

Chapter 2

Probabilistic Analysis I: Concentration Inequalities

We will study a central topic in probability theory and probabilistic analysis in this chapter: *concentration inequalities*. These are analytical tools that allow us to bound the probability of random variables deviating from their expected value. The reader may want to consult with [Appendix A](#) for a quick refresher on probabilistic background.

2.1 Balls and Bins Experiments

We shall use the following simple running example in this chapter to review “power” of different concentration inequalities. We are given a collection of n balls and n bins. We throw each ball *independently* to one of the bins chosen *uniformly at random*. We define the **load** of a bin as the number of balls thrown to that bin in this experiment.

Problem 2.1. What is the asymptotically maximum load of a bin with probability at least $2/3$, when throwing n balls independently and uniformly into n bins?

More precisely, find the smallest possible function $T(n)$, asymptotically, such that

$$\Pr(\text{the load of every bin is at most } T(n)) \geq \frac{2}{3}.$$

Let us first see how we can simplify our analysis. The original question involves a quantifier *for-all*—as in, *every* bin should have at most $T(n)$ balls. But, can we turn this into a *for-each* quantifier—for a *fixed* bin, say, bin one, upper bound the load? It is easy to see that if we simply bound the load of a fixed bin with constant probability, it cannot tell us anything meaningful about the maximum load across all bins¹. However, if we can bound the load of each bin with a much higher probability, then we can extend this to bound the maximum load as well. In particular, for a fixed choice of $k \in \mathbb{N}$, define the following events:

- E_i for $i \in [n]$: the event that the load of bin i is more than k ;
- $E_{\max}$: the event that the maximum load is more than k .

¹It is an easy calculation to see that with probability $(1 - 1/n)^n \approx 1/e$, the first bin receives zero balls but of course the maximum load is not zero!

We have

$$\Pr(E_{\max}) = \Pr\left(\bigvee_{i=1}^n E_i\right).$$

But we can now use a very simple—but extremely helpful—inequality, the **union bound**:

Fact 2.2. *For any two events A and B , $\Pr(A \vee B) \leq \Pr(A) + \Pr(B)$.*

So, applying union bound to the above terms iteratively gives us that

$$\Pr(E_{\max}) \leq \sum_{i=1}^n \Pr(E_i) = n \cdot \Pr(E_1);$$

in the last equality, we used the fact that the distribution of load of all bins is the same.

Consequently, if we define L to be the random variable for the load of the first bin—or any other fixed bin—then, to answer [Problem 2.1](#), we can instead find a function $T(n)$ such that

$$\Pr(L > T(n)) \leq \frac{1}{3n}. \quad (2.1)$$

Thus, for the rest of the proof, we will focus on solving this problem. In effect, we have now moved from a for-all quantifier in the original problem to a for-each quantifier, as desired.

Let us define indicator random variables $B_1, \dots, B_n$ where for any $i \in [n]$, $B_i = 1$ iff the i -th ball is sent to the first bin and is thus counted in L . As such, we have $L = \sum_{i=1}^n B_i$ and hence, using the linearity of expectation,

$$\mathbb{E}[L] = \mathbb{E}\left[\sum_{i=1}^n B_i\right] = \sum_{i=1}^n \mathbb{E}[B_i] = \sum_{i=1}^n \Pr(\text{ball } i \text{ is sent to bin one}) = \sum_{i=1}^n \frac{1}{n} = 1.$$

Hence, we know that in expectation, load of every bin is 1. But what can this tell us about [Eq \(2.1\)](#), namely, the probability that load of a bin is far away from its expectation? For this, we need to take a quick detour.

2.1.1 Concentration Inequalities

When working with random variables, perhaps the easiest way to “summarize” the variable is to focus on its expected value. This is for instance what we did in the previous chapter with the running time of our (Las Vegas) randomized algorithm.

However, expected value on its own can often be misleading: for instance, consider a random variable which is 0 with probability $1/2$ and is 1 with the remaining probability. Expected value of X is $1/2$ but of course we do not ‘expect’ X to ever take the value of $1/2$. This, and many other examples, suggest that summarizing a random variable just down to its expectation may lose too much information.

