This manual includes GOAL
Customizable, versatile online exam question bank.
Thousands of questions!

Access your exclusive ***Study**Plus*$^+$ bonus content:
GOAL | Flashcards | Formula sheet

a/s/m
Exam STAM Study Manual

New! Video Solutions
Now includes FREE video solutions by Abe Weishaus

Abraham Weishaus, Ph.D., F.S.A., C.F.A., M.A.A.A.

a/s/m
Actuarial Study Materials
Learning Made Easier

Adaptable, customizable, online actuarial exam prep tool.

- 10,000+ exam style problems with detailed solutions
- Adaptive quizzes
- 3 learning modes
- 3 difficulty modes
- Practice. Quiz. Test. Pass!

Free with your ACTEX or ASM Study Manual

- GOAL is an e-Learning test prep tool for students to practice skills learned in class or from independent study
- The online platform offers a database of SOA & CAS exam-style problems with detailed solutions
- GOAL offers instructor support as well as performance reporting and tracking to monitor your progress
- Don't forget to check out GOAL Score to see how ready you are for the exam

Available now for P, FM, IFM, LTAM, STAM, SRM, MAS-I, MAS-II, CAS 5, CAS 6U, & CAS 6C

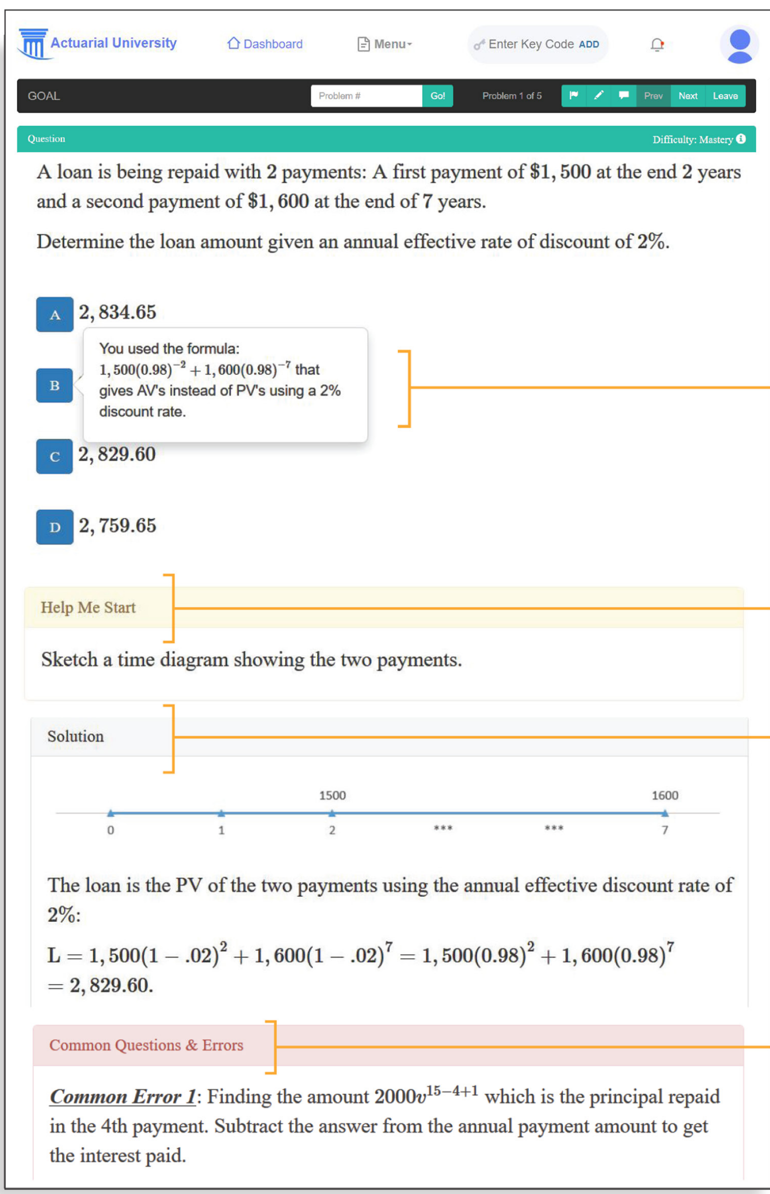

Flag for review, record notes & email the professor

Monitor difficulty level

Got it wrong? Often there is a simple reason why

Helpful strategies to get you started

Graphs and other solution techniques demonstrated when applicable

Commonly encountered errors

Start Your Exam Prep Today!

Preface

Exam STAM discusses the mathematics of short-term insurance. Short-term insurance is insurance for periods of one year or less. Even if the insurance is renewed from year to year, premium rates are usually updated every year. There is no cost for switching insurers from one year to the next. Contrast this with life insurance.

As a candidate for ASA and FSA, you probably won't be working in a property/casualty insurance company. And the typical insurance discussed in this course is auto insurance. But some of the methods we discuss are useful for medical and dental insurance, products sold by life insurance companies. And some concepts, such as credibility, are useful for mortality studies.

Prerequisites for most of the material are few beyond knowing probability (and calculus of course). Some elementary statistics will be helpful, especially for Part II of the manual.

This manual

The exercises in this manual

I've provided lots of my own exercises, as well as relevant exercises from old exams. Though the style of exam questions has changed a little, these are still very useful practice exercises which cover the same material—don't dismiss them as obsolete!

