Chapter 2

Expectations

2.1 Overview

Recall from Chapter 1 that a random variable is just a function that attaches a number to each item in the sample space. Less formally, a random variable corresponds to a numerical quantity whose value is determined by some chance mechanism.

Just as events have (conditional) probabilities attached to them, with possible interpretation as a long-run frequency, so too do random variables have a number interpretable as a long-run average attached to them. Given a particular piece of information (info), the symbol

$$\mathbb{E}\,(X \mid \text{info})$$

denotes the ***(conditional) expected value*** or ***(conditional) expectation*** of the random variable X (given that information). When the information is taken as understood, the expected value is abbreviated to $\mathbb{E}X$.

Expected values are not restricted to lie in the range from zero to one. For example, if the info forces a random variable X to always take values larger than 16 then $\mathbb{E}\,(X \mid \text{info})$ will be larger than 16.

As with conditional probabilities, there are convenient abbreviations when the conditioning information includes something like {event F has occurred}:

$$\mathbb{E}\,(X \mid \text{info and ``}F\text{ has occurred''})$$
$$\mathbb{E}\,(X \mid \text{info}, F)$$

Unlike many authors, I will take the expected value as a primitive concept, not one to be derived from other concepts. All of the methods that those authors use to *define* expected values will be *derived* from a small number of basic rules. I will provide an interpretation for just one of the rules, using long-run averages of values generated by independent repetitions of random experiments. You should provide analogous interpretations for the other rules.

> **Remark.** See the Appendix to this Chapter for another interpretation, which does not depend on a preliminary concept of independent repetitions of an experiment. The expected value $\mathbb{E}X$ can be interpreted as a "fair price" to pay up-front, in exchange for a random return X later—something like an insurance premium.

Rules for (conditional) expectations

Let X and Y be random variables, c and d be constants, and F_1, F_2, ... be events. Then:

(E1) $\mathbb{E}\left(cX + dY \mid \text{info}\right) = c\mathbb{E}\left(X \mid \text{info}\right) + d\mathbb{E}\left(Y \mid \text{info}\right)$;

(E2) if X can only take the constant value c under the given "info" then $\mathbb{E}\left(X \mid \text{info}\right) = c$;

(E3) if the given "info" forces $X \le Y$ then $\mathbb{E}\left(X \mid \text{info}\right) \le \mathbb{E}\left(Y \mid \text{info}\right)$;

(E4) if the events $F_1, F_2, \dots$ are disjoint and have union equal to the whole sample space then

$$\mathbb{E}\left(X \mid \text{info}\right) = \sum\nolimits_i \mathbb{E}\left(X \mid F_i, \text{info}\right) \mathbb{P}\left(F_i \mid \text{info}\right).$$

Rule (E4) combines the power of both rules (P4) and (P5) for conditional probabilities. Here is the frequency interpretation for the case of two disjoint events F_1 and F_2 with union equal to the whole sample space: Repeat the experiment (independently) a very large number (n) of times, each time with the same conditioning info, noting for each repetition the value taken by X and which of F_1 or F_2 occurs.

	1	2	3	4	...				$n-1$	n	total
F_1 occurs	✓	✓		✓	...				✓	✓	n_1
F_2 occurs			✓		...	✓	✓	✓			n_2
X	x_1	x_2	x_3	x_4	...				x_{n-1}	x_n	

By the frequency interpretation of probabilities, $\mathbb{P}(F_1 \mid \text{info}) \approx n_1/n$ and $\mathbb{P}(F_2 \mid \text{info}) \approx n_2/n$. Those trials where F_1 occurs correspond to conditioning on F_1:

$$\mathbb{E}(X \mid F_1, \text{info}) \approx \frac{1}{n_1} \sum_{F_1 \text{ occurs}} x_i.$$

Similarly,

$$\mathbb{E}(X \mid F_2, \text{info}) \approx \frac{1}{n_2} \sum_{F_2 \text{ occurs}} x_i$$

Thus

$$\begin{aligned}
&\mathbb{E}(X \mid F_1, \text{info})\,\mathbb{P}(F_1 \mid \text{info}) + \mathbb{E}(X \mid F_2, \text{info})\,\mathbb{P}(F_2 \mid \text{info}) \\
&\approx \left(\frac{1}{n_1} \sum_{F_1 \text{ occurs}} x_i\right)\left(\frac{n_1}{n}\right) + \left(\frac{1}{n_2} \sum_{F_2 \text{ occurs}} x_i\right)\left(\frac{n_2}{n}\right) \\
&= \frac{1}{n}\sum_{i=1}^{n} x_i \\
&\approx \mathbb{E}(X \mid \text{info}).
\end{aligned}$$

As n gets larger and larger all approximations are supposed to get better and better, and so on.

Modulo some fine print regarding convergence of infinite series, rule (E1) extends to sums of infinite sequences of random variables,

$$(E1)' \qquad \mathbb{E}(X_1 + X_2 + \dots) = \mathbb{E}(X_1) + \mathbb{E}(X_2) + \dots$$

(For mathematical purists: the asserted equality holds if $\sum_i \mathbb{E}|X_i| < \infty$.)

Remark. The rules for conditional expectations actually include all the rules for conditional probabilities as special cases. This delightfully convenient fact can be established by systematic use of particularly simple random variables. For each event A the ***indicator function of*** A is defined by

$$\mathbb{I}_A = \begin{cases} 1 & \text{if the event } A \text{ occurs,} \\ 0 & \text{if the event } A^c \text{ occurs.} \end{cases}$$

Each $\mathbb{I}_A$ is a random variable.

Rule (E4) with $F_1 = A$ and $F_2 = A^c$ gives

$$\begin{aligned}
\mathbb{E}(\mathbb{I}_A \mid \text{info}) &= \mathbb{E}(\mathbb{I}_A \mid A, \text{ info})\,\mathbb{P}(A \mid \text{info}) + \\
&\quad + \mathbb{E}(\mathbb{I}_A \mid A^c, \text{ info})\,\mathbb{P}(A^c \mid \text{info}) \\
&= 1 \times \mathbb{P}(A \mid \text{info}) + 0 \times \mathbb{P}(A^c \mid \text{info}) \qquad \text{by (E2).}
\end{aligned}$$

That is, $\mathbb{E}(\mathbb{I}_A \mid \text{info}) = \mathbb{P}(A \mid \text{info})$.

If an event A is a disjoint union of events $A_1, A_2, \ldots$ then $\mathbb{I}_A = \mathbb{I}_{A_1} + \mathbb{I}_{A_2} + \ldots$. (Why?) Taking expectations then invoking the version of (E1) for infinite sums we get rule (P4).

As an exercise, you might try to derive the other probability rules, but don't spend much time on the task or worry about it too much. Just keep buried somewhere in the back of your mind the idea that you can do more with expectations than with probabilities alone.

You will find it useful to remember that $\mathbb{E}(\mathbb{I}_A \mid \text{info}) = \mathbb{P}(A \mid \text{info})$, a result that is easy to recall from the fact that the long-run frequency of occurrence of an event, over many repetitions, is just the long-run average of its indicator function.

