Lecture Notes

Gambling, Markov Chains, Martingales and h-transform
Stochastic Processes Winter 2022/23
University of Potsdam

Work in progress – use with care

Dr. Peter Keller

Last Update: October 13, 2023

Preface

These lecture notes contain a version of the topics taught in the first half of the lecture "stochastic processes" held at the University of Potsdam in teh winter of 2022.

We motivate the theory of stochastic processes using gambling examples, especially the well-known *Gambler's Ruin Problem*. We continue the discussion with conditional expectation in the version for discrete random variables and use this to introduce basics of martingale theory. The following remarks on Markov Chains are then connected with the martingale theory by the Lévy-Martingale and usage of Doob's h-transform utilising the Markov Operator. In this framework we discuss the maximum principle and the related Dirichlet-problem to characterize basic properties of absorption problems in Markov Chains.

We assume throughout the lecture basics of probability theory with a bit of measure theory. The audience for this course are advanced Bachelor students or Master students of mathematics.

During the lecture we had several exercises. Some of them made it into the script.

I do thank *Julian Kern* for finding many interesting exercises and working with the students. Thanks goes also to *Jessiva Havemann* correcting the exercises.

Main resources: [Eth10], [Bré20], [Pri18] and [Wil91]. Some notes were taken form the computer science course script [LLM10] as taught in 2010 at MIT.

Contents

1 Notation and short recap

We assume throughout that the reader is familiar with the basics of measure theory. We will need measure theory only in the case of discrete random variables. However, by skipping some of the proofs, the contents can be understood and applied through many examples. After each section, some exercises are provided.

Let Ω be the event space and assume that Ω is either finite or countable. The sigma-Algebra $\mathcal{A}$ is the event algebra describing all observable events and is subset of the power set of Ω, i.e. $\mathcal{A} \subset 2^{\Omega}$. If Ω is finite we can set $\mathcal{A} = 2^{\Omega}$ in most relevant cases.

Let $\mathbb{P}$ be a probability measure with the usual properties given by Kolmogoroff's axioms:

- $\mathbb{P} : \mathcal{A} \to [0, 1]$
- $\mathbb{P}(\Omega) = 1$
- For any sequence of events $(B_n)_{n \geq 0}$, with $B_i \neq B_j$, $i \neq j$,

$$\mathbb{P}\Big(\bigcup_n B_n\Big) = \sum_n \mathbb{P}(B_n).$$

A (numeric) *random variable* X is a measurable mapping $X : \mathcal{A} \to \mathbb{R}$ and we can assign probabilities to subsets of the image of X by $\mathbb{P}_X = \mathbb{P} \circ X^{-1}$. In many cases, we will drop the index. In the case of discrete random variables, we sometimes call $\mathbb{P}_X(\{x_k\}) = \mathbb{P}(X = x_k) =: p_k$ the *probability mass function* or *counting measure* of X.

For a given random variable X, we denote with $\sigma(X)$ the smallest σ-algebra to which X is measurable. In particular, the only zero set in $\sigma(X)$ is the empty set $\emptyset$.

The expectation of a discrete random variable is defined as usual:

$$\mathbb{E}[X] = \sum_{k \in X(\Omega)} k \, \mathbb{P}(X = k).$$

In the discrete case, the expectation of a random variable X concatenated with a measurable function g can be computed using the LOTUS[1] principle:

$$\mathbb{E}[g(X)] := \sum_{k \in X(\Omega)} g(k) \, \mathbb{P}(X = k).$$

The notation $X(\Omega)$ is referring to the image of X. In many cases $X(\Omega) = \{1, 2, \ldots, n\}$ for some $n \in \mathbb{N}$.

[1] Law Of The Unthinking Statistician.

2 Discrete Conditional Expectation

Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space. Let X be a random variable on this space. Then we can define the *conditional probability* with respect to an event $A \in \mathcal{A}$ with $\mathbb{P}(A) \neq 0$ via

$$\mathbb{P}(X = k \mid A) := \frac{\mathbb{P}(\{X = k\} \cap A)}{\mathbb{P}(A)}.$$

The event $A \subset \Omega$ acts in that case as a new event space, where possible observations are restricted to observations in A. One checks the Kolmogoroff axioms, to show that $\mathbb{P}(X = . \mid A)$ is a probability measure for any fixed $A \in \mathcal{A}$ with $\mathbb{P}(A) > 0$. If A is an impossible event, i.e. $\mathbb{P}(A) = 0$, we can set the conditional probability to zero.

With this new measure, it is possible to define the expectation of X conditional to A, when we have only partial information. In that sense, $(A, \sigma(A), \mathbb{P}(. \mid A)$ defines a new discrete probability space. The operator σ applied to a set returns the smallest sub-σ-algebra of $\mathcal{A}$ that contains A, i.e. $\sigma(A) \subset \mathcal{A}$.

Definition 2.1 (conditional expectation)

For an event $A \in \mathcal{A}$ and a discrete random variable X defined on $(\Omega, \mathcal{A}, \mathbb{P})$, the *conditional expectation* of X given A is

$$\mathbb{E}[X \mid A] := \sum_{k \in X(\Omega)} k\, \mathbb{P}(X = k \mid A).$$

As an example, we throw a six sided fair die once. Then the corresponding stochastic model is $\Omega = \{1, 2, 3, 4, 5, 6\}$, $\mathcal{A} = 2^{\Omega}$ and $\mathbb{P}$ is the uniform measure on Ω. Thus for any $A \in \mathcal{A}$

$$\mathbb{P}(A) = \frac{\#A}{\#\Omega} = \frac{\#A}{6}.$$

Let $X : \Omega \to \Omega$ be the result of a throw, i.e. $X(\omega) = \omega$. The expectation of X is

$$\mathbb{E}[X] = \sum_{k \in X(\Omega)} k\, \mathbb{P}(X = k) = \sum_{k=1}^{6} k\, \mathbb{P}(\{k\}) = \frac{1}{6}\,(1 + 2 + 3 + 4 + 5 + 6) = \frac{7}{2}.$$

Therefore, if we repeat the experiment of throwing a die multiple times, we can expect that the average of the outcomes is about 3.5.

Suppose a friend is reporting the result of throwing the die to you and you can not see the original result. But s/he only reports whether the die showed an even number or not. The information available to you is now contained in the following σ-algebra

$$\mathcal{A} = \{\emptyset, \{1, 3, 5\}, \{2, 4, 6\}, \Omega\}.$$

Let now $A \in \mathcal{A}$ be the event where your friend reported an even number, then $A = \{2, 4, 6\}$. The probability of A in the chosen stochastic model is $\mathbb{P}(A) = 3/6 = 1/2$.