All SOA or joint exam questions in this manual from exams given in 2000 and later, with solutions, are also available on the web from the SOA. When the 2000 syllabus was established in 1999, sample exams 3 and 4 were created, consisting partially of questions from older exams and partially of new questions, not all multiple choice. These sample exams were not real exams, and some questions were inappropriate or defective. These exams are no longer posted on the web. I have included appropriate questions, labeled "1999 C3 Sample" or "1999 C4 Sample". *These refer to these 1999 sample exams, not to the 328 sample questions currently posted on the web, which are discussed later in this introduction.*

Questions from old exams are marked xxx:yy, where xxx is the time the exam was given, with S for spring and F for fall followed by a 2-digit year, and yy is the question number. Sometimes xxx is preceded with SOA or CAS to indicate the sponsoring organization. From about 1986 to 2000, SOA exams had 3-digit numbers (like 160) and CAS exams were a number and a letter (like 4B). From 2000 to Spring 2003, exam 3 was jointly sponsored, so I do not indicate "SOA" or "CAS" for exam 3 questions from that period. There was a period in the 1990's when the SOA, while releasing old exam questions, did not indicate which exam they came from. As a result, I sometimes cannot identify the source exam for questions from this period. In such a case, I mark the question aaa-bb-cc:yy, where aaa-bb-cc is the study note number and yy is the question number. Generally aaa is the exam number (like 160), and cc is the 2-digit year the study note was published.

Other Useful Features of This Manual

The SOA site has a set of 328 sample questions and solutions.[1] Almost all of these questions are from released exams that are readily available; nevertheless many students prefer to use this list since non-syllabus material has been removed. Appendix B has a complete cross reference between these questions and the exams they come from, as well as the page in this manual having the question.

This manual has an index. Whenever you remember some topic in this manual but can't remember where you saw it, check the index. If it isn't in the index but you're sure it's in the manual and an index listing would be appropriate, contact the author.

[1] Actually less than 328. Some questions were mistakenly put on the original list and deleted, and questions pertaining to topics not on the STAM syllabus were deleted.

New for The Third Edition

The CAS redesigned their website. As part of this redesign they removed pre-2011 exams. This edition adds questions from CAS-Exams 3 given in 2005 and 2006. Previous editions only included solutions to these questions.

Tables

Download the tables you will be given on the exam. They will often be needed for the examples and the exercises; I take the information in these tables for granted. If you see something in the text like "That distribution is a Pareto and therefore the variance is . . . ", you should know that I am getting this from the tables; you are not expected to know this by heart. So please download the tables.

The direct address of the tables at this writing is

https://www.soa.org/Files/Edu/2019/2019-02-exam-stam-tables.pdf

The tables include distribution tables and the following statistical tables: the normal distribution function and chi-square critical values. They also include 5 formulas.

The distribution tables are an abbreviated version of the *Loss Models* appendix. Whenever this manual refers to the tables from the *Loss Models* appendix in this manual, the abbreviated version will be sufficient.

The tables (on the second page) specify rules for using the normal distribution table that is supplied: *Do not interpolate in the table. Simply use the nearest value.* If you are looking for $\Phi(0.0244)$, use $\Phi(0.02)$. If you are given the cumulative probability $\Phi(x) = 0.8860$ and need x, use 1.21, the nearest x available. The examples, exercises, and quizzes in this manual use this rounding method. On real exams, they will try to avoid ambiguous situations, so borderline situations won't occur, but my interpretation of the rules (used for problems in this manual) is that if the third place is 5, round up the absolute value. So I round 0.125 to 0.13 and -0.125 to -0.13.

Flashcards

Many students find flashcards a useful tool for learning key formulas and concepts. ASM flashcards, available from the same distributors that sell this manual, contain the formulas and concepts from this manual in a convenient deck of cards. The cards have crossreferences, usually by page, to the manual.

Notes About the Exam

Topic weights

The syllabus breaks the course down into 9 topics. Their weights and the lessons in the manual that cover them are:

Topic	Weight	Lessons
1. Severity Models	2.5–7.5%	1–4
2. Frequency Models	2.5–7.5%	19–21
3. Aggregate Models	2.5–7.5%	22–27
4. Coverage Modifications	2.5–7.5%	11–13, 17–18
5. Risk Measures	2.5–7.5%	16
6. Construction and Selection of Parametric Models	20–30%	29–36
7. Credibility	20–25%	38–54
8. Insurance and Reinsurance Coverage	5–10%	5–6, 15
9. Pricing and Reserving for Short-Term Insurance Coverages	15–25%	7–10, 14–15

Since the exam has 35 questions, each 2.5% is a little less than one question.

Guessing penalty

There is no guessing penalty on this exam. So fill in every answer—you may be lucky! Leave yourself a couple of seconds to do this.

Table 1: Twelve Week Study Schedule for Exam STAM

Week	Subject	Lessons	Rarely Tested
1	Probability basics	1–4	1.4, 2.2, *4.1.3*
2	Short term insurances and loss reserves	5–8	
3	Ratemaking	9–10	
4	Severity modifications	12–14	
5	Reinsurance, risk measures, and frequency	15–21	16.4,*18*,19.2,*21.1*
6	Aggregate loss	22–27	*25*,*27.2*
7	Maximum likelihood	29–30	
8	Maximum likelihood (continued) and hypothesis testing	31–36	*31.1.3*,*32.4*
9	Classical and discrete Bayesian credibility	38–41	39
10	Continuous Bayesian credibility	42–46	44
11	Bühlmann credibility	47–50	
12	Bühlmann credibility	51–54	52

Calculators

A wide variety of calculators are permitted: the TI-30Xa, TI-30X II battery or solar, TI-30X MultiView battery or solar, the BA-35 (battery or solar), and the BA-II Plus (or BA II Plus Professional Edition). You may bring several calculators into the exam. The MultiView calculator is considered the best one, due to its data tables which allow fast statistical calculations. The data table is a very restricted spreadsheet. Despite its limitations, it is useful.