Rules (E2) and (E4) can be used to calculate expectations from probabilities, for random variables that take values in "discrete" set. Consider the case of a random variable Y expressible as a function $g(X)$ of another random variable, X, which takes on only a discrete set of values $c_1, c_2, \ldots$. Let F_i be the subset of S on which $X = c_i$, that is, $F_i = \{X = c_i\}$. Then by E2,

$$\mathbb{E}(Y \mid F_i, \text{info}) = g(c_i),$$

and by E5,

$$\mathbb{E}(Y \mid \text{info}) = \sum\nolimits_i g(c_i)\mathbb{P}(F_i \mid \text{info}).$$

More succinctly,

$$(E5) \qquad \mathbb{E}(g(X) \mid \text{info}) = \sum\nolimits_i g(c_i)\mathbb{P}(X = c_i \mid \text{info}).$$

In particular,

$$(E5)' \qquad \mathbb{E}(X \mid \text{info}) = \sum\nolimits_i c_i\mathbb{P}(X = c_i \mid \text{info}).$$

Both (E5) and (E5)' apply to random variables X that take values in the "discrete set" $\{c_1, c_2, \ldots\}$.

Remark. For random variables that take a continuous range of values an approximation argument (see Chapter 6) will provide us with an analog of (E5) with the sum replaced by an integral.

You will find it helpful to remember expectations for a few standard mechanisms, such as coin tossing, rather than have to rederive them repeatedly.

Example <2.1> Expected value for the geometric(p) distribution is $1/p$.

The calculation of an expectation is often a good way to get a rough feel for the behaviour of a random process, but it doesn't tell the whole story.

Example <2.2> Expected number of tosses to get tthh is 16 with fair coin.

Compare with the next Example.

Example <2.3> Expected number of tosses to get hhh is 14 with fair coin.

Don't the last two results seem strange? On average it takes longer to reach tthh than hhh, but also on average the pattern tthh appears first.

Remark. You should also be able to show that the expected number of tosses for the completion of the game with competition between hhh and tthh is $9\,1/3$. Notice that the expected value for the game with competition is smaller than the minimum of the expected values for the two games. Why must it be smaller?

Probabilists study standard mechanisms, and establish basic results for them, partly in the hope that they will recognize those same mechanisms buried in other problems. In that way, unnecessary calculation can be avoided, making it easier to solve more complex problems. It can, however, take some work to find the hidden simplification.

Example <2.4> [Coupon collector problem] In order to encourage consumers to buy many packets of cereal, a manufacurer includes a Famous Probabilist card in each packet. There are 10 different types of card: Chung, Feller, Lévy, Kolmogorov, ..., Doob. Suppose that I am seized by the desire to own at least one card of each type. What is the expected number of packets that I need to buy in order to achieve my goal?

For the coupon collectors problem I assumed large numbers of cards of each type, in order to justify the analogy with coin tossing. Without that assumption the depletion of cards from the population would have a noticeable effect on the proportions of each type remaining after each purchase. The next example illustrates the effects of sampling from a finite population without replacement, when the population size is not assumed very large.

The example will also provides an illustration of the ***method of indicators***, whereby a random variable is expressed as a sum of indicator variables $\mathbb{I}_{A_1} + \mathbb{I}_{A_2} + \dots$, in order to reduce calculation of an expected value to separate calculation of probabilities $\mathbb{P}A_1$, $\mathbb{P}A_2$, ... via the formula

$$
\begin{aligned}
&\mathbb{E}\left(\mathbb{I}_{A_1} + \mathbb{I}_{A_2} + \dots \quad | \text{ info }\right) \\
&\qquad = \mathbb{E}\left(\mathbb{I}_{A_1} \mid \text{info}\right) + \mathbb{E}\left(\mathbb{I}_{A_2} \;\mid\; \text{info }\right) + \dots \\
&\qquad = \mathbb{P}\left(A_1 \mid \text{info}\right) + \mathbb{P}\left(A_2 \mid \text{info}\right) + \mathbb{P}\left(A_2 \mid \text{info}\right) + \dots
\end{aligned}
$$

Example <2.5> Suppose an urn contains r red balls and b black balls, all identical except for color. Suppose you remove one ball at a time, without replacement, at each step selecting at random from the urn: if k balls remain then each has probability $1/k$ of being chosen.Show that the expected number of red balls removed before the firstblack ball equals $r/(b+1)$.

Compare the solution $r/(b+1)$ with the result for sampling with replacement, where the number of draws required to get the first black would have a geometric$(b/(r+b))$ distribution. With replacement, the expected number of reds removed before the first black would be

$$(b/(r+b))^{-1} - 1 = r/b.$$

Replacement of balls after each draw increases the expected value slightly. Does that make sense?

The conditioning property (E5) can be used in a subtle way to solve the classical gambler's ruin problem. The method of solution invented by Abraham de Moivre, over two hundred years ago, has grown into one of the main technical tools of modern probability.

Example <2.6> Suppose two players, Alf and Betamax, bet on the tosses of a fair coin: for a head, Alf pays Betamax one dollar; for a tail, Betamax pays Alf one dollar. The stop playing when one player runs out of money. If Alf starts with α dollar bills, and Betamax starts with β dollars bills (both α and β whole numbers), what is the probability that Alf ends up with all the money?

De Moivre's method also works with biased coins, if we count profits in a different way—an even more elegant application of conditional expectations. The next Example provides the details. You could safely skip it if you understand the tricky idea behind Example <2.6>.

Example <2.7> Same problem as in Example <2.6>, except that the coin they toss has probability $p \neq 1/2$ of landing heads. (Could be skipped.)

You could also safely skip the final Example. It contains a discussion of a tricky little problem, that can be solved by conditioning or by an elegant symmetry argument.

Example <2.8> Big pills, little pills. (Tricky. Should be skipped.)

2.2 Things to remember

- Expectations (and conditional expectations) are linear (E1), increasing (E3) functions of random variables, which can be calculated as weighted averages of conditional expectations,

$$\mathbb{E}(X \mid \text{info}) = \sum\nolimits_i \mathbb{E}(X \mid F_i, \text{info})\, \mathbb{P}(F_i \mid \text{info}),$$

  where the disjoint events $F_1, F_2, \dots$ cover all possibilities (the weights sum to one).

- The indicator function of an event A is the random variable defined by

$$\mathbb{I}_A = \begin{cases} 1 & \text{if the event } A \text{ occurs,} \\ 0 & \text{if the event } A^c \text{ occurs.} \end{cases}$$

  The expected value of an indicator variable, $\mathbb{E}(\mathbb{I}_A \mid \text{info})$, is the same as the probability of the corresponding event, $\mathbb{P}(A \mid \text{info})$.

- As a consequence of the rules,

$$\mathbb{E}(g(X) \mid \text{info}) = \sum\nolimits_i g(c_i) \mathbb{P}(X = c_i \mid \text{info}),$$

  if X can take only values $c_1, c_2, \dots$.

2.3 The examples

<2.1> **Example.** For independent coin tossing, what is the expected value of X, the number of tosses to get the first head?

Suppose the coin has probability $p > 0$ of landing heads. (So we are actually calculating the expected value for the geometric(p) distribution.) I will present two methods.