On the other extreme, a random variable on a domain Ω can be uniquely identified by its probability distribution:

$$\mathbb{P}_X : \Omega \rightarrow [0, 1] \quad \text{such that} \quad \mathbb{P}_X(a) = \Pr(X = a).$$

Yet, the distribution of even very simple random variables can be quite cumbersome to work with. Consider, for instance, the simple example of throwing a fair coin 100 times and defining X to be the number of heads. Here, for every $a \in \{0, \dots, 100\}$,

$$\Pr(X = a) = \binom{100}{a} \cdot 2^{-100},$$

which, even in this simple form, is not that easy to work with.

Concentration inequalities are a saving grace between these two extremes: morally speaking (but not strictly speaking true), they allow us to extract “important” information about the distribution of our random variables, without getting to compute the very precise distribution itself. More accurately, they allow us to bound the probability of *deviation* of a random variable from its expectation (as a function of its distance from the expectation).




Figure 2.1: An illustration of three different random variables and their probability distributions. Here, all these variables have the same expected value. Moreover, the first two variables show roughly the same concentration for the particular choice of t ; i.e., for the first two variables, the probability that their values are more than t away from the expectation is almost the same (the probability is the part shaded in red) even though the distributions are quite different; the third variable is much more concentrated.

We will see several concentration inequalities in this book as they arise quite frequently in the analysis of randomized algorithms (and way beyond). This chapter includes three of the most basic and highly applicable ones.

2.1.2 Concentration Inequalities for Balls and Bins

Let us now go back to our original balls and bins question. Equipped with the definition of concentration inequalities, it is easy to see that bounding Eq (2.1) is precisely the topic of concentration inequalities.

Markov Bound

The simplest and most basic variant of concentration results is **Markov’s inequality** or **Markov bound**.

Proposition 2.3 (Markov Bound). *For a non-negative random variable X and $t > 0$,*

$$\Pr(X \geq t \cdot \mathbb{E}[X]) \leq \frac{1}{t}.$$

Proof. Let $\mu := \mathbb{E}[X]$. We can use the law of total expectation to have:

$$\begin{aligned} \mathbb{E}[X] &= \mathbb{E}[X \mid X \geq t \cdot \mu] \cdot \Pr(X \geq t \cdot \mu) + \mathbb{E}[X \mid X < t \cdot \mu] \cdot \Pr(X < t \cdot \mu) \\ &\geq t \cdot \mu \cdot \Pr(X \geq t \cdot \mu) + 0. \end{aligned}$$

(the first term since we conditioned on $X \geq t \cdot \mu$ and the second term since X is non-negative)

Thus, $\Pr(X \geq t \cdot \mathbb{E}[X]) = \Pr(X \geq t \cdot \mu) \leq 1/t$, otherwise the RHS above will be larger than the LHS. $\square$

Markov bound only bounds the *upper tail* of the distribution²: the probability that a random variable takes value t times larger than its expectation is at most $1/t$. This is a basic but

²Although one can use it to bound the lower tail in special cases as well, e.g., for $X \in [n]$ to bound $\Pr(X < t)$, we can instead define $Y := n - X$ and bound $\Pr(Y > n - t)$ which can be done using Markov bound. However, such bounds are generally quite weak.

extremely useful property. One can alternatively state the Markov bound as follows (by simply picking $t = b/\mathbb{E}[X]$ in [Proposition 2.3](#)).

Corollary 2.4. *For a non-negative random variable X and $b > 0$,*

$$\Pr(X \geq b) \leq \frac{\mathbb{E}[X]}{b}.$$

Note that a random variable X , which, with probability $\mathbb{E}[X]/b$ takes the value b and otherwise is 0 will be a tight example for Markov bound; i.e., one cannot expect to improve Markov bound in general.

Remark. Even though Markov bound may sound almost trivial (and it is straightforward indeed), it is the basis for proving all other concentration inequalities that we use in this course. Moreover, Markov bound is used one way or another in analysis of almost every randomized algorithm.

For our application to the balls and bins experiment, we can thus have for any $k \geq 1$,

$$\Pr(L \geq k) \leq \frac{\mathbb{E}[L]}{k} = \frac{1}{k},$$

and thus to bound this probability with $1/3n$, we need to take $k = 3n$. But of course this is entirely useless because we only have n balls to begin with and so $L \leq n$ happens with probability 1 already!

