6 possible consecutive differences: 0, 1, ..., 5, and they appear among any pair of dice with probability $6/36, 10/36, 8/36, 6/36, 4/36, 2/36$ respectively. We can simplify the problem by modelling these as independent random events (and adding 1 onto our answer at the end, since the first roll doesn't generate a consecutive difference). Now the question becomes: given these independent random variables, what is the expected number of rolls until we attain them all?

Getting the first consecutive difference clearly only takes 1 roll. The most likely first consecutive difference to get is the one with probability $10/36$, then the expected number of rolls until getting a new consecutive difference would be $36/(36 - 10) = 36/26$. We can proceed in this way, getting an estimate of $1 + (36/36) + (36/26) + (36/18) + (36/12) + (36/6) + (36/2) \approx 32.38$. This is likely to be an overestimate, because any other order in which we find the consecutive differences will give us a lower sum - for instance, finding them in reverse likelihood order gives us an expected value of 11.36. So we expect a value somewhere between these two, but much closer to 32.38, because finding the consecutive differences in order of decreasing likelihood is the most likely order to find them. Sure enough, the true value turns out to be about 25.85.

[Question](#)

13. There will be at least N rolls iff the first N dice are distinct, and in ascending order. The probability of the $(m + 1)$ th dice being distinct from the first m is $(6 - m)/6$, so the probability they are all distinct is:

$$\frac{5}{6} \times \frac{4}{6} \times \dots \times \frac{6 - (N - 1)}{6}$$

If they are all distinct, the probability of them being in ascending order is $1/N!$ (because there are $N!$ possible arrangements, and only one of them is ascending). From this, we get a formula for $P(X \geq n)$:

$$\mathbb{P}(X \geq N) = \frac{5}{6} \times \frac{4}{6} \times \dots \times \frac{6 - (N - 1)}{6} \times \frac{1}{N!}$$

Using the formula $\mathbb{E}[X] = \sum_{N=1}^{\infty} \mathbb{P}(X \geq N)$, we get the expected value as

$$\mathbb{E}[X] = 1 + \frac{5}{12} + \frac{5}{54} + \dots$$

We could sum all 6 terms terms, but notice that the terms are decreasing in size, and the third is already very small. If we truncate at the third term, we expect to get a very good approximation for the true value. In fact, we get $163/108 = 1.509$, which is very close to the true value of 1.522.

For the sum of rolls, first suppose we rolled a 6, then expected sum is 6 (because 6 must be our only roll). Now suppose our first roll was a 5, then our expected sum is $5 + 6p$, where p is the probability that our next roll is a 6. $p = 1/6$, so we get $5 + 6(1/6) = 6$. Continuing in this way, we can see that our expected sum will be 6 regardless of what our first roll is.

Now consider the case for general N . By a similar argument, we find that the expected sum of rolls is N . Using the same formula for expected number of rolls, we can see that for very large N , the factorial term in the previous equation will decay much faster than the product of fractions at the start, so the expected value will tend to $\sum_{N=1}^{\infty} 1/N!$. This is the [Taylor series](#) for e , with the first term missing, so the limit is $e - 1 \approx 1.718$.

Question

14. First, we consider the case of no snake or ladder. You have to move at least 24 squares, and the expected amount by which you move on each roll is 3.5. You might be tempted to say that the expected number of rolls is $24/3.5$, and this is a pretty good approximation. However, there is an additional factor. You don't stop when you get to exactly square 25, you stop when you get at least 25. A useful result to know is that, if you roll a d6 and stop when you get cumulative sum at least X , then for large values of X , your expected stopping value tends to $X + 5/3 \approx X + 1.67$. To intuitively understand this, note that if your average roll is 3.5, then we should expect the distance between X and the first number you hit that is at least X to be approx. uniformly distributed between 0 and 3.5, so expected value should be about $3.5/2 = 1.75$. Using the true value of $5/3$, the expected number of rolls until finishing is:

$$E_1 = \left(\frac{24 + 5/3}{3.5} \right) = \frac{22}{3} \approx 7.33$$

Now consider the snake. The average roll is 3.5, so the probability of landing on any individual square is $1/3.5$ (note this clearly doesn't hold in small cases - e.g. your probability of landing on square 2 is only $1/6$ - but for larger numbers it holds with extremely high accuracy). So the number of attempts until we get past the snake is geometric with probability $(1 - 1/3.5) = 2.5/3.5$, meaning the expected number of times we hit the snake is $(3.5/2.5) - 1 = 0.4$. Each time we hit the snake, it adds $15/3.5$ to our expected total number of rolls. This means the expected number of rolls in the snake game is:

$$E_2 = \frac{22}{3} + \left(0.4 \times \frac{15}{3.5} \right) = \frac{190}{21} \approx 9.05$$

Finally, we now consider the ladder. Once again the probability of hitting it is $1/3.5$, and if we hit it then we subtract $5/3.5$ rolls from the expected number of rolls in our game, so we get:

$$E_3 = \left(\frac{22}{3} - \frac{1}{3.5} \times \frac{5}{3.5} \right) + \left(0.4 \times \left(\frac{15}{3.5} - \frac{1}{3.5} \times \frac{5}{3.5} \right) \right) = \frac{178}{21} \approx 8.47$$

Sanity check: the snake is three times longer than the ladder, so we expect E_3 to be between E_1 and E_2 , but much closer to E_2 - this is indeed what we have found.