How does the knowledge of this event change the observed, expected value of X? We get by direct computation using definition 2.1:

$$\mathbb{E}[X \mid A] = \sum_{k\in\Omega} k\,\mathbb{P}(X = k \mid A) = \sum_{k\in\{2,4,6\}} k\,\frac{\mathbb{P}(X = k)}{\mathbb{P}(A)} = 2\,\frac{1}{6}(2+4+6) = 4.$$

This expectation is now higher than the "unconditional" expectation $\mathbb{E}[X] = 3.5$.[2] This is intuitive, since we, in a way, censor the outcomes that are odd. With an analogue computation we can condition the expectation on the event that the die showed an uneven number, thus define $B := \{1,3,5\} = A^c$. The probability of B is also $1/2$ and the conditional expectation of X given B is $\mathbb{E}[X \mid B] = 3$.

We notice that in our case $A \cap B = \emptyset$ and $A \cup B = \Omega$ and

$$\frac{7}{2} = \mathbb{E}[X] = \mathbb{E}[X \mid A]\,\mathbb{P}(A) + \mathbb{E}[X \mid B]\,\mathbb{P}(B) = 4 \cdot \frac{1}{2} + 3 \cdot \frac{1}{2}.$$

This is no coincidence, as the following theorem shows. In fact, this is analogue to the law of total probability.

Theorem 2.1 (Law Of Total Expectation)

Let X be a discrete, integrable random variable on the probability space $(\Omega, \mathcal{A}, \mathbb{P})$. Let $(B_i)_{i\in\mathbb{N}}$ a sequence of pairwise disjoint events in $\mathcal{A}$ that form a partition of Ω and assume $\mathbb{E}[X \mid B_i]$ exist for all $i \in \mathbb{N}$. Then the *law of total expectation* holds:

$$\mathbb{E}[X] = \sum_{i=1}^{\infty} \mathbb{E}[X|B_i]\,\mathbb{P}(B_i).$$

Proof. We first rewrite the conditional expectation with respect to an arbitrary event $A \in \mathcal{A}, A \neq \emptyset$:

$$\mathbb{E}[X \mid A] = \sum_{k\in X(A)} k\,\frac{\mathbb{P}(\{X = k\} \cap A)}{\mathbb{P}(A)} = \sum_{k\in\Omega} k\,\mathbb{1}_A \frac{\mathbb{P}(X = k)}{\mathbb{P}(A)} = \frac{1}{\mathbb{P}(A)}\mathbb{E}[X\mathbb{1}_A]. \tag{2.1}$$

Then

$$\sum_{i=1}^{n} \mathbb{E}[X \mid B_i]\mathbb{P}(B_i) = \sum_{i=1}^{n} \mathbb{E}[X\mathbb{1}_{B_i}] = \mathbb{E}[X\mathbb{1}_{\bigcup_{i=1}^n B_i}].$$

We notice $X = \lim_{n\to\infty} X\mathbb{1}_{\bigcup_{i=1}^n B_i}$ almost surely and $X\mathbb{1}_{\bigcup_{i=1}^n B_i} \leq X$. By the dominated convergence theorem, we have

$$\mathbb{E}[X] = \lim_{n\to\infty} \mathbb{E}[X\mathbb{1}_{\bigcup_{i=1}^n B_i}],$$

which concludes the proof. □

[2] We can think of $\mathbb{E}[X]$ as expectation conditional to the whole event space, i.e. $\mathbb{E}[X \mid \Omega]$.

Equation 2.2 will be useful in the following, as it allows to express the conditional expectation in terms of the ordinary expectation. We could have used this as the definition, but it is also useful to give a definition with regards to the conditional probability.

We continue the previous example about throwing a fair six sided die. Let us assume the existence of another random variable Y that gives information of whether or not the die showed a number strictly larger than three. Y is clearly the indicator variable of the event $\{4, 5, 6\}$ and is one with probability $1/2$. We also have

$$\begin{aligned}
&\{Y = 1\} = \{\omega \in \{1, 2, 3, 4, 5, 6\} : Y(\omega) = 1\} = \{X \in \{4, 5, 6\}\} = \{4, 5, 6\}, \\
&\{Y = 0\} = \{Y = 1\}^c = \{1, 2, 3\}.
\end{aligned}$$

We can now compute the expectation of X conditional to the two outcomes of Y using equation 2.2:

$$\begin{aligned}
\mathbb{E}[X \,|\, Y = 1] &= \frac{1}{\mathbb{P}(Y = 1)} \mathbb{E}[X \mathbb{1}_{\{Y=1\}}] = 2 \cdot \frac{4 + 5 + 6}{6} = 5, \\
\mathbb{E}[X \,|\, Y = 0] &= 2 \cdot \frac{1 + 2 + 3}{6} = 2.
\end{aligned}$$

Naturally, Theorem 2.1 is fulfilled, since

$$\mathbb{E}[X \,|\, Y = 0]\mathbb{P}(Y = 0) + \mathbb{E}[X \,|\, Y = 1]\mathbb{P}(Y = 1) = \frac{7}{2} = \mathbb{E}[X].$$

We formalise the previous in the following definition.

Definition 2.2 (Conditional Expectation of X given Y)

Let X and Y be random variables on the probability space $(\Omega, \mathcal{A}, \mathbb{P})$. Then we define $\mathbb{E}[X \,|\, Y] : \sigma(Y) \to \mathbb{R}$ by

$$\mathbb{E}[X \,|\, Y](\omega) := \mathbb{E}[X \,|\, Y = Y(\omega)]$$

and call $\mathbb{E}[X \,|\, Y]$ the *conditional expectation of X given Y*.

Note, that $\mathbb{E}[X \,|\, Y]$ is a random variable by definition. The random variable is determined by the values $\mathbb{E}[X \,|\, Y = y]$, as we will see in the next theorem.

But first, we return to the examplary die throw. We can state for $y \in Y(\Omega)$

$$\mathbb{E}[X \,|\, Y = y] = 2 + 3y =: g(y).$$

Indeed, we can just plugin the random variable Y and get

$$g(Y) = 2 + 3Y.$$

Taking the expectation of the random variable $g(Y) = \mathbb{E}[X \,|\, Y]$, we get

$$\mathbb{E}[\mathbb{E}[X \,|\, Y]] = \mathbb{E}[2 + 3Y] = 2 + 3 \cdot \mathbb{E}[Y] = \frac{7}{2} = \mathbb{E}[X].$$

This property is called the *tower property*. We prove this and some other basic properties in the following theorem.