Chebyshev's Inequality

We consider our second concentration inequality: **Chebyshev's inequality**. Unlike Markov bound that only required a knowledge of the expected value of the random variable to bound its deviation probability, Chebyshev's inequality applies to the settings in which we can additionally bound the *variance* of the random variable as well.³

Recall that for a random variable X , **variance** of X is:

$$\text{Var}[X] := \mathbb{E}[(X - \mathbb{E}[X])^2] = \mathbb{E}[X^2] - (\mathbb{E}[X])^2.$$

Chebyshev's inequality allows us to bound deviation of a random variable based on its variance.

Proposition 2.5 (Chebyshev's Inequality). *For any random variable X and $t > 0$,*

$$\Pr(|X - \mathbb{E}[X]| \geq t \cdot \mathbb{E}[X]) \leq \frac{\text{Var}[X]}{\mathbb{E}[X]^2 \cdot t^2}.$$

Proof. Define a new random variable $Y := (X - \mathbb{E}[X])^2$. Clearly, Y is non-negative. Moreover, $|X - \mathbb{E}[X]| \geq t \cdot \mathbb{E}[X]$ if and only if $Y = (X - \mathbb{E}[X])^2 \geq t^2 \cdot \mathbb{E}[X]^2$. Hence,

$$\Pr(|X - \mathbb{E}[X]| \geq t \cdot \mathbb{E}[X]) = \Pr(Y \geq t^2 \cdot \mathbb{E}[X]^2) \leq \frac{\mathbb{E}[Y]}{\mathbb{E}[X]^2 \cdot t^2}$$

(by Markov bound of [Corollary 2.4](#))

³As we showed earlier, Markov bound *can* be tight for certain random variables; thus, naturally whenever we need a stronger bound we should show that our variable satisfies additional guarantees than what is only required by Markov bound.

$$= \frac{\text{Var}[X]}{\mathbb{E}[X]^2 \cdot t^2}.$$

(as $\mathbb{E}[Y] = \mathbb{E}[(X - \mathbb{E}[X])^2] = \text{Var}[X]$ by definition)

□

A useful variant of Chebyshev's inequality is the following (by simply picking $t = b/\mathbb{E}[X]$ in [Proposition 2.5](#)).

Corollary 2.6. *For a random variable X and $b > 0$,*

$$\Pr(|X - \mathbb{E}[X]| \geq b) \leq \frac{\text{Var}[X]}{b^2}.$$

For our balls and bins experiment, to apply Chebyshev's inequality, we need to bound $\text{Var}[L]$. We have,

$$\text{Var}[L] = \text{Var}\left[\sum_{i=1}^n B_i\right] = \sum_{i=1}^n \text{Var}[B_i] \leq \sum_{i=1}^n \mathbb{E}[B_i] = \mathbb{E}[L] = 1,$$

where in the second equality, we used the fact that $B_1, \dots, B_n$ are *independent* and hence variance of their sum is equal to their sum of variances, as proven in the following [fact](#)⁴.

Fact 2.7. *For any two independent random variables X, Y :*

$$\text{Var}[X + Y] = \text{Var}[X] + \text{Var}[Y].$$

Proof. We have,

$$\begin{aligned} \text{Var}[X + Y] &= \mathbb{E}[(X + Y)^2] - (\mathbb{E}[X + Y])^2 \\ &= \mathbb{E}[X^2] + \mathbb{E}[Y^2] + 2 \cdot \mathbb{E}[X \cdot Y] - (\mathbb{E}[X])^2 - (\mathbb{E}[Y])^2 - 2 \cdot \mathbb{E}[X] \cdot \mathbb{E}[Y] \\ &\quad \text{(by expanding the terms)} \\ &= \mathbb{E}[X^2] - (\mathbb{E}[X])^2 + \mathbb{E}[Y^2] - (\mathbb{E}[Y])^2 \\ &\quad \text{(as } \mathbb{E}[X \cdot Y] = \mathbb{E}[X] \cdot \mathbb{E}[Y] \text{ when } X \perp Y) \\ &= \text{Var}[X] + \text{Var}[Y]. \end{aligned}$$