[Question](#)

15. If x_n was just a randomly chosen integer, then for any fixed A_n , we have $\mathbb{P}(A_n \text{ divides } x_n) = 1/A_n$, so the expected value of $x_{n+1} - x_n$ is:

$$\begin{aligned} \mathbb{E}[x_{n+1} - x_n | A_n] &= \frac{1}{A_n} \times A_n + \left(1 - \frac{1}{A_n} \right) \times (-A_n) \\ &= -A_n + 2 \end{aligned}$$

We then calculate the expected value of $x_{n+1} - x_n$ for random A_n :

$$\begin{aligned} \mathbb{E}[x_{n+1} - x_n] &= \sum_{A_n=1}^{10} \frac{1}{10} \times (-A_n + 2) \\ &= \frac{1}{10} \times (-0.5 \times 10 \times 11 + 20) \\ &= -3.5 \end{aligned}$$

This can be assumed to hold approximately in all cases, except $n = 0$ (because 0 is divisible by every number). The expected value of x_1 is equal to the expected value of A_n , i.e. 5.5. This gives us an answer of:

$$\mathbb{E}[x_{100}] = 5.5 + 99 \times (-3.5) = -341$$

This is pretty close to the true answer, which is about -341.6 (attained by random simulation). The true answer being slightly smaller than our estimate may be due to effects at the start (e.g. $x_1 = A_0$, so if $A_0 < A_1$, then $\mathbb{P}(A_1 \text{ divides } x_1) = 0$, so A_1 will certainly be subtracted).

[Question](#)

16. As described in the hints, we will call the three states (0), (1) and (100). If you are in (0), then $X \rightarrow 1/X$ takes you to (100), and $X \rightarrow X+1$ takes you to (1). So if some level of the probability tree has A nodes in state (0), this will result in A nodes being in states (1) and (2) respectively on the next level of the tree. Similar analysis for the other two states shows us that if you have (A, B, C) in states (0, 1, 100) on some level of the tree, then the next level has (C, A+2B, A+C). Below is a table of these values:

N	0	1	2	3	4	5	6	7	8
A	0	1	1	2	3	5	8	13	21
B	0	0	1	3	8	19	43	94	201
C	1	1	2	3	5	8	13	21	34

Note, this is a lot easier to fill in if you recognise that A and C are just the Fibonacci series! I will now show 2 possible ways we can finish: one very fast, and one slower but slightly less accurate.

METHOD 1: The most direct way is to multiply the probabilities in the table by 0, 1 and 100 respectively. The approximate answer we get is:

$$\begin{aligned}
 E &= \frac{21}{256}(0) + \frac{201}{256}(1) + \frac{34}{256}(100) \\
 &= \frac{201 + 3400}{256} \\
 &= \frac{3601}{256} \\
 &= 10 + \frac{1041}{256} \\
 &\approx 14
 \end{aligned}$$

METHOD 2: To get a more accurate answer, it's worth digging a bit further into our (0, 1, 100) simplification. It's a pretty good model, but breaks down in a few places.

- (0)-states will always be strictly greater than 0, although they are so small this doesn't mess up our calculations much.
- (1)-states might be larger or smaller than 1, e.g. $1 \rightarrow 2 \rightarrow 3 \rightarrow 1/3$ all count as (1)-states in our calculations. Clearly the larger numbers will have a bigger impact on the expected value than the smaller numbers, so an expected value of 2 for the (1)-states seems like a more reasonable guess.
- (100)-states can be any odd number from 101 to 107 (can you see why this is?), so a 104 is a reasonable guess for their expected value.

Now we can calculate a better expected value:

$$\begin{aligned}
 E &= \frac{21}{256}(0) + \frac{201}{256}(2) + \frac{34}{256}(104) \\
 &= \frac{402 + (3400 + 4 \times 34)}{256} \\
 &= \frac{3938}{256} \\
 &= 10 + \frac{1378}{256} \\
 &= 15 + \frac{98}{256} \\
 &\approx 15 + \frac{100}{250} \\
 &= 15.4
 \end{aligned}$$

The true value turns out to be 15.21, so both methods worked pretty well. In an interview, you probably wouldn't be expected to get an answer as accurate as method 2; the first method would be fine.