Theorem 2.2 (Conditional Expectation – Properties)

Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space and X, Y discrete, integrable random variables on that space. Then

1. $\mathbb{E}[X \,|\, Y]$ is a random variable.
2. $\mathbb{E}[X] = \mathbb{E}[\mathbb{E}[X \,|\, Y]]$ (tower property).
3. If X and Y are independent
$$\mathbb{E}[X \,|\, Y] = \mathbb{E}[X].$$
4. For any bounded, measurable function $f : Y(\Omega) \to \mathbb{R}$
$$\mathbb{E}[X\, f(Y) \,|\, Y] = f(Y)\mathbb{E}[X \,|\, Y].$$

Proof. The proof is straight forward.

1. We first note, that $\sigma(Y)$ is generated by the events $(\{Y = k\})_{k\in Y(\Omega)}$. Use equation 2.1 to write for all $\omega \in \Omega$
$$\mathbb{E}[X \,|\, Y = Y(\omega)] = \frac{1}{\mathbb{P}(Y = Y(\omega))}\mathbb{E}[X\mathbb{1}_{\{Y=Y(\omega)\}}].$$
This is a concatenation of measurable functions and thus measurable. Therefore $\mathbb{E}[X \,|\, Y]$ is a random variable, since measurability on the generating set is enough to conclude measurability.
2. Follows directly from the law of total expectation, see Theorem 2.1.
3. Using independence and equation 2.1 we get for $y \in Y(\Omega)$
$$\mathbb{E}[X \,|\, Y = y] = \frac{1}{\mathbb{P}(Y = y)}\mathbb{E}[X]\,\mathbb{E}[\mathbb{1}_{\{Y=y\}}] = \mathbb{E}[X].$$
Therefore $\mathbb{E}[X \,|\, Y]$ is equal to the constant $\mathbb{E}[X]$.
4. We compute by LOTUS-principle
$$\mathbb{E}[X\, f(Y) \,|\, Y = y] = f(y) \cdot \sum_{x\in X(\Omega)} x\,\mathbb{P}(X = x \,|\, Y = y).$$

□

Example 2.1 (Random Sum)

In applications in insurance it is interesting to model the (random) number of, for example, car crashes during a certain time period. Depending on the severity of the crash, the company is only interested in those crashes, that trigger a payment claim. The probability of that has a certain probability $p \in (0,1)$. Naturally, the company is interested in the expected number of such claims, if something is known or assumed about the frequency of car crashes. A common way is to assume that the number of accidents follows a Poisson distribution. We model this in the following.

Let $\Omega := \{0,1\}^{\mathbb{N}_0}$, the space of all infinite sequences of zeros and ones, where one codes for a claim. We note that the cardinality of Ω is equal to the cardinality of the interval $[0,1]$, since each $\omega \in \Omega$ codes for a number $x \in [0,1]$ in binary. This implies the problem, that now the σ-algebra $\mathcal{A}$ can not be chosen arbitrarily. The power set is too big: it is impossible to define a probability measure (see Vitali's Theorem). Therefore, we take the sigma algebra generated from the cylinder sets that fix the first n coordinates of an infinite sequence in Ω.

$$A_{\omega_0,\omega_1,\ldots,\omega_n} := \{\omega \in \Omega : \omega = (\omega_0, \omega_1, \ldots, \omega_n, \ldots)\}.$$

For fixed n, the events $A_{\omega_0,\omega_1,\ldots,\omega_n}$ are pairwise disjoint, but contain Accordingly, there is only one possible probability measure:

$$\mathbb{P}(A_{\omega_1,\ldots,\omega_n}) := (1-p)^{\#\{\omega_i=0\}} p^{\#\{\omega_i=1\}}$$

for a given $p \in (0,1)$.

Instead of some fixed n, we want to look at a sequence of random length. Therefore, let N be a Poisson random variable with intensity $\lambda \in \mathbb{R}^+$ and $(X_i)_{i\in\mathbb{N}_0}$ with $X_i(\omega) = \omega_i$ (the coordinate projections) be a sequence of Bernoulli random variables with parameter $p \in (0,1)$. Assume that all random variables are independent and indentically distributed.[3]

We define a new random variable $S : \Omega \to \mathbb{N}_0$, the random number of claims, for $\omega \in \Omega$ via

$$S(\omega) := \sum_{i=0}^{N(\omega)} X_i(\omega).$$

We can compute the expectation of S using the tower property directly. However, more is possible: We can actually compute the distribution of S using generating functions. If we knew, there were exactly n car crashes, we could compute using independence and $z \in (-1,1)$:

$$\mathbb{E}[z^S \mid N = n] = \mathbb{E}[z^{X_0}]^n = ((1-p) + pz)^n.$$

This is the generating function of a binomial random variable. We conclude, that the generating function of S has the following form

$$G_S(z) := \mathbb{E}[z^S] = \mathbb{E}[((1-p)+pz)^N] = e^{-\lambda} \sum_{n=0}^{\infty} ((1-p)+pz)^n \frac{\lambda^n}{n!} = e^{\lambda p(z-1)}.$$

[3]It is possible to prove independence and identical distribution of the X_i from the definition, but it is too technical for now.

From this particular result, we conclude that $S \sim Poi(\lambda p)$, which means $\mathbb{E}[S] = \lambda p$. It is noteworthy, that the result can be generalized, see exercise 2.

Exercises

1. (4P) Let X and Y be independent and uniformly distributed on $\{1, 2, 3, 4, 5, 6\}$. What is the expectation of $X \cdot Y$ under the events $\{X < Y\}$, $\{X = Y\}$ and $\{X > Y\}$?

2. Generalise example 2.1 in the following sense. Let N be an arbitrary discrete $\mathbb{N}_0$-valued random variable with generating function G_N. Let $(X_i)_{i \in \mathbb{N}_0}$ be a iid sequence of random variables independent of N with common generating function G_X. Define the random sum $S := \sum_{i=0}^{N} X_i$ and show

$$G_S(z) = G_N(G_X(z)).$$

   Use G_S to show

$$\begin{aligned} \mathbb{E}[S] &= \mathbb{E}[N]\,\mathbb{E}[X_0], \\ Var[S] &= Var[N]\mathbb{E}[X_0^2] + \mathbb{E}[N]Var[X_0]. \end{aligned}$$

3. (4P) Let X and Y be independent discrete random variables on a common probability space. Compute $\mathbb{P}(X = k|X+Y = N)$, $k \leq N$, and identify the distribution $\mathbb{P}(X = .|X + Y = N)$ for fixed N, when

   (a) $X, Y \sim Bin(n, p)$, $n \in \mathbb{N}$, $p \in (0, 1)$.