□

Thus,

$$\Pr(L \geq k) \leq \Pr(|L - \mathbb{E}[L]| \geq k - 1) \leq \frac{\text{Var}[L]}{(k - 1)^2} = \frac{1}{(k - 1)^2}.$$

This time, we will get something non-trivial: by setting $k = \sqrt{3n} + 1$, we have

$$\Pr(L \geq k) \leq \frac{1}{3n},$$

which is the type of the bound we wanted. In other words, using Chebyshev's inequality allows us to bound $T(n) = \sqrt{3n} + 1 = O(\sqrt{n})$. While this is non-trivial (unlike what we had with Markov's bound), as we shall see this is still very far from the *right* bound for the problem.

⁴Recall that in general, we do not have linearity of variance, only linearity of expectation.

Chernoff Bound

We now switch to one of the strongest concentration inequalities that we will encounter in this book, the **Chernoff bound**. To be able to use Chernoff bound for bounding deviation of a random variable, we need a much stronger knowledge than variance in Chebyshev's inequality and expectation in Markov bound; we now need to know that the random variable X is a *sum of bounded-value independent random variables*.

Proposition 2.8 (Chernoff Bound). *Suppose $X_1, \dots, X_n$ are independent random variables in $[0, 1]$ and $X = \sum_i X_i$. Then, for any $t \geq 1$,*

$$\Pr(|X - \mathbb{E}[X]| \geq t \cdot \mathbb{E}[X]) \leq 2 \cdot \exp\left(-\frac{t \cdot \mathbb{E}[X]}{3}\right),$$

and, for any $\varepsilon \in (0, 1]$,

$$\Pr(|X - \mathbb{E}[X]| \geq \varepsilon \cdot \mathbb{E}[X]) \leq 2 \cdot \exp\left(-\frac{\varepsilon^2 \cdot \mathbb{E}[X]}{3}\right).$$

Proposition 2.8 is often called the *multiplicative* Chernoff bound as opposed to the *additive* Chernoff bound as follows⁵. The additive Chernoff bound is weaker than the multiplicative one but sometimes it can be easier to work with.

Proposition 2.9 (Additive Chernoff Bound). *Suppose $X_1, \dots, X_n$ are independent random variables in $[0, 1]$ and $X = \sum_i X_i$. Then, for any $b \geq 1$,*

$$\Pr(|X - \mathbb{E}[X]| \geq b) \leq 2 \cdot \exp\left(-\frac{2b^2}{n}\right).$$

Proving **Proposition 2.8** or **Proposition 2.9** is not as easy as those of Markov bound or Chebyshev's inequality, but it is also not particularly hard. For the interested reader (and to showcase the main ideas), we present a proof of a simpler form in **Section 2.3**.

Let us now apply Chernoff bound to our problem. We have that $L = B_1 + \dots + B_n$, each $B_i \in [0, 1]$, and B_i 's are independent of each other, as each ball is thrown independently. So, we can indeed apply Chernoff bound to L and obtain that

$$\Pr(L \geq k) \leq \Pr(|L - \mathbb{E}[L]| \geq k - 1) \leq 2 \cdot \exp\left(-\frac{(k - 1) \cdot \mathbb{E}[L]}{3}\right) = 2 \cdot \exp\left(-\frac{(k - 1)}{3}\right).$$

By setting $k = 3 \cdot \ln(6n) + 1$, we have,

$$\Pr(L \geq k) \leq 2 \cdot \exp(-\ln(6n)) = 2 \cdot \frac{1}{6n} = \frac{1}{3n}.$$

Thus, Chernoff bound allows us to set $T(n) = 3 \cdot \ln(6n) + 1 = O(\log n)$, exponentially stronger than the bound obtained via Chebyshev's inequality.

2.1.3 Optimal Bounds for Balls and Bins?

The bound of $O(\log n)$ on maximum load obtained via Chernoff bound is quite close to the optimal solution but it is *not* the right answer still. While we can use a stronger version of

⁵There is also a more general version of Chernoff bound which is rather unwieldy to use in practice (at least initially), but is a bit stronger than the above bounds (especially for “small” values of $\mathbb{E}[X]$); however, the above bounds are quite more convenient for us to use and as such we will ignore the stronger bound for now.