[Question](#)

4.3 Strategy Game Problems

EASY

1. Random Dice Game #1

You roll a fair d6. You can either win the number on the dice, or you can pay £1 to roll again, in which case you have to take whatever number you get. What is the best strategy, and what is the expected value?

Note - this question has a harder follow-up [here](#).

[Hint](#)
[Solution](#)

MEDIUM

2. Weighted Dice

You and the interviewer both play a game: you both pick a number (doesn't need to be an integer), then roll two weighted d6 dice. The person who gets closest to the sum wins. The dice are non-uniform: there is a 40% chance of getting a sum of 12, and a uniform distribution for all other possible sums. What is your strategy, both in the cases where you go first, and where the interviewer goes first?

[Hint](#)
[Solution](#)

3. Dice Sum Game - Greater Than 8

You roll a d8, and keep track of the cumulative sum. You can stop and take the sum at any time, but if the sum is ever greater than 8 then the game ends, and you get a payout of zero. What is the best strategy, and what is the expected value?

[Hint](#)
[Solution](#)

4. Random Dice Game #2

This is a follow-on question from [Random Dice Game #1](#). In this case, you roll a d100, and you can re-roll as many times as you want (but you have to pay £1 each time). What is the best strategy and expected value in this case?

Note - it is possible to get the exact right answer with a bit of effort, but for this question it is recommended to try to find a quicker and more approximate method.

Now suppose the numbers were between 1 and 1000. Would you expect the value to be less than, equal to, or greater than ten times the original game?

[Hint](#)
[Solution](#)

5. Product of Dice

You roll 2 d4s, and your payout is their product. You can pay £2 to re-roll either of the dice, as many times as you like. What is the best strategy, and what is the optimum value? What happens when the cost of re-rolling gets very large, or very small?

[Hint](#)
[Solution](#)

HARD

6. Betting On Colour

A deck of cards has 26 red and 26 black cards. It is shuffled, and each card is dealt face up, one after the other. Before each card is dealt, you may bet any amount of money (up to your current pot size) on it being red or black. If you are correct then you double your bet; if you are incorrect then you lose your bet. You start with £1. What strategy maximises your expected profit, and what is this profit?

[Hint](#)
[Solution](#)

7. 10 Cards Game

There are 10 cards, face down, numbered $\{1, 2, \dots, 10\}$. If you pay £1, you can point to any 2 cards and I will tell you which card is smaller and what its value is. At any point you can finish the game by choosing a card, and you win the value of that card. What is the expected value of this game?

[Hint #1](#)
[Hint #2](#)
[Solution](#)

8. Threes

In this game, threes count as zero, while the other faces count normally. The goal is to get as low a sum as possible. You start by rolling 5 d6s, and on each roll at least once dice must be kept. Any dice that are kept are added to your sum, and you keep re-rolling the remaining dice until all have been added to your sum. The game lasts at most five rolls, and the score can be anywhere from 0 to 30. For instance, a game might go like this:

- $\{2, 3, 3, 4, 6\} \rightarrow$ keep both threes,
- $\{1, 5, 5\} \rightarrow$ keep the one,
- $\{1, 2\} \rightarrow$ keep both, final score is $0 + 0 + 1 + 1 + 2 = 4$.

What is the expected value of this game?

[Hint #1](#)
[Hint #2](#)
[Solution](#)

9. Dice Sum Game - Multiples Of 10

You roll a d6, and keep track of the cumulative sum. You can stop at any point and take your current sum, but if you ever land on a multiple of 10 then you stop and win nothing. What is the best strategy, and what is the expected value of this strategy?

[Hint](#)
[Solution](#)

10. Decreasing Sequences

9 cards (numbered from 1 to 9) are shuffled and then dealt, one after the other. You can stop at any point, and receive a payout equal to the sum of the last decreasing subsequence. For example: say the cards were arranged 752194836, then you could have stopped after the 1 for a payout of $7 + 5 + 2 + 1 = 15$, or after the 4 for a payout of $9 + 4 = 13$. If you waited until the end, you'd have to take a payout of 6. How much would you pay to play this game?

[Hint #1](#)
[Hint #2](#)
[Solution](#)

Hints

1. Suppose you choose to re-roll, what is the expected value of the number you get? What does this imply about the numbers you should choose to re-roll with?

[Question](#)
[Solution](#)

2. First think about the case where you choose the number second. If the interviewer chooses X , would you ever choose $X+2$ instead of $X+1$?

[Question](#)
[Solution](#)

3. Consider strategies of the form “re-roll until you get at least X ”. Your payout is $A - B$, where A is the value you stop at, and B is the number of extra rolls it took. This is a useful framing for calculating expected value.