   (b) $X \sim Poi(\lambda)$, $Y \sim Poi(\mu)$, $\lambda, \mu \in \mathbb{R}^+$.

4. (4P) Let X and Y be independent and identically distributed discrete random variables defined on the same probability space with countable number of outcomes. Show

$$\mathbb{E}[X|X + Y = k] = \frac{k}{2}.$$

5. (4P) Let $(\Omega, 2^\Omega, \mathbb{P})$ be a probability space where $\Omega = \{\omega_1, \omega_2, \omega_3\}$ and $\mathbb{P}$ is the uniform measure on Ω. The random variables X and Y are defined as

$$X(\omega_i) \mapsto \begin{cases} +1 & \text{, if } i = 1 \\ 0 & \text{, if } i = 2 \\ -1 & \text{, if } i = 3 \end{cases} \qquad\qquad Y(\omega_i) \mapsto \begin{cases} -1 & \text{, if } i = 1 \\ +1 & \text{, if } i = 2 \\ 0 & \text{, if } i = 3 \end{cases}$$

   (a) Compute the joint distribution of X and Y and determine whether X and Y are independent.

   (b) Determine the distribution of $X + Y$.

   (c) Compute $\mathbb{E}[X|X + Y = k]$ and $\mathbb{E}[Y|X + Y = k]$ for $k \in \{-1, 0, 1\}$.

3 Basic Gambling Problems

In this section, we look at classical gambling problems. The Problem of Points was stated already in the 14^{th} century, and was solved correctly by Pascal and Fermat in 1654, for the first time. Shortly after, many other mathematicians generalised the problem. The main method is the solution of certain recursions, that today would be derived from conditional expectations. In a sense, the 17^{th} century solutions used the concept way before it had been introduced formally. A variant is the famous Matchox Problem by Banach.

The second problem is the *Gambler's Ruin*, which is a variant of the Problem of Points, where the players continue to play, until one of them is bankrupt. The random duration and ruin probability are of interest here. Although this constitutes a Markov Chain, one can solve basic questions without this theory. We continue the problems later using Markov Chains and Martingales. However, the ruin probability for arbitrary initial asset has been computed as early as 1656 by Pascal and Fermat, and also by other famous probabilists in the 18^{th} century.

3.1 The Problem of Points

Let us assume the following situation:

> The problem of points concerns a game of chance with two players, Ann (A) and Bob (B). The players contribute equally to a prize pot, and agree in advance that the first player to have won a certain number $n \geq 1$ of rounds will collect the entire prize C. The probability that Ann wins is $p \in (0, 1)$ and the probability that Bob wins is $1 - p$.
>
> Now suppose that the game is interrupted by external circumstances before either player has achieved victory. How does one then divide the pot fairly?[4]

There are some suggestions to solve the *problem of points* as early as in the 15^{th} century by Italian mathematicians Paciolo[5] and Tartaglia.[6] However, the first complete solution of the problem has been stated in the famous exchange of letters between Fermat[7] and Pascal.[8]

We want to model this problem and take on Ann's perspective. Let $\Omega = \{0, 1\}^{2n-1}$, $n \geq 1$, be the event space of all sequences of possible games (in the worst case the number of games is $2n - 1$), where 1 decodes a win by Ann and 0 a defeat. Let $X_i : \Omega \to \{0, 1\}$ be the indicator whether or not Ann wins the n^{th} round, i.e. $X_i(\omega) = \omega_i$. Clearly, $X_i \sim Ber(p)$. We can set $\mathbb{P}$ as the product measure, i.e.

$$\mathbb{P}(\{(\omega_1, \dots, \omega_{2n-1})\}) = \prod_{i=1}^{2n-1} \mathbb{P}(X_i = \omega_i).$$

[4] From Wikipedia, Problem of Points, adapted, last visit: 2023/03/22, 2pm.

[5] Luca Pacioli, 1447-1517.

[6] Niccolò Fontana Tartaglia, 1499-1557.

[7] Pierre de Fermat, 1607-1665

[8] Blaise Pascal, 1623-1662

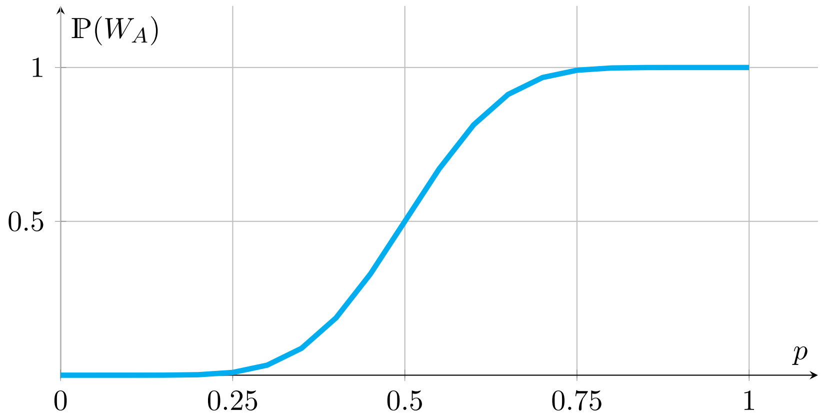

Figure 1: The probability of W_A for different $p \in (0,1)$ and 10 points to win. Small deviations from the fair case ($p = 1/2$) has large influence on the winning probability. For example if $p = 0.75$, it is almost certain that Ann wins the game.

Since Ω is finite, we can choose $(\Omega, 2^\Omega, \mathbb{P})$ as a stochastic model for this game of chance.[9]

We want to analyse the probability of the following two events

$$S_{n:j} := \left\{ \sum_{k=1}^{n-1+j} X_k = n-1 \right\} \cap \{X_{n+j} = 1\}, \qquad W_A := \bigcup_{j=0}^{n-1} S_{n:j}.$$

The event $S_{n:j}$ is the event when Ann wins and Bob has j points.[10] The probability W_A is the event when Ann wins, regardless of the number of points that Bob achieved.