Chernoff bound to obtain the optimal bound via the use of a “generic” concentration inequality, we are going to instead use a more direct approach: by calculating the probability of the event explicitly ourselves (as somewhat of a last resort).

We have,

$$\Pr(L \geq k) = \sum_{b=k}^n \Pr(L = b) = \sum_{b=k}^n \binom{n}{b} \cdot \left(\frac{1}{n}\right)^b \cdot \left(1 - \frac{1}{n}\right)^{n-b};$$

the RHS is because for $L = b$, we need to have a set of b different B_i 's (chosen from $\binom{n}{b}$ choices), all set to 1 (which happens with probability n^{-b}), and all the remaining B_i 's set to 0 (which happens with probability $(1 - 1/n)^{n-b}$).

Notice that in this approach, we already have a rather unwieldy calculation to work with. So, even in this case, it would help to use some inequalities. In particular, we are going to calculate the above probability rather differently by upper bounding it first.

$$\Pr(L \geq k) = \Pr(\text{there exists some } k \text{ balls that are mapped to the bin})$$

(a sufficient and necessary condition for $L \geq k$ is that some k balls are mapped to this bin)

$$\leq \sum_{\text{balls } a_1, \dots, a_k} \Pr(a_1, \dots, a_k \text{ are mapped to the bin})$$

(by union bound where the sum ranges over all possible choices of picking k balls out of n)

$$\begin{aligned} &= \sum_{\text{balls } a_1, \dots, a_k} \left(\frac{1}{n}\right)^k \\ &\quad \text{(as the probability of a fixed } k \text{ balls mapping to the bin is } n^{-k}) \\ &= \binom{n}{k} \cdot \left(\frac{1}{n}\right)^k. \quad \text{(the number of ways of picking } k \text{ balls out of } n \text{ is } \binom{n}{k}) \end{aligned}$$

So far, in these calculations, we already used one inequality where we applied the union bound. We are going to apply yet another one now⁶.

Fact 2.10. For any integers $a \geq b$,

$$\left(\frac{a}{b}\right)^b \leq \binom{a}{b} \leq \left(\frac{e \cdot a}{b}\right)^b,$$

where $e \approx 2.73\dots$ is the Euler's number (the base of the natural logarithm).

Proof. We have,

$$\binom{a}{b} = \frac{a!}{b! \cdot (a-b)!} = \frac{a}{b} \cdot \underbrace{\frac{(a-1)}{(b-1)} \cdot \frac{(a-2)}{(b-2)} \dots \frac{(a-b+1)}{1}}_{b \text{ terms in total}} \geq \left(\frac{a}{b}\right)^b,$$

because each of the b terms is at least a/b . This proves the LHS. For the RHS, we have,

$$\binom{a}{b} \leq \frac{a^b}{b!},$$

⁶The reason we are pointing out the inequalities explicitly is that we need to be extra careful with them if our goal is to compute an asymptotically optimal bound; any loose inequality may negate all our effort of calculating the probability explicitly.

because each of the terms in the numerator of the first equation is at most a . Furthermore, by writing the Taylor expansion of e^b , we have that

$$e^b = \sum_{k=0}^{\infty} \frac{b^k}{k!} \geq \frac{b^b}{b!}. \quad (\text{as this is just one term, } k = b, \text{ in the whole series})$$

Plugging in $b! \geq (b/e)^b$ in the above inequality, we get the RHS as well. $\square$

Continuing our calculations from before, we now have,

$$\begin{aligned} \Pr(L \geq k) &\leq \binom{n}{k} \cdot \frac{1}{n^k} && (\text{as calculated above}) \\ &\leq \left(\frac{e \cdot n}{k}\right)^k \cdot \frac{1}{n^k} = \left(\frac{e}{k}\right)^k && (\text{by Fact 2.10}) \\ &= \frac{e^k}{e^{k \cdot \ln k}} && (\ln k \text{ is the log of } k \text{ in base } e) \\ &= \exp(-k \cdot (\ln k - 1)) = \exp(-k \cdot \ln(k/e)). \end{aligned}$$