[Question](#)
[Solution](#)

4. Suppose you roll the dice, and get values $\{X, Y\}$, with $X < Y$. If you choose to re-roll, which of the two dice should you re-roll? What does that tell you about the form that the optimal strategy must have?

[Question](#)
[Solution](#)

5. Clearly, you don't want to re-roll if you have a sum of 8, but you do want to re-roll with a sum of 0. There must be some maximum value that you re-roll on. Try finding that value, then find the expected value by working downwards.

[Question](#)
[Solution](#)

6. The [Kelly criterion](#) is a useful tool here; in particular the section on *even money bets*.

[Question](#)
[Solution](#)

7. HINT #1

You always have the option to just guess at random; this is your baseline strategy. What is its expected value? How does the information you get the first time you pay £1 increase the expected value?

HINT #2

Say you pay £1, point at 2 cards, and are shown the minimum one, and told its value is X . Which of the 10 cards has the largest expected value now, and what is that value in terms of X ? What does this tell you about the expected value of the game after you see X , given that X is random?

[Question](#)
[Solution](#)

8. **HINT #1**

With pen and paper, this problem is tractable with 1 or 2 dice, but very difficult beyond that. Start with these cases, then see if you can extrapolate to guess a value for the 5-dice case. If you are struggling with the 2-dice case, see the next hint.

HINT #2

Hopefully you found that the expected value in the 1-dice case is 3. Suppose you roll 2 dice, then you have to keep the smaller one. What is the condition on whether you should keep the bigger one?

[Question](#)
[Solution](#)

9. **HINT #1**

Prompt questions:

- Should you stop if you are at 9? Why/why not?
- How about at 99?

For an answer to these questions, see the start of the solution. For more prompt questions, see hint #2.

HINT #2

Prompt questions:

- For any given multiple of 10, what is the probability you “get past it”?
- What is the probability you get past the first N multiples of 10?
- What is the expected value of the game if you get past the first N multiples of 10, and then stop before you are at risk of hitting the next multiple of 10?

[Question](#)
[Solution](#)

10. **HINT #1**

Calculating the expected value of the longest decreasing subsequence is very difficult. However, it’s also unnecessary, because you can’t guarantee to choose the longest decreasing subsequence - sometimes it will be a better choice (in terms of expected value) to stop earlier. Can you come up with a strategy that tells you when to stop, and normally ends on a subsequence with a high sum?

For a more specific hint about strategies you could try, see hint #2.

HINT #2

One very simple strategy would be to wait until you draw a 9, then stop soon after that (e.g. as soon as there is a risk of the next card being higher than the preceding one). This works pretty well, but there are better strategies available. Try to think about the situations in which this strategy works poorly. Can you use these observations to come up with a better strategy?

[Question](#)
[Solution](#)

Solutions

1. The expected value of a single roll is 3.5. If you pay to re-roll, the expected value of the game is $3.5 - 1 = 2.5$. Thus, if you roll a 1 or 2 on your first try, you will be increasing the expected value by paying to re-roll, and if you roll a 3 or more you will be decreasing your expected value. The value of the game is: $(2.5 + 2.5 + 3 + 4 + 5 + 6)/6 = 23/6 = 3.83$.

Question

2. First, note there are 10 possible sums other than 12 (from 2 to 11 inclusive), so each has probability 6%. We can split into 3 cases:

- $X > 10$. The probability of the number being in the range 2-10 is 54%, so you should choose just below X , and your probability of winning is at least 54%.
- $X = 10$. If you choose just below 10, your probability of winning is 48%. If you choose just above, your probability is 46%. So you should choose just below.
- $X < 10$. You should choose just above X , and your probability of winning is at least 52%.

From this analysis, we can see that the best situation for the interviewer is if he chooses $X = 10$, because this is the only way he can get better than 50% odds. So if you are first to choose, this should be your strategy.

Question

3. If you re-roll until you get at least X , then your payout is $A - B$, where A is the value you stop at, and B is the number of extra rolls it took. We can work out the expected value of both A and B as a function of X , and then differentiate that expression with respect to X to find out the stationary point, i.e. the optimum value of X . However, for a faster method, we can try plugging in values.

Let $\mathbb{E}_X[A]$, $\mathbb{E}_X[B]$ = expected values of A & B with the strategy “stop when we get at least X ”. Then:

- $\mathbb{E}_{50}[A] = 75$, since the stopping value is uniform between 50 and 100.
- $\mathbb{E}_{50}[B] = 1$, since the number of rolls until you get at least 50 is a geometric random variable with probability $51/100 \approx 1/2$, so expected value 2, and we get the first roll for free.