Method A: a Markov argument without the picture
Condition on whether the first toss lands heads (H_1) or tails (T_1).

$$\begin{aligned}\mathbb{E}X &= \mathbb{E}(X \mid H_1)\mathbb{P}H_1 + \mathbb{E}(X \mid T_1)\mathbb{P}T_1 \\ &= (1)p + (1 + \mathbb{E}X)(1 - p).\end{aligned}$$

The reasoning behind the equality

$$\mathbb{E}(X \mid T_1) = 1 + \mathbb{E}X$$

is: After a tail we are back where we started, still counting the number of tosses until a head, except that the first tail must be included in that count.

Solving the equation for $\mathbb{E}X$ we get

$$\mathbb{E}X = 1/p.$$

Does this answer seem reasonable? (Is it always at least 1? Does it decrease as p increases? What happens as p tends to zero or one?)

Method B
By the formula (E5),

$$\mathbb{E}X = \sum\nolimits_{k=1}^{\infty} k(1-p)^{k-1}p.$$

There are several cute ways to sum this series. Here is my favorite. Write q for $1 - p$. Write the kth summand as a a column of k terms pq^{k-1}, then sum by rows:

$$\begin{aligned}\mathbb{E}X = p + pq &+ pq^2 + pq^3 + \dots \\ +pq &+ pq^2 + pq^3 + \dots \\ &+ pq^2 + pq^3 + \dots \\ &\qquad\;\; + pq^3 + \dots \\ &\qquad\quad \vdots\end{aligned}$$

Each row is a geometric series.

$$\begin{aligned}\mathbb{E}X &= p/(1-q) + pq/(1-q) + pq^2/(1-q) + \dots \\ &= 1 + q + q^2 + \dots \\ &= 1/(1-q) \\ &= 1/p,\end{aligned}$$

same as before. $\square$

<2.2> **Example.** The "HHH versus TTHH" Example in Chapter 1 solved the following problem:

> Imagine that I have a fair coin, which I toss repeatedly. Two players, M and R, observe the sequence of tosses, each waiting for a particular pattern on consecutive tosses: M waits for hhh, and R waits for tthh. The one whose pattern appears first is the winner. What is the probability that M wins?

The answer—that M has probability 5/12 of winning—is slightly surprising, because, at first sight, a pattern of four appears harder to achieve than a pattern of three.

A calculation of expected values will add to the puzzlement. As you will see, if the game is continued until each player sees his pattern, it takes tthh longer (on average) to appear than it takes hhh to appear. However, when the two patterns are competing, the tthh pattern is more likely to appear first. How can that be?

For the moment forget about the competing hhh pattern: calculate the expected number of tosses needed before the pattern tthh is obtained with four successive tosses. That is, if we let X denote the number of tosses required then the problem asks for the expected value $\mathbb{E}X$.

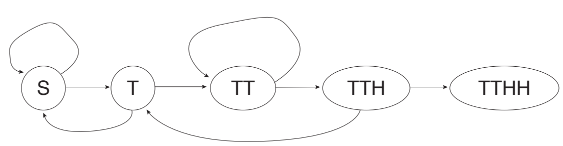

The Markov chain diagram keeps track of the progress from the starting state (labelled S) to the state TTHH where the pattern is achieved. Each arrow in the diagram corresponds to a transition between states with probability 1/2. The corresponding transition matrix is:

$$P = \begin{array}{c} \\ \boxed{\text{S}} \\ \boxed{\text{T}} \\ \boxed{\text{TT}} \\ \boxed{\text{TTH}} \\ \boxed{\text{TTHH}} \end{array} \begin{array}{c} \begin{array}{ccccc} \boxed{\text{S}} & \boxed{\text{T}} & \boxed{\text{TT}} & \boxed{\text{TTH}} & \boxed{\text{TTHH}} \end{array} \\ \begin{pmatrix} 1/2 & 1/2 & 0 & 0 & 0 \\ 1/2 & 0 & 1/2 & 0 & 0 \\ 0 & 0 & 1/2 & 1/2 & 0 \\ 0 & 1/2 & 0 & 0 & 1/2 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} \end{array}.$$

Once again it is easier to solve not just the original problem, but a set of problems, one for each starting state. Let

$$\begin{aligned} \mathcal{E}_S &= \mathbb{E}(X \mid \text{start at S}) \\ \mathcal{E}_H &= \mathbb{E}(X \mid \text{start at H}) \\ &\vdots \end{aligned}$$

Then the original problem is asking for the value of $\mathcal{E}_S$.

To solve gthe problem, condition on the outcome of the first toss, writing $\mathcal{H}$ for the event {first toss lands heads} and $\mathcal{T}$ for the event {first toss lands tails}. From rule E4 for expectations,

$$\begin{aligned} \mathcal{E}_S = \mathbb{E}(X \mid \text{start at S}, \mathcal{T})\mathbb{P}(\mathcal{T} \mid \text{start at S}) \\ + \mathbb{E}(X \mid \text{start at S}, \mathcal{H})\mathbb{P}(\mathcal{H} \mid \text{start at S}) \end{aligned}$$

Both the conditional probabilities equal 1/2 ("fair coin"; probability does not depend on the state). For the first of the conditional expectations, count 1 for the first toss, then recognize that the remaining tosses are just those needed to reach TTHH starting from the state T:

$$\mathbb{E}(X \mid \text{start at S}, \mathcal{T}) = 1 + \mathbb{E}(X \mid \text{start at T})$$

Don't forget to count the first toss. An analogous argument leads to an analogous expression for the second conditional expectation. Substitution into the expression for $\mathcal{E}_S$ then gives

$$\mathcal{E}_S = 1/2(1 + \mathcal{E}_T) + 1/2(1 + \mathcal{E}_S)$$

Similarly,

$$\begin{aligned} \mathcal{E}_T &= 1/2(1 + \mathcal{E}_{TT}) + 1/2(1 + \mathcal{E}_S) \\ \mathcal{E}_{TT} &= 1/2(1 + \mathcal{E}_{TT}) + 1/2(1 + \mathcal{E}_{TTH}) \\ \mathcal{E}_{TTH} &= 1/2(1 + 0) + 1/2(1 + \mathcal{E}_T) \end{aligned}$$

What does the zero in the last equation represent?

The four linear equations in four unknowns have the solution $\mathcal{E}_S = 16$, $\mathcal{E}_T = 14$, $\mathcal{E}_{TT} = 10$, $\mathcal{E}_{TTH} = 8$. Thus, the solution to the original problem is that the expected number of tosses to achieve the tthh pattern is 16. □

<2.3> **Example.** *Expected number of tosses to get hhh with fair coin is 14.*

I could use the same method as for the tthh problem but I want to show you a variation on the method that is easier to generalize. It involves a lot more notation, but it captures better the recursive nature of the problem.

First relabel the states: $S_0 = S$, $S_1 = H$, $S_2 = HH$, and $S_3 = HHH$. Then write X_k for the number of steps to reach state S_k and define

$$\tau_k = \mathbb{E}(X_k \mid \text{start at } S_0).$$

Clearly X_{k+1} is bigger than X_k, so the random variable $Y_{k+1} = X_{k+1} - X_k$ is nonnegative. The variation works by calculating $\mathbb{E}(Y_{k+1} \mid \text{start at } S_0)$.