We can now pick $k = \frac{e \cdot \ln n}{\ln \ln n}$ to have,

$$\begin{aligned} k \cdot \ln(k/e) &= \left(\frac{e \cdot \ln n}{\ln \ln n}\right) \cdot \ln\left(\frac{\ln n}{\ln \ln n}\right) \\ &= \left(\frac{e \cdot \ln n}{\ln \ln n}\right) \cdot (\ln \ln n - \ln \ln \ln n) \\ &\geq \left(\frac{e \cdot \ln n}{\ln \ln n}\right) \cdot \frac{1}{2} \cdot \ln \ln n && (\text{as for sufficiently large } n, \ln \ln n > 2 \cdot \ln \ln \ln n) \\ &= \frac{e}{2} \cdot \ln n \geq \ln(3n). && (\text{for sufficiently large } n) \end{aligned}$$

Plugging in all these—honestly, rather tedious⁷—calculations back into our bounds, we get that

$$\Pr\left(L \geq \frac{e \cdot \ln n}{\ln \ln n}\right) \leq \frac{1}{3n},$$

for sufficiently large n . This allows us to say $T(n) = O(\log n / \log \log n)$ also holds for our problem in Eq (2.1).

Before moving on, let us also mention that the bound of $O(\log n / \log \log n)$ we proved for Problem 2.1 is also asymptotically optimal. The proof is simply as follows:

$$\begin{aligned} \Pr(L \geq k) &\geq \Pr(L = k) \\ &= \binom{n}{k} \cdot \frac{1}{n^k} \cdot \left(1 - \frac{1}{n}\right)^{n-k} \\ &\geq \binom{n}{k} \cdot \frac{1}{n^k} \cdot \frac{1}{4} && (\text{as } (1 - 1/n)^m \geq (1 - 1/n)^n \geq 1/4 \text{ for all } m \leq n \text{ and } n \geq 2) \\ &\geq \left(\frac{n}{k}\right)^k \cdot \frac{1}{n^k} \cdot \frac{1}{4} = \frac{1}{4} \cdot \frac{1}{k^k}. && (\text{by Fact 2.10}) \end{aligned}$$

Thus, if we set $k = \frac{1}{2} \cdot \frac{\ln n}{\ln \ln n}$, we get that $\Pr(L \geq k) \geq \frac{1}{4 \cdot \sqrt{n}}$; this means that $T(n)$ is also $\Omega(\log n / \log \log n)$, and hence the correct answer asymptotically is $\Theta(\log n / \log \log n)$.

⁷Well, this was kind of the whole point of showing how generic concentration inequalities are much easier to work with than an ad-hoc argument.

Remark. Think of working directly with the probability distribution of a random variable as the “ultimate solution” – technically speaking, this approach *always* gives you the optimal answer as you will be calculating the exact probability. But in most cases, these probabilities are just too cumbersome or downright impossible to calculate, hence the need for more general concentration inequalities.

Markov, Chebyshev, and Chernoff bounds are among the most used concentration inequalities in the analysis of the algorithms. But, they are only the tip of the iceberg among the vast array of concentration inequalities known. See the excellent book of

Devdatt P. Dubhashi and Alessandro Panconesi, Concentration of Measure for the Analysis of Randomized Algorithms.

that provides a more detailed overview of concentration inequalities in this context.

2.2 Success Amplification of Randomized Algorithms

Let us see an application of the bounds we saw in this chapter: *boosting* the probability of success of a randomized algorithm to any desired bound (by repeating it a certain number of times in parallel).

We now formalize the problem.

Problem 2.11. Suppose A is a randomized algorithm for some problem P with probability of success at least $2/3$. The problem P can be either a decision problem which has a *Yes/No* answer, or an estimation problem which requires outputting a number in a certain range. How can we boost its probability of success to $1 - \delta$ for a given $\delta \in (0, 1)$ by running A repeatedly?

Majority trick for decision problems: If P is a decision problem and thus A returns a *Yes/No* answer, we can do the following: run A independently for $k := 18 \ln(2/\delta)$ times and return the *majority* answer.

