

For newer version of some of these notes, cf. ‘under construction’ [online ‘book’](#)

1 First: Overview of Sampling Distributions

1.1 Examples of Sampling Distributions

Distribution	Statistic whose sampling variability it describes
Binomial(n, π)	no. (or prop'n, p) of 1's in s.r.s. of n from infinite-sized universe containing proportion π of 1's & $(1 - \pi)$ of 0's. — can think of p as a <i>mean</i> of n (Bernoulli) 1's and 0's.
Hypergeometric (N_1 1's, N_0 0's)	no. (or prop'n, p) of 1's in s.r.s. of size n from universe of 1's and 0's, but universe size ($N = N_1 + N_0$) is <i>finite</i> .
Poisson(μ)	no. of 1's in s.r.s. of n from infinite-sized universe of 0's and 1's, but where π is small, and n is large, so that $np(1 - p) \approx np = \mu$ [limiting case of Binomial]. No. of ‘events’ in a sampled volume of experience (conditions apply ! – see later).
Gaussian	mean, proportion, count, difference, etc. (n large)
Student's t	$(\bar{y} - \mu) / (s_y / \sqrt{n})$; $y \sim N(\mu, \sigma)$; $s_y^2 = \sum(y - \bar{y})^2 / (n - 1)$.
F	ratio of sample variances (used for ANOVA)
???	order statistic as estimate of quantile

1.2 Ways of calculating sampling variability

- i. directly from the relevant discrete distribution, by adding probabilities of the variations in question, e.g. :
 - $0.010 + 0.001 = 0.011$ Binomial prob. of ≥ 9 1's in $n = 10$ if $\pi = 0.5$.
 - 2.5% probability of (Poisson) count ≥ 5 if $\mu = 1.624$
 - 2.5% probability of (Poisson) count ≤ 5 if $\mu = 11.668$
- ii. from specially-worked out distributions for more complex statistics calculated from continuous or rank data –
 - Student's t , F ratio, χ^2 , distribution of Wilcoxon statistic.

iii. (**very common**) from the **Gaussian approximation** to the relevant discrete or continuous distribution – by using (an estimate of) the standard deviation of the sampling variation in question and assuming the variation is reasonably symmetric and bell-shaped [every sampling distribution has a standard deviation – but it isn't very useful if the distribution is quite skewed or heavy-tailed]. **We give a special name (standard error) to the standard deviation of a sampling distribution in order to distinguish it from the measure of variability of individuals.**¹ Interestingly, we haven't given a special name to the square of the SD of a statistic – we use Variance to denote both SE^2 and SD^2 .

iv. **jack-knife** (only for variance), or **bootstrap** (more detailed picture)

1.3 Standard Error (SE) of a sample statistic

What it is

An estimate of the SD of the *different possible values* of the sample statistic one would obtain in different random samples² of a given size n .

Since we observe only one of the many possible different random samples of a given size, the SD of the sample statistic is not directly measurable: is merely **conceptual**. The concept of ‘repeated sampling’ is not very helpful.

In this course, in computer simulations, and in mathematical statistics courses, we have the luxury of knowing the relevant information about each element in the population and thus the probabilities of all the possible sample statistics. Thus, for example, we can say that **if** individual Y 's are such that $Y \sim N(\mu, \sigma)$, **then** the different possible \bar{y} 's will vary from μ in a certain

¹Note: Up to Ch 5, M&M use the same notation for the SD of a *mean* or a difference of means as they do for the SD of *individuals* – they use ‘SD’ for both. Many texts distinguish the two by using SE (Standard Error) when dealing with the SD of a mean or proportion or other *statistic*, and SD when dealing with *individual* variation. Moore & McCabe in page 500 of Ch 7 say “when the SD of a statistic is *estimated* from the data, the result is called the SE of the statistic.” This is a more restricted definition than many authors use. JH's advice: always say what SD or SE one is referring to: the SD or SE of a mean, SD or SE of a median, the SD or SE of a proportion, the SD or SE of a slope, the SD of individual measurements etc. If one sees a SD on its own i.e., without reference to a specific statistic, one would suspect (but cannot be sure) that it is the SD of individuals. *However a SE is never in relation to individuals; it is always in relation to a statistic.*

²Notice how JH says ‘different possible samples’ rather than ‘in repeated sampling’. Try to avoid this ‘repeated sampling’ notion, and instead think of *what might have been* if the sample has been based on a different starting seed for the random number generator, or starting at a different place in a sampling frame, etc. In real life there is only one sample: one spends one's entire budget on it, and there are no possibilities of ‘repeating’ it or winding back the clock and getting another one instead.

known way. In real life, we don't know the value of μ and are interested in estimating it *using the one sample we are allowed to observe*. Thus the **SE of a statistic** is usually an estimate or a projection of the variation in a *conceptual distribution* i.e. *the SD of all the "might-have-been" statistics*.

Use

If n is large enough, the *different possible values* of the statistic **would** have a Gaussian distribution with a spread of 2-3 SE's on each side of the "true" parameter value [note the "would have"]

So, one can calculate the chance of various deviations from the true value.

Thus, we can assess the range of parameter values for which the observed statistic would or would not be an extreme observation. Note the convoluted legal wording: this is not as satisfactory as we would wish, but under the frequentist paradigm, it is the best we can do. Under the Bayesian paradigm, we would speak directly about the parameter, and where (now that we have new data) we think it is. Under the frequentist approach, we speak about the behaviour of the (new) data: there are no prior data, and there is just a hypothetical value (or a range of hypothetical values) for the parameter.

e.g.

if statistic is \bar{y} , we talk of SE of the mean (SEM)

$SE(\bar{y})$ describes variation of \bar{y} from μ ;

$SD(y)$ describes variation of y from μ (or from \bar{y}).

DIGRESSION: How old is the (statistical) concept of a 'mean' ? ³

The delightful 2016 book *The Seven Pillars of Statistical Wisdom*, by Stephen Stigler, historian of statistics, [and available online at the McGill library] addresses this in his first Chapter (the "Aggregation" pillar).

But when did the scientific analysis of statistical data begin? **When did the use of the arithmetic mean become a formal part of a statistical analysis?** Was it really not much before the seventeenth century? Why was the mean not used to combine observations in some earlier era – in astronomy, surveying, or economics? The mathematics of a mean was certainly known in antiquity. The Pythagoreans knew already in 280 BCE of three kinds of means: the arithmetic, the geometric, and

the harmonic. And by 1000 CE the philosopher Boethius had raised that number to at least ten, including the Pythagorean three. To be sure, these means were deployed in philosophical senses, in discussing proportions between line segments, and in music, not as data summaries.

Surely we might expect the Greeks or the Romans or the Egyptians to have taken means of data in day-to-day life more than two millennia ago. Or, if they did not, surely the mean may be found in the superb astronomical studies of Arabic science a thousand years ago. **But diligent and far-ranging searches for even one well-documented example from those sources have come up empty.**

The most determined seeker of an early use of the mean was the indefatigable researcher Churchill Eisenhart, who spent most of his professional career at the National Bureau of Standards. Over several decades he pursued historical uses of the mean, and he summarized his researches in his presidential address to the American Statistical Association in 1971. His enthusiasm carried the address to nearly two hours, but for all that effort, the earliest documented uses of the mean he found were those I have already mentioned by D. B. and Gellibrand. Eisenhart found that Hipparchus (ca. 150 BCE) and Ptolemy (ca. 150 CE) were silent on their statistical methods; al-Biruni (ca. 1000 CE) gave nothing closer to a mean than using the number produced by splitting the difference between a minimum and a maximum. The mean occurred in practical geometry in India quite early; Brahmagupta, in a tract on mensuration written in 628 CE, suggested approximating the volume of irregular excavations using that of the rectangular solid with the excavation's mean dimensions.

Over all these years, the historical record shows that data of many types were collected. In some cases, inevitably, summaries would be needed; if the mean was not used, what did people do in order to summarize, to settle on a single figure to report? Perhaps we can get a better idea of how this question was seen in pre-statistical times by looking at a few examples where something similar to a mean was employed.

One story told by Thucydides involved siege ladders and dates to 428 BCE: Ladders were made to match the height of the enemy's wall, which they measured by the layers of bricks, the side turned towards them not being thoroughly whitewashed. These were counted by many persons at once; and though some might miss the right calculation, most would hit upon it, particularly as they counted over and over again, and were no great way from the wall, but could see it easily enough for their purpose. The length required for the ladders was thus obtained, being calculated from the breadth of the brick. Thucydides described the use of what we can recognize as the mode – the most frequently reported value. With the expected lack of independence among counts, the mode is not particularly accurate, but if the reports were tightly clustered it was likely as good as any other summary. Thucydides did not give the data.

Another, much later example dates from the early 1500s and is reported

³cf The Early History of Average Values and Implications for Education by A. Bakker in J Stat. Educ. 11(1)2003.

by Jacob Köbel in a finely illustrated book on surveying. As Köbel tells us, the basic unit of land measure in those times was the rod, defined as sixteen feet long. And in those days a foot meant a real foot, but whose foot? Surely not the king's foot, or each change of monarch would require a renegotiation of land contracts. The solution Köbel reports was simple and elegant:

sixteen representative citizens (all male in those days) would be recruited after a church service and asked to stand in a line, toe to heel, and the sixteen-foot rod would be the length of that line. Köbel's picture, etched by himself, is a masterpiece of explanatory art (see Figure 1.7).

It was truly a community rod! And, after the rod was determined, it was subdivided into sixteen equal sections, each representing the measure of a single foot, taken from the communal rod. *Functionally, this was the arithmetic mean of the sixteen individual feet, but nowhere was the mean mentioned.*

Wikipedia [Foot (unit)] also has the picture, and the wording:

Stand at the door of a church on a Sunday and bid 16 men to stop, tall ones and small ones, as they happen to pass out when the service is finished; then make them put their left feet one behind the other, and the length thus obtained shall be a right and lawful rood to measure and survey the land with, and the 16th part of it shall be the right and lawful foot.



2 Sampling Distribution of \bar{y} : Expectation / SD / Shape

- Quantitative variable (characteristic) of interest : Y
- N (effectively) infinite (or sampling with replacement)
- Mean of all Y values in population: μ
- Variance of all Y values in population: σ^2
- **Shape of distribution of Y 's: Unknown/Unspecified**
- Sample of size n ; (i.i.d.) observations y_1, \dots, y_n
- Sample mean: $\bar{y} = (1/n) \sum y_i$

Statistic	E(Statistic)	SD(Statistic)
\bar{y}	μ_y	σ_y / \sqrt{n}

2.1 ?? Shape of the sampling distribution of \bar{y} ??

The sampling distribution is the frequency distribution (e.g. in the form of a histogram or other depiction) we **would** get if we could observe the mean (or any other calculated statistic) of *each of the (infinite number of) different possible random samples of a given size*. It quantifies probabilistically how the different possible values of the statistic *would* vary around some central value. The sampling distribution is *strictly conceptual* (except, for illustration purposes, in toy classroom exercises where we can actually do the 'what if' exercise for all possible samples from some made-up universe of known values).

Relevance of knowing shape of a sampling distribution:

We will only observe the mean in the one sample we chose; however we can, with certain assumptions, mathematically (beforehand) calculate how far the mean (\bar{y}) of a randomly selected sample is likely to be from the mean (μ) of the population. Thus we can say with a specified probability (95% for example) that the \bar{y} that we are about to observe will be no more than Q (some constant, depending on whether we use 90%, 95%, 99%, ...) units from μ . In 'frequentist' inference, we say that in *95% of the applications of our procedure*, our estimate will come within the stated distance of the target, and so we

can have this much ‘confidence’ in the *procedure*. The probability statement associated with the confidence interval for μ is really about the stochastic behaviour of \bar{y} in relation to μ .⁴ We also use the sampling distribution to assess the (probabilistic) distance of a sample mean from some “test” or “Null Hypothesis” value in statistical tests.

2.1.1 Example of the distribution of a sample mean:

When summing (or averaging) n 0’s and 1’s (i.e numbers measured on a 2-point scale), there are only $n+1$ unique possibilities for the result $(0, 1, \dots, n)$. However, if we were studying a variable, e.g. cholesterol or income, that was measured on a continuous scale, the numbers of possible sample means would be very large and not easy to enumerate. For the sake of illustration, we instead take a simpler variable, that is measured on a *discrete* integer scale with a very limited range. However, the principle is the same as for a truly continuous variable.

Imagine we are interested in the average number of cars per household μ in a city area with a large number (N) of households. With an estimate of the average number per household and the total number of households we can then estimate the total number of cars $N \times \mu$. It is not easy to get data on every single one of the N , so we draw a random sample, with replacement, of size n . [The sampling with replacement is simply for the sake of simplicity in this example – we would use sampling without replacement in practice].

How much sampling variation can there be in the estimates we might obtain from the sample? What will the degree of “error” or “noise” depend on? Can we anticipate the magnitude of possible error and the *pattern* of the errors in estimation caused by use of a finite sample?