So $\mathbb{E}_{50}[A - B] = 74$. Using a similar method, we get:

- $\mathbb{E}_{80}[A - B] = 90 - 4 = 86$
- $\mathbb{E}_{90}[A - B] = 95 - 9 = 86$
- $\mathbb{E}_{100}[A - B] = 100 - 99 = 1$

From this, we can guess the expected value curve has a peak somewhere between $X = 80$ and $X = 90$. We could also guess that the peak is slightly higher than 85, since the curve must fall away very steeply as X moves up past 90 (if this isn't clear, try making a sketch). In fact, the peak is somewhere between 86 and 87, and the expected value of using this strategy (i.e. stopping when you get at least 87) is about 87.3.

For the game with numbers 1-1000, we expect the value to be *more* than ten times the previous game. An easy way to see this is by observing that the 1-1000 game is similar to playing the previous game with a re-roll cost of £0.10 rather than £1, and then multiplying your net profit by 10. The smaller re-roll cost relative to the size of numbers you are getting makes this game much more attractive. In fact, the best strategy in the second game turns out to be re-rolling until you get at least 956, and the expected payout of this strategy is about 956.8.

Question

4. Suppose the dice values are X and Y , with $X < Y$. If you choose to re-roll, you should clearly re-roll X , because this will increase their expected product more than re-rolling Y . From this, we can deduce that any optimal strategy must be of the form:

“keep re-rolling both dice, stop as soon as the minimum of the two dice is at least N ”

where N could be either 1, 2, 3, 4. We can calculate the expected value in each of these cases.

- **$N = 1$**

The expected value of each die is $(1 + 4)/2 = 2.5$, so the expected product is $2.5^2 = 6.25$ (because $E[XY] = E[X]E[Y]$ when X and Y are independent). The expected number of extra rolls it takes to get all dice ≥ 1 is 0, so the expected value of this strategy is 6.25.

- **$N = 2$**

If we re-roll both until both dice are 2, 3 or 4, the expected stopping value of each die is $(2 + 4)/2 = 3$, so expected product is 9. The probability of any single die being at least 2 is $3/4$, so the expected number of rolls it takes is $4/3$, and the expected amount you pay for re-rolling both dice is $2 \times 2 \times (4/3 - 1) = 4/3$, meaning expected value of strategy is $9 - 4/3 = 7.67$.

- **$N = 3$**

Expected value is $3.5^2 - 2 \times 2 \times (4/2 - 1) = 8.25$

- **$N = 4$**

Expected value is $4^2 - 2 \times 2 \times (4/1 - 1) = 4$

So we deduce the best strategy is to re-roll both dice until they are both at least 3, and the expected value of the game in this case is 8.25.

If the cost of re-rolling gets very large, we are more likely to keep the values we get at the first roll, so the expected value will tend to $2.5^2 = 6.25$. If the cost gets very small, we are more likely to keep re-rolling until we get both dice equal to 4 (since this won't cost much in expectation), so the expected value will tend to 16.

[Question](#)

5. We can consider the case where you re-roll on certain numbers, starting from 4 and then working downwards.

- If you re-roll on 4, your expected value is $(5 + 6 + 7 + 8)/8 = 3.25$. This is less than 4, so you shouldn't re-roll on 4 (or by extension, on any higher numbers).
- If you re-roll on 3, your expected value is $(4 + 5 + 6 + 7 + 8)/8 = 3.75$. This is greater than 3, so you should re-roll on 3 (and by extension, any smaller numbers). This also tells us that if your first roll is a 3, the expected value of the game at this point is 3.75.
- If you re-roll on 2, your expected value is $(3.75 + 4 + 5 + 6 + 7 + 8)/8 = 135/32 = 4 + \frac{7}{32} \approx 4.25$.
- If you re-roll on 1, your expected value is $(4.25 + 3.75 + 4 + 5 + 6 + 7 + 8)/8 = 4.75$.

Finally, your expected value at the start of the game is approximately:

$$(4.75 + 4.25 + 3.75 + 4 + 5 + 6 + 7 + 8)/8 = 5.34$$

An interesting observation about the game: the worst number to roll is a 3. Intuitively, this is because a higher number is good enough to stick with, whereas a lower number is unlikely to “go bust” if you re-roll.

[Question](#)

6. The [Kelly criterion](#) gives the mathematically correct²² amount to wager on an even-money bet, to maximise the expected rate at which your wealth grows. If you have probability $p > 0.5$ of winning the bet, then you should wager a proportion $(2p - 1)$ of your total wealth on the outcome. Mathematically, this maximises the expected value of the natural log of your wealth. One intuitive reason for why you should choose to maximise your log wealth rather than your actual wealth is that $\log(x)$ is an increasing function of x , but tends to $-\infty$ as $x \rightarrow 0$, in other words going bust has an infinite penalty.