I've provided several examples of using the data table of the Multiview calculator to speed up calculations.

Another feature of the Multiview is storage of previous calculations. They can be recalled and edited.

Other features which may be of use are the K constant and the table feature, which allows calculation of a function at selected values or at values in an arithmetic progression.

Financial calculations do not occur on this exam; interest is almost never considered. You will not miss the lack of financial functions on the Multiview.

Study Schedule

Different students will have different speeds and different constraints, so it's hard to create a study schedule useful for everybody. However, I offer a sample 12-week study schedule, Table 1, as a g uide. The last column lists rarely tested materials so you can skip those if you are behind in your schedule. Italicized sections in this column are, in my opinion, extremely unlikely exam topics.

The Future of This Exam

The last administration of this exam will be in June 2022. The SOA hasn't yet released final syllabi for the exams under the new syllabus, but probably some of the material will move to Exam FAM and some will move to Exam ASTAM. Exam ASTAM will be optional, in the sense that students will have the option to take either ASTAM or ALTAM.

Errata

Please report any errors you find. Reports may be sent to the publisher (`mail@studymanuals.com`) or directly to me (`errata@aceyourexams.net`).

An errata list will be posted at `http://errata.aceyourexams.net`

Acknowledgements

I wish to thank the Society of Actuaries and the Casualty Actuarial Society for permission to use their old exam questions. These questions are the backbone of this manual.

I wish to thank Donald Knuth, the creator of TeX, Leslie Lamport, the creator of LaTeX, and the many package writers and maintainers, for providing a typesetting system that allows such beautiful typesetting of mathematics and figures. I hope you agree, after looking at mathematical material apparently typed with Word (e.g., the Dean study note) that there's no comparison in appearance.

I wish to thank Geoff Tims for proofreading the new STAM material in the body of the manual and Michael Bean for proofreading the new-for-STAM practice exam questions.

I wish to thank the many readers who have sent in errata, or who have reported them anonymously at the Actuarial Discussion Forum, for this manual. A partial list of readers who sent in errata is: German Altgelt, Samuel Anang, Elizabeth Arment, Rachid Belhachemi, Rey Chou, Vo Duy Cuong, Dylan Eilertson, Gavin Ferguson, Matthew Gorham, Brett Harmelink, Minh Ho, Dennis Jerry, Jeffery Lam, Chad Larson, Jonathan Lim, Yin Sze Lim, Karine Prud'homme, Jason Rosenzweig, Leilei Shao, Yuting Sun, Kyle Tenke, Leong Joon Wah, Steven Yap.

I thank Professor Homer White for help with Example 31D.

Lesson 2

Parametric Distributions

Reading: *Loss Models* Fourth Edition or *Loss Models* Fifth Edition 4, 5.1, 5.2.1–5.2.3, 5.3—5.4

A *parametric distribution* is one that is defined by a fixed number of parameters. Examples of parametric distributions are the exponential distribution (parameter θ) and the Pareto distribution (parameters α, θ). Any distribution listed in the *Loss Models* appendix is parametric.

The alternative to a parametric distribution is a *data-dependent distribution*. A data-dependent distribution is one where the specification requires at least as many "parameters" as the number of data points in the sample used to create it; the bigger the sample, the more "parameters".

It is traditional to use parametric distributions for claim counts (frequency) and loss size (severity). Parametric distributions have many advantages. One of the advantages of parametric distributions which makes them so useful for severity is that they handle inflation easily.

2.1 Scaling

A parametric distribution is a member of a *scale family* if any positive multiple of the random variable has the same form. In other words, the distribution function of cX, for c a positive constant, is of the same form as the distribution function of X, but with different values for the parameters. Sometimes the distribution can be parametrized in such a way that only one parameter of cX has a value different from the parameters of X. If the distribution is parametrized in this fashion, so that the only parameter of cX having a different value from X is θ, and the value of θ for cX is c times the value of θ for X, then θ is called a *scale parameter*.

All of the continuous distributions in the tables (Appendix A) are scale families. The parametrizations given in the tables are often different from those you would find in other sources, such as your probability textbook. They are parametrized so that θ is the scale parameter. Thus when you are given that a random variable has any distribution in the appendix and you are given the parameters, it is easy to determine the distribution of a multiple of the random variable.

The only distributions not parametrized with a scale parameter are the lognormal and the inverse Gaussian. Even though the inverse Gaussian has θ as a parameter, it is not a scale parameter. The parametrization for the lognormal given in the tables is the traditional one. *If you need to scale a lognormal, proceed as follows: if X is lognormal with parameters (μ, σ), then cX is lognormal with parameters $(\mu + \ln c, \sigma)$.*

To scale a random variable not in the tables, you'd reason as follows. Let $Y = cX$, $c > 0$. Then

$$F_Y(y) = \Pr(Y \le y) = \Pr(cX \le y) = \Pr\left(X \le \frac{y}{c}\right) = F_X\left(\frac{y}{c}\right)$$

One use of scaling is in handling inflation. In fact, handling inflation is the only topic in this lesson that is commonly tested directly. If loss sizes are inflated by $100r\%$, the inflated loss variable Y will be $(1+r)X$, where X is the pre-inflation loss variable. For a scale family with a scale parameter, you just multiply θ by $(1+r)$ to obtain the new distribution.

Example 2A Claim sizes expressed in dollars follow a two-parameter Pareto distribution with parameters $\alpha = 5$ and $\theta = 90$. A euro is worth \$1.50.

Calculate the probability that a claim will be for 20 euros or less.