Suppose that:

⁴Ideally any description of the CI should involve sentences in which \bar{y} is the subject; μ should not be the subject of the sentence. In the ‘frequentist’ approach, we are not allowed to say before (or after) the fact that there is a 95% probability that the target will be (is) within the stated distance of where the estimate lands. If one is pretty sure that a particular location is within 15 Km of downtown Montreal, then it is *mathematically* correct to say that one is pretty sure that downtown Montreal is within 15 Km of the location in question. In the frequentist approach, however, it is not ‘*statistically correct*’ to turn this type of statement around and to say that there therefore is a 95% chance that the population mean (μ , the quantity we would like to make inferences about) will not be more than Q units away from the sample mean (\bar{y}) we (are about to) observe. The reason has to do with the different (asymmetric) logical status of each of the 2 quantities: even though it is unknown, μ is treated as a *fixed* point, while \bar{y} is treated as the *stochastic* element. Thus, for example, if μ were the speed of light, and \bar{y} was a future estimate of it, we cannot speak of μ ‘falling’ randomly somewhere near \bar{y} : instead. In Bayesian inference, it is permitted to speak of the pre-sample and thus the post-sample uncertainty concerning μ .

- 50% have 0 cars,
- 30% have 1 car,
- 20% have 2 cars.

i.e. in all, there are $0.5 \times N$ 0’s, $0.30 \times N$ 1’s, and $0.20 \times N$ 2’s.

You would be correct to object “but how can we know this - this is the point of sampling”; however, this is a purely conceptual or “what if” exercise; the relevance will become clear later.

The **mean** of the entire set of Y ’s is

$$\mu_Y = 0 \times 0.5 + 1 \times 0.3 + 2 \times 0.2 = \mathbf{0.7}.$$

The **variance** of the Y ’s is

$$\begin{aligned} \sigma_Y^2 &= (0 - 0.7)^2 \times 0.5 + (1 - 0.7)^2 \times 0.3 + (2 - 0.7)^2 \times 0.2 \\ &= 0.49 \times 0.5 + 0.09 \times 0.3 + 1.69 \times 0.2 \\ &= \mathbf{0.61}. \end{aligned}$$

[Thus, the **SD**, $\sigma = \sqrt{0.61} = \mathbf{0.78}$ is slightly larger than μ].

We take a s.r.s. of $n = 2$ houses, obtain y_1 and y_2 , and use $\bar{y} = (y_1 + y_2)/2$ as $\hat{\mu}_Y$. What estimates might we obtain?

The distribution of all possible \bar{y} ’s when $n = 2$ is:

Probability (frequency)	$\hat{\mu}$ [i.e., \bar{y}]	error [$\bar{y} - \mu$]	% error [% of μ]
25%	$\frac{0}{2} = 0.0$	-0.7	-100
30%	$\frac{1}{2} = 0.5$	-0.2	-29
29%	$\frac{2}{2} = 1.0$	+0.3	+43
12%	$\frac{3}{2} = 1.5$	+0.8	+114
4%	$\frac{4}{2} = 2.0$	+1.3	+186

Most of the possible estimates of μ from samples of size 2 will be “off the target” by quite serious amounts. It’s not much good saying that “on average, over all possible samples” the sample will produce the correct estimate.

Check:

$$\begin{aligned} \text{Average}[\bar{y}] &= 0 \times 0.25 + 0.5 \times 0.30 + 1.0 \times 0.29 + 1.5 \times 0.12 + 2.0 \times 0.04 \\ &= 0.7 \\ &= \mu \end{aligned}$$

$$\begin{aligned} \text{Variance}[\bar{y}] &= (-0.7)^2 \times 0.25 + \dots (1.3)^2 \times 0.04 \\ &= 0.305 \\ &= \sigma^2/2 \end{aligned}$$

The distribution of all possible \bar{y} 's when $n = 3$ can be found by convoluting the one for $n = 2$ with the one for $n = 1$:

Probability (frequency)	$\hat{\mu}$ [i.e., \bar{y}]	error [$\bar{y} - \mu$]	% error [% of μ]
12.5%	$\frac{0}{3} = 0.00$	-0.7	-100
22.5%	$\frac{1}{3} = 0.33$	-0.37	-53
28.5%	$\frac{2}{3} = 0.67$	-0.03	-4
20.7%	$\frac{3}{3} = 1.00$	+0.30	+43
11.4%	$\frac{4}{3} = 1.33$	+0.63	+90
3.6%	$\frac{5}{3} = 1.66$	+0.96	+137
0.8%	$\frac{6}{3} = 2.00$	+1.30	+186

Check that: Mean $[\bar{y}] = 0.7$ & Var $[\bar{y}] = \sigma^2/3$.

A sample of size $n = 4$ would give less variable estimates. The distribution of the $3^n = 81$ possible sample configurations, and their corresponding estimates of μ can be enumerated manually as:

Probability (frequency)	$\hat{\mu}$ [i.e., \bar{y}]	error [$\bar{y} - \mu$]	% error [% of μ]
6.25%	$\frac{0}{4} = 0.00$	-0.70	-100
15.00%	$\frac{1}{4} = 0.25$	-0.45	-64
23.50%	$\frac{2}{4} = 0.50$	-0.20	-29
23.4%	$\frac{3}{4} = 0.75$	+0.05	+7
17.61%	$\frac{4}{4} = 1.00$	+0.30	+43
9.36%	$\frac{5}{4} = 1.25$	+0.55	+79
3.76%	$\frac{6}{4} = 1.50$	+0.80	+114
0.96%	$\frac{7}{4} = 1.75$	+1.05	+150
0.16%	$\frac{8}{4} = 2.00$	+1.30	+186

Of course, there is still a good chance that the estimate will be a long way from the correct value of $\mu = 0.7$. But the variance or scatter of the possible estimates is less than it would have been had one used $n = 2$.

Check:

$$\begin{aligned} \text{Average}[\bar{y}] &= 0 \times 0.0625 + 0.25 \times 0.15 + \dots + 2.0 \times 0.0016 \\ &= 0.7 \\ &= \mu \end{aligned}$$

$$\begin{aligned} \text{Variance}[\bar{y}] &= (-0.7)^2 \times 0.0625 + (-0.45)^2 \times 0.15 + \dots \\ &= 0.1525 \\ &= \sigma^2/4 \end{aligned}$$

If we are happy with an estimate that is not more than 50% in error, then the above table says that with a sample of $n = 4$, there is a $23.50 + 23.40 + 17.61$ or $\approx 65\%$ chance that our sample will result in an “acceptable” estimate (i.e. within $\pm 50\%$ of μ). In other words, we can be 65% confident that our sample will yield an estimate within 50% of the population parameter μ .

For a given n , we can trade a larger % error for a larger degree of confidence and vice versa e.g. if $n = 4$, we can be 89% confident that our sample will result in an estimate within 80% of or be 25% confident that our sample will result in an estimate within 10% of μ .

If we use a bigger n , we can increase the degree of confidence, or narrow the margin of error (or a mix of the two), since with a larger sample size, the distribution of possible estimates is tighter around μ . With $n = 100$, we can associate a 20% error with a statement of 90% confidence or a 10% error with a statement of 65% confidence.

But one could argue that there are two problems with these calculations: first, **they assumed that we knew both μ and the distribution of the individual Y 's** before we start; second, they used manual enumeration of the possible configurations for a small n and Y 's with a small number (3) of possible integer values.

2.1.2 What about real situations with a sample of 10 or 100 from an unknown distribution of Y on a continuous scale?

The answer can be seen by examining the sampling distributions as a function of n in the ‘cars per household’ example, and in other examples dealing with Y 's with a more continuous distribution (see Colton p103-108, A&B p80-83 and M&M 403-404). All the examples show the following:

- i. As expected, the variation of possible sample means about the (in practice, unknown) target μ is less in larger samples. We can use the variance or SD of \bar{y} to measure this scatter. The SD (scatter) in the possible \bar{y} 's from samples of size n is σ/\sqrt{n} , where σ is the SD of the individual Y 's.

This is true no matter what the shape of the distribution of the individual Y 's.

- ii. If the individual Y 's **DO HAVE** a Gaussian distribution, then the distribution of all possible \bar{y} 's will be Gaussian, even for n 's as low as 2.

BUT...

even if the individual Y 's DO NOT a Gaussian distribution...

the larger the n [and the more symmetric and unimodal the distribution of the individual Y 's is to start with], the more the distribution of possible \bar{y} 's resembles a Gaussian distribution. And for many distributions, this approximation is already quite good for samples of $n = 30$ or fewer.

The sampling distribution of \bar{y} [or of a sample proportion, or slope or correlation, or other statistic created by aggregation of individual observations ..] is, for a large enough n [and under other conditions⁵], close to Gaussian in shape no matter what the shape of the distribution of individual Y values. This phenomenon is referred to as the **CENTRAL LIMIT THEOREM.**

We use the notation $Y \sim \text{Distribution}(\mu_y, \sigma_y)$ as shorthand to say that “ Y has a certain type of distribution with mean μ_y and standard deviation σ_y ”.

In this notation, the Central Limit Theorem says that

$$\text{if } Y \sim ???(\mu_Y, \sigma_Y), \text{ then}$$

$$\bar{y} \sim N(\mu_Y, \sigma_Y/\sqrt{n}), \text{ if } n \text{ is large enough and ...}$$

The Gaussian approximation to certain Binomial distributions is an example of the Central Limit Theorem in action: Individual (Bernoulli) Y 's have a 2-point distribution: a proportion $(1 - \pi)$ have the value $Y = 0$ and the remaining proportion π have $Y = 1$.

The mean (μ) of all (0, 1) Y values in the population is π .

The (unit) variance (σ^2) of all Y values in population

$$\sigma^2 = (0 - \pi)^2 \times (1 - \pi) + (1 - \pi)^2 \times \pi = \pi(1 - \pi)$$

From a sample of size n :

observations y_1, y_2, \dots, y_n (sequence of n 0's and 1's)

$$\text{sample mean } \bar{y} = \frac{\sum y_i}{n} = \frac{\text{number of 1's}}{n} = p.$$

⁵On the degree of symmetry and dispersion of the distribution of the individual Y 's.

CLT ...

If $Y \sim \text{Bernoulli}(\mu = \pi, \sigma = \sqrt{\pi[1 - \pi]})$, then

$p = \bar{y} \sim N(\pi, \sqrt{\pi[1 - \pi]}/\sqrt{n})$ if n is sufficiently ‘large’ and π is not extreme.⁶

Returning to example on estimating $\mu_{\text{cars/household}}$.

If $n = 100$, then the SD of possible \bar{y} ’s from samples of size $n = 100$ is $\sigma/\sqrt{100} = 0.78/10 = 0.078$. Thus, we can approximate the distribution of possible \bar{y} ’s by a Gaussian distribution with a mean of 0.7 and a standard deviation of 0.078, to get ...

		Interval	Prob.	% Error
$\mu \pm 1.00SD(\bar{y})$	0.7 ± 0.078	0.62 to 0.77	68%	$\pm 11\%$
$\mu \pm 1.50SD(\bar{y})$	0.7 ± 0.117	0.58 to 0.81	87%	$\pm 17\%$
$\mu \pm 1.96SD(\bar{y})$	0.7 ± 0.143	0.55 to 0.84	95%	$\pm 20\%$
$\mu \pm 3.00SD(\bar{y})$	0.7 ± 0.234	0.46 to 0.93	99.7%	$\pm 33\%$

[The Gaussian-based intervals are only slightly different from the results of a computer simulation in which we drew samples of size 100 from the above Y distribution]

If this variability in the possible estimates is still not acceptable and we use a sample size of $n = 200$, the standard deviation of the possible \bar{y} ’s is not halved (divided by 2) but rather divided by $\sqrt{2} = 1.4$. We would need to go to $n = 400$ to cut the s.d. down to half of what it is with $n = 100$.

[Notice that in all of this (as long as we sample with replacement, so that the n members are drawn independently of each other), the size of the population (N) didn’t enter into the calculations at all. The errors of our estimates (i.e. how different we are from μ on randomly selected samples) vary directly with σ and inversely with \sqrt{n} . However, if we were interested in estimating $N\mu$ rather than μ , the absolute error would be N times larger, *although the relative error would be the same in the two scales.*]

Message from diagram on next page:

⁶E[no. ‘positive’ = numerator = $\sum y_i$] needs to be sufficiently far ‘inland’ from 0 and from n , and n needs to be large enough that Binomial(n, π) distribution does not have much probability mass on 0 or n , i.e., so that the Gaussian approximation to it does not spill over onto, and thus place substantial probability mass on, silly values such as $\dots -3, -2, -1$ or on $n+1, n+2, \dots$. One Rule of Thumb for when the Gaussian approximation provides a reasonable accurate approximation is that both $n \times \pi \geq 5$ and $n \times (1 - \pi) \geq 5$, i.e. the expected number of ‘positives’ should be ‘inland’ or ‘away from the edge’ by at least 5 from both boundaries.

The variation of the possible sample means is closer to Gaussian than the variation of the individual observations (the panel where we have a mean of $n = 1$ values can be taken as the distribution of individual Y ’s), and the bigger the sample size, the closer to Gaussian: with large enough n , you could not tell from the sampling distribution of the means what the shape of the distribution of the individual ‘parent’ observations was. Averages of $n = 16$ are “effectively” Gaussian in this example. **How ‘fast’ or ‘slowly’ the CLT will ‘kick in’ is a function of how symmetric, or how asymmetric and ‘CLT-unfriendly,’ the distribution of Y is.**

2.1.3 Another example of central limit theorem at work: word lengths

The distribution of the lengths of words has a long right tail (see ‘ $n=1$ ’ panel in Fig 2), but the (sampling) distribution of the possible values of the sample mean when $n = 2$ has less of a long right tail, and the distribution of $\bar{y}_{n=4}$ is less asymmetric and closer to Gaussian, and that of $\bar{y}_{n=16}$ even more so.