Now we consider the Kelly criterion applied to this game. Suppose there are R red and B black cards left to be dealt, and $R > B$, then the probability of the next card being red is $\frac{R}{R+B}$, so you should bet $(2 \times \frac{R}{R+B}) - 1 = \frac{R-B}{R+B}$ of your wealth on red. You should do this at every stage. One way to calculate the expected returns from this strategy is as follows: if your bet is correct, your wealth is multiplied by $\frac{2R}{R+B}$, and R reduces by 1. If your bet is incorrect then your wealth is multiplied by $\frac{2B}{R+B}$, and B reduces by 1. By considering the effect of every single bet, and multiplying all these factors together, we can see that your final wealth will be:

$$\frac{((2 \times 26) \times (2 \times 25) \times \dots \times (2 \times 1))^2}{52 \times 51 \times \dots \times 1} = \frac{2^{52} \times 26!^2}{52!} = 2^{52} \left/ \binom{52}{26} \right.$$

It might seem strange that there is no variance, i.e. you always end up with this wealth regardless of which order the cards come out. However, there is actually a very nice interpretation of this strategy that explains this confusing feature. The strategy is actually equivalent to betting $1/\binom{52}{26}$ of your wealth at the start on every single one of the $\binom{52}{26}$ possible combinations of red and black cards. All of these combinations will be wrong eventually, except for exactly one which will be correct on every card, and so your bet doubles 52 times.

[Question](#)

7. If you guess at random straight away, your expected value is $(1 + 10)/2 = 5.5$, so this is our lower bound. Now suppose you point to the first two cards, and are told the min of them is X . The max of those two cards now has a uniform distribution between $X+1$ and 10, so expected value is $((X + 1) + 10)/2 = X/2 + 5.5$. Choosing this card is better than choosing a different random card (unless the min of the first 2 cards is 1, in which case we are indifferent). Now, the question is: what is the expected value of $X/2 + 5.5$? We might say it is 8.25, by plugging in $\mathbb{E}[X] = 5.5$, but this is actually not correct. X isn't uniformly random; it's the min of a randomly chosen pair. To calculate $\mathbb{E}[X]$, we note that there are $\binom{10}{2} = 45$ possible equally likely pairs, and the number that have minimum $(1, 2, \dots, 9)$ are $(9, 8, \dots, 1)$ respectively. This means the expected value of the min of 2 cards is:

$$\frac{1 \times 9 + 2 \times 8 + \dots + 9 \times 1}{45} = 11/3$$

So we deduce $\mathbb{E}[X/2 + 5.5] = 11/6 + 5.5 = 7.33$. Subtracting the cost of pointing to 2 cards, we get 6.33. This is our new lower bound on the value of the game; it is the value we get from using the strategy:

“Pay £1 to look at the min of 2 cards, then choose the larger one.”

However, this might not be the optimal strategy. If the X is large then we should probably stop there and pick the larger one (e.g. if $X = 9$ then the larger one must be 10), but if X is very low then we might want to keep playing (e.g. if $X = 1$ then we are in a similar position to where we started; playing a version of the game with 9 cards rather than 10). The more cards we look at, the less extra information we get from seeing more, and so the more unlikely it is that it will be worth keeping playing. It is a reasonable guess that the value of the game isn't much larger than the value of 6.33 (when we are only allowed to look at one card), so a guess of between 6.5 and 7 seems reasonable. The true value (as verified by a friend's Python code) turns out to be approximately 6.54.

²²Mathematically correct in the sense that, if you had to make N bets, then as $N \rightarrow \infty$, this strategy will eventually outperform any other strategy almost surely.

It's hard to craft an explicit strategy for this game, because there are quite a few edge cases to consider. However, we can see that the best strategy is going to be some kind of refinement on the following:

*“Pay £1 to look at the min of 2 cards.
If the min is sufficiently large, stop and choose the max.
If the min is small, pay again to compare the max to a different card. Repeat.”*

where “sufficiently large” could mean something like ≥ 3 at the start, but might change slightly depending on which cards have already come up.

Question

8. For ease of notation, I will refer to the dice as having 6 sides: 0, 1, 2, 4, 5, 6. Let E_N denote the expected value of the game with N dice.

The 1-dice case is simple; you have to stick with whatever you roll. The expected value of your roll is $E_1 = (0 + 1 + 2 + 4 + 5 + 6)/6 = 3$.

For the 2-dice case, we have to keep the smallest dice. The relevant strategic question is when to re-roll the largest dice, and when to keep it. From our analysis of the 1-dice case, we deduce that we should re-roll only if its expected value is more than 3 (i.e. it is 4, 5 or 6). I will show 2 possible ways you could arrive at the answer (in reality these methods are essentially the same, although they offer slightly different perspectives).