Solution: If claim sizes in dollars are X, then claim sizes in euros are $Y = X/1.5$. The resulting euro-based random variable Y for claim size will be Pareto with $\alpha = 5$, $\theta = 90/1.5 = 60$. The probability that a claim will be no more than 20 euros is

$$\Pr(Y \le 20) = F_Y(20) = 1 - \left(\frac{60}{60+20}\right)^5 = \boxed{\mathbf{0.7627}}$$

□

Example 2B Claim sizes in 2019 follow a lognormal distribution with parameters $\mu = 4.5$ and $\sigma = 2$. Claim sizes grow at 6% uniform inflation during 2020 and 2021.

Calculate $f(1000)$, the probability density function at 1000, of the claim size distribution in 2021.

Solution: If X is the claim size random variable in 2019, then $Y = 1.06^2 X$ is the revised variable in 2021. The revised lognormal distribution of Y has parameters $\mu = 4.5 + 2\ln 1.06$ and $\sigma = 2$. The probability density function at 1000 is

$$\begin{aligned} f_Y(1000) &= \frac{1}{\sigma(1000)\sqrt{2\pi}} e^{-(\ln 1000-\mu)^2/2\sigma^2} \\ &= \frac{1}{(2)(1000)\sqrt{2\pi}} e^{-[\ln 1000-(4.5+2\ln 1.06)]^2/2(2^2)} \\ &= (0.000199471)(0.518814) = \boxed{\mathbf{0.0001035}} \end{aligned}$$

□

Example 2C Claim sizes expressed in dollars follow a lognormal distribution with parameters $\mu = 3$ and $\sigma = 2$. A euro is worth \$1.50.

Calculate the probability that a claim will be for 100 euros or less.

Solution: If claim sizes in dollars are X, then claim sizes in euros are $Y = X/1.5$. As discussed above, the distribution of claim sizes in euros is lognormal with parameters $\mu = 3 - \ln 1.5$ and $\sigma = 2$. Then

$$F_Y(y) = \Phi\left(\frac{\ln 100 - 3 + \ln 1.5}{2}\right) = \Phi(1.01) = \boxed{\mathbf{0.8438}}$$

□

Example 2D Claim sizes X initially follow a distribution with distribution function:

$$F_X(x) = 1 - \frac{1}{e^{0.01x}(1+0.01x)} \qquad x > 0$$

Claim sizes are inflated by 50% uniformly.

Calculate the probability that a claim will be for 60 or less after inflation.

Solution: Let Y be the increased claim size. Then $Y = 1.5X$, so $\Pr(Y \le 60) = \Pr(X \le 60/1.5) = F_X(40)$.

$$F_X(40) = 1 - \frac{1}{1.4e^{0.4}} = \boxed{\mathbf{0.5212}}$$

□

2.2 Transformations

Students report that there have been questions on transformations of random variables on recent exams. However, you only need to know the simplest case, how to transform a single random variable using a monotonic function.

If $Y = g(X)$, with $g(x)$ a one-to-one monotonically increasing function, then

$$F_Y(y) = \Pr(Y \le y) = \Pr\big(X \le g^{-1}(y)\big) = F_X\big(g^{-1}(y)\big) \tag{2.1}$$

and differentiating,

$$f_Y(y) = f_X\big(g^{-1}(y)\big)\frac{\mathrm{d}g^{-1}(y)}{\mathrm{d}y}$$

If $g(x)$ is one-to-one monotonically decreasing, then

$$F_Y(y) = \Pr(Y \le y) = \Pr\bigl(X \ge g^{-1}(y)\bigr) = S_X\bigl(g^{-1}(y)\bigr) \tag{2.2}$$

and differentiating,

$$f_Y(y) = -f_X\bigl(g^{-1}(y)\bigr)\frac{\mathrm{d}g^{-1}(y)}{\mathrm{d}y}$$

Putting both cases (monotonically increasing and monotonically decreasing) together:

$$f_Y(y) = f_X\bigl(g^{-1}(y)\bigr)\left|\frac{\mathrm{d}g^{-1}(y)}{\mathrm{d}y}\right| \tag{2.3}$$

Example 2E X follows a two-parameter Pareto distribution with parameters α and θ. You are given

$$Y = \ln\left(\frac{X}{\theta} + 1\right)$$

Determine the distribution of Y.

Solution:

$$\begin{aligned}
y &= \ln\left(\frac{x}{\theta} + 1\right)\\
e^y - 1 &= \frac{x}{\theta}\\
x &= \theta(e^y - 1)\\
F_Y(y) &= F_X\bigl(\theta(e^y - 1)\bigr)\\
&= 1 - \left(\frac{\theta}{\theta + \theta(e^y - 1)}\right)^\alpha\\
&= 1 - \left(\frac{\theta}{\theta e^y}\right)^\alpha\\
&= 1 - e^{-\alpha y}
\end{aligned}$$

So Y's distribution is exponential with parameter $\theta = 1/\alpha$. □

We see in this example that an exponential can be obtained by transforming a Pareto.

There are a few specific transformations that are used to create distributions:

1. If the transformation $Y = X^\tau$ is applied to a random variable X, with τ a positive real number, then the distribution of Y is called *transformed*. Thus when we talk about transforming a distribution we may be talking about any transformation, but if we talk about a transformed Pareto, say, then we are talking specifically about raising the random variable to a positive power.

2. If the transformation $Y = X^{-1}$ is applied to a random variable X, then the distribution of Y is prefaced with the word *inverse*. Some examples you will find in the tables are inverse exponential, inverse Weibull, and inverse Pareto.

3. If the transformation $Y = X^\tau$ is applied to a random variable X, with τ a negative real number, then the distribution of Y is called *inverse transformed*.

4. If the transformation $Y = e^X$ is applied to a random variable X, we name Y with the name of X preceded with "log". The lognormal distribution is an example.