You can think of the effects of increasing n as two-fold:

- It makes for a ‘finer’ measuring scale (just as with a ruler with finer gradations). For example, if the Y ’s are recorded with a ‘bin-width’ of δY (integers in our two examples), then the sample mean has a ‘bin-width’ of $\delta Y \div n$.
- (With independence between elements) extreme sums, and thus extreme means, are less likely: with large enough n , there are enough extremes from each end of the distribution that they will tend to cancel each other.

2.1.4 Other examples of central limit theorem in action:

- Lengths of n -th generation copies of the 1-metre bar:

Suppose we use a piece of string (or a large photocopier) to make 2 copies of the 1-metre prototype bar <http://en.wikipedia.org/wiki/Metre>. But suppose that in doing so, we make independent errors of either +1cm or -1cm. From each of these 2, we make 2 second generation copies, again with independent +/- errors of 1cm, and from these 8 third generation copies, etc.. What would be the distribution of the lengths of the 2^{16} 16-th generation copies? They will have a binomial-shape distribution, ranging from 84cm to 116cm, and centered on 100cm. A




Figure 1: Illustration of Central Limit Theorem in the case where Y is the number of cars per household




Figure 2: Another illustration of Central Limit Theorem (harder?) at work

plot this (using

```
plot(100+seq(0,16),dbinom(seq(0,16),16,0.5),type="h")
```

say, in R) you will see that it has the shape of what Gauss called the Law or Errors. Much earlier, de Moivre worked out a normal approximation directly from the binomial in his 1733 pamphlet ‘A Method of approximating the Sum of the Terms of the Binomial $(a + b)^n$ expanded to a Series from whence are deduced some practical Rules to estimate the Degree of Assent which is to be given to Experiments.’ If you make the errors smaller, but have more of them, the variation will be effectively on a continuous scale. One way to establish the Normal density $\phi(y, \mu, \sigma)$ is (as de Moivre did more generally for $\pi \neq 0.5$) to apply Stirling’s formula (http://en.wikipedia.org/wiki/Stirling's_approximation) to the Binomial probabilities in the case of a large n and “success” probability $\pi = 0.5$.

- Generating random numbers from an (approx.) Gaussian distribution:

Since Φ^{-1} , the inverse of the cdf of a $N(0,1)$ random variable, does not have a closed form, the inverse cdf method of obtaining Gaussian random numbers has to rely on an approximation involving powers.⁷ Another way to produce values that have close to a $N(0,1)$ distribution is by summing $n = 12$ realizations from a $U(0,1)$ distribution and subtracting 6 from the sum.

```
# sum of 12 random numbers from U(0,1)
```

```
r = function(dummy) sum(runif(12))-6 ;
sum.12.uniforms = sapply(1:50000,r);
hist(sum.12.uniforms,breaks=50)
```

- There is also a CTL that applies to sums of *independent* but *not identically distributed* random variables. The key element is the *independence*. See the **cartoon “The Central Limit Theorem in Action** (courtesy Lawrence Joseph)” in the [Resources page](#). If the components were correlated, say because of weather, then it would impede the cancellation of extremes.

```
days=2000;
walk.to.bus = rnorm(days,mean=4,sd=1);
wait.for.bus = runif(days,4,16);
bus.ride = rnorm(days,mean=20,sd=2);
walk.up.hill = rgamma(days,scale=2,shape=3/2);
```

⁷For an exact method, see http://en.wikipedia.org/wiki/Box-Muller_transform

```
hist(walk.to.bus); summary(walk.to.bus);
hist(wait.for.bus); summary(wait.for.bus);
hist(bus.ride); summary(bus.ride);
hist(walk.up.hill); summary(walk.up.hill);
```

```
total.time = walk.to.bus + wait.for.bus + bus.ride +
walk.up.hill;
summary(total.time);
c(mean(total.time),sd(total.time),var(total.time))
hist(total.time)
boxplot(total.time)
```

Other examples can be found in §’s [12.7.3](#), [16.3](#), and [18.2.4](#) of ‘online ’book’.

3 Standard Error (SE) of combination or weighted average of estimates

$$\begin{aligned}
 SE(\sum estimates) &= \sqrt{\sum ([SE \text{ of each estimate}]^2)} \\
 SE(constant \times estimate) &= constant \times SE(estimate) \\
 SE(constant + estimate) &= SE(estimate) \\
 SE(\sum w_i \times estimate_i) &= \sqrt{\sum (w_i^2 \times [SE estimate_i]^2)} \quad (1)
 \end{aligned}$$

This last one is important for combining estimates from stratified samples, and for meta-analysis.

In an estimate for the overall population, derived from a stratified sample, the weights are chosen so that the overall estimate is unbiased for the weighted average of the stratum-specific parameters i.e. the w ’s are the relative sizes of the segments (strata) of the overall population (see “combining estimates ... entire population” below). The parameter values usually differ between strata: this is why stratified sampling helps. The *estimate* for this weighted average of the stratum-specific parameters is formed as a weighted average of the age-specific parameter estimates, and so one has no choice in the choice of weights: they must reflect the proportions of population in the various strata.

If instead, one had several estimates of a **single** parameter value (a big assumption in the ‘usual’ approach to meta-analyses), but each estimate had a different uncertainty (precision), one should take a weighted average of them,

but with the weights inversely proportional to the amount of uncertainty in each. From the formula above one can verify by algebra or trial and error that the smallest variance for the weighted average is obtained by using weights proportional to the inverse of the variance (squared standard error) of each estimate. If there is variation in the parameter value, a ‘fixed effects’ SE is too small. The ‘random effects’ approach to meta-analyses weights each estimate in inverse relation to an amalgam of (i) each SE and (ii) the ‘greater-than-random’ variation between estimates [it allows for the possibility that the parameter estimates from each study would not be the same, even if each study used huge n ’s). The SE of this weighted average is larger than that using the simpler (called fixed effects) model; as a result, CI’s are also wider.

3.1 Combining Estimates from Subpopulations to form an Estimate for the Entire Population

Suppose several (say k) sub-populations or “strata” of sizes N_1, N_2, \dots, N_k , form one entire population of size $\sum N_k = N$. Suppose we are interested in the average level of a quantitative characteristic, or the prevalence of a qualitative characteristic in the **entire** population. Denote this numerical or binary characteristic in each individual by Y , and an average or proportion (or total) across all individuals in the population by θ . It could stand for a mean (μ), a total ($T_{amount} = N \times \mu$), a proportion (π), a percentage ($\% = 100\pi$) or a total count ($T_c = N \times \pi$).

Examples:

If Y is a measured variable (i.e. “numerical”)

- μ : the annual (per capita) consumption of cigarettes
- T_{amount} : the total undeclared yearly income
($T_{amount} = N \times \mu$ and conversely $\mu = T_{amount}/N$)

If Y is a binary variable (i.e. “yes / no”)

- π : the proportion of persons who exercise regularly
- $100\pi\%$: the percentage of children who have been fully vaccinated
- $N\pi$: the total number of persons who need R_x for hypertension
($T_c = N\pi$; $\pi = T_c/N$)

The sub-populations might be age groups, the 2 sexes, occupations, provinces, etc. There is a corresponding θ_i for the i -th of the k sub-populations. Rather than study every individual each each stratum, one might instead measure Y in a sample from each stratum.

3.2 Estimate of overall $\mu, \pi, \text{ or } \pi\%$, by combining estimates:

Sub Popln	Size	Relative Size $W_i = N_i/N$	Sample Size	Estimate of θ_i	SE of estimate
1	N_1	W_1	n_1	e_1	$SE(e_1)$
...
...
k	N_k	W_k	n_k	e_k	$SE(e_k)$
Total	$\sum N_i = N$	$\sum W_i = 1$	$\sum n_i = n$	$\sum W_i e_i$	$\sum W_i^2 [SE(e_i)]^2$

Note 1 To estimate T_{amount} or T_c , use weights $W_i = N_i$;

Note 2 If any **sampling fraction** $f_i = n_i/N_i$ is **substantial**, the SE of the e_i should be scaled down i.e. it should be multiplied by $\sqrt{1 - f_i}$.

Note 3 If variability in Y within a stratum is smaller than across strata, the smaller SE obtained from the SE’s of the individual stratum specific estimates more accurately reflects the uncertainty in the overall estimate. Largest gain over SRS is when large inter-stratum variability.

4 (FREQUENTIST) Inference for μ – small n : Student’s t distribution

Use: when we replace σ by s (an estimate of σ) when forming CI’s, or carrying out statistics tests, using the sample mean and the standard error of the mean.⁸ We proceed in the usual way – expressing the distance of \bar{y} from μ in terms of multiples of $SE_{\bar{y}} = s/\sqrt{n}$ – except that we use a different ‘reference’ distribution than the usual Z (Gaussian) one. The percentiles of this new distribution are further from 0 than the familiar 0.84, 1.28, 1.645, 1.96, etc, of the Z distribution: how much further depends on the amount of data, i.e., the $(n - 1)$ used to estimate σ^2 .

To paraphrase, and quote from, Student’s 1908 paper... (*italics* by JH)

(Until now) “the usual method of calculating the probability that “ μ is within a given distance of \bar{x} ”⁹ is to assume $\mu \sim N(\bar{x}, s/\sqrt{n})$, where s is the standard deviation of the sample, and to use the tables of the (Normal) probability integral.” *But, with smaller n , the value of s “becomes itself subject to increasing error.”* In some instances, we can use a more reliable value of s from earlier experiments, but “in some chemical, many biological, and most agricultural and large scale experiments,” we are forced to “judge of the uncertainty of the results from a small sample, which itself affords the only indication of the variability.” Inferential methods for such small-scale experiments had “hitherto been outside the range of statistical enquiry.”

Rather than complain, Gosset [pen-name ‘Student’] did something about it.

Although it is well known that the method of using *the normal curve is only trustworthy when the sample is “large”, no one has yet told us very clearly where the limit between “large” and “small” samples is to be drawn.* The aim of the present paper is to *determine the point at which we may use the tables of the (Normal) probability integral* in judging of the significance of the *mean* of a series of experiments, and to furnish *alternative tables for use when the number of experiments is too few.*

⁸it is also used in a wider context, where we have a ratio of a Gaussian random variable, and the square root of an independent random variable that has a chi-squared distribution.

⁹This way of writing, i.e., of making μ the subject of the sentence, was commonplace in 1908; it is not politically or statistically correct today, unless one adopts a Bayesian viewpoint, where the focus is directly on the pre- and post-data uncertainty concerning μ . [JH]

Student assumed that the Y values are normally distributed, so that \bar{y} has a Gaussian sampling distribution.¹⁰

“Student’s” t distribution is the (conceptual) distribution one would get if one...

- took (an infinite number of) samples, of a given size n , from a $\text{Normal}(\mu, \sigma)$ distribution
- formed the ratio $t = (\bar{y} - \mu) / (s/\sqrt{n})$ from each sample
- compiled a histogram of the ratios.

In fact, to check that his derivation was correct, Gosset¹¹ actually performed a simulation in which he followed the above process:

Before I had succeeded in solving my problem analytically, I had endeavored to do so empirically. The material I used was a ... table containing the height and left middle finger measurements of 3000 criminals.... The measurements were written out on 3000 pieces of cardboard, which were then very thoroughly shuffled and drawn at random... each consecutive set of 4 was taken as a sample... [i.e. $n = 4$ above]... and the mean [and] standard deviation of each sample determined.... This provides us with two sets of... 750 (ratios) on which to test the theoretical results arrived at. The height and left middle finger... table was chosen because the distribution of both was approximately normal...”

Sampling distribution of t

- is symmetric around 0 (just like Z distribution)
- has shape like that of the Z distribution, but with SD slightly larger than unity i.e. slightly flatter & more wide-tailed; $\text{Var}[t] = df/(df - 2)$.
- its shape becomes indistinguishable from that of Z distribution as $n \rightarrow \infty$ (in fact as n goes much beyond 30.)

¹⁰Even if the Y ’s were not normally distributed, but n was sufficiently large, the Central Limit Theorem would guarantee that the distribution of all possible \bar{y} ’s is close to a Gaussian distribution – but with large enough n , one would have sufficient degrees of freedom to estimate σ quite precisely, and so the problem would disappear.

¹¹Student. The probable error of a mean, *Biometrika* 1908. See JH’s website for his 2008 anniversary’ paper and ‘related material, including 1907 & 2007 simulations. The web page containing images from the book containing Gosset’s 750 samples of size $n = 4$ also contains a clue as to **why Gosset may have chosen the pen-name ‘Student.’**

- Instead of $\pm 1.96\sigma/\sqrt{n}$ for 95% confidence, we need

Multiple	n	Degrees of freedom ('df')
± 3.182	4	3
...
± 2.228	11	10
...
± 2.086	21	20
...
± 2.042	31	30
...
± 1.980	121	120
...
± 1.960	∞	∞

Test of $\mu = \mu_0$	Confidence Interval (CI) for μ
$t \text{ ratio} = (\bar{y} - \mu_0)/(s/\sqrt{n})$	$\bar{y} \pm t \times s/\sqrt{n}$

Data:¹²

Subject	Hours of Sleep †		Difference: Drug minus Placebo
	Drug	Placebo	
			d
1	6.1	5.2	0.9
2	7.0	7.9	-0.9
3	8.2	3.9	4.3
4	•	•	2.9
5	•	•	1.2
6	•	•	3.0
7	•	•	2.7
8	•	•	0.6
9	•	•	3.6
10	•	•	-0.5
			$\bar{d} = 1.78$
			$SD[d] = 1.77$

Test statistic: $t = (1.78 - 0)/(1.77/\sqrt{10}) = 3.18.$

Critical Value: $|t_9| = 2.26$

Since $3.18 > |t_9|$, we “reject” H_0 .