METHOD 1

Perhaps the most direct method is to write out all the possible combinations of two dice, which I have done in the table below. The entries of the table are the value of the game for the corresponding dice values. For instance, if you roll $\{1, 2\}$ then you keep both of them (because the larger of the two dice is less than 3), so expected value is $1 + 2 = 3$, but if you roll $\{1, 4\}$ then you keep 1 and re-roll 4, so expected value is $1 + 3 = 4$. Note, we only have to fill in the upper corner of the table, because it is symmetric. Summing over all these entries, we get:

	0	1	2	4	5	6
0	0	1	2	3	3	3
1	-	2	3	4	4	4
2	-	-	4	5	5	5
4	-	-	-	7	7	7
5	-	-	-	-	8	8
6	-	-	-	-	-	9

$$E_2 = \frac{(0 + 2 \times (1 + 2 + 3 + 3 + 3)) + (2 + 2 \times (3 + 4 + 4 + 4)) + \dots + (8 + 2 \times 8) + 9}{36} = \frac{158}{36} \approx 4.39$$

METHOD 2

An alternative, slightly slicker method is to consider that your score is the sum of both dice, minus $\max\{X - 3, 0\}$, where X is the value of the larger dice. This is positive if $X = 4, 5$ or 6 , which happens with probability $7/36, 9/36, 11/36$ respectively. This gives us an expected value of:

$$E_2 = (3 + 3) - \left(\frac{7}{36} \times 1 + \frac{9}{36} \times 2 + \frac{11}{36} \times 3 \right) = \frac{158}{36} \approx 4.39$$

Now, consider the case of 3 dice. The strategy can be deduced directly from the previous cases. Let the dice be X, Y, Z where $X \leq Y \leq Z$, then the strategy is:

- $\min\{X + Y + Z, X + Y + 3, X + 4.39\} = X + Y + Z \implies$ keep all three dice
- $\min\{X + Y + Z, X + Y + 3, X + 4.39\} = X + Y + 3 \implies$ keep X and Y , re-roll Z
- $\min\{X + Y + Z, X + Y + 3, X + 4.39\} = X + 4.39 \implies$ keep X , re-roll Y and Z

So the expected value is $\mathbb{E}[\min\{X + Y + Z, X + Y + 3, X + 4.39\}]$. Working out this expected value is pretty hard. One way we could approach it is to calculate the expected value of the much simpler strategy “always keep X , re-roll Y and Z ”. This turns out to be 5.56, which is an upper bound on the expected value of the 3-dice game. However, doing this takes a while, and in a trading interview we should be more willing to try and approximate. We know that adding more dice will increase the expected value, but it seems like this marginal increase in expected value will get smaller as we add more dice (e.g. in an extreme case with 100 dice, we’re almost certain to roll at least one zero, so the expected value is almost identical to the 99-dice game). At this point, I encourage any readers to try and estimate the value of the 3, 4 and 5-dice games. The true values can be found on pages 81-82 of the [Dice Problems pdf](#).

An interesting point to make about the 5-dice game is that even though we don’t know its exact value, we can actually have enough information to play the 5-dice game with a strategy that is quite close to optimal. We know the optimal strategy will have the same form as the one described in the 3-dice game, except there will be 5 possible minimum expressions instead of 3 (corresponding to the 5 possible choices of how many dice to keep), and two of these expressions will include the values of the 3-dice and 4-dice games. Small errors in estimating these values won’t change which expression is the smallest much of the time, and in the cases where it does, the difference in expected value between the option you chose and the best option is probably small.

[Question](#)

9. SOLUTION TO HINT #1

If you re-roll at 9, then you have a $1/6$ chance of hitting 10, and a $5/6$ chance of getting past 10. If you get past 10, then you will get at least 14 (because if you are below 14, you can keep rolling without fear of hitting 20). In fact, as discussed in an earlier problem, your expected stopping value if you keep re-rolling until you get to at least 14 is $14 + 5/3 = 15.67$. If you are on a number X , the value of the game must be at least X , because you have the choice of stopping there. So your expected value from re-rolling at 9 is at least $(5/6) \times 15.67$, which is much greater than 9. On the other hand, if you re-roll at 99, you have a $1/6$ chance of hitting 100, and a $5/6$ chance of passing it and getting at most an extra 10 before you risk hitting another multiple of 10. So re-rolling at 99 will give you an expected value of less than 99, i.e. you shouldn’t re-roll here.

SOLUTION TO HINT #2

The average value of a dice roll is 3.5, so the probability of hitting any number is approximately $1/3.5$. This means the probability of “getting past” some multiple of 10 is $2.5/3.5$. The probability of getting past the first N multiples of 10 is $(2.5/3.5)^N$.

SOLUTION TO FULL QUESTION

From these observations, we can now guess what the optimal strategy might look like. We keep rolling; for the lower multiples of 10 we keep rolling until we get past them (we are happy to risk landing on them, because re-rolling has positive expectancy), but for the larger multiples of 10 the expected value of continuing decreases.