As an example, let's develop the distribution and density functions of an inverse exponential. Start with an exponential with parameter θ:

$$\begin{aligned} F(x) &= 1 - e^{-x/\theta} \\ f(x) &= \frac{e^{-x/\theta}}{\theta} \end{aligned}$$

and let $y = 1/x$. Notice that this is a one-to-one monotonically decreasing transformation, so when transforming the density function, we will multiply by the negative of the derivative. Then

$$\begin{aligned} F_Y(y) &= \Pr(Y \le y) = \Pr(X \ge 1/y) = S_X(1/y) = e^{-1/(y\theta)} \\ f_y(y) &= f_x(1/y)\left|\frac{\mathrm{d}x}{\mathrm{d}y}\right| = \frac{e^{-1/(y\theta)}}{\theta y^2} \end{aligned}$$

However, θ is no longer a scale parameter after this transformation. Therefore, the tables in the appendix use the reciprocal of θ as the parameter and call it θ:

$$\begin{aligned} F_Y(y) &= e^{-\theta/y} \\ f_y(y) &= \frac{\theta e^{-\theta/y}}{y^2} \end{aligned}$$

As a result of the change in parametrization, the negative moments of the inverse exponential, as listed in the tables, are different from the corresponding positive moments of the exponential. Even though $Y = X^{-1}$, the formula for $\mathbf{E}[Y^{-1}]$ is different from the one for $\mathbf{E}[X]$ because the θ's are not the same.

To preserve the scale parameters,[1] the transformation should be done after the random variable is divided by its scale parameter. In other words

1. Set $Y/\theta = (X/\theta)^\tau$ for a transformed random variable.
2. Set $Y/\theta = (X/\theta)^{-1}$ for an inverse random variable.
3. Set $Y/\theta = (X/\theta)^{-\tau}$ for an inverse transformed random variable.
4. Set $Y/\theta = e^{X/\theta}$ for a logged random variable.

Let's redo the inverse exponential example this way.

$$\begin{aligned} \frac{Y}{\theta} &= \left(\frac{X}{\theta}\right)^{-1} \\ F_Y(y) &= \Pr(Y \le y) \\ &= \Pr\left(\frac{Y}{\theta} \le \frac{y}{\theta}\right) \\ &= \Pr\left(\frac{X}{\theta} \ge \frac{\theta}{y}\right) \\ &= \Pr\left(X \ge \frac{\theta^2}{y}\right) \\ &= e^{-\theta^2/y\theta} \\ &= e^{-\theta/y} \end{aligned}$$

[1]This method was shown to me by Ken Burton

Table 2.1: Summary of Scaling and Transformation Concepts

- If a distribution has a scale parameter θ and X has that distribution with parameter θ, then cX has the same distribution with parameter $c\theta$.
- All continuous distributions in the exam tables has scale parameter θ except for lognormal and inverse Gaussian.
- If X is lognormal with parameters μ and σ, then cX is lognormal with parameters $\mu + \ln c$ and σ.
- If $Y = g(X)$ and g is monotonically increasing, then

$$F_Y(y) = \Pr(Y \le y) = \Pr\bigl(X \le g^{-1}(y)\bigr) = F_X\bigl(g^{-1}(y)\bigr) \tag{2.1}$$

- If $Y = g(X)$ and g is monotonically decreasing, then

$$F_Y(y) = \Pr(Y \le y) = \Pr\bigl(X \ge g^{-1}(y)\bigr) = S_X\bigl(g^{-1}(y)\bigr) \tag{2.2}$$

- If $Y = g(X)$ and g is monotonically increasing or decreasing, then

$$f_Y(y) = f_X\bigl(g^{-1}(y)\bigr)\left|\frac{\mathrm{d}g^{-1}(y)}{\mathrm{d}y}\right| \tag{2.3}$$

2.3 Common parametric distributions

The tables provide a lot of information about the distributions, but if you don't recognize the distribution, you won't know to use the table. Therefore, it is a good idea to be familiar with the common distributions.

You should familiarize yourself with the *form* of each distribution, but not necessarily the constants. The constant is forced so that the density function will integrate to 1. If you know which distribution you are dealing with, you can figure out the constant. To emphasize this point, in the following discussion, we will use the letter c for constants rather than spelling out what the constants are. You are not trying to recognize the constant; you are trying to recognize the form.

We will mention the means and variances or second moments of the distributions. You need not memorize any of these. The tables give you the raw moments. You can calculate the variance as $\mathbf{E}[X^2] - \mathbf{E}[X]^2$. However, for frequently used distributions, you may want to memorize the mean and variance to save yourself some time when working out questions.

We will graph the distributions. You are not responsible for graphs, but they may help you understand the distributions.

The tables occasionally use the gamma function $\Gamma(x)$ in the formulas for the moments. You should have a basic knowledge of the gamma function; if you are not familiar with this function, see the sidebar. The tables also use the incomplete gamma and beta functions, and define them, but you can get by without knowing them.

2.3.1 Uniform

A uniform distribution has a constant density on $[d, u]$:

$$f(x; d, u) = \frac{1}{u-d} \qquad d \le x \le u$$

$$F(x; d, u) = \begin{cases} 0 & x \le d \\ \dfrac{x-d}{u-d} & d \le x \le u \\ 1 & x \ge u \end{cases}$$

The gamma function

The gamma function $\Gamma(x)$ is a generalization to real numbers of the factorial function, defined by

$$\Gamma(x) = \int_0^\infty u^{x-1} e^{-u} \mathrm{d}u$$

For positive integers n,

$$\Gamma(n) = (n-1)!$$

The most important relationship for $\Gamma(x)$ that you should know is

$$\Gamma(x+1) = x\Gamma(x)$$

for any real number x.