95% CI for μ_D : $1.78 \pm t_9 \times SE_{\bar{d}}$
 $1.78 \pm 2.26 \times (1.77/\sqrt{10})$
 1.78 ± 1.26
 0.5 to 3.0 hours

¹² The table deliberately omits the full data on the drug and placebo conditions: this is to emphasize that all we need for the analysis are the 10 differences. Incidentally, Stephen Senn has traced these classic data and found that Student, and Fisher after him, did not describe them correctly: the experiment was a bit more complicated than was initially described.

4.1 WORKED EXAMPLE: CI and Test of Significance

Response of interest: D: Increase (D) in hours of sleep with a test medication.

Test:

$$\begin{aligned} \mu_D &= 0 & H_0 \\ &\neq 0 & H_{alt} \\ \alpha &= 0.05 & 2 \text{ sided} \end{aligned}$$

4.2 Another worked Example, with graphic:

Posture, blood flow, and prophylaxis of venous thromboembolism.

CPG Barker, *The Lancet* Vol 345. April 22, 1995, p. 1047.

Sir-Ashby and colleagues (Feb 18, p 419) report adverse effects of posture on femoral venous blood flow. They noted a moderate reduction velocity when a patient was sitting propped up at 35° in a hospital bed posture and a further pronounced reduction when the patient was sitting with legs dependent. Patients recovering from operations are often asked to sit in a chair with their feet elevated on a footrest. The footrests used in most hospitals, while raising the feet, compress the posterior aspect of the calf. Such compression may be important in the aetiology of venous thrombo-embolism. We investigated the effect of a footrest on blood flow in the deep veins of the calf by dynamic radionuclide venography.

Calf venous blood flow was measured in fifteen young (18-31 years) healthy male volunteers. 88 MBq technetium-99m-labelled pertechnetate in 1 mL saline was injected into the lateral dorsal vein of each foot, with ankle tourniquets inflated to 40 mm Hg, and the time the bolus took to reach the lower border of the patella was measured (Sophy DSX Rectangular Gamma Camera). Each subject had one foot elevated with the calf resting on the footrest and the other plantigrade on the floor as a control. The mean transit time of the bolus to the knee was 24.6 s (SE 2.2) for elevated feet and 14.8 s (SE 2.2) for control feet [see figure 3]. The mean delay was 9.9 s (95% CI 7.8-12.0).

Simple leg elevation without hip flexion increases leg venous drainage and femoral venous blood flow. The footrest used in this study raises the foot by extension at the knee with no change in the hip position. Ashby and colleagues' findings suggest that such elevation without calf compression would produce an increase in blood flow. Direct pressure of the posterior aspect of the calf therefore seems to be the most likely reason for the reduction in flow we observed. Sitting cross-legged also reduced calf venous blood flow, probably by a similar mechanism. If venous stasis is important in the aetiology of venous thrombosis, the practice of nursing patients with their feet elevated on footrests may need to be reviewed.

[Data abstracted from diagram; calculations won't match exactly those in text]

$$\bar{d}(SD) = 9.8(4.1); t = (9.8 - 0)/(4.1/\sqrt{15}) = 9.8/1.0 = 9.8$$

Critical ratio: $t_{14,0.05} = 2.145$. So, the observed difference is 'off the t -scale'. This corroborates the impression gained from visual display of the data.

95% CI for μ_D : $9.8 \pm 2.145 \times 1.0$ i.e., 7.7s to 11.9s.




Figure 3: Raw data, and summary statistics. JH encourages a display like this for showing paired data, and for letting the data themselves tell the story. Here there is no need for a formal statistical test: the data easily pass the 'Inter Ocular Traumatic Test': Plot the data. if the result hits you between the eyes, then it's (statistically – but of course, not necessarily clinically) significant.

Remarks:

Whereas the mean, 9.8, of the 15 within-person between-condition differences is *arithmetically* equal to the difference of the 2 means of 15, the SE of the mean of these 15 differences is not the same as the SE of the difference of two independent means. In general...

$$Var(\bar{y}_1 - \bar{y}_2) = Var(\bar{y}_1) + Var(\bar{y}_2) - 2 \times Covariance(\bar{y}_1, \bar{y}_2)$$

Double-check that one can arrive at the SE of 1.1 for the mean delay by subtracting twice the covariance from the sum of the two variances, and then taking the square root of this.

Indeed, the effect of pairing is to *remove* the intrinsic between-person variance, and focus the within-person differences. **Applying an inefficient statistical analysis to data collected by an efficient statistical design is a common ‘Type III’ error!**

Authors continue to report the SE of each of the 2 means, but the 2 separate SEs are of little use here, since we are not interested in the difference of means, but in the mean difference.

Calculating

$$Var(\bar{y}_1 - \bar{y}_2) = Var(\bar{y}_1) + Var(\bar{y}_2) = 2.2^2 + 2.2^2 = 9.7$$

so that the $SE_{diff. in means}$ is $\sqrt{9.7} = \sqrt{2} \times 2.2 = 3.1$ assumes that we used *one set of 15 subjects* for the No FootRest condition, and *a different set of 15* for the FootRest condition, a much noisier contrast.

Fortunately, it turned out that in this study the signal is much greater than the ‘noise’. Thus, even the inefficient (2-independent samples) analysis, based on a $SE_{\bar{y}_1 - \bar{y}_0} = 3.1$, would have produced a statistically significant 2-sample *t*-ratio of $9.8/3.1 = 3.2$.

See article (in *jh*’s catalogued collection) on display of data from pairs.

4.3 Sample Size for CI’s and test involving μ

4.3.1 n required for a (2 sided) CI with margin of error (ME) at confidence level $1 - \alpha$

<-- Margin of Error(ME) -- • -- Margin of Error(ME) -->
 <-----> • <----->

- large-sample CI: • $\pm ME = \bar{y} \pm Z_{\alpha/2} SE(\bar{y})$
- $SE(\bar{y}) = \sigma/\sqrt{n}$, so solving for $n...$
- $n = (\sigma^2 \times Z_{\alpha/2}^2) / ME^2$.
- If n small, replace $Z_{\alpha/2}$ by $t_{\alpha/2}$

Typically we do not know σ , so we use use a pre-study estimate of it.

In planning n for example just discussed, authors might have had pilot data on inter leg differences in transit time – with both legs in the No FootRest position. Sometimes, one has to ‘ask around’ as to what the SD of the d ’s will be. Always safer to assume a higher SD than might turn out to be the case.

4.3.2 n required to have power $1 - \beta$ when testing $H_0 : \mu = \mu_0$, if unknown mean, μ , is Δ units from μ_0 , i.e., if $\mu_{alt} - \mu_0 = \Delta$, and if test is carried out with Probability[type I error] = α .

[cf. Fig 4, as well as Colton p. 142, and CRC table on next page. See also ‘Sample Size, Precision and Power Calculations: A Unified Approach’ by Hanley and Moodie]

- Assume that the ‘unit variability’, σ_Y , is the same under H_0 and H_{alt} , so that

$$SE_0[\bar{y}] = SE_{alt}[\bar{y}] = \sigma_Y/\sqrt{n}.$$

- Need

$$Z_{\alpha/2} \times SE_0[\bar{y}] + Z_{\beta} \times SE_{a;lt}[\bar{y}] \geq \Delta.$$

- Substitute $SE[\bar{y}] = \sigma_Y/\sqrt{n}$.
- Solve for n :

$$n \geq [Z_{\alpha/2} + Z_{\beta}]^2 \times [\sigma_Y/\Delta]^2 \quad \sigma_Y/\Delta \text{ is the “noise-to-signal” ratio.}$$




Figure 4: Link between test size (α), sample size, n , power ($1 - \beta$) and Δ .

Notes:

- To make life simpler, JH has made the diagram and formula in terms of the absolute values of $Z_{\alpha/2}$ and Z_β . Thus, be careful with the sign of Z_β : If $\mu_{alt} > \mu$, then the alternative distribution of \bar{y} is to the right of μ_0 (as in diagram), so that a power of more than 50% means that technically, Z_β is negative. e.g. :

$$\alpha = 0.05 \ \& \ \beta = 0.2 \ \Rightarrow \ Z_{\alpha/2} = 1.96 \ \& \ Z_\beta = -0.84.$$

If back-solving for Z_β (and thus β) in terms of n , Δ and σ_Y , be especially careful as to the sign of Z_β : always draw a diagram.

- While it can be α or $\alpha/2$, its always $1 - \beta$, never $1 - \beta/2$!
- Technically, if n is small, should use the more conservative $t_{\alpha/2}$ and t_β : see table on the following page. Since the required n is a function of $t_{\alpha/2}$ and t_β and vice versa, arriving at this table takes some iteration.
- The question of what Δ to use is not a matter of statistics or samples, or what the last researcher found in a study, but rather the “difference that would make a difference” i.e., it is a clinical judgement, and includes the impact, cost, alternatives, etc... JH thinks of it as the Δ that IF IT WERE TRUE would lead to a difference in management or a substantial risk, or ...

4.4 Sign Test for median

Test:

$$\begin{aligned} \text{Median}_D &= 0 & H_0 \\ &\neq 0 & H_{alt} \end{aligned}$$

$$\alpha = 0.05 \quad \text{2 sided}$$

Reference: Binomial [$n = 10$; $\pi(+) = 0.5$]. See also Sign Test Table which I have provided in Chapter on Distribution-free Methods.

Data:

DIFFERENCE	SIGN
Drug-Placebo	
0.9	+
-0.9	-
4.3	+
2.9	+
1.2	+
3.0	+
2.7	+
0.6	+
3.6	+
-0.5	-
Σ 8+, 2-	

Upper-tail: $\text{Prob}[\geq 8 + | \pi = 0.5] = 0.0439 + 0.0098 + 0.0010 = 0.0547$.

2-tails: $P = 0.0547 + 0.0547 = 0.1094$. $P > 0.05$ (2-sided) ...less powerful than t -test.

In above example on Blood Flow, fact that all 15/15 had delays makes any formal test unnecessary... the “**Inter Ocular Traumatic Test**” says it all.

[Q: could it be that the investigators always raised the left leg, and blood flow is less in the left leg? JH doubts it, but asks the question just to point out that just because we find a numerical difference doesn’t necessarily mean that we know what **caused** the difference! *]

* A famous scientist, begins by removing one leg from an insect and, in an accent I cannot reproduce on paper, says “quick march”. The insect walks briskly. The scientist removes another leg, and again on being told “quick march” the insect walks along... This continues until the last leg has been removed, and the insect no longer walks. Whereupon the scientist, again in an accent I cannot convey here, pronounces “There! it goes to prove my theory: when you remove the legs from an insect, it cannot hear you anymore!”.

4.4.1 Number of Observations to ensure specified power β if use 1-sample or paired t -test concerning μ_Y or μ_d

Required n for test where $\alpha = 0.005$ (1-sided) or $\alpha = 0.01$ (2-sided)

Δ/σ	β	0.01	0.05	0.10	0.20	0.50
	Power	99%	95%	90%	80%	50%
0.2						
0.3					134	78
0.4			115	97	77	45
0.5		100	75	63	51	30
0.6		71	53	45	36	22
0.7		53	40	34	28	17
0.8		41	32	27	22	14
0.9		34	26	22	18	12
1.0		28	22	19	16	10
1.2		21	16	14	12	8
1.4		16	13	12	10	7
1.6		13	11	10	8	6
1.8		12	10	9	8	6
2.0		10	8	8	7	5
2.5		8	7	6	6	
3.0		7	6	6	5	

Notes:

- $\Delta/\sigma = (\mu - \mu_0)/\sigma =$ “Signal” / “Noise”
- Table entries transcribed from Table IV.3 of CRC Tables of Probability and Statistics. Table IV.3 tabulates the n 's for the Signal/Noise ratio increments of 0.1, and also includes entries for $\alpha = 0.01$ (1-sided) / 0.02 (2-sided). See also Colton, page 142, or the ‘Sample Size, Precision and Power Calculations: A Unified Approach’ Hanley-Moodie article.
- Sample sizes based on t -distribution, and so slightly larger (and more realistic, when n small) than those given by Z -based formula:
 $n = (Z_\alpha + Z_\beta)^2 \times (\sigma/\Delta)^2$.