For each integer m write $\mathcal{H}_m$ for the event that the mth toss results in a head and $\mathcal{T}_m$ for $\mathcal{H}_m^c$. I claim that

$$\mathbb{E}(Y_{k+1} \mid X_k = m, \mathcal{H}_{m+1}) = 1$$
$$\mathbb{E}(Y_{k+1} \mid X_k = m, \mathcal{T}_{m+1}) = 1 + \tau_{k+1}.$$

The second equality reflects the fact that the tail sends us right back to the start.

By rule (E4),

$$\begin{aligned}\mathbb{E}(Y_{k+1} \mid \text{start at } S_0)&\\ &= \sum\nolimits_m \mathbb{E}(Y_{k+1} \mid X_k = m, \mathcal{H}_{m+1}, \text{start at } S_0)\times\\ &\qquad \mathbb{P}(X_k = m, \mathcal{H}_{m+1} \mid \text{start at } S_0)\\ &\quad + \sum\nolimits_m \mathbb{E}(Y_{k+1} \mid X_k = m, \mathcal{T}_{m+1}, \text{start at } S_0)\times\\ &\qquad \mathbb{P}(X_k = m, \mathcal{T}_{m+1} \mid \text{start at } S_0)\\ &= \sum\nolimits_m \mathbb{P}(X_k = m)\left((1/2 \times 1) + 1/2 \times (1 + \tau_{k+1})\right)\end{aligned}$$

Here I have used conditional independence of the events $\{X_k = m\}$ and $\mathcal{H}_{m+1}$ (or $\mathcal{T}_{m+1}$) given "start at S_0". Assuming that X_k is finite with probability one (given "start at S_0"), the probabilities $\mathbb{P}(X_k = m)$ sum to one, leaving

$$\mathbb{E}(Y_{k+1} \mid \text{start at } S_0) = 1 + \tau_{k+1}/2.$$

We now have a recursive equation

$$\tau_{k+1} = \mathbb{E}(X_k + Y_{k+1} \mid \text{start at } S_0) = \tau_k + 1 + \tau_{k+1}/2$$

or $\tau_{k+1} = 2\tau_k + 2$. If you are very brave you might use $\tau_0 = 0$, otherwise you would appeal to Example <2.1> to get $\tau_1 = 2$. The recursive equality then gives $\tau_2 = 6$ and $\tau_3 = 14$, as asserted. □

<2.4> **Example.** *In order to encourage consumers to buy many packets of cereal, a manufacurer includes a Famous Probabilist card in each packet. There are 10 different types of card: Chung, Feller, Lévy, Kolmogorov, ..., Doob. Suppose that I am seized by the desire to own at least one card of each type. What is the expected number of packets that I need to buy in order to achieve my goal?*

Assume that the manufacturer has produced enormous numbers of cards, the same number for each type. (If you have ever tried to collect objects of this type, you might doubt the assumption about equal numbers. But, without it, the problem becomes exceedingly difficult.) The assumption ensures, to a good approximation, that the cards in different packets are independent, with probability 1/10 for a Chung, probability 1/10 for a Feller, and so on.

The high points in my life occur at random "times" T_1, $T_1 + T_2$, ..., $T_1 + T_2 + \cdots + T_{10}$, when I add a new type of card to my collection: After one card (that is, $T_1 = 1$) I have my first type; after another T_2 cards I will get something different from the first card; after another T_3 cards I will get a third type; and so on.

The question asks for $\mathbb{E}(T_1 + T_2 + \cdots + T_{10})$, which rule E1 (applied repeatedly) reexpresses as $\mathbb{E}T_1 + \mathbb{E}T_2 + \cdots + \mathbb{E}T_{10}$.

The calculation for $\mathbb{E}T_1$ is trivial because T_1 must equal 1: we get $\mathbb{E}T_1 = 1$ by rule (E2). Consider the mechanism controlling T_2. For concreteness suppose the first card was a Doob. Each packet after the first is like a coin toss with probability 9/10 of getting a head (= a nonDoob), with T_2 like the number of tosses needed to get the first head. Thus

$$T_2 \text{ has a geometric}(9/10) \text{ distribution.}$$

Deduce from Example <2.1> that $\mathbb{E}T_2 = 10/9$, a value slightly larger than 1.

Now consider the mechanism controlling T_3. Condition on everything that was observed up to time $T_1 + T_2$. Under the assumption of equal abundance and enormous numbers of cards, most of this conditioning information is acually irrelevent; the mechanism controlling T_3 is independent of the past information. (Hard question: Why would the T_2 and T_3 mechanisms not be independent if the cards were not equally abundant?) So what is that T_3 mechanism? I am waiting for any one of the 8 types I have not yet collected. It is like coin tossing with probability 8/10 of heads:

$$T_3 \text{ has geometric } (8/10) \text{ distribution,}$$

and thus $\mathbb{E}T_3 = 10/8$.

Remark. More precisely, T_3 is independent of T_2 with conditional probability distribution geometric (8/10). That is, with $p = 8/10$,

$$\mathbb{P}\{T_3 = k \mid T_2 = \ell\} = (1-p)^{k-1}p \qquad \text{for } k = 1, 2, \ldots$$

for every possible ℓ.

And so on, leading to

$$\mathbb{E}T_1 + \mathbb{E}T_2 + \cdots + \mathbb{E}T_{10} = 1 + 10/9 + 10/8 + ... + 10/1 \approx 29.3.$$

I should expect to buy about 29.3 packets to collect all ten cards. □

Note: The independence between packets was **not** needed to justify the appeal to rule (E1), to break the expected value of the sum into a sum of expected values. It did allow me to recognize the various geometric distributions without having to sort through possible effects of large T_2 on the behavior of T_3, and so on.

You might appreciate better the role of independence if you try to solve a similar (but much harder) problem with just two sorts of card, not in equal proportions.

<2.5> **Example.** *Suppose an urn contains r red balls and b black balls, all identical except for color. Suppose you remove one ball at a time, without replacement, at each step selecting at random from the urn: if k balls remain then each has probability* $1/k$ *of being chosen.Show that the expected number of red balls removed before the firstblack ball equals* $r/(b+1)$.

The problem might at first appear to require nothing more than a simple application of rule (E5)' for expectations. We shall see. Let T be the number of reds removed before the first black. Find the distribution of T, then appeal to E5′ to get

$$\mathbb{E}T = \sum\nolimits_k k\mathbb{P}\{T = k\}.$$

Sounds easy enough. We have only to calculate the probabilities $\mathbb{P}\{T = k\}$.

Define $R_i = \{i\text{th ball red}\}$ and $B_i = \{i\text{th ball black}\}$. The possible values for T are $0, 1, \dots, r$. For k in this range,

$$\begin{aligned}
\mathbb{P}\{T = k\} &= \mathbb{P}\{\text{first k balls red, (k+1)st ball is black}\} \\
&= \mathbb{P}(R_1 R_2 \dots R_k B_{k+1}) \\
&= (\mathbb{P}R_1)\mathbb{P}(R_2 \mid R_1)\mathbb{P}(R_3 \mid R_1 R_2) \dots \mathbb{P}(B_{k+1} \mid R_1 \dots R_k) \\
&= \frac{r}{r+b} . \frac{r-1}{r+b-1} \dots \frac{b}{r+b-k}.
\end{aligned}$$

The dependence on k is fearsome. I wouldn't like to try multiplying by k and summing. If you are into pain you might try to continue this line of argument. Good luck.