Therefore for every $j < n$:

$$\mathbb{P}(S_{n:j}) = \binom{n-1+j}{j} p^n \, (1-p)^j.$$

Since $S_{n:i} \cap S_{n:j} = \emptyset$ for any $i \neq j$, we deduce

$$\mathbb{P}(W_A) = \sum_{j=0}^{n-1} \mathbb{P}(S_{n:j}) = \sum_{j=0}^{n-1} \binom{n-1+j}{j} p^n \, (1-p)^j.$$

It is intuitive, that W_A has a negative binomial probability, as we are waiting for the first time Ann has won n rounds. See figure 1 for a numerical analysis of the winning probability.

If the game is interrupted at score $(n-i) : (n-j)$, Ann needs to win i rounds before Bob wins j rounds. If we want to compute the probability that Ann wins, Bob can win

[9]We have made similar assumptions in Section 1, example 2.1, but for infinite sequences.

[10]For $p = 1/2$ we get *Banach's matchbox problem*, where j matches are left in the non-empty matchbox, after one of the boxes is found empty. In this case $\mathbb{P}(W_A) = 1/2$, so it does not matter which pocket we actually look at.

(5-i):(5-j)	$j = 1$	2	3	4
$i = 1$	0.500	0.637	0.773	0.891
2	0.363	0.500	0.656	0.813
3	0.227	0.344	0.500	0.688
4	0.109	0.188	0.313	0.500

Table 1: Winning probability of Ann at score $(5 - i) : (5 - j)$, or equivalently the proportion of the prize Ann should get when the game is interrupted. For example, if the game is interrupted at score $4 : 1$ ($i = 1$, $j = 4$) Ann should get about 89% of the prize.

at most $j - 1$ games. Thus

$$\mathbb{P}(W_A \,|\, \text{score } (n-i) : (n-j)) = \sum_{k=0}^{j-1} \mathbb{P}(S_{i:k}) = \sum_{k=0}^{j-1} \binom{i-1+k}{k} p^i \, (1-p)^k.$$

This probability also has a negative binomial form. For $n = 5$, and $p = 1/2$, table 1 gives the proportion that Ann should get from the prize C.

Remark 3.1

Historically, another method was used, that we will later call *first step analysis*. However the method is quite intuitive. Using conditional expectation and the law of total expectation we get

$$\mathbb{P}(W_A | S_{i:j}) = \mathbb{E}[\mathbb{1}_{W_A} | S_{i:j}] = p \, \mathbb{E}[\mathbb{1}_{W_A} | S_{i+1:j}] + (1-p) \, \mathbb{E}[\mathbb{1}_{W_A} | S_{i:j+1}].$$

This approach implies a recurrence equation, with some boundary conditions. The solution is the result we computed earlier.

One can use this approach for a three or more player game. Such solutions were found as early as the 18^{th} century by Lagrange. The article [Gor14] gives a very complete overview over the methods and history of the solutions for the Problem of Points.

3.2 Gambler's Ruin

In the previous section Ann and Bob were playing a game that ends after one of them wins a certain fixed number of rounds. The game does not last forever, but ends after at most $2n - 1$ games.

We want to alterate the game as follows: Ann has a coins, and Bob has b coins, $a, b \in \mathbb{N}$. At the beginning of each round each player bets a coin. Either Ann wins with probability $p \in (0, 1)$ or Bob wins.[11] The winner gets both coins, so the net win is one coin. For example, if Ann wins, she has $a + 1$ coins and Bob has $b - 1$. We want to analyse, what is the probability that Ann is ruined at some point, i.e. has no coins

[11] We do not consider draws.

left, and Bob has all $N := a + b$ coins. We will later also compute the expected game duration and similar quantities.

Schematically, we can depict this dynamic as follows:

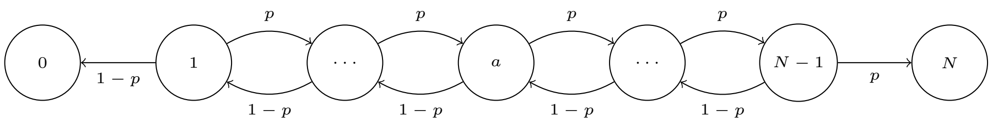

We model the development of Ann's capital in the following way. Similar to previous models, we keep track of the movement in Ann's capital and set $\Omega := \{-1, 1\}^{\mathbb{N}}$. Note, that we will fix the initial capital a later. It could also be random, but we will not consider this case explicitly.

For $\omega \in \Omega$ define the coordinate projections

$$X_i(\omega) = \omega_i.$$

The sequence $(X_i)_{i\in\mathbb{N}}$ is iid and $X_i \sim (1-p)\,\delta_{-1} + p\,\delta_{+1}$. We set $X_0 := a \in \mathbb{Z}$, so the initial X_0 is a constant almost surely.

The σ-algebra and probability measure are defined analogue to example 2.1. However, it is convenient to consider a filtrated probability space.

Definition 3.1 (Filtration and filtered probability space)

A sequence $(\mathcal{F}_n)_{n\in\mathbb{N}_0}$ of sub-sigma-algebras of $\mathcal{A}$ in a probability space $(\Omega, \mathcal{A}, \mathbb{P})$, with

$$\mathcal{F}_0 \subset \mathcal{F}_1 \subset \ldots \subset \mathcal{F}_n \subset \ldots \subset \mathcal{A},$$

is called *filtration.*

The quadruple $(\Omega, \mathcal{A}, (\mathcal{F}_n)_{n\in\mathbb{N}_0}, \mathbb{P})$ is called a *filtered probability space.*

In the case of the gambler's ruin we can use the so called *canonical σ-algebra.* Provided X_0 is a constant, we set

$$\mathcal{F}_0 := \{\emptyset, \Omega\}, \quad \mathcal{F}_n := \sigma(X_0, \ldots, X_n).$$

This filtration contains all the information about the movements of Ann's capital up to time n. Her capital at time n is then defined via

$$S_0 := X_0 = a, \qquad\qquad S_n := a + \sum_{k=1}^{n} X_n.$$

See figure 2 for some sample trajectories. Note, that we have not introduced any restrictions on whether Ann and Bob can go into debts. Therefore, $(S_n)_{n\in\mathbb{N}_0}$ can take values in $\mathbb{Z}$, not just between 0 and $N = a + b$. In general $(S_n)_{n\in\mathbb{N}_0}$ is called *Random Walk on $\mathbb{Z}$.*

To model the condition of stopping, whenever 0 or N is reached, we introduce another random variable. Let

$$T := \inf\{n \geq 0 : X_n \in \{0, N\}\}$$

be the *hitting time* or, in our case, be the *game duration.* The random variable T is a special type of random variable, that we will need later.