Required n for test where $\alpha = 0.025$ (1-sided) or $\alpha = 0.05$ (2-sided)

Δ/σ	β	0.01	0.05	0.10	0.20	0.50
	Power	99%	95%	90%	80%	50%
0.2						99
0.3				119	90	45
0.4		117	84	68	51	26
0.5		76	54	44	34	18
0.6		53	38	32	24	13
0.7		40	29	24	19	10
0.8		31	22	19	15	9
0.9		25	19	16	12	7
1.0		21	16	13	10	6
1.2		15	12	10	8	5
1.4		12	9	8	7	
1.6		10	8	7	6	
1.8		8	7	6		
2.0		7	6	5		
2.5		6				
3.0		5				

Required n for test where $\alpha = 0.05$ (1-sided) or $\alpha = 0.1$ (2-sided)

Δ/σ	β	0.01	0.05	0.10	0.20	0.50
	Power	99%	95%	90%	80%	50%
0.2						70
0.3			122	97	71	32
0.4		101	70	55	40	19
0.5		65	45	36	27	13
0.6		46	32	26	19	9
0.7		34	24	19	15	8
0.8		27	19	15	12	6
0.9		21	15	13	10	5
1.0		18	13	11	8	5
1.2		13	10	8	6	
1.4		10	8	7	5	
1.6		8	6	6		
1.8		7	6			
2.0		6				
2.5						
3.0						

4.5 “Definitive Negative” Studies

Starch Blockers – their effect on calorie absorption from a high-starch meal

Abstract: It has been known for more than 25 years that certain plant foods, such as kidney beans and wheat, contain a substance that inhibits the activity of salivary and pancreatic amylase. More recently, this anti-amylase has been purified and marketed for use in weight control under the generic name “starch blockers.” Although this approach to weight control is highly popular, it has never been shown whether starch-blocker tablets actually reduce the absorption of calories from starch. Using a one-day calorie-balance technique and a high-starch (100 g) meal (spaghetti, tomato sauce, and bread), we measured the excretion of fecal calories after normal subjects had taken either placebo or starch-blocker tablets. **If the starch-blocker tablets had prevented the digestion of starch, fecal calorie excretion should have increased by 400 kcal. However, fecal calorie excretion was the same on the two test days (mean ± S.E.M., 80 ± 4 as compared with 78 ± 2). We conclude that starch-blocker tablets do not inhibit the digestion and absorption of starch calories in human beings.**

Bo-Linn GW. et al New England J of Medicine. 307(23):1413-6, 1982 Dec 2.

Overview of Methods: The one-day calorie-balance technique begins with a preparatory washout in which the entire gastrointestinal tract is cleansed of all food and fecal material by lavage with a special calorie-free, electrolyte-containing solution. The subject then eats the test meal, which includes ⁵¹CrCl₃ as a non absorbable marker. After 14 hours, the intestine is cleansed again by a final washout. The rectal effluent is combined with any stool (usually none) that has been excreted since the meal was eaten. The energy content of the ingested meal and of the rectal effluent is determined by bomb calorimetry. The completeness of stool collection is evaluated by recovery of the non absorbable marker.]

See Powell-Tuck J. A defence of the small clinical trial: evaluation of three gastroenterological studies. Br Med J Clinical Research Ed..292(6520): 599-602, 1986 Mar 1. for a good paper on ‘negative’ studies,

Table 1: Standard Test Meal

Ingredients		Dietary constituents*	
Spaghetti (dry weight)**	100 g	Protein	19 g
Tomato sauce	112 g	Fat	14 g
White bread	50 g	Carbohydrate (starch)	108 g (97 g)
Margarine	10 g		
Water	250 g		
⁵¹ CrCl ₃	4μCi		

*Determined by adding food-table contents of each item.

**Boiled for seven minutes in 1 liter of water.

Table 2. Results in Five Normal Subjects on Days of Placebo and Starch-Blocker Tests.

	Placebo Test Day			Starch-Blocker Test Day		
	Duplicate Test Meal*	Rectal Effluent	Marker Recovery	Duplicate Test Meal*	Rectal Effluent	Marker Recovery
subject	kcal	kcal	%	kcal	kcal	%
1	664	81	97.8	665	76	96.6
2	675	84	95.2	672	84	98.3
3	682	80	97.4	681	73	94.4
4	686	67	95.5	675	75	103.6
5	676	89	96.3	687	83	106.9
Means ± S.E.M.	677 ±4	80 ±4	96.4 ±0.5	676 ±4	78 ±2	100 ±2

Does not include calories contained in three placebo tablets (each tablet, 1.2±0.1 kcal) or in three Carbo-Lite tablets (each tablet, 2.8±0.1 kcal) that were ingested with each test meal.

Is this a Definitive Negative Study?


---0-----100-----200-----300----- | <- Company’s Claim: **400 kcal**
 ---***-----100-----200-----300----- |

---0-----100-----200-----300-----400-- kcal blocked

*** 95% CI estimate from study

0 Exercises

0.1 Age-structures of Populations, then and now



- Comment on the shape, and on any other remarkable features. Give possible reasons for them.
- Visually estimate the mean age, the SD, and the quartiles: Q_{25} , Q_{50} (median) and Q_{75}
- How CLT-friendly is this distribution? i.e., at what n might the sampling distribution of the mean age of n randomly sampled persons be close enough ('for government work'¹³) to Gaussian? Hint: see [code](#) & Q13.

0.2 Clinical Research in General Medical Journals: A 30-Year Perspective

Table from Fletcher et al., NEJM 301: 180-183, 1979

Fletcher et al. studied the features of 612 randomly selected articles published in the NEJM, JAMA and Lancet since 1946. Two features were the number of authors per article and the number of subjects studied in each article.

Year	No. articles examined	No. authors Mean (SD)	No. Subjects Median
1946	151	2.0 (1.4)	25
1956	149	2.3 (1.6)	36
1966	157	2.8 (1.2)	16
1976	155	4.9 (7.3)	30

- From the mean and SD, *roughly* reconstruct the actual frequency distribution of the no. of authors per article for 1946 [Excel or R can help]
- Can the 1976 SD of 7.3 really be larger than the mean of 4.9? Explain.

0.3 Statistics in the U.K. Parliament

Hansard 29 November 1991

Mr. Arbuthnot: the Labour party's suggestion of a minimum wage is in itself rather obscure and bizarre. As I understand it, it is tied to the average and would therefore not only be relatively high at £3.40 but would increase as the average wage itself increased. With each increase in the average rate of pay, the minimum wage itself would have to go up and it would be forever chasing its own tail.

Mr. Tony Lloyd: Perhaps I can help the hon. Gentleman. It will be tied to the median, which is not the same as the average. It is simply a mid-point on the range and would not be affected by changes in the minimum wage.

Mr. Arbuthnot: From what I understand, even an amount tied to the median would be affected because if the lowest wage were increased to £3.40 per hour, the median would have to rise.

Mr. Tony Lloyd: I shall put the matter in simple terms. The median, the mid-point in a series of numbers such as 2.2, 5.6 and 7, is defined as being the difference between 2 and 7, which is 3.5. If we alter the figures 2 and 2 to

¹³<http://www.yourdictionary.com/close-enough-for-government-work>

3.5, the middle figure of 5 would remain unaltered because it is independent of the bottom figures.

Mr. Arbuthnot: I do not understand the hon. Gentleman's mathematics and I slightly doubt whether he does.

Mr. Matthew Carrington (Fulham): I am extremely confused. I studied mathematics for some years at school and I have not totally forgotten all of them. The median is not the mid-point between the first number and the last. It is where the largest number of items in a sample comes to, whereas the average is obviously the sample multiplied by the number of items. The hon. Member for Stretford (Mr Lloyd) is obviously extremely confused. The median has a precise mathematical definition which is absolutely right, and my hon. Friend is correct in saying that the median is bound to alter if the number at the bottom on the scale is changed. That will alter the average as well in a different way, but it is bound to alter the median. Perhaps the hon. Member for Stretford wishes to define median in a non mathematical sense.

Mr. Arbuthnot: I am grateful to my hon. Friend for sorting out at least the hon. Gentleman's mathematics with obvious skill and knowledge.

Exercise: Correct the honourable Gentlemen

0.4 Can it be that 60% of people live longer than 'average' ?

The following file https://www.prhdh.umontreal.ca/BDLC/data/pdfs/TablesMortalite_SC_2009-2011_EN.pdf contains in Table 1b (p6-7) the Complete life table for females in Canada from 2009 to 2011. For a given age x , column l_x gives the number alive (the - fictitious - cohort starts with 100,000 people) and column d_x gives the number of deaths (this column adds up to 100,000 people). The mean age of death is 83.6 years.

- i. In this lifetable, what percentage of people live longer than the mean?
- ii. Make up a news story headline that makes a big deal of this 'new' finding, and a tweet that calls it 'fake news' and 'Canadian propaganda and BAD statistics.'

0.5 Maternal Lead Levels after Alterations to Water Supply. The Lancet 1981

<https://www-sciencedirect-com.proxy3.library.mcgill.ca/science/article/pii/S0140673681903846>

[Still an issue! Cf. Blood Lead Levels of Children in **Flint, Michigan: 2006-2016** [https://www.jpeds.com/article/S0022-3476\(17\)31758-4/fulltext](https://www.jpeds.com/article/S0022-3476(17)31758-4/fulltext) and **EpiMonitor** piece: <https://www.epimonitor.net/Flint-Water.htm>]

- i. What blood lead concentrations fall into the different tabulated intervals used to construct the two frequency distributions? What are the midpoints of these true intervals?¹⁴
- ii. Make histograms for the 1977 and 1980 distributions, putting them both on the same graph to allow them to be easily compared. Pay attention to these grouped data with unequal intervals, and make sure the areas add to 1. Assume, arbitrarily, that the last category ends at 2.00.

If you prefer to show side by side boxplots or dot diagrams (rather than histograms) to display the distributions, you can first use R to simulate what the raw (unbinned) data might have looked like, i.e., randomly assign values within each bin.¹⁵

- iii. Paraphrase the statement "the median value was $0.8\mu\text{mol/l}$ " (para. 2).
- iv. Calculate the (arithmetic) mean and standard deviation of the grouped 1977 data (use a semi-arbitrary "midpoint" for the 14 values in the last interval). Is the SD very useful? What might be a better summary?

¹⁴This issue of where the real boundaries were was a much bigger issue back in the years 'BC' (before computers!) when, to save on manual calculations, researchers grouped the data, and put all the values in a bin at the mid-point of the bin. They also had to check that the rules for allocating values to bins were unambiguous, and that no data points were missed, or duplicated. They also tended to make histograms by hand, and so had to worry, if they used unequal bin widths, about giving a distorted impression - it was important to have each area represent the frequency. These boundary issues and unequal bin widths are much less of an issue nowadays, when 'computers' are no longer persons. See <https://interestingliterature.com/2020/02/origin-word-computer-etymology/>,

¹⁵We will fit parametric (e.g. log-normal and gamma) distributions to these 'interval-censored' data when we come to Chapter 3 (Likelihood) in Clayton and Hills.

0.6 Are all head sizes alike?

Stephen Jay Gould’s book [“The Mismeasure of Man”](#) discusses a table from a 1978 article by Epstein. Gould read the original article and found that “a glance at E A. Hooton’s original table, reproduced below,¹⁶ reveals that the *SE* column had been copied and re-labelled *SD*” Then, using this *SD*, and the *n*, to compute a much smaller-than-it-should-be *SE*, Epstein was able to “show” that the *CI*’s for mean head circumference for people of varied vocational statuses did not overlap, and thus that there were “statistically significant” inter-group differences.

Vocational Status	N	Mean (in mm)	“S.D.”
Professional	25	569.9	1.9
Semiprofessional	61	566.5	1.5
Clerical	107	566.2	1.1
Trades	194	565.7	0.8
Public service	25	564.1	2.5
Skilled trades	351	562.9	0.6
Personal services	262	562.7	0.7
Laborers	647	560.7	0.3

- Explain why - if they really were *SD*s –the “*SD*s” in the table should not decrease with increasing *n*, i.e., why the *SD* from a smaller *n* is as likely to be greater than the *SD* from a bigger *n* as it is to be smaller. If *SD*’s were smaller (some argue larger) in larger samples, then the *SD* of the diameters of red blood cells should be different for a larger adult than a smaller adult!
- Also, from what you have seen of hat-sizes, what would be a reasonable *SD*, and thus a reasonable *CV*, for inter-individual headsizes? Does this fit with the various *SD*s you can back-calculate from the *n* and ‘fake-*SD*’ (in reality, standard error of the mean, $SEM = SD/\sqrt{n}$) columns?
- (A more focused version of part i.) You plan to use **R** to compute the *sd* of 10 random numbers from a $N(0,1)$ distribution: `sd(rnorm(10))`. Your classmate plans to obtain the *sd* of 20 (other) random numbers from that same distribution, `sd(rnorm(20))`. She is convinced that (a) the *sd* of the values in her sample will (‘almost surely’) be smaller than your *sd* based on 10. Another classmate is convinced that (b) the *sd* based on 20 will (‘almost

¹⁶Table VIII-17 “Mean and standard deviation of head circumference for people of varied vocational statuses” , The American Criminal, v. 1, Harvard U. Press, 1939,

surely’) be bigger than the one based on 10.

You, on the other hand, think (c) there is a 50% chance that `sd(rnorm(20)) > sd(rnorm(10))` and a 50% chance that `sd(rnorm(20)) < sd(rnorm(10))`.

A mathematical statistician claims that (d) (to 2 decimal places) there is a 53% chance that `sd(rnorm(20)) > sd(rnorm(10))`¹⁷

Who is correct, and why?

0.7 Births after The Great Blackout of 1966

On November 9, 1965, the electric power went out in New York City, and it stayed out for a day – The Great Blackout. Nine months later, newspapers suggested that New York was experiencing a baby boom. The table shows the number of babies born every day during a twenty-five day period, centered nine months and ten days after The Great Blackout.