There is a much easier way to calculate the expectation, by breaking T into a sum of much simpler random variables for which (E5)' is trivial to apply. This approach is sometimes called the ***method of indicators***.

Suppose the red balls are labelled $1, \dots, r$. Let T_i equal 1 if red ball number i is sampled before the first black ball, zero otherwise. That is, T_i is the indicator for the event

$$\{\text{red ball number } i \text{ is removed before any of the black balls}\}.$$

(Be careful here. The black balls are not thought of as numbered. The first black ball is not a ball bearing the number 1; it might be any of the b black balls in the urn.) Then $T = T_1 + \dots + T_r$. By symmetry—it is assumed that the numbers have no influence on the order in which red balls are selected—each T_i has the same expectation. Thus

$$\mathbb{E}T = \mathbb{E}T_1 + \dots + \mathbb{E}T_r = r\mathbb{E}T_1.$$

For the calculation of $\mathbb{E}T_1$ we can ignore most of the red balls. The event $\{T_1 = 1\}$ occurs if and only if red ball number 1 is drawn before all b of the black balls. By symmetry, the event has probability $1/(b+1)$. (If $b+1$ objects are arranged in random order, each object has probability $1/(1+b)$ of appearing first in the order.)

Remark. If you are not convinced by the appeal to symmetry, you might find it helpful to consider a thought experiment where all $r+b$ balls are numbered and they are removed at random from the urn. That is, treat all the balls as distinguishable and sample until the urn is empty. (You might find it easier to follow the argument in a particular case, such as all $120 = 5!$ orderings for five distinguishable balls, 2 red and 3 black.) The sample space consists of all permutations

> of the numbers 1 to $r+b$. Each permutation is equally likely. For each permutation in which red 1 precedes all the black balls there is another equally likely permutation, obtained by interchanging the red ball with the first of the black balls chosen; and there is an equally likely permutation in which it appears after two black balls, obtained by interchanging the red ball with the second of the black balls chosen; and so on. Formally, we are partitioning the whole sample space into equally likely events, each determined by a relative ordering of red 1 and all the black balls. There are $b+1$ such equally likely events, and their probabilities sum to one.

Now it is easy to calculate the expected value for red 1.

$$\mathbb{E}T_1 = 0\,\mathbb{P}\{T_1 = 0\} + 1\,\mathbb{P}\{T_1 = 1\} = 1/(b+1)$$

The expected number of red balls removed before the first black ball is equal to $r/(b+1)$. □

<2.6> **Example.** *Suppose two players, Alf (A for short) and Betamax (B for short), bet on the tosses of a fair coin: for a head, Alf pays Betamax one dollar; for a tail, Betamax pays Alf one dollar. They stop playing when one player runs out of money. If Alf starts with α dollar bills, and Betamax starts with β dollars bills (both α and β whole numbers), what is the probability that Alf ends up with all the money?*

Write X_n for the number of dollars held by A after n tosses. (Of course, once the game ends the value of X_n stays fixed from then on, at either $a+b$ or 0, depending on whether A won or not.) It is a random variable taking values in the range $\{0, 1, 2, \ldots, a+b\}$. We start with $X_0 = \alpha$. To solve the problem, calculate $\mathbb{E}X_n$, for very large n in two ways, then equate the answers. We need to solve for the unknown $\theta = \mathbb{P}\{\text{A wins}\}$.

First calculation

Invoke rule (E4) with the sample space broken into three pieces,

$$\begin{aligned}
A_n &= \{\text{A wins at, or before, the } n\text{th toss}\},\\
B_n &= \{\text{B wins at, or before, the } n\text{th toss}\},\\
C_n &= \{\text{game still going after the } n\text{th toss}\}.
\end{aligned}$$

For very large n the game is almost sure to be finished, with $\mathbb{P}A_n \approx \theta$, $\mathbb{P}B_n \approx 1-\theta$, and $\mathbb{P}C_n \approx 0$. Thus

$$\begin{aligned}
\mathbb{E}X_n &= \mathbb{E}(X_n \mid A_n)\mathbb{P}A_n + \mathbb{E}(X_n \mid B_n)\mathbb{P}B_n + \mathbb{E}(X_n \mid C_n)\mathbb{P}C_n\\
&\approx \big((\alpha+\beta)\times\theta\big) + \big(0\times(1-\theta)\big) + \big((\text{something})\times 0\big).
\end{aligned}$$

The error in the approximation goes to zero as n goes to infinity.

Second calculation

Calculate conditionally on the value of X_{n-1}. That is, split the sample space into disjoint events $F_k = \{X_{n-1} = k\}$, for $k = 0, 1, \ldots, a+b$, then work towards another appeal to rule (E4). For $k = 0$ or $k = a+b$, the game will be over, and X_n must take the same value as X_{n-1}. That is,

$$\mathbb{E}(X_n \mid F_0) = 0 \qquad \text{AND} \qquad \mathbb{E}(X_n \mid F_{\alpha+\beta}) = \alpha + \beta.$$

For values of k between the extremes, the game is still in progress. With the next toss, A's fortune will either increase by one dollar (with probability 1/2) or decrease by one dollar (with probability 1/2). That is, for $k = 1, 2, \ldots, \alpha + \pm 1$,

$$\mathbb{E}(X_n \mid F_k) = {}^1\!/\!_2(k+1) + {}^1\!/\!_2(k-1) = k.$$

Now invoke (E4).

$$E(X_n) = (0 \times \mathbb{P}F_0) + (1 \times \mathbb{P}F_1) + \cdots + (\alpha+\beta)\mathbb{P}F_{\alpha+\beta}.$$

Compare with the direct application of (E5)' to the calculation of EX_{n-1}:

$$\begin{aligned}\mathbb{E}(X_{n-1}) = \big(0 \times \mathbb{P}\{X_{n-1} = 0\}\big) + \big(1 \times \mathbb{P}\{X_{n-1} = 1\}\big) + \\ \cdots + \big((\alpha+\beta) \times \mathbb{P}\{X_{n-1} = \alpha+\beta\}\big),\end{aligned}$$

which is just another way of writing the sum for $\mathbb{E}X_n$ derived above. Thus we have

$$\mathbb{E}X_n = \mathbb{E}X_{n-1}$$

The expected value doesn't change from one toss to the next.

Follow this fact back through all the previous tosses to get

$$\mathbb{E}X_n = \mathbb{E}X_{n-1} = \mathbb{E}X_{n-2} = \cdots = \mathbb{E}X_2 = \mathbb{E}X_1 = \mathbb{E}X_0.$$

But X_0 is equal to α, for certain, which forces $\mathbb{E}X_0 = \alpha$.