Suppose we used the strategy:

“keep rolling until you get past $10N$, then stop when you get at least $10N+4$ ”

Our probability of getting past $10N$ is $(2.5/3.5)^N$, and our expected stopping value if we do get that far is $10N + 5.67$. We can try to find the value of N that maximises $(2.5/3.5)^N \times (10N + 5.67)$. Trying $N=1, 2, 3, \dots$ we get 11.193, 13.097, 12.999, 11.888, 10.351, \dots , so it seems the value of the strategy peaks early, then deteriorates slowly. The optimum is $N = 2$, which corresponds to the strategy of rolling until you get past 10, then stopping once you get to at least 14. We can assume that our strategy has missed some finer points (e.g. if you get past 10, land on 13, then roll a 6 bringing you to 19, it is actually better to keep rolling than to stop). This means our strategy isn't optimal, so we should expect to underestimate the true value of the game. However, it turns out our strategy gets us very close to the true optimal value of 13.217, which is achieved by using the following (slightly odd-looking) strategy:

“stop on 24, 25, or when you get at least 34”

It's interesting how close we were able to get to the optimal value, even with a strategy that looks pretty different to the optimal one. This is because, much like in the previous question, the cases where our strategy diverges from the optimal strategy are all very marginal choices, and only result in very small losses in expected value.

Note - I haven't discussed this here because it's not really relevant to quant trading interviews, but the way you can computationally verify the optimal strategy and expected value for these types of questions is called [dynamic programming](#).

[Question](#)

10. For simplicity, I will refer to getting a card higher than your previous card as *“going bust”*.

Perhaps the simplest, lowest-risk strategy you could use is the following:

“Keep drawing cards until you get a 9, then keep drawing until you have a chance of going bust.”

Most of the time, this will mean stopping on the next card after 9. Since the cards from 1 to 8 have average value 4.5, we could guess that this strategy will have expected value about $9 + 4.5 = 13.5$. However, this strategy leaves a lot to be desired. For instance, if the first cards were (8, 7), you have the option of stopping here at a sum of 15, which is better than your expected value if you keep drawing until you get a 9. This suggests we try a strategy of the form:

“Keep drawing cards until you get to a sum of at least X , then stop as soon as you have a chance of going bust. If you never get to X , then just stop at the end.”

It's very difficult to make any precise statements about the value of strategy, but by comparing it to the first one, we can guess X should be at least 13, and the value of this strategy should be larger than 13.5. In fact, the best value of X turns out to be 13, and the value of this strategy is about 14.14. Interestingly, this is pretty close to the expected value of the sum of the longest decreasing sequence (which is 16.22), so our strategy is capturing most of the value of the game.

[Question](#)

4.4 Fermi Problems

The questions in this section are structured slightly differently. There will just be a list of Fermi problems for you to practice, with answers. Some of them require the estimation of several uncertain quantities so are challenging to estimate in their own right. Others have just one uncertain quantity, so are more suitable for giving confidence intervals.

Note, most of these answers were found using Google, I've tried to source them reliably, but please take them with a pinch of salt!

1. What percentage of the global population is less than 15 years old (as of 2018)? [Solution](#)
2. How much higher is the average global life expectancy in 2020 than in 1900? [Solution](#)
3. What is the maximum recorded temperature change over the US in one day? [Solution](#)
4. How many US businesses went bankrupt in 2014? [Solution](#)
5. How many cups of tea are drunk in the UK per day? [Solution](#)
6. What is the average weight of an adult giraffe? [Solution](#)
7. How many people are airborne over the US at any given moment? [Solution](#)
8. How much TV will the average British person watch in their lifetime? [Solution](#)
9. How many pennies would you need to stack on top of each other, until your stack was as high as the total annual rainfall in the UK (as of 2019)? [Solution](#)
10. If you filled an average Olympic swimming pool (of recommended depth) with Coca-Cola, how many kilograms of sugar would it contain? [Solution](#)

Solutions

1. 26% [Question](#)

2. 41 years (32 to 73) [Question](#)

3. 57.2°C, or 103°F [Question](#)

4. 28,319 [Question](#)

5. 100 million [Question](#)

6. 1192kg [Question](#)

7. 61,000 [Question](#)

8. Time watching TV per week: 30 hours
Average life expectancy: 81 years
Total time watching TV: $30 \times 52 \times 81 = 126,360$ hours (or about 14.4 years!) [Question](#)

9. Rainfall in 2019: 1,416.6mm
Thickness of penny: 1.52mm
Number of pennies: $1,416.6/1.52 = 932$ [Question](#)

10. Dimensions of pool: $50\text{m} \times 25\text{m} \times 3\text{m}$
Coke: 10.6g sugar per 100ml
Mass of sugar: $(25 \times 50 \times 3) \times (10.6 \times 10^{-3}) \times 10^4 = 397,500\text{kg}$ [Question](#)