Number of births in New York, Monday August 1-Thursday August 25, 1966.

Mon	Tue	Wed	Thu	Fri	Sat	Sun
451	468	429	448	466	377	344
448	438	455	468	462	405	377
451	497	458	429	434	410	351
467	508	432	426			

These numbers average 436. This turns out to be not unusually high for New York. But there is an interesting twist: the 3 Sundays only average 357.

- How likely is it that the average of three days chosen at random from the table will be 357 or less? What do you infer? Hint: The *SD* of the 25 numbers in the table is about 40. Formulate the null hypothesis; the normal approximation can be used.
- The above question and the following footnote come from the (very readable) *Statistics* text by Freedman et al. [link](#)
 ”Apparently, the New York Times sent a reporter around to a few hospitals on Monday August 8, and Tuesday August 9, nine months after the blackout. The hospitals reported that their obstetric wards were busier

¹⁷and that it would be 61% if the sample sizes were 4 and 2, and 51% of they were 100 and 50. As it happens, the form (Normal, Uniform, Gamma, ...) of the distribution does not *greatly* affect the answer, but you happened to select one (*N*) where one can work out the probabilities from math-stat theory, rather than by simulation.

than usual – apparently because of the general pattern that weekends are slow, Mondays and Tuesdays are busy. These “findings” were published in a front-page article on Wednesday, August 10, 1966, under the headline “Births Up 9 Months After the Blackout.” This seems to be the origin of the baby-boom myth.”

Exercise: Suggest a better plan for estimating the impact, if any, of the Blackout on the number of births.

Credits for story & questions: Freeman, Pisani and Purves and their book *Statistics*

- iii. (Still on the subject of births, but now in Québec). In an effort to bolster the sagging birth rate, the Québec government in its budget of March 1988 implemented a cash bonus of \$4,500 to parents who had a third child. Suggest a method of measuring the impact of this incentive scheme – be both precise and concise.

0.8 Planning ahead: 1

One has to travel a distance of 7,500 Km by 4-wheel jeep, over very rough terrain, with no possibility of repairing a tire that becomes ruptured. Suppose one starts with 14 intact tires (the 4, plus 10 spares). It is known that on average, tires rupture at the rate of 1 per 5,000 tire-Kms (the mean interval between ruptures is 5,000 tire-Kms). Assume ruptures occur independently of the tire position or the distance already driven with the tire (i.e., the sources of failure are purely *external*). Also, ignore the possibility of multiple failures from a single source, e.g. a short bad section of the trail. [This R [code](#) is a crude replacement for the the java animation that no longer runs]¹⁸

Calculate the probability of completing the trip, using the . . .

- i. Poisson Distribution for the *number* of ruptures.
- ii. Exact distribution of a sum of distances i.e. of a (fixed) number of ‘*distance*’ random variables.
- iii. Central Limit Theorem to approximate the distribution in ii.
- iv. Central Limit Theorem to approximate¹⁹ the distribution in i.
- v. Random number fns. in R/SAS to simulate intervals between ruptures.

¹⁸Since the tires are exchangeable, you could think ‘serial’ rather than ‘parallel’, i.e., you could represent it as 1 line of length $4 \times 7,500 = 30,000$ tire-Kms.

¹⁹Also comment on why the approximations in iii and iv. work as well as they do.

0.9 A random selection?

A colony of laboratory mice consisted of several hundred animals. Their average weight was about 40 grams, with an SD of about 5 grams. As part of an experiment, graduate students were instructed to choose 25 animals haphazardly, without any definite method. The average weight of these 25 sampled animals was 43 grams. Is choosing animals haphazardly the same as drawing them at random? Assess this by calculating the probability, under strict random selection, of obtaining an average of 43 grams or greater.

0.10 Planning ahead: 2

On the average, conventioners weigh about 150 pounds; the SD is 25 pounds.

- i. If a large elevator for a convention centre is designed to lift a maximum of 15,500 pounds, the chance it will be overloaded by a random group of 100 conventioners is closest to which of the following: 0.1 of 1%, 2%, 5%, 50%, 95%, 98%, 99.9% ? Explain your reasoning.
- ii. The weights of conventioners are unlikely to have a Gaussian (“Normal”) distribution. Are you still comfortable using the Normal distribution for your calculations in part i? Explain carefully. Explain why the ‘random’ is key to being able to answer part i. (see ‘16.3.3 Example 2’ [here](#)) and what impact it would have if it is not the case (see 16.4).

0.11 An unexpected pattern: or is it?

Suppose the pathology reports from biopsies taken during a given week are equally likely to be received on any one of the weekdays from Monday August 5 to Friday August 9.²⁰ Women with positive biopsies then have their surgery scheduled for one of the weekdays of the last full week of August, i.e., Monday August 26 to Friday August 30, equally likely to be any one of these 5 weekdays, and unrelated to which day of the week the biopsy result was received.

- i. Derive and plot the probability distribution of the wait (i.e., no. of days) from when the *biopsy was received* until the woman had the *surgery*.
- ii. Calculate the mean and standard deviation of the distribution.

²⁰Exercise freely adapted from patterns seen in study by Mayo et al, [Waiting time for breast cancer surgery in Quebec](#). CMAJ. 2001 Apr 17:1133-8.

- iii. Comment on its shape, and why it is this shape, and what would happen if there were several stages, not just 2.
- iv. What if we had a 7-day work-week rather than 5?

0.12 A snail's pace

A snail (escargot) starts out to climb a very high wall. During the day it moves upwards an average of 22 cm (SD 4 cm); during the night, independently of how well it does during the day, it slips back down an average of 12 cm (SD 3 cm). The forward and backward movements on one day/night are also independent of those on another day/night.

- i. After 16 days and 16 nights, how much vertical progress will it have made? Answer in terms of a mean and SD. Note that – contrary to what many students in previous years calculated – the SD of the total progress made is not 80 cm; show that it is in fact 20 cm.
- ii. What is the chance that, after 16 days and 16 nights, it will have progressed more than 150 cm?
- iii. “Independence” was ‘given’. Did you have to make strong [and possibly unjustified] distributional assumptions in order to answer part b? Explain carefully.

0.13 How CLT-friendly are Triangular-Shaped Distributions?

The **uniform distributions** are regarded as quite ‘CLT-friendly.’ Indeed, in the early days of computing, when it took a lot of effort to compute $\Phi^{-1}[U(0,1)]$ as a way to generate a random draw from a $N(0,1)$ distribution, it was common practice, and faster, to generate it by summing 12 independent $U(0,1)$ random numbers and subtracting 6. If $Y = (\sum_1^{12} U_i) - 6$, then $E[Y] = 0$, and $\text{Var}[Y] = 12 \times (1/12) = 1$. Many textbooks show the near-Gaussian shape of this distribution, and it is easy to simulate it today.

Incidentally, Laplace has worked out a formula for the exact distribution of a sum of i.i. U 's a long time ago, but lacked the computing power to evaluate the very long sums of powers as n got large. (see question 21 below)

Clearly, the Law of Cancellation of Extremes works best if there are no extremes to begin with, and a lot of the mass is centralized, with a mode that is well in from both extremes, and (ideally) symmetry.

Many textbooks also show the much less CLT friendly case of the **(negative) exponential distribution**, with a lot of mass stacked up against zero, and a long right tail, and how big an n it takes to wipe out the extreme shape. They usually do it by simulation, but one could also appreciate it by taking advantage of the fact that a sum of exponential random variables has a gamma distribution whose shape parameter is the number of variables being summed. So, for example, you can examine how close the shape of a $\text{gamma}(10, \cdot)$ distribution comes to a Gaussian distribution. This gamma distribution is the sampling distribution of the sum or mean of a sample of 10 from the negative exponential distribution.

The **pyramid- or triangular-shaped (age) distributions** seen in ‘the old days’ or in less developed countries today can be used to appreciate how CLT-friendly is a distribution that is ‘in between’ the uniform and the negative exponential distributions. It has finite support, say $(0, \text{Max})$, and a less steep descent from zero.

Can the sampling distribution of the mean or sum of n triangular r.v.'s be described by a nice closed-form expression? Does the fact that a $\text{uniform}(0, \text{Max})$ random variable has a (scaled) beta distribution with parameters $\alpha = 1$ and $\beta = 2$ help any? Could it be that the sum of two such (i.i.d.) r.v.'s again has a beta distribution with parameters $\alpha = 2$ and $\beta = 4$?

The answer is left to you to figure out. JH suspects that you will use the moment generating or characteristic function to do so.

JH expects that if we had the chance to ask Laplace, he could have come up with a (maybe even nastier) formula than he did for the sum of Uniforms, which are beta distributions with $\alpha = 1$ and $\beta = 1$?

Fortunately, we can get somewhere if we use the **discrete but not- quite-triangular-shaped (age) distribution** shown in Question 0.01.

To follow along in **R**, the frequencies of the 110 reported ‘integer’ ages, from 0 to 109 (males and females combined) can be downloaded from here (note that there are no persons in some age bins beyond 100).

```
path = "https://jhanley.biostat.mcgill.ca/bios601/AgeFrequencies.txt"
age.distrn = read.table(path,header=TRUE)
```

First, convert the frequencies to proportions/probability masses: $\sum_0^{109} p_a = 1$.

The trick is to (recursively) use convolutions. To obtain the sampling distribution of the mean age of all possible samples of size $n = 2$, sum the $p_{a'} \times p_{a''}$ products in each of the 219 diagonals and off-diagonals of the $110 \times$

110 grid.²¹ You can make this grid with the `outer` function in R, and ignore the fact that this implies sampling with rather than without replacement – a negligible issue when sampling so few people from so many.

The sums from the 219 diagonals (where the 2 ages sum to 0, 1, . . . , 218 – or the mean ages are 0, 1/2, 1, 1 and 1/2, 2, . . . , 109) provide the sampling distribution (of length 219) for the mean age of $n = 2$ randomly selected persons.

This R code: `sapply(split(M, col(M) + row(M)), sum)`, courtesy of a Google search, sums the diagonal entries in a matrix M.

Now, to get to the case of $n = 3$, compute the outer product of the original p vector of size 110 and the new one of size 219, to make a 110×219 grid of products, and again obtain the sums from the (now) 328 diagonals. These (where the 3 ages sum to 0, 1, 2, . . . , 317 – or the mean ages are 0, 1/3, 2/3, 1, 1 and 1/3, 1 and 2/3, . . . , 109) provide the sampling distribution (of length 328) for the total or mean age of $n = 3$ randomly selected persons.

Continue on like this, each time using the outer product of the original p ('1') distribution and the 'sum/mean of n ' distribution to get the sampling distribution of the sum or mean for a sample of size $n + 1$. Each time, the spacings on the 'ruler' (think of a sample as yielding a measurement on a ruler) become finer, so that with a samples of 10 and 20, the mean age in the sample has increments (granularity/resolution) of 0.1 and 0.05 years, respectively.

- i. Carry out 9 successive convolutions, so as to get to the sampling distribution of the mean age of 10 randomly sampled persons
- ii. Plot the 10 distributions using discrete probability *masses* or spikes (`type="h"` in the `plot` or `lines`). This way the masses add to 1.
- iii. Plot the 10 distributions using probability *density* functions so that the areas under them add to 1. To plot the 10 densities (probabilities per year of age) on the same graph, you will need to rescale the probability masses.
- iv. Comment on the shapes, and on whether they are closer to or further from Gaussian than the corresponding gamma ones obtained by summing/averaging exponential random variables. Why do you think this is?

²¹Cf. 6×6 grid of (all-equal) probabilities $\frac{1}{6} \times \frac{1}{6} = \frac{1}{36}$ for 2 randomly thrown dice. Within each of its 11 diagonals, the sum of the numbers of dots ('pips', 1:6) is a different constant, ranging from 2 to 12. The distribution of the sum/mean is a symmetric triangle.

- v. How quickly do the 'extra' masses at multiples of 10 years (and presumably the deficits in mass elsewhere) get 'polished/smoothed away' by the aggregation? For example, are they still visible at $n = 2$? 3? 4?

0.14 Student's t -distribution: beyond $n = 10$

"Student"'s table was for $z = (\bar{y} - \mu_0)/s$, not the $t = (\bar{y} - \mu_0)/(s/\sqrt{n})$ tabulated and used today [Also, the s in Student's z was obtained by $\div n$, not $\div(n-1)$].

Moreover, the 1908 table only went up to $n = 10$. For $n > 10$ he suggested using $z = (\bar{y} - \mu_0)/(s/\sqrt{n-3})$ and obtaining the (approximate) p-value by using the Normal table to finding the tail area corresponding to this z value.

His first e.g.'s had $n = 10, 6$ and 2 , he "conclude(d) with an example which comes beyond the range of the tables, there being eleven experiments."