Putting the two answers together

We have two results: $\mathbb{E}X_n = \alpha$, no matter how large n is; and $\mathbb{E}X_n$ gets arbitrarily close to $\theta(\alpha+\beta)$ as n gets larger. We must have $\alpha = \theta(\alpha+\beta)$. That is, Alf has probability $\alpha/(\alpha+\beta)$ of eventually winning all the money. □

Remark. Twice I referred to the sample space, without actually having to describe it explicitly. It mattered only that several conditional probabilities were determined by the wording of the problem.

Danger: The next two Examples are harder. They can be skipped.

<2.7> **Example.** Same problem as in Example <2.6>, except that the coin they toss has probability $p \neq 1/2$ of landing heads.

The cases $p = 0$ and $p = 1$ are trivial. So let us assume that $0 < p < 1$ (and $p \neq 1/2$). Essentially De Moivre's idea was that we could use almost the same method as in Example <2.6> if we kept track of A's fortune on a geometrically expanding scaled. For some number s, to be specified soon, consider a new random variable $Z_n = s^{X_n}$.

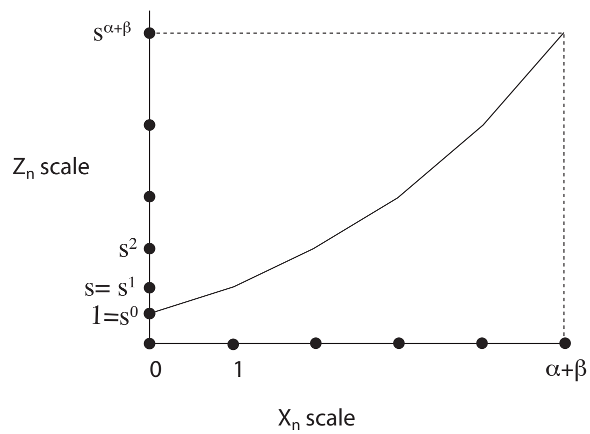

Once again write θ for $\mathbb{P}\{\text{A wins}\}$, and give the events A_n, B_n, and C_n the same meaning as in Example <2.6>.

As in the first calculation for the other Example, we have

$$\begin{aligned}\mathbb{E}Z_n &= \mathbb{E}(s^{X_n} \mid A_n)\mathbb{P}A_n + \mathbb{E}(s^{X_n} \mid B_n)\mathbb{P}B_n + \mathbb{E}(s^{X_n} \mid C_n)\mathbb{P}C_n \\ &\approx \left(s^{\alpha+\beta} \times \theta\right) + \left(s^0 \times (1-\theta)\right) + \left((\text{something}) \times 0\right)\end{aligned}$$

if n is very large.

For the analog of the second calculation, in the cases where the game has ended by at or before the $(n-1)$st toss we have

$$\mathbb{E}(Z_n \mid X_{n-1} = 0) = s^0 \qquad \text{AND} \qquad \mathbb{E}(Z_n \mid X_{n-1} = \alpha+\beta) = s^{\alpha+\beta}.$$

For $0 < k < \alpha + \beta$, the result of the calculation is slightly different.

$$\mathbb{E}(Z_n \mid X_{n-1} = k) = ps^{k+1} + (1-p)s^{k-1} = \left(ps + (1-p)s^{-1}\right) s^k.$$

If we choose $s = (1-p)/p$, the factor $\left(ps + (1-p)s^{-1}\right)$ becomes 1. Invoking rule E4 we then get

$$\begin{aligned}\mathbb{E}Z_n &= \mathbb{E}(Z_n \mid X_{n-1} = 0) \times \mathbb{P}(X_{n-1} = 0) + \mathbb{E}(Z_n \mid X_{n-1} = 1) \times \mathbb{P}(X_{n-1} = 1)\\ &\qquad + \dots + \mathbb{E}(Z_n \mid X_{n-1} = \alpha+\beta) \times \mathbb{P}(X_{n-1} = \alpha+\beta)\\ &= s^0 \times \mathbb{P}(X_{n-1} = 0) + s^1 \times \mathbb{P}(X_{n-1} = 1)\\ &\quad + \dots + s^{\alpha+\beta} \times \mathbb{P}(X_{n-1} = \alpha+\beta)\end{aligned}$$

Compare with the calculation of $\mathbb{E}Z_{n-1}$ via (E5).

$$\begin{aligned}\mathbb{E}Z_{n-1} &= \mathbb{E}(s^{X_{n-1}} \mid X_{n-1} = 0) \times \mathbb{P}(X_{n-1} = 0)\\ &\qquad + \mathbb{E}(s^{X_{n-1}} \mid X_{n-1} = 1) \times \mathbb{P}(X_{n-1} = 1)\\ &\qquad + \dots + \mathbb{E}(s^{X_{n-1}} \mid X_{n-1} = \alpha+\beta) \times \mathbb{P}(X_{n-1} = \alpha+\beta)\\ &= s^0 \times \mathbb{P}(X_{n-1} = 0) + s^1 \times \mathbb{P}(X_{n-1} = 1) + \dots\\ &\qquad + s^{\alpha+\beta} \times \mathbb{P}(X_{n-1} = \alpha+\beta)\end{aligned}$$

Once again we have a situation where $\mathbb{E}Z_n$ stays fixed at the initial value $\mathbb{E}Z_0 = s^\alpha$, but, with very large n, it can be made arbitrarily close to $\theta s^{\alpha+\beta} + (1-\theta)s^0$. Equating the two values, we deduce that

$$\mathbb{P}\{\text{Alf wins}\} = \theta = \frac{1-s^\alpha}{1-s^{\alpha+\beta}} \qquad \text{where } s = (1-p)/p.$$

What goes wrong with this calculation if $p = 1/2$? As a check we could let p tend to $1/2$, getting

$$\begin{aligned}\frac{1-s^\alpha}{1-s^{\alpha+\beta}} &= \frac{(1-s)(1+s+\dots+s^{\alpha-1})}{(1-s)(1+s+\dots+s^{\alpha+\beta-1})} \qquad \text{for } s \neq 1\\ &= \frac{1+s+\dots+s^{\alpha-1}}{1+s+\dots+s^{\alpha+\beta-1}}\\ &\to \frac{\alpha}{\alpha+\beta} \qquad \text{as } s \to 1.\end{aligned}$$

Comforted? □

<2.8> **Example.** My interest in the calculations in Example <2.5> was kindled by a problem that appeared in the August-September 1992 issue of the American Mathematical Monthly. My solution to the problem—the one I first

came up with by application of a straightforward conditioning argument—reduces the calculation to several applications of the result from the previous Example. The solution offered by two readers of the Monthly was slicker. The following brown paragraphs are taken hyper-verbatim from the Monthly; I was seeing how closely LaTeX could reproduce the original text.

E 3429 [1991, 264]. *Proposed by Donald E. Knuth and John McCarthy, Stanford University, Stanford, CA.*

A certain pill bottle contains m large pills and n small pills initially, where each large pill is equivalent to two small ones. Each day the patient chooses a pill at random; if a small pill is selected, (s)he eats it; otherwise (s)he breaks the selected pill and eats one half, replacing the other half, which thenceforth is considered to be a small pill.