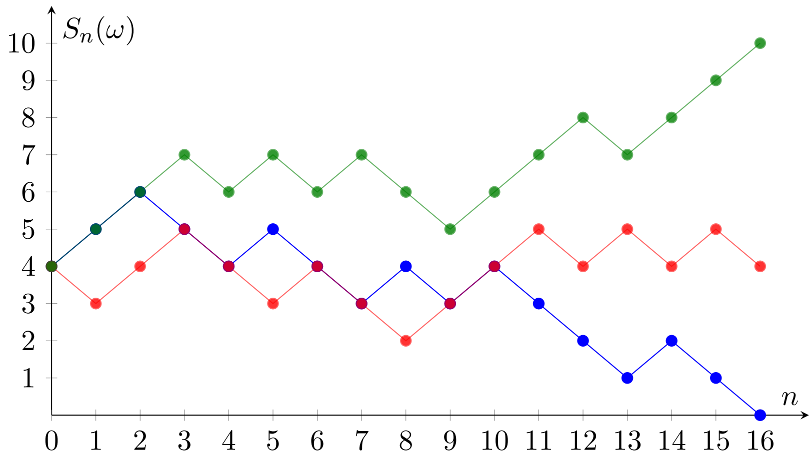

Figure 2: Sample trajectory for $N = 10$ and $a = 4$.

Definition 3.2 (Stopping time)

Let $(\Omega, \mathcal{A}, (\mathcal{F}_n)_{n\in\mathbb{N}_0}, \mathbb{P})$ be a filtered probability space. The random variable

$$T : \Omega \to \mathbb{N}_0 \cup \{+\infty\}$$

is called a *stopping time*, if

$$\{T \leq n\} \in \mathcal{F}_n.$$

Indeed, the game duration is a stopping time with respect to the canonical filtration defined above, since

$$\{T \leq n\} \in \mathcal{F}_n = \sigma(X_0, X_1, \dots, X_n) = \sigma(S_n).$$

The equality $\sigma(X_0, X_1, \dots, X_n) = \sigma(S_n)$ is due to the fact that S_n is a measurable function (sum) of $X_0, \dots, X_n$.

3.3 Ann's Ruin Probability

With the notion of stopping time, we can give sense to the expression S_T (Ann's capital at the end of the game) via

$$(S_T)(\omega) := S_{T(\omega)}(\omega)$$

and the stopped process $(S_{n\wedge T})_{n\in\mathbb{N}_0}$ is the stopped Random Walk or in our case the *Gambler's Ruin model.*[12]

Clearly, $S_T \in \{0, N\}$, where $N = a + b$ is the combined fortune of Ann and Bob. Ann's ruin probability is then $\mathbb{P}(S_T = 0 \,|\, S_0 = a)$. Before we actually compute this probability, we need some intuition. Let $q := 1 - p$, the probability that Ann loses

[12] $x \wedge y := \inf\{x, y\}$, for $x, y \in \mathbb{R}$.

a round. We look at the *bias* $r := q/p$. For a fair game this bias is equal to one, it is equally like for Ann and Bob to win a game. If the odds are in favour of Bob ($q > p$), then $r > 1$. We might expect, that Ann is drifting towards ruin and thus will be eventually ruined with higher probability than in the equal odds case. The opposite should hold, if Ann has the better odds or skill. We will see that even the slightest deviation from the fair case will have a dramatic influence on the ruin probability.

Theorem 3.1 (Ann's Ruin Probability)

For $a \in \mathbb{N}$ let

$$S_n := a + \sum_{k=1}^{n} X_k$$

be the sum of iid random variables with $X_k \sim q\,\delta_{-1} + p\,\delta_{+1}$ for $p \in (0,1)$ and $q := 1-p$. Assume $b \in \mathbb{N}$ and define

$$T := \inf\{n \geq 0 : S_n \in \{0, a+b\}\}.$$

Then, for $r := q/p$,

1. T is almost surely finite, i.e. $\mathbb{P}(T < +\infty) = 1$.
2. The probability to hit zero (ruin) is

$$\mathbb{P}(S_T = 0 \,|\, S_0 = a) = \begin{cases} 1 - \frac{a}{a+b} & \text{if } r = 1, \\ r^a \frac{1-r^b}{1-r^{a+b}} & \text{if } r \neq 1. \end{cases}$$

Proof. Since $\{T < +\infty\} = \{S_T = 0\} \sqcup \{S_T = a+b\}$, it is enough to show that $\mathbb{P}(S_T = 0) + \mathbb{P}(S_T = a+b) = 1$.

We notice first, that reaching 0 means there is at least one index n such $S_n = 0$, so

$$\{S_T = 0\} = \bigcup_{n=0}^{\infty} \{S_n = 0\}.$$

We deduce using total probability and the independence and identical distribution of the X_i

$$\begin{aligned}
\mathbb{P}(S_T = 0 \,|\, S_0 = a) &= \mathbb{P}\Big(\bigcup_{n=0}^{\infty}\{S_n = 0\} \,|\, S_0 = a\Big) \\
&= \mathbb{P}\Big(\bigcup_{n=0}^{\infty}\{S_n = 0\} \cap \{X_1 = -1\} \,|\, S_0 = a\Big) + \mathbb{P}\Big(\bigcup_{n=0}^{\infty}\{S_n = 0\} \cap \{X_1 = +1\} \,|\, S_0 = a\Big) \\
&= \mathbb{P}\Big(\bigcup_{n=2}^{\infty}\{S_n = 0\} \cap \{S_1 = a-1\} \,|\, S_0 = a\Big) + \mathbb{P}\Big(\bigcup_{n=2}^{\infty}\{S_n = 0\} \cap \{S_1 = a+1\} \,|\, S_0 = a\Big) \\
&= q\,\mathbb{P}(S_T = 0 \,|\, S_0 = a-1) + p\,\mathbb{P}(S_T = 0 \,|\, S_0 = a+1).
\end{aligned}$$

Let $f(a) := \mathbb{P}(S_T = 0 \,|\, S_0 = a)$. Then the above result can be written as $f(0) = 1$, $f(a+b) = 0$ and for $k \in \{1, 2, \ldots, a+b-1\}$

$$f(k) = q\, f(k-1) + p\, f(k+1).$$

In [Pri18], p. 46, a solution method similar to standard methods of differential equations is presented. The idea is to assume that solutions have the form $f(k) = c\, x^k$. Plugging this into the recursion, gives the so called *characteristic equation*

$$cx^k = qcx^{k-1} + pcx^{k+1} \quad \Leftrightarrow \quad x^2 - \frac{1}{p}x + r = 0.$$

The solution(s) to this quadratic equation are 1 and r.