For this, he uses the approximation $\Delta \sim N(\bar{d}, s/\sqrt{n-3})$ to arrive at the statement that there is a 0.934 probability "that kiln-dried barley seed gives a higher barley yield than non-kiln-dried seed." [i.e. that $\Delta > 0$ – see below]

- i. Use today's packages/functions (e.g. the `pt` function in R or `tdist` function in Excel, or `probt` in SAS) to check how accurate his approximation was in this case.²² Note that he calculated each SD as $\{(1/11) \times \sum(\text{Diff} - \overline{\text{Diff}})^2\}^{1/2}$.
- ii. Do likewise with his other 3 p values (notice the typo in the mean difference in crop value in the last column).
- iii. Today, how 'politically correct' is it to interpret the p value of $1-0.934 = 0.066$ as 'the odds are about 14 to 1 that kiln-dried corn gives the higher yield.' ? (*It sounds like a Bayesian statement, but starting from no opinion whatever, a bit strange given that he had been aware of past finding and talking to his friends*)
- iv. In the clearest language you can, give your own more politically correct interpretation of this p-value, acceptable to practitioners of all statistical religions.

²²Others had to wait for his extended z table published in 1917, in order to obtain the exact probability.

Excerpts from section IX of Student’s 1908 paper...²³

To test whether it is of advantage to kiln-dry barley seed before sowing, seven varieties of barley were sown (both kiln-dried [KD] and not kiln-dried [NKD]) in 1899 and four in 1900; the results are given in the table. (corn price is in shillings per quarter and the value of the crop is in shillings per acre).

It will be noticed that the kiln-dried seed gave on an average the larger yield of corn and straw, but that the quality was almost always inferior. At first sight this might be supposed to be due to superior germinating power in the kiln-dried seed, but my farming friends tell me that the effect of this would be that the kiln-dried seed would produce the better quality barley. Dr Voelcker draws the conclusion: “In such seasons as 1899 and 1900 there is no particular advantage in kiln-drying before mowing.” Our examination completely justifies this and adds “and the quality of the resulting barley is inferior though the yield may be greater.”

In this case I propose to use the approximation given by the normal curve with standard deviation $s/\sqrt{n-3}$ and therefore use Sheppard’s (Normal) tables, looking up the difference divided by $s/\sqrt{8}$. The probability in the case of yield of corn per acre is given by looking up $33.7/22.3 = 1.51$ in Sheppard’s tables. **This gives $p = 0.934$** , or the odds are about 14 to 1 that kiln-dried corn gives the higher yield.

Similarly $0.91/0.28 = 3.25$, corresponding to $p = 0.9994$ ²⁴ so that the odds are very great that kiln-dried seed gives barley of a worse quality than seed which has not been kiln-dried.

Similarly, it is about 11 to 1 that kiln-dried seed gives more straw and about 2 to 1 that the total value of the crop is less with kiln-dried seed.

Year	lbs. head corn per acre			price head corn			cwts. straw per acre			value of crop per acre		
	NKD	KD	Diff	NKD	KD	Diff	NKD	KD	Diff	NKD	KD	Diff
1899	1903	2009	+106	26.5	26.5	0	19.25	25	+5.75	140.5	152	+11.5
1899	1935	1915	-20	28	26.5	-1.5	22.75	24	+1.25	152.5	145	-7.5
1899	1910	2011	+101	29.5	28.5	-1	23	24	+1	158.5	161	+2.5
1899	2496	2463	-33	30	29	-1	23	28	+5	204.5	199.5	-5
1899	2108	2180	+72	27.5	27	-0.5	22.5	22.5	0	162	164	+2
1899	1961	1925	-36	26	26	0	19.75	19.5	-0.25	142	139.5	-2.5
1899	2060	2122	+62	29	26	-3	24.5	22.25	-2.25	168	155	-13
1900	1444	1482	+38	29.5	28.5	-1	15.5	16	+0.5	118	117.5	-0.5
1900	1612	1542	-70	28.5	28	-0.5	18	17.25	-0.75	128.5	121	-7.5
1900	1316	1443	+127	30	29	-1	14.25	15.75	+1.5	109.5	116.5	+7
1900	1511	1535	+24	28.5	28	-0.5	17	17.25	+0.25	120	120.5	+0.5
Ave.	1841.5	1875.2	+33.7	28.45	27.55	-0.91	19.95	21.05	+1.10	145.82	144.68	+1.14
SD	-	-	63.1	-	-	0.79	-	-	2.25	-	-	6.67
SD/ $\sqrt{8}$	-	-	22.3	-	-	0.28	-	-	0.80	-	-	2.40

²³Full 1908 article: ‘The Probable Error of a Mean’ by ‘Student’ <https://jhanley.biostat.mcgill.ca/bios601/Mean-Quantile/Student-t.pdf>

²⁴As pointed out in §V, the normal curve gives too large a value for p when the probability is large. I find the true value in this case to be $p = 0.9976$. It matters little, however, to a conclusion of this kind whether the odds in its favour are 1660 to 1 or merely 416 to 1.




Figure 5: The cartoon, from the textbook *Statistics* by Freedman, Pisani and Purves, refers to switching from the ratio $(\bar{y} - \mu_Y)/(\sigma/\sqrt{n})$ (where σ is known) to the ratio $(\bar{y} - \mu_Y)/(s/\sqrt{n})$ (where s is an estimate of the unknown σ). Ironically, there is another z as well: in 1908 Student derived and tabulated the distribution of the ratio: $z = (\bar{y} - \mu_Y)/s^*$, with s^* obtained using a divisor of n . Later, in the mid 1920s, Fisher got him to switch to the ratio $(\bar{y} - \mu_Y)/(s/\sqrt{n})$, with s obtained using a divisor of $n - 1$. It appears that Student was the one who made the name change from *Student's z* to *Student's t*, and Fisher who did the heavy math lifting, and who saw the much wider applicability of the t distribution. Fisher saw a t r.v. as (proportional to) the ratio of a Gaussian r.v. to the square root of an independent r.v. with a chi-squared distribution, and the centrality of the concept of 'degrees of freedom'. For more, see 2008 article by JH MJ and EM under Resources.



'Student' in 1908

Figure 6: from <http://www.york.ac.uk/depts/math/histstat/people/>. See also the presentations on JH's webpage devoted to Student (Gosset).

0.15 Experiments to Determine the Density of the Earth. By Henry Cavendish, Esq. F.R.S. and A. S.

The 29 measurements (cf. an earlier assignment sheet) are repeated here:

5.5 5.61 4.88 5.07 5.26 5.55 5.36 5.29 5.58 5.65 5.57 5.53 5.62 5.29 5.44 5.34
5.79 5.1 5.27 5.39 5.42 5.47 5.63 5.34 5.46 5.3 5.75 5.68 5.85

The following is from pp 521-522 of his report.

From this table it appears, that though the experiments agree pretty well together, yet the difference between them, both in the quantity of motion of the arm and in the time of vibration, is greater than can proceed merely from the error of observation. As to the difference in the motion of the arm, it may very well be accounted for, from the current of air produced by the difference of temperature; but, whether this can be accounted for the difference in the time of vibration, is doubtful. If the current of air was regular, and of the same swiftness in all parts of the vibration of the ball, I think it could not; but, as there will most likely be much irregularity in the current, it may very likely be sufficient to account for the difference.

By a mean of the experiments made with the wire first used, the density of the earth comes out **5.48** times greater than that of water; and by a mean of those made with the latter wire, it comes out the same; and the extreme difference of the results of the 23 observations made with this wire, is only **.75**; so that the extreme results do not differ from the mean by more than **.38**, or $\frac{1}{14}$ of the whole, and therefore the density should be determined hereby, to great exactness.

It, indeed, may be objected, that as the result appears to be influenced by the current of air, or some other cause, the laws of which we are not well acquainted with, this cause may perhaps act always, or commonly, in the same direction, and thereby make a considerable error in the result. But yet, as the experiments were tried in various weathers, and with considerable variety in the difference of temperature of the weights and air, and with the arm resting at different distances from the sides of the case, it seems very unlikely that this cause should act so uniformly in the same way, as to make the error of the mean result nearly equal to the difference between this and the extreme; and, therefore, it seems very unlikely that the density of the earth should differ from 5.48 by so much as $\frac{1}{14}$ of the whole.

Another objection, perhaps, may be made to these experiments, namely, that it is uncertain whether, in these small distances, the force of gravity follows exactly the same law as in greater distances. There is no reason, however, to think that any irregularity of this kind takes place, until the bodies come within the action of what is called the attraction of cohesion, and which seems to extend only to very minute distances. With a view to see whether the result could be affected by this attraction, I made the 9th, 10th, 11th and 15th experiments, in which the balls were made to rest as close to the sides of the case as they could; but there is no difference to be depended on, between the results under that circumstance, and when the balls are placed in any other part of the case.

According to the experiments made by Dr. MASKELYNE on the attraction of the hill Schehallien, the density of the earth is $4\frac{1}{2}$ times that of water; which differs rather more from the preceding determination than I should have expected. But I forbear entering into any consideration of which determination is most to be depended on, till I have examined more carefully how much the preceding determination is affected by irregularities whose quantity I cannot measure.

Exercise

- i. Find the mean of the 29 values.
- ii. Calculate a 95% CI to accompany it.
- iii. What ‘fraction of the whole’ (i.e., what fraction of the point estimate) does the margin of error in this CI represent?
- iv. Would a “trimmed mean” be useful here?
(see http://en.wikipedia.org/wiki/Truncated_mean).

0.16 Power and sample size calculations: 1

Refer to the exercises in Section 6.4 of Moore and McCabe’s text

- i. Exercise 6.82.
- ii. Exercise 6.84. In addition, calculate the n that yields a power of 80% against a shortfall of 5, 2, 1 and 0.1 cc, respectively. Make a plot of the power as a function of the shortfall, (Δ)
- iii. Exercise 6.86. In addition, calculate the n that yields a power of 50%, and a power of 80%, against a mean of 130, 135, and 140 respectively.

0.17 Power and sample size calculations: 2

Is this milk watered down?

(Adapted from Q 15.17 from Moore and McCabe, 4th Edition) A cheese maker buys milk from several suppliers. It suspects that some producers are adding water to their milk to increase their profits. Excess water can be detected by measuring the freezing point of the milk. The freezing temperature of natural milk varies according to a Gaussian distribution, with mean $\mu = -0.540^\circ$ Celsius (C) and standard deviation $\sigma = 0.008^\circ$ C. Added water raises the freezing temperature toward 0° C, the freezing point of water. The laboratory manager measures the freezing temperature of five consecutive lots of milk from one producer. The mean of these 5 measurements is -0.533° C. Is this good evidence that the producer is adding water to the milk?

Moore and McCabe asked students to ‘State hypotheses, carry out the test, give the P -value, and state your conclusion.’ *In this course*, we will go further and ask how much water a farmer could add to the milk before (s)he has a 10% , 50%, 80% chance of getting caught (of the buyer ‘detecting’ the cheating). Assume the buyer continues to use an n of 5, and the same $\sigma = 0.008^\circ$ C, and bases the boundary for rejecting/accepting the product on a 5% ‘ α ’, and a 1-sided test, i.e, $z = 1.645$.²⁵, i.e., the buyer sets the cutoff²⁶ at $-0.540 + (\text{qnorm}(0.95) = 1.645) \times 0.008/\sqrt{5} = -0.534^\circ$ C.

Assume that mixtures of M% milk and W% water would freeze at a mean of $\mu = (M/100) \times -0.540^\circ\text{C} + (W/100) \times 0^\circ\text{C}$ and that the σ would remain unchanged. Thus, mixtures of 99% milk and 1% water would freeze at a mean of $\mu = (99/100) \times -0.540^\circ\text{C} + (1/100) \times 0^\circ\text{C} = -0.5346^\circ\text{C}$. Mixtures of 98:2 would freeze at $\mu = -0.5292^\circ\text{C}$. [Hint: *drawing overlapping distributions will help.*]

²⁵Given the extensive experience with the variability of freezing point measurements, and the practice of using a ‘known’ σ in quality control, the use of a z multiple of 1.645 (rather than a t_4 multiple of 2.13) to establish the ‘critical value’ makes sense. It does not make sense to use an s based on a sample of 5 to re-estimate σ .

²⁶ http://ansci.illinois.edu/static/ansc438/Milkcompsynth/milkcomp_freezing.html states that ‘The current official freezing point limit (-0.525 degrees Horvet or -0.505 degrees C; see Sherbon 1988 for discussion of Horvet vs Centigrade) was designed for whole-herd, bulk-tank samples or processed milk samples, and not for samples from individual cows or individual quarters. The value of -0.525 degrees Horvet is considered the upper limit which statistically is suppose to be a cut-off for most, but not absolutely all, samples to be considered “water-free” (that is, no added water). <https://van.physics.illinois.edu/qa/listing.php?id=1606> states that ‘the exact freezing point of milk (also called the melting point) varies slightly according to the individual cow, the breed, the time of day / season that the milk is collected, the type of feed that the cow receives, etc. According to ... , the majority of cows produce milk with a natural freezing point of -0.5250 to -0.5650 ° C, with an average of about -0.5400 C.’

0.18 Power and sample size calculations: 3

Suppose we wished to assess, via a formal statistical test, whether [at an *population*, rather than an individual, level] a step-counting device or app is unbiased (H_0) or under-counts (H_1). Suppose we will do so the way Case et al. did, but measuring n persons just once each. We observe the device count when the true count on the treadmill reaches 500. The statistical test will declare the test ‘positive’ (the departure from 500 is *statistically significant*) if the mean of the n observations is *below* $500 - 1.96 \times s/n^{1/2}$, where s is the SD of the n observations, and the 1.96 ²⁷ is chosen to give the test a type I error of 5% ²⁸ Since each person is only being measured once, we will not be able to distinguish the genuine between-person variance, σ_B^2 from the within-person variance, σ_W^2 .²⁹ Thus the sample variance, s^2 will be an estimate of $\sigma^2 = \sigma_B^2 + \sigma_W^2$.