Example 2F Evaluate $\dfrac{\Gamma(8.5)}{\Gamma(6.5)}$.

Solution:

$$\frac{\Gamma(8.5)}{\Gamma(6.5)} = \left(\frac{\Gamma(8.5)}{\Gamma(7.5)}\right)\left(\frac{\Gamma(7.5)}{\Gamma(6.5)}\right) = (7.5)(6.5) = \boxed{\mathbf{48.75}}$$

□

You recognize a uniform distribution both by its finite support[2] and by the lack of an x in the density function.

Its moments are

$$\mathbf{E}[X] = \frac{d+u}{2}$$

$$\mathrm{Var}(X) = \frac{(u-d)^2}{12}$$

Its mean, median, and midrange are equal. The best way to calculate the second moment is to add up the variance and the square of the mean. However, some students prefer to use the following easy-to-derive formula:

$$\mathbf{E}[X^2] = \frac{1}{u-d}\int_d^u x^2\,\mathrm{d}x = \frac{u^3-d^3}{3(u-d)} = \frac{u^2+ud+d^2}{3} \tag{2.4}$$

If $d = 0$, then the formula reduces to $u^2/3$.

The uniform distribution is not directly in the tables, so I recommend you memorize the formulas for mean and variance. However, if $d = 0$, then the uniform distribution is a special case of a beta distribution with $\theta = u$, $a = 1$, $b = 1$.

2.3.2 Beta

The probability density function of a beta distribution with $\theta = 1$ has the form

$$f(x; a, b) = cx^{a-1}(1-x)^{b-1} \qquad 0 \le x \le 1$$

The parameters a and b must be positive. They may equal 1, in which case the corresponding factor is missing from the density function. Thus if $a = b = 1$, the beta distribution is a uniform distribution.

You recognize a beta distribution both by its finite support—it's the only common distribution for which the density is nonzero only on a finite range of values—and by factors with x and $1-x$ raised to powers and no other use of x in the density function.

[2]"Support" is the range of values for which the probability density function is nonzero.

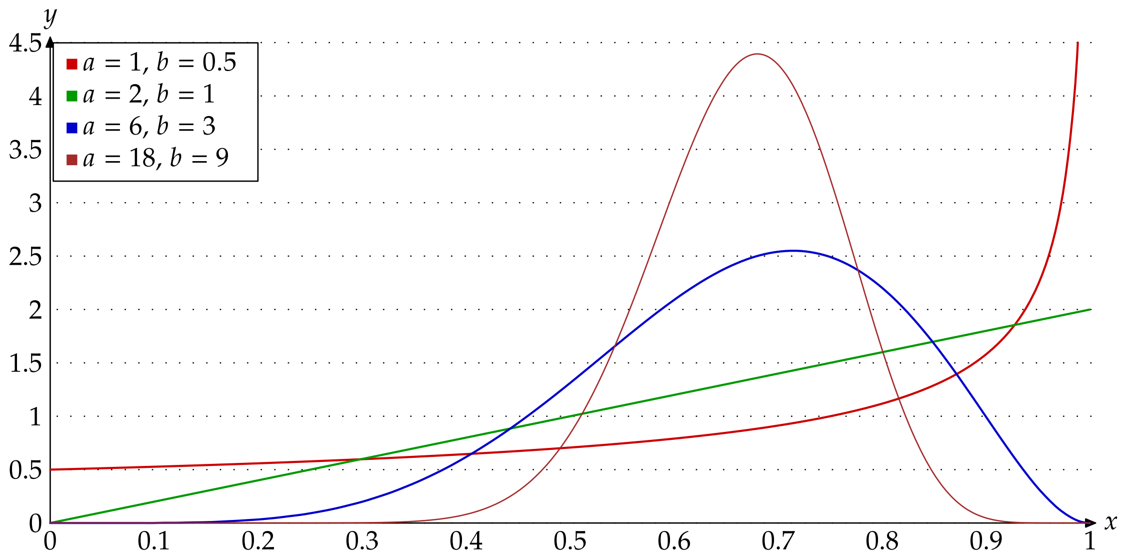

Figure 2.1: Probability density function of four beta distributions with $\theta = 1$ and mean 2/3

If θ is arbitrary, then the form of the probability density function is

$$f(x; a, b, \theta) = cx^{a-1}(\theta - x)^{b-1} \qquad 0 \le x \le \theta$$

The distribution function can be evaluated if a or b is an integer. The moments are

$$\mathbf{E}[X] = \frac{\theta a}{a+b}$$
$$\mathrm{Var}(X) = \frac{\theta^2 ab}{(a+b)^2(a+b+1)}$$

The mode is $\theta(a-1)/(a+b-2)$ when a and b are both greater than 1, but you are not responsible for this fact.

Figure 2.1 graphs four beta distributions with $\theta = 1$ all having mean 2/3. You can see how the distribution becomes more peaked and normal looking as a and b increase.

2.3.3 Exponential

The probability density function of an exponential distribution has the form

$$f(x; \theta) = ce^{-x/\theta} \qquad x \ge 0$$

θ must be positive.

You recognize an exponential distribution when the density function has e raised to a multiple of x, and no other use of x.

The distribution function is easily evaluated. The moments are:

$$\mathbf{E}[X] = \theta$$
$$\mathrm{Var}(X) = \theta^2$$

Figure 2.2 graphs three exponential distributions. The higher the parameter, the more weight placed on higher numbers.