(a) What is the expected number of small pills remaining when the last large pill is selected?

(b) On which day can we expect the last large pill to be selected?

Solution from AMM:

Composite solution by Walter Stromquist, Daniel H. Wagner, Associates, Paoli, PA and Tim Hesterberg, Franklin & Marshall College, Lancaster, PA. The answers are (a) $n/(m+1)+\sum_{k=1}^{m}(1/k)$, and (b) $2m+n-(n/(m+1))-\sum_{k=1}^{m}(1/k)$. The answer to (a) assumes that the small pill created by breaking the last large pill is to be counted. A small pill present initially remains when the last large pill is selected if and only if it is chosen last from among the $m+1$ element set consisting of itself and the large pills—an event of probability $1/(m+1)$. Thus the expected number of survivors from the original small pills is $n/(m+1)$. Similarly, when the kth large pill is selected $(k=1,2,\dots,m)$, the resulting small pill will outlast the remaining large pills with probability $1/(m-k+1)$, so the expected number of created small pills remaining at the end is $\sum_{k=1}^{m}(1/k)$. Hence the answer to (a) is as above. The bottle will last $2m+n$ days, so the answer to (b) is just $2m+n$ minus the answer to (a), as above.

I offer two alternative methods of solution for the problem. The first method uses a conditioning argument to set up a recurrence formula for the expected numbers of small pills remaining in the bottle after each return of half a big pill. The equations are easy to solve by repeated substitution. The second method uses indicator functions to spell out the Hesterberg-Stromquist method in more detail. Apparently the slicker method was not as obvious to most readers of the Monthly (and me):

> *Editorial comment.* Most solvers derived a recurrence relation, guessed the answer, and verified it by induction. Several commented on the origins of the problem. Robert High saw a version of it in the MIT Technology Review of April, 1990. Helmut Prodinger reports that he proposed it in the Canary Islands in 1982. Daniel Moran attributes the problem to Charles MacCluer of Michigan State University, where it has been known for some time.

Solved by 38 readers (including those cited) and the proposer. One incorrect solution was received.

Conditioning method.

Invent random variables to describe the depletion of the pills. Initially there are $L_0 = n$ small pills in the bottle. Let S_1 small pills be consumed before the first large pill is broken. After the small half is returned to the bottle let there be L_1 small pills left. Then let S_2 small pills be consumed before the next big pill is split, leaving L_2 small pills in the bottle. And so on.

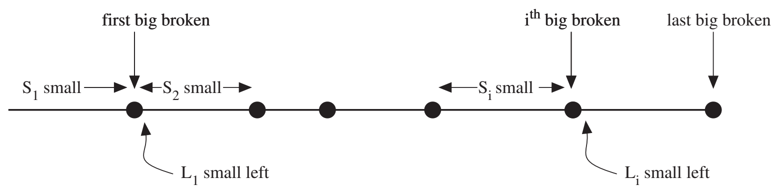

With this notation, part (a) is asking for $\mathbb{E}L_m$. Part (b) is asking for $2m+n-\mathbb{E}L_m$: If the last big pill is selected on day X then it takes $X+L_m$ days to consume the $2m+n$ small pill equivalents, so $\mathbb{E}X+\mathbb{E}L_m = 2m+n$.

The random variables are connected by the equation

$$L_i = L_{i-1} - S_i + 1,$$

the $-S_i$ representing the small pills consumed between the breaking of the $(i-1)$st and ith big pill, and the $+1$ representing the half of the big pill that is returned to the bottle. Taking expectations we get

<2.9> $$\mathbb{E}L_i = \mathbb{E}L_{i-1} - \mathbb{E}S_i + 1.$$

The result from Example <2.5> will let us calculate $\mathbb{E}S_i$ in terms of $\mathbb{E}L_{i-1}$, thereby producing the recurrence formula for $\mathbb{E}L_i$.

Condition on the pill history up to the $(i-1)$st breaking of big pill (and the return of the unconsumed half to the bottle). At that point there are

L_{i-1} small pills and $m-(i-1)$ big pills in the bottle. The mechanism controlling S_i is just like the urn problem of Example <2.5>, with

$$
\begin{aligned}
&r = L_{i-1} \text{ red balls (= small pills)}\\
&b = m-(i-1) \text{ black balls (= big pills)}.
\end{aligned}
$$

From that Example,

$$\mathbb{E}\{S_i \mid \text{history to } (i-1)\text{st breaking of a big pill}\} = L_{i-1}1 + m - (i-1).$$

To calculate $\mathbb{E}S_i$ we would need to average out using weights equal to the probability of each particular history:

$$\mathbb{E}S_i = \frac{1}{1+m-(i-1)} \sum_{\text{histories}} \mathbb{P}\{\text{history}\}(\text{value of } L_{i-1} \text{ for that history}).$$

The sum on the right-hand side is exactly the sum we would get if we calculated $\mathbb{E}L_{i-1}$ using rule E4, partitioning the sample space according to possible histories up to the $(i-1)$st breaking of a big pill. Thus

$$\mathbb{E}S_i = \frac{1}{2+m-i}\mathbb{E}L_{i-1}.$$

Now we can eliminate $\mathbb{E}S_i$ from equality <2.9> to get the recurrence formula for the $\mathbb{E}L_i$ values:

$$\mathbb{E}L_i = \left(1 - \frac{1}{2+m-i}\right)\mathbb{E}L_{i-1} + 1.$$

If we define $\theta_i = \mathbb{E}L_i/(1+m-i)$ the equation becomes

$$\theta_i = \theta_{i-1} + \frac{1}{1+m-i} \qquad \text{for } i = 1, 2, \dots, m,$$

with initial condition $\theta_0 = \mathbb{E}L_0/(1+m) = n/(1+m)$. Repeated substitution gives

$$
\begin{aligned}
\theta_1 &= \theta_0 + \frac{1}{m}\\
\theta_2 &= \theta_1 + \frac{1}{m-1} = \theta_0 + \frac{1}{m} + \frac{1}{m-1}\\
\theta_3 &= \theta_2 + \frac{1}{m-2} = \theta_0 + \frac{1}{m} + \frac{1}{m-1} + \frac{1}{m-2}\\
&\vdots\\
\theta_m &= \cdots = \theta_0 + \frac{1}{m} + \frac{1}{m-1} + \cdots + \frac{1}{2} + \frac{1}{1}.
\end{aligned}
$$

That is, the expected number of small pills left after the last big pill is broken equals

$$\begin{aligned}\mathbb{E}L_m &= (1+m-m)\theta_m \\ &= \frac{n}{1+m} + 1 + \frac{1}{2} + \cdots + \frac{1}{m}.\end{aligned}$$

Rewrite of the Stromquist-Hesterberg solution.
Think in terms of half pills, some originally part of big pills. Number the original half pills $1,\dots,n$. Define

$$H_i = \begin{cases} +1 & \text{if original half pill } i \text{ survives beyond last big pill} \\ 0 & \text{otherwise.}\end{cases}$$

Number the big pills $1,\dots,m$. Use the same numbers to refer to the half pills that are created when a big pill is broken. Define

$$B_j = \begin{cases} +1 & \text{if created half pill } j \text{ survives beyond last big pill} \\ 0 & \text{otherwise.}\end{cases}$$