The analysis is as follows. Let $X_1, \dots, X_k$ be k indicator random variables where $X_i = 1$ iff the i -th run of A returns the correct answer. Let $X := \sum_{i=1}^k X_i$. Firstly, we have

$$\mathbb{E}[X] = \sum_{i=1}^k \mathbb{E}[X_i] \geq \sum_{i=1}^k \frac{2}{3} = k \cdot \frac{2}{3}$$

as each $X_i = 1$ with probability at least $2/3$. Secondly, for the majority answer to be wrong, we need $X \leq k/2$. Thus,

$$\begin{aligned} \Pr(\text{majority answer is wrong}) &\leq \Pr(|X - \mathbb{E}[X]| \geq k/6) && (\text{as } 2k/3 - k/2 = k/6) \\ &\leq 2 \cdot \exp\left(-\frac{2 \cdot k^2}{36 \cdot k}\right) \end{aligned}$$

by [Proposition 2.9](#) for $b = k/6$ and $n = k$ and since X is a sum of independent random variables in $[0, 1]$. Continuing, we have,

$$= 2 \cdot \exp(-\ln(2/\delta)) = 2 \cdot (\delta/2) = \delta.$$

Median trick for estimation problems: What if P was an estimation problem and thus A returns a number that is in the correct range $[x : y]$ with probability at least $2/3$? In this case, returning the majority will not work because it is possible that none of the numbers returned by the different copies of A are even the same. The solution however is almost the same with a simple tweak: run A independently for $k := 18 \ln(2/\delta)$ times and return the *median* answer⁸.

The analysis is also similar to before. Let $X_1, \dots, X_k$ be k indicator random variables where $X_i = 1$ iff the i -th run of A returns an answer in the correct range $[x : y]$. Let $X := \sum_{i=1}^k X_i$. We again have that $\mathbb{E}[X] \geq k \cdot 2/3$. Moreover, for the median to be a wrong answer, *at least* half of X_i 's have to be zero (this is a necessary condition for failing but is not sufficient). So,

$$\Pr(\text{median answer is wrong}) \leq \Pr(|X - \mathbb{E}[X]| \geq k/6) \leq \delta,$$

where the rest of the calculation is exactly as done for the majority trick.

Remark. It is worth emphasizing that both the majority trick and the median trick completely treated the algorithm (and analysis) of their input algorithm in a black-box way. One can use these tricks to boost the probability of success of *any* algorithm (in many different settings) – this is indeed the reason that for most algorithms, the dependence of resources on δ is almost always an $O(\ln(1/\delta))$ factor^a.

^aWe emphasize that this is O -notation and *not* Θ ; there are a good number of cases also that one can do considerably better than this bound also using more sophisticated (and often ad-hoc) approaches.

2.3 Proof of (a Special Case of) Chernoff Bound

We are *not* going to give the entire proof of the Chernoff bound as it is beyond the scope of this book. However, to provide enough intuition, we will prove a simpler variant of Chernoff (which is actually sufficient for our balls in bins argument and many other settings as well).

Let us assume that each X_i is a Bernoulli random variable⁹ with mean p_i (instead of arbitrary random variable in $[0, 1]$). For simplicity, we are also going to only prove the *upper tail* of the deviation bound instead of both tails (but the lower tail can be proven symmetrically). In particular, we prove the following result.

Proposition 2.12 (A weaker form of Chernoff bound). *Let $X_1, \dots, X_n$ be independent Bernoulli random variables in $\{0, 1\}$ with mean $p_1, \dots, p_n$ and $X = \sum_i X_i$. Then, for any $\varepsilon > 0$,*

$$\Pr(X \geq (1 + \varepsilon) \cdot \mathbb{E}[X]) \leq \exp\left(-\frac{\varepsilon^2}{2 + \varepsilon} \cdot \mathbb{E}[X]\right).$$

Notice that in this result, ε is not limited to be less than 1 and thus we can recover the bounds of [Proposition 2.8](#) by considering $\varepsilon < 1$ and $\varepsilon \geq 1$ in the RHS separately.

Proof. Fix $\alpha > 0$. Note that $X \geq (1 + \varepsilon) \mathbb{E}[X]$ if and only if

$$\exp(\alpha \cdot X) \geq \exp(\alpha \cdot (1 + \varepsilon) \cdot \mathbb{E}[X]).$$

Define a random variable $Y = \exp(\alpha \cdot X)$. Using Markov bound on random variable Y ,

$$\Pr(X \geq (1 + \varepsilon) \cdot \mathbb{E}[X]) = \Pr(Y \geq \exp(\alpha \cdot (1 + \varepsilon) \cdot \mathbb{E}[X])) \leq \frac{\mathbb{E}[Y]}{\exp(\alpha \cdot (1 + \varepsilon) \cdot \mathbb{E}[X])}. \quad (2.2)$$

⁸You are strongly encouraged to think about why returning the *average* fails (in certain cases miserably).