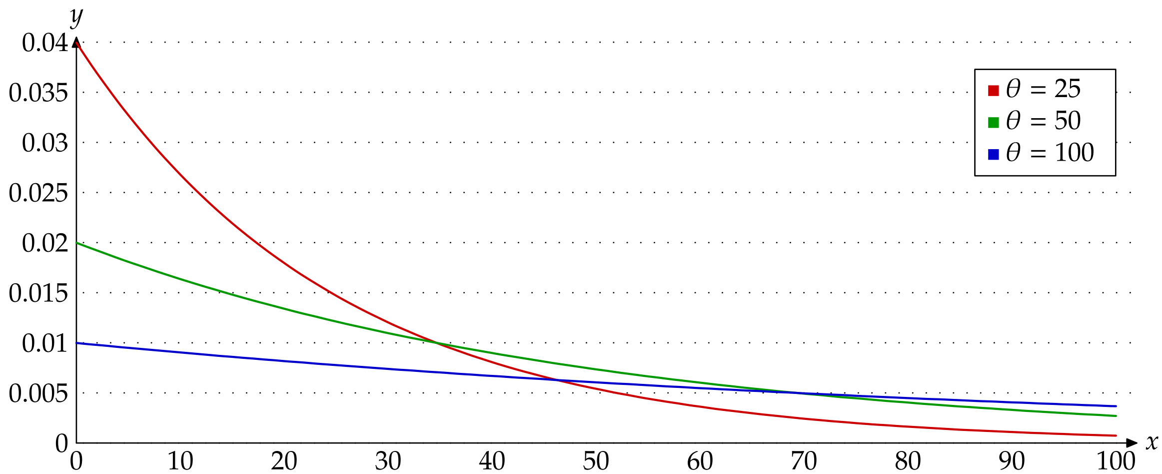

Figure 2.2: Probability density function of three exponential distributions

2.3.4 Weibull

A Weibull distribution is a transformed exponential distribution. If Y is exponential with mean μ, then $X = Y^{1/\tau}$ is Weibull with parameters $\theta = \mu^{1/\tau}$ and τ. An exponential is a special case of a Weibull with $\tau = 1$.

The form of the density function is

$$f(x;\tau,\theta) = cx^{\tau-1}e^{-(x/\theta)^{\tau}} \qquad x \geq 0$$

Both parameters must be positive.

You recognize a Weibull distribution when the density function has e raised to a multiple of a power of x, and in addition has a corresponding power of x, one lower than the power in the exponential, as a factor.

The distribution function is easily evaluated, but the moments require evaluating the gamma function, which usually requires numerical techniques. The moments are

$$\mathbf{E}[X] = \theta\Gamma(1 + 1/\tau)$$
$$\mathbf{E}[X^2] = \theta^2\Gamma(1 + 2/\tau)$$

Figure 2.3 graphs three Weibull distributions with mean 50. The distribution has a non-zero mode when $\tau > 1$. Notice that the distribution with $\tau = 0.5$ puts a lot of weight on small numbers. To make up for this, it will also have to put higher weight than the other two distributions on very large numbers, so although it's not shown, its graph will cross the other two graphs for high x

2.3.5 Gamma

The form of the density function of a gamma distribution is

$$f(x;\alpha,\theta) = cx^{\alpha-1}e^{-x/\theta} \qquad x \geq 0$$

Both parameters must be positive.

When α is an integer, a gamma random variable with parameters α and θ is the sum of α independent exponential random variables with parameter θ. In particular, when $\alpha = 1$, the gamma random variable is exponential. The gamma distribution is called an Erlang distribution when α is an integer. We'll discuss this more in Subsection 27.1.2.

You recognize a gamma distribution when the density function has e raised to a multiple of x, and in addition has x raised to a power. Contrast this with a Weibull, where e is raised to a multiple of a *power* of x.

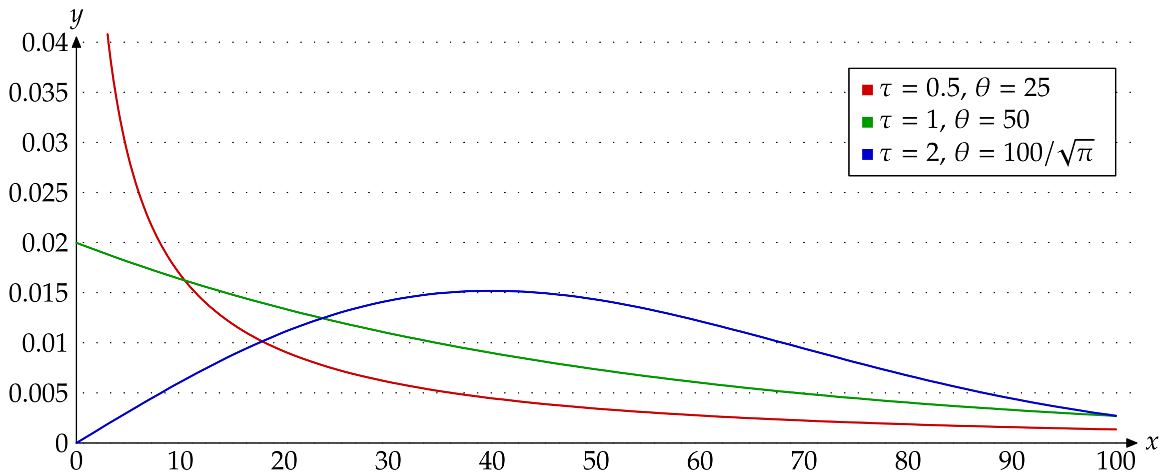

Figure 2.3: Probability density function of three Weibull distributions with mean 50

The distribution function may be evaluated if α is an integer; otherwise numerical techniques are needed. However, the moments are easily evaluated:

$$\mathbf{E}[X] = \alpha\theta$$
$$\mathrm{Var}(X) = \alpha\theta^2$$

Figure 2.4 graphs three gamma distributions with mean 50. As α goes to infinity, the graph's peak narrows and the distribution converges to a normal distribution.