Therefore, the solution of the recursion has the form

$$f(k) = c_1 + c_2\, r^k.$$

The boundary conditions lead to the following system of linear equaitons

$$\begin{aligned} 1 &= f(0) = c_1 + c_2 \\ 0 &= f(a+b) = c_1 + c_2\, r^{a+b}. \end{aligned}$$

The unique solutions are

$$c_1 = \frac{r^{a+b}}{r^{a+b}-1}, \qquad\qquad c_2 = \frac{1}{1-r^{a+b}}.$$

From this, we can deduce

$$f(a) = r^a \frac{1-r^b}{1-r^{a+b}}.$$

We have skipped over the case $r = 1$, but a simple limit computation using L'Hospital, yields

$$\lim_{r \to 1} \frac{1-r^b}{1-r^{a+b}} = \lim_{r \to 1} \frac{-b\, r^{b-1}}{-(a+b)r^{a+b-1}} = \frac{b}{a+b} = 1 - \frac{a}{a+b}.$$

For the computation of $\mathbb{P}(S_T = a+b)$, we just need to start the same computation with $X_0 := b$.[13] Then

$$f(a) + f(b) = \frac{1}{1-r^{a+b}} \Big(r^a(1-r^b) + r^b(1-r^a) \Big) = 1,$$

which establishes $\mathbb{P}(T < \infty) = 1$. $\square$

Theorem 3.1 allows to plot the ruin probability depending on p and Bob's initial capital, see figure 3. The plot demonstrates a rather sharp "cutoff" at $p = 1/2$. A slight deviation from the fair game changes Ann's fortune dramatically. In the limiting case $b \to \infty$, Ann's ruin is certain, if she is at an disadvantage ($r < 1$):

$$\lim_{b \to \infty} r^a \frac{1-r^b}{1-r^{a+b}} = 1$$

[13] Ruin probability of Bob

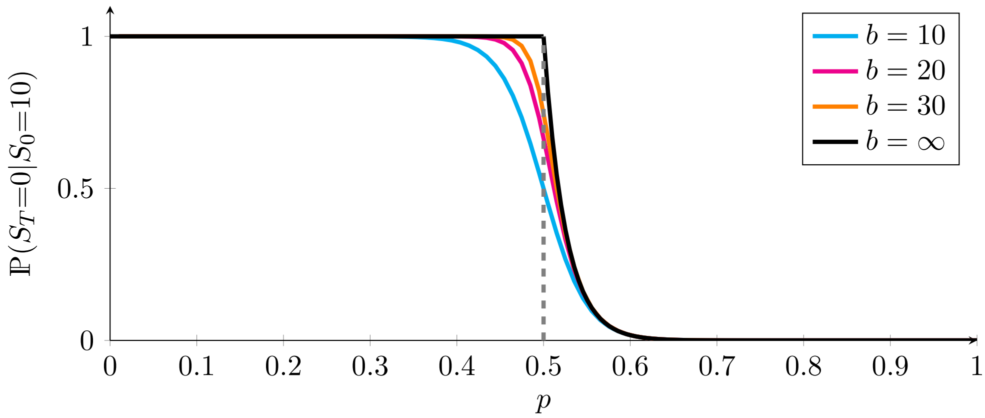

Figure 3: Ann's ruin probability as a function of p with $a = 10$.

For $r > 1$, we have

$$\lim_{b\to\infty} r^a \frac{1 - r^b}{1 - r^{a+b}} = r^a,$$

which means that ruin probability decays exponentially fast with increasing skill of Ann (i.e. increasing r).

So, if Bob is somewhat richer than Ann, her chances of winning against Bob are slim, if she is at a disadvantage. This is the reason why opening a casino is much more profitable than gambling.[14] However, many gamblers do think that a small disadvantage should asymptotically be negligible.[15] That this is not the case shows the following argumentation, see [LLM10], 20.2. The probability to reach a capital $a + c$, where c is some constant smaller than b, is by theorem 3.1

$$\mathbb{P}(S_T = a + c \,|\, S_0 = a) = \frac{1 - r^a}{1 - r^{a+c}}.$$

In the case of a disadvantage, i.e. $r > 1$, and a large, we can estimate

$$\mathbb{P}(S_T = a + c \,|\, S_0 = a) \sim \frac{1}{r^c}.$$

As an example on how small this probability in real games of chance really is, we take American roulette. It is equally likely that the ball on the wheel hits any one of the 38 numbers, but only 18 of them are red (there are 18 black numbers, and two green ones, 0 and 00). So betting on red is successful with a probability of $18/38 = 9/19 \approx 0.474$, which seems not too unfair. Then $1/r = p/(1 - p) = 9/10$. However, the probability to win 100 dollars, if hypothetically it was allowed to bet just one dollar[16], is for a sufficiently large initial capital a just

$$\mathbb{P}(S_T = a + 100 \,|\, S_0 = a) \sim \left(\frac{9}{10}\right)^{100} \approx 2.66 \times 10^{-5}.$$

[14]The author recommends neither option: The first one is unethical, knowing what we know now. The other will ruin you with certainty.

[15]In German we say “Ausgleichende Gerechtigkeit” (retributive justice).

[16]Minimum bet is often upwards of 10 dollars.

The conclusion must be, that even with a great initial fortune, the chance of winning 100 dollars is close enough to zero to never see that happen in any ones lifetime. In a sense, this implies that the "general drift towards ruin dominates any streak of luck".[17]

Please remember that in roulette one can bet on other numbers or combinations that have a higher payout than just 1:1. Unfortunately, that only implies that it takes a little longer to hit certain ruin.

3.4 Game Duration

So far, we have looked at the probability that the game ends with ruin for Ann. But the actual game duration is not yet known. We do know it is finite, by theorem 3.1, but what game duration can we expect on average?