The number of small pills surviving beyond the last big pill equals

$$H_1 + \cdots + H_n + B_1 + \cdots + B_m.$$

By symmetry, each H_i has the same expected value, as does each B_j. The expected value asked for by part (a) equals

<2.10> $$n\mathbb{E}H_1 + m\mathbb{E}B_1 = n\mathbb{P}\{H_1 = 1\} + m\mathbb{P}\{B_1 = 1\}.$$

For the calculation of $\mathbb{P}\{H_1 = +1\}$ we can ignore all except the relative ordering of the m big pills and the half pill described by H_1. By symmetry, the half pill has probability $1/(m+1)$ of appearing in each of the $m+1$ possible positions in the relative ordering. In particular,

$$\mathbb{P}\{H_1 = +1\} = \frac{1}{m+1}.$$

For the created half pills the argument is slightly more complicated. If we are given that big pill number 1 the kth amongst the big pills to be broken, the created half then has to survive beyond the remaining $m-k$ big pills. Arguing again by symmetry amongst the $(m-k+1)$ orderings we get

$$\mathbb{P}\{B_1 = +1 \mid \text{big number 1 chosen as kth big}\} = \frac{1}{m-k+1}.$$

Also by symmetry,

$$\mathbb{P}\{\text{big 1 chosen as kth big}\} = \frac{1}{m}.$$

Average out using the conditioning rule E4 to deduce

$$\mathbb{P}\{B_1 = +1\} = \frac{1}{m}\sum_{k=1}^{m}\frac{1}{m-k+1}.$$

Notice that the summands run through the values $1/1$ to $1/m$ in reversed order.

When the values for $\mathbb{P}\{H_1 = +1\}$ and $\mathbb{P}\{B_1 = +1\}$ are substituted into <2.10>, the asserted answer to part (a) results. □

2.4 Appendix: The fair price interpretation of expectations

Consider a situation—a bet if you will—where you stand to receive an uncertain return X. You could think of X as a random variable, a real-valued function on a sample space S. For the moment forget about any probabilities on the sample space S. Suppose you consider $p(X)$ the fair price to pay in order to receive X. What properties must $p(\cdot)$ have?

Your net return will be the random quantity $X - p(X)$, which you should consider to be a ***fair return***. Unless you start worrying about the utility of money you should find the following properties reasonable.

(i) ***fair + fair = fair***. That is, if you consider $p(X)$ fair for X and $p(Y)$ fair for Y then you should be prepared to make both bets, paying $p(X) + p(Y)$ to receive $X + Y$.

(ii) ***constant × fair = fair***. That is, you shouldn't object if I suggest you pay $2p(X)$ to receive $2X$ (actually, that particular example is a special case of (i)) or $3.76p(X)$ to receive $3.76X$, or $-p(X)$ to receive $-X$. The last example corresponds to willingness to take either side of a fair bet. In general, to receive cX you should pay $cp(X)$, for constant c.

(iii) There is no fair bet whose return $X - p(X)$ is always ≥ 0 (except for the trivial situation where $X - p(X)$ is certain to be zero).

If you were to declare a bet with return $X - p(X) \geq 0$ under all circumstances to be fair, I would be delighted to offer you the opportunity to receive the "fair" return $-C\,(X - p(X))$, for an arbitrarily large positive constant C. I couldn't lose.

Fact 1: *Properties (i), (ii), and (iii) imply that $p(\alpha X + \beta Y) = \alpha p(X) + \beta p(Y)$ for all random variables X and Y, and all constants α and β.*

Consider the combined effect of the following fair bets:

you pay me $\alpha p(X)$ to receive αX
you pay me $\beta p(Y)$ to receive βY
I pay you $p(\alpha X + \beta Y)$ to receive $(\alpha X + \beta Y)$.

Your net return is a constant,

$$c = p(\alpha X + \beta Y) - \alpha p(X) - \beta p(Y).$$

If $c > 0$ you violate (iii); if $c < 0$ take the other side of the bet to violate (iii). The asserted equality follows.

Fact 2:*Properties (i), (ii), and (iii) imply that $p(Y) \leq p(X)$ if the random variable Y is always $\leq$ the random variable X.*

If you claim that $p(X) < p(Y)$ then I would be happy for you to accept the bet that delivers

$$(Y - p(Y)) - (X - p(X)) = -(X - Y) - (p(Y) - p(X)),$$

which is always < 0.

The two Facts are analogous to rules E1 and E3 for expectations. You should be able to deduce the analog of E2 from (iii).

As a special case, consider the bet that returns 1 if an event F occurs, and 0 otherwise. If you identify the event F with the random variable taking the value 1 on F and 0 on F^c (that is, the indicator of the event F), then it follows directly from Fact 1 that $p(\cdot)$ is additive: $p(F_1 \cup F_2) = p(F_1) + p(F_2)$ for disjoint events F_1 and F_2, an analog of rule P4 for probabilities.

Contingent bets

Things become much more interesting if you are prepared to make a bet to receive an amount X, but only when some event F occurs. That is, the bet is made ***contingent*** on the occurrence of F. Typically, knowledge of the occurrence of F should change the fair price, which we could denote by $p(X \mid F)$. Let me write Z for the indicator function of the event F, that is,

$$Z = \begin{cases} 1 & \text{if event } F \text{ occurs} \\ 0 & \text{otherwise} \end{cases}$$

Then the net return from the contingent bet is $(X - p(X \mid F))\, Z$. The indicator function Z ensures that money changes hands only when F occurs.

By combining various bets and contingent bets, we can deduce that an analog of rule E4 for expectations: if S is partitioned into disjoint events $F_1, \ldots, F_k$, then

$$p(X) = \sum_{i=1}^{k} p(F_i) p(X \mid F_i).$$

Make the following bets. Write c_i for $p(X \mid F_i)$.

(a) For each i, pay $c_i p(F_i)$ in order to receive c_i if F_i occurs.

ritem[(b)] Pay $-p(X)$ in order to receive $-X$.

(c) For each i, make a bet contingent on F_i: pay c_i (if F_i occurs) to receive X.

If event F_k occurs, your net profit will be

$$-\sum_i c_i p(F_i) + c_k + p(X) - X - c_k + X = p(X) - \sum_i c_i p(F_i),$$

which does not depend on k. Your profit is always the same constant value. If the constant were nonzero, requirement (iii) for fair bets would be violated.

If you rewrite $p(X)$ as the expected value $\mathbb{E}X$, and $p(F)$ as $\mathbb{P}F$ for an event F, and $\mathbb{E}(X \mid F)$ for $p(X \mid F)$, you will see that the properties of fair prices are completely analogous to the rules for probabilities and expectations. Some authors take the bold step of interpreting probability theory as a calculus of fair prices. The interpretation has the virtue that it makes sense in some situations where there is no reasonable way to imagine an unlimited sequence of repetions from which to calculate a long-run frequency or average.

See de Finetti (1974) for a detailed discussion of expectations as fair prices.

References

de Finetti, B. (1974). *Theory of Probability*, Volume 1. New York: Wiley.