⁹Recall that a Bernoulli random variable Z with mean p gets value 1 w.p. p and 0 w.p. $1 - p$.

As such, to bound the probability of deviation of X from $\mathbb{E}[X]$, we only need to bound $\mathbb{E}[Y]$ and then we can apply [Eq \(2.2\)](#). We can now upper bound $\mathbb{E}[Y]$ as follows:

$$\begin{aligned}
\mathbb{E}[Y] &= \mathbb{E}[\exp(\alpha \cdot X)] = \mathbb{E}\left[\exp\left(\alpha \cdot \sum_{i=1}^n X_i\right)\right] \\
&= \mathbb{E}\left[\prod_{i=1}^n \exp(\alpha \cdot X_i)\right] \\
&= \prod_{i=1}^n \mathbb{E}[\exp(\alpha \cdot X_i)] \quad (\text{for independent random variables } A, B: \mathbb{E}[AB] = \mathbb{E}[A] \cdot \mathbb{E}[B]) \\
&= \prod_{i=1}^n (1 - p_i + p_i \cdot e^\alpha) \quad (\text{by the assumption that } X_i \text{ is Bernoulli with mean } p_i) \\
&= \prod_{i=1}^n (1 + p_i \cdot (e^\alpha - 1)) \leq \prod_{i=1}^n \exp(p_i \cdot (e^\alpha - 1)) \quad (1 + x \leq e^x \text{ for all } x) \\
&= \exp\left((e^\alpha - 1) \cdot \sum_{i=1}^n p_i\right) = \exp((e^\alpha - 1) \cdot \mathbb{E}[X]).
\end{aligned}$$

Let us now set $\alpha = \ln(1 + \varepsilon)$ and use [Eq \(2.2\)](#) to obtain that:

$$\begin{aligned}
\Pr(X \geq (1 + \varepsilon) \cdot \mathbb{E}[X]) &\leq \exp((e^\alpha - 1) \cdot \mathbb{E}[X] - \alpha \cdot (1 + \varepsilon) \cdot \mathbb{E}[X]) \\
&= \exp(\varepsilon \cdot \mathbb{E}[X] - (1 + \varepsilon) \cdot \ln(1 + \varepsilon) \cdot \mathbb{E}[X]) \\
&= \left(\frac{e^\varepsilon}{(1 + \varepsilon)^{(1 + \varepsilon)}}\right)^{\mathbb{E}[X]}. \tag{2.3}
\end{aligned}$$

The bound above is already a *very strong form* of Chernoff bound (for Bernoulli random variables)¹⁰ – note also that in this equation, we do *not* need ε to be in $(0, 1)$. We can simplify this bound to get the bounds we want in the statement of the proposition (this will weaken the bound slightly; very rarely this weakening can be problematic and we may need to use the stronger bound above directly).

We are going to use this inequality (the proof is omitted) to simplify [Eq \(2.3\)](#): For $x > 0$,

$$1 + x \geq \frac{e^x}{1 + x/2}.$$

By applying this, we have (the proof is again omitted),

$$\left(\frac{e^\varepsilon}{(1 + \varepsilon)^{(1 + \varepsilon)}}\right) = \exp(\varepsilon - (1 + \varepsilon) \cdot \ln(1 + \varepsilon)) \leq \exp\left(-\frac{\varepsilon^2}{2 + \varepsilon}\right).$$

And thus by [Eq \(2.3\)](#),

$$\Pr(X \geq (1 + \varepsilon) \cdot \mathbb{E}[X]) \leq \exp\left(-\frac{\varepsilon^2}{2 + \varepsilon} \cdot \mathbb{E}[X]\right).$$

This concludes the proof. □

¹⁰And we could have recovered the bound of $T(n) = O(\log n / \log \log n)$ for balls-and-bins experiments by using this stronger Chernoff bound; you are encouraged to try this on your own.