Theorem 3.2 (Average Game Duration)

For $a \in \mathbb{N}$ let

$$S_n := a + \sum_{k=1}^{n} X_k$$

be the sum of iid random variables with $X_k \sim q\,\delta_{-1} + p\,\delta_{+1}$ for $p \in (0,1)$ and $q := 1-p$. Assume $b \in \mathbb{N}$ and define

$$T := \inf\{n \geq 0 : S_n \in \{0, a+b\}\}.$$

Then, for $r := q/p$,

$$\mathbb{E}[T \mid S_0 = a] = \begin{cases} ab & \text{if } r = 1, \\ \frac{1}{1-2p}\left(a - (a+b)\frac{1-r^a}{1-r^{a+b}}\right) & \text{if } r \neq 1. \end{cases}$$

Proof. The proof is very similar to the proof of theorem 3.1. We define

$$g(k) := \mathbb{E}[T \mid S_0 = k]$$

and notice the boundary conditions $g(0) = g(a+b) = 0$. For $k \in \{1, 2, \dots, a+b-1\}$ we condition on the outcome of X_1. By law of total expectation, see Theorem 2.1, we have:

$$g(k) = \mathbb{E}[T \mid X_1 = +1, S_0 = k]\mathbb{P}(X_1 = +1) + \mathbb{E}[T \mid X_1 = -1, S_0 = k]\mathbb{P}(X_1 = -1)$$

We can now rethink the situation: instead of conditioning on the outcome of X_1 and $S_0 = k$, we know one step with certainty and find ourselves in the situation of starting in $S_0 = k + X_1$ instead. Therefore, we get the recursion

$$g(k) = (g(k+1) + 1)\,p + (g(k-1) - 1)\,q = g(k+1)\,p + g(k-1)\,q + 1.$$

[17] As stated similarly in [LLM10], Section 20.2., Caption of figure 20.2, p. 546.

The recursion is similar to the one in theorem 3.1, but with a constant term. Following Privault, see [Pri18] p. 57, we first search for a particular solution of the form $c\,k$. We get

$$c\,k = c\,(k+1)\,p + c\,(k-1)\,q + 1 = c\,k\,(p+q) + c\,(p-q) + 1,$$

and thus

$$c = \frac{1}{1-2p}.$$

For now, we shall ignore the case $r = 1$. Now, the general solution has the form

$$g(k) = c_1 + c_2 r^k + \frac{1}{2-p}k.$$

The boundary conditions give two equations

$$\begin{aligned} g(0) &= 0 = c_1 + c_2 \\ g(a+b) &= 0 = c_1 + c_2 r^{a+b} + \frac{1}{1-2p}(a+b). \end{aligned}$$

We deduce

$$c_1 = -c_2 = \frac{1}{1-2p}\frac{a+b}{1-r^{a+b}}.$$

Plugging these constants into the general solution gives the result for $r \neq 1$. For the case $r = 1$, we no longer find a particular solution with the ansatz we made previously. However, a simple limit computation for $r \to 1$ gives the desired result. $\square$

Theorems 3.1 and 3.2 imply that the game duration is finite with probability one. Thus it is certain that one of the players will be ruined in finite time. It is intuitive that in the case of equal skill, $p = 1/2$, the game duration is maximised.

The situation changes somewhat, when Bob has a lot more money to play with than Ann. For the case that Ann is at an advantage, the game does go on for very long, as Bob is not easily ruined. In the case of equal odds, $r = 1$, the game could also go on for quite a while, since Ann and Bob might win and lose back and forth, but Ann might be ruined by a streak of bad luck, but only early on when Ann does not have won many games. Thus the game might be arbitrary long. This intuition is made clear by taking the limit $b \to \infty$:

$$\lim_{b\to\infty} \mathbb{E}[T \mid S_0 = a] = \begin{cases} +\infty & r \leq 1 \\ \frac{a}{1-2p} & r > 1. \end{cases}$$

It might be surprising that for $r = 1$ the time to ruin is almost surely finite, but has infinite expectation. It just means that the game can endure for an arbitrary length of time, before it eventually terminates.

A plot of the game duration as p changes is given in figure 4.

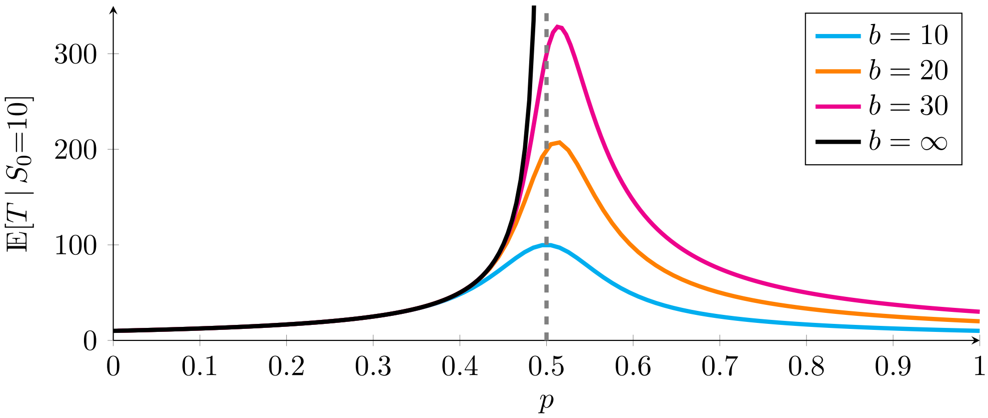

Figure 4: Average Game Duration as a function p with $a = 10$.

References

[Bré20] P. Brémaud. *Markov chains. Gibbs fields, Monte Carlo simulation and queues.* English. 2nd revised and augmented edition. Vol. 31. Texts Appl. Math. Cham: Springer, 2020. DOI: `10.1007/978-3-030-45982-6`.

[Eth10] S. N. Ethier. *The doctrine of chances. Probabilistic aspects of gambling.* English. Probab. Appl. Berlin: Springer, 2010. DOI: `10.1007/978-3-540-78783-9`.

[Gor14] P. Gorroochurn. "Thirteen Correct Solutions to the "Problem of Points" and Their Histories". In: *Math Intelligencer* 36 (2014), pp. 56–64. DOI: `https://doi.org/10.1007/s00283-014-9461-5`.

[LLM10] E. Lehmann, F. T. Leighton, and A. R. Meyer. *Mathematics for Computer Science.* `https://ocw.mit.edu/courses/6-042j-mathematics-for-computer-science-fall-2010/resources/mit6_042jf10_notes/`. [Online; latest access: October 11th, 2023]. 2010.

[Pri18] N. Privault. *Understanding Markov chains. Examples and applications.* English. 2nd edition. Springer Undergrad. Math. Ser. Singapore: Springer, 2018. DOI: `10.1007/978-981-13-0659-4`.

[Wil91] D. Williams. *Probability with martingales.* English. Cambridge etc.: Cambridge University Press, 1991. DOI: `10.1017/CBO9780511813658`.