The gamma distribution is one of the few for which the moment generating function has a closed form. In particular, the moment generating function of an exponential has a closed form. The only other distributions in the tables with closed form moment generating functions are the normal distribution (not actually in the tables, but the formula for the lognormal moments is the MGF of a normal) and the inverse Gaussian.

2.3.6 Pareto

When we say "Pareto", we mean a *two-parameter* Pareto. On recent exams, they write out "two-parameter" to make it clear, but on older exams, you will often find the word "Pareto" with no qualifier. It always refers to a two-parameter Pareto, not a single-parameter Pareto.

The form of the density function of a two-parameter Pareto is

$$f(x) = \frac{c}{(\theta + x)^{\alpha+1}} \qquad x \geq 0$$

Both parameters must be positive.

You recognize a Pareto when the density function has a denominator with x plus a constant raised to a power.

The distribution function is easily evaluated. The moments are

$$\mathbf{E}[X] = \frac{\theta}{\alpha - 1} \qquad \alpha > 1$$
$$\mathbf{E}[X^2] = \frac{2\theta^2}{(\alpha-1)(\alpha-2)} \qquad \alpha > 2$$

When α does not satisfy these conditions, the corresponding moments don't exist.

A shortcut formula for the variance of a Pareto is

$$\mathrm{Var}(X) = \mathbf{E}[X]^2\left(\frac{\alpha}{\alpha - 2}\right)$$

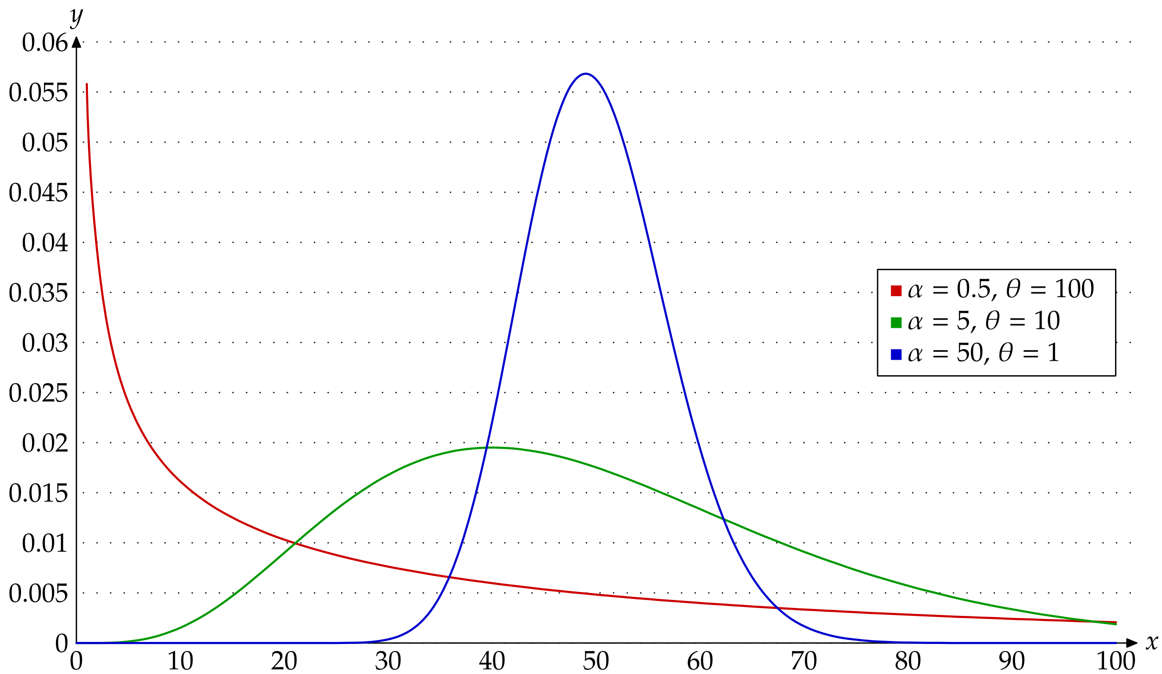

Figure 2.4: Probability density function of three gamma distributions with mean 50

Figure 2.5 graphs three Pareto distributions, one with $\alpha < 1$ and the other two with mean 50. Although the one with $\alpha = 0.5$ puts higher weight on small numbers than the other two, its mean is infinite; it puts higher weight on large numbers than the other two, and its graph eventually crosses the other two as $x \to \infty$.

2.3.7 Single-parameter Pareto

The form of the density function of a single-parameter Pareto is

$$f(x) = \frac{c}{x^{\alpha+1}} \qquad x \geq \theta$$

α must be positive. θ is not considered a parameter since it must be selected in advance, based on what you want the range to be.

You recognize a single-parameter Pareto by the range of nonzero values for its density function—unlike most other distributions, this range does not start at 0—and by the form of the density function, which has a denominator with x raised to a power. A beta distribution may also have x raised to a negative power, but its density function is 0 above a finite number.

A single-parameter Pareto X is a two-parameter Pareto Y shifted by θ: $X = Y + \theta$. Thus it has the same variance, and the mean is θ greater than the mean of a two-parameter Pareto with the same parameters.

$$\mathbf{E}[X] = \frac{\alpha\theta}{\alpha - 1} \qquad \alpha > 1$$

$$\mathbf{E}[X^2] = \frac{\alpha\theta^2}{\alpha - 2} \qquad \alpha > 2$$

2.3.8 Lognormal

The form of the density function of a lognormal distribution is

$$f(x) = \frac{ce^{-(\ln x - \mu)^2/2\sigma^2}}{x} \qquad x > 0$$