- i. Using a planned sample size of $n = 25$, and $\sigma = 60$ steps as a pre-study best-guess as to the s that might be observed in them, calculate the critical value $500 - 1.96 \times s/n^{1/2}$.
- ii. Now imagine that in an infinite sample, the mean would not be the null 500, but $\mu = 470$. Calculate the probability that the mean in the sample of 25 will be less than this critical value. [Use the same s for the alternative that you used for the null.]
- iii. By trial and error, or better still by deriving a general formula, adjust the n until this probability (i.e., the power) is 80%. Show the 2 probabilities in a diagram, as was done in Figure 4 in section 4.3.2 of the notes.³⁰
- iv. Show that this n satisfies the equality

$$1.96 \times SE_{null}[\bar{y}] + 0.84 \times SE_{alt}[\bar{y}] = 30.$$

²⁷Assume the ultimate n will be large enough that the t and z distributions are interchangeable.

²⁸The way the task is phrased, it looks like a 1-sided alternative, but – even if interest is only in 1 direction – grant review panels and journal editors often insist on a 2-tailed p-value. This insistence on a 2-tailed test might be appropriate if, even though the interest is in a drug *reducing* the death rate, an *increased* death rate would be of concern and not considered the same as a null ‘no difference’. If you find this made-up example a bit un-natural, imagine instead that you were a regulator who had to ensure that a company’s claim (H_0) that its soft-drink bottles contained (at least) 500 ml was in fact true [versus the alternative (H_1) that it was *under-filling* them].

²⁹With 2 replicates on each subject, Case et al. *could* have obtained separate estimates of the two variance components.

³⁰Note that there, the concern was with a μ that is *greater* than μ_0 .

0.19 Unsure if a z-SE's confidence interval is appropriate: bootstrap!

See [Efron & Diaconis article in Scientific American](#)

- i. Select a (.csv) dataset from the list, at the bottom of this [Oceanography webpage](#). Each one consists of the ocean depths (or land heights, i.e., altitudes) at 200 randomly selected locations on the earth. First, calculate a z-based 95% CI for μ , the mean depth of the ocean. Then, use the R code from the next question (or 'roll your own') to investigate (via bootstrap³¹) whether your n of ocean-depth measurements is large enough, and the 'parent' distribution well-behaved (CLT-friendly) enough to assure that a Gaussian- and SE-based CI for μ (the mean ocean depth) is reasonably accurate.
- ii. Refer to the On Time data for the first 1/2 million U.S. airline flights in 2013, along with the R code supplied. These can be found under **Data / Data Analysis** in this webpage. Use the R code (or 'roll your own') to investigate (via bootstrap) whether $n = 30$ is large enough, and the 'parent' distribution is sufficiently CLT-friendly to assure that a Gaussian- and SE-based CI for μ (the mean delay) is reasonably accurate. Even though you might be tempted in this case to 'peek' at the universe of (1/2 million) observations, in practice you won't have this luxury (you will have spent your entire budget getting the n observations).

What if you had a budget for a sample of $n = 200$?

Yes! we know, it is just the cost of a fraction of a second of your time, and that of R, in this toy example, and a minimal battery or electricity charge for the extra milliseconds of computer time; but imagine that each observation – as it might in a medical trial of how much a certain treatment delays some event – costs several thousand dollars? or this is how many new eligible patients in your city develop the condition of interest in a 2-year window!

The basis for the bootstrap: A few years ago, one student investigated question i by taking many samples of the same size (n was 140) from the oceanography database – and found that indeed the sampling variability of the sample mean was close to Gaussian. BUT, in practice, one would not have this luxury. So, instead of sampling from the actual universe of N , Efron decided to 'photocopy' or 'clone' the n many many many times,

and pretended that this was like having the universe of all N – from which he could then sample as often as he wanted, just like that student did.


Of course, if we have a computer, we don't need to make physical copies. We can sample n **with replacement** from the n , and do so enough times until the histogram of sample means (or sample medians, or whatever statistic) becomes smooth enough to trust. Then we can take the central (say) 95% portion of this as the 95% CI for μ (or whatever parameter – generically θ – we are interested in.

On average, any specific one of the n observed values will appear in the bootstrap sample 1 time... sometimes it won't appear at all, sometimes the sample will have 1 copy of it, sometimes 2 copies, etc.

You might be interested in a trick we came up with for not having to make physical copies.. see [this paper](#) by Hanley JA and Macgibbon B. Creating non-parametric bootstrap samples using Poisson frequencies. Computer Methods and Programs in Biomedicine. 2006 Jul;83(1):57-62. Epub 2006 May 30.

Here is an excellent *expository piece* on the Bootstrap by Diaconis and Efron in 1983 in *Scientific American*. <http://folk.uio.no/deilerts/phd/docs/Efron-scientificamerican0583-116.pdf>.

There is also also-excellent 2016 *book* by Efron and Hastie. It is divided into 3 parts: **Part I Classic Statistical Inference** (Algorithms and Inference; Frequentist Inference; Bayesian Inference; Fisherian Inference and Maximum Likelihood Estimation; Parametric Models and Exponential Families) **Part II Early Computer-Age Methods** (Empirical Bayes; James-Stein Estimation and Ridge Regression; Generalized Linear Models and Regression Trees; Survival Analysis and the EM Algorithm; **The Jackknife and the Bootstrap; Bootstrap Confidence Intervals**; Cross-Validation and Cp Estimates of Prediction Error; Objective Bayes Inference and MCMC; Postwar Statistical Inference and Methodology) and **Part III Twenty-First-Century Topics**(Large-Scale Hypothesis Testing and FDRs; Sparse Modeling and the Lasso; Random Forests and Boosting; Neural Networks and Deep Learning; Support-Vector Machines and Kernel Methods; Inference After Model Selection; Empirical Bayes Estimation Strategies).



<https://www-cambridge-org.proxy3.library.mcgill.ca/core/books/computer-age-statistical-inference/E32C1911ED937D75CE159BBD21684D37>

³¹Thanks to your fellow student, with the same initials B.S. we sometimes use for bootstrap, for suggesting this.

0.20 Planning ahead – the (2015) sequel

A few years ago, when EBOH offices were in Puvis Hall at the top of Peel Street, JH has an appointment for 2pm, one door away from the Vendôme Metro station, and consulted the STM website <http://www.stm.info/en/info/networks/metro> to find out how much time it would take if he took the Metro from the Peel to the Vendôme Metro station. It said 20 minutes.

How far ahead should he be at Peel metro station in order to be on time? (as it turned out, he had to wait over an hour once he got to his destination!)


He doesn't know when the next Green Line train will be, but (for this exercise) assume that they also arrive every 10 minutes at this time of day.

Lionel Groulx station, where he will switch to the Orange Line, is the 3rd stop (count on 1 minute per stop). Again, he does not know, when planning ahead, when the trains depart from there, but does know (lets assume) that they come every 12 minutes at this time of day. Vendôme station is the second stop.

- i. What is the probability that if he arrives at the Peel Metro platform at 1.40 pm, he will be on the Vendome platform by 1.57 at the latest? State any assumptions made.³²

Hint: Last year, students spent many pages deriving the answer by integral calculus, only to arrive at a fraction involving two integers. So, this year, you might want to sketch the diagram and calculate the probability geometrically, using only the formulae for the areas of rectangles and triangles! (If that fails, obtain an estimate by simulation!)

32



This problem is similar to one by Laplace, discussed in Chapter 5 of a book we will meet in the next question.


- ii. What if the trip involved 3 Lines, involving 9 stops?

0.21 BC – Before [electronic] Computers

McGill library has the ebook A History of the Central Limit Theorem: From Classical to Modern Probability by Hans Fischer. Section 2.1.1 (Sums of Independent Random Variables) appears in the chapter entitled The Central Limit Theorem from Laplace to Cauchy, and page 19 is reproduced overleaf.

If you Google ‘Laplace comets central limit theorem’ you will find Chapter 4 in an older book The Life and Times of the Central Limit Theorem, by William J. Adams (McGill Library has a physical copy)

And you can use Google to find Chapter 5 ‘Distribution of Comets’ in Pierre-Simon Laplace, 1749-1827: A Life in Exact Science by Charles Coulston Gillispie and I. Grattan-Guinness. This chapter discusses Laplace’s treatment of the **sum of 2 or 3 Uniform random variables**. As the authors explained in a p41 footnote, while they say they got the algebra correct, they drew parts of Figure 6 (the ‘sum of 3’ case) with ‘curvature the reverse of that shown.’ The fixed-up Figure 6 is shown here.



Laplace then turned to the case of three bodies (or comets) M, N, and P, where $AB = a$ (Fig. 6) is divided into three equal parts, Aa , ab , bB . In order to find the probability that the mean inclination is equal to any abscissa AY , he draws the probability curve $AmMnB$. The reason-

All three books tell of how Laplace tried to submit to statistical “analysis

the probabilities that the mean inclination of the comets (...) will fall within certain limits, under the assumption that they have been orbited by chance.” **Their sum (or mean) involves the angles of each of 63 comets, measured against the ecliptic, and assumed to be distributed randomly (independently) according to a uniform distribution between 0 and 90 degrees, and stochastically independent of each other.**

Earlier in 2022, when I interviewed just-retired McGill Math-Stat professor [David Wolfson](#), and the topic turned to teaching, he mentioned that seeing the diagrams as one went from $n = 1$ to $n = 2$ or $n = 3$ (as in the previous page) was sufficient for him to believe in the Central Limit Theorem. He found these better than simulations!

2.1.1 Sums of Independent Random Variables

Sums of independent random variables had played an important role in Laplace’s probabilistic work from the very beginning.³ In this context, the problem of calculating the probability distribution of the sum of angles of inclination, which were assumed to be determined randomly, as well as the related problem of calculating the probabilities of the deviations between the arithmetic mean of data which were afflicted by observational errors and the underlying “true value,” became especially important. In one of his first published papers, [Laplace \[1776\]](#) had already set out to determine the probability that the sum of the angles of inclination of comet orbits (or the arithmetic mean of these angles respectively) is within given limits. He assumed that all angles, which had to be measured against the ecliptic, were distributed randomly according to a uniform distribution between 0° and 90° (and also tacitly presupposed that all angles were stochastically independent). Laplace succeeded in calculating these probabilities for an arbitrary number of comets via induction (with a minor mistake which was subsequently corrected in [\[Laplace 1781\]](#)). In this 1781 paper, Laplace even introduced a general—however very intricate—method, based on convolutions of density functions, in order to exactly determine the probability that a sum of independent random variables (“quantités variables,” as Laplace put it) was within given limits.⁴ In the most simple case, each of the n variables had the same rectangular distribution between 0 and h . For the probability P that the sum of those variables was between a and b with $0 \leq a < b \leq nh$, Laplace obtained (in modern notation)

$$P = \frac{1}{h^n n!} \left(\sum_{i=0}^N \binom{n}{i} (-1)^i (b - ih)^n - \sum_{i=0}^M \binom{n}{i} (-1)^i (a - ih)^n \right), \quad (2.1)$$

where $N = \min(n, \lfloor \frac{b}{h} \rfloor)$ and $M = \min(n, \lfloor \frac{a}{h} \rfloor)$. Formulae of this kind were too complicated for a direct numerical evaluation if the number of random variables exceeded a relatively small value. The arithmetic mean of the actual observed angles of inclination of the then known 63 comets was $46^\circ 16'$. Through the use of (2.1) alone, Laplace was unable to address the hypothesis that the comets’ planes of motion resulted at “random.” At this stage of his mathematical work, however, Laplace could not develop usable approximations.

³ For a comprehensive biography also dealing with Laplace’s probabilistic work, see [\[Gillispie 1997\]](#). Detailed discussions of Laplace’s contributions to probability and statistics can be found in [\[Sheynin 1976; 1977; 2005b; Stigler 1986; Hald 1998\]](#). The web site already referred to in footnote 1 contains English translations of most works in probability theory by Laplace.

⁴ See [\[Sheynin 1973, 219 f.\]](#) and [\[Hald 1998, 56–60\]](#) for descriptions of this method.

- i. Suggest a replacement (admittedly approximate) for the at-that-time-not-very-practical Laplace formula (2.1 in Fischer's book). Use it to judge whether the arithmetic mean of the observed inclination of the then known 63 planets (46 degrees, 16 minutes), did indeed, as Laplace interpreted it, 'deviate(d) but little' from 45 degrees.
- ii. (Before checking) How accurate do you think your formula is? Then check it against formula 2.1 or against Monte Carlo simulation (whichever methods is faster/easier), in the $n = 63$ example.
- iii. At the time, Laplace left it to others the calculation for $n = 63$ because 'le calcul serait pénible' (see line 3, section IX on page 301 from his *Des Orbites des Comètes*, reproduced on the following page). What about the correctness of the calculation with the 'supposed' $n = 12$ he was content to report there? Again, check it against formula 2.1, or by Monte Carlo simulation. Also check further out in the tail.

The article 'A note on the sum of uniform random variables by' Buonocore, et al. in *Statistics & Probability Letters*, Volume 79, Issue 19, 1 Oct 2009, pp 2092-97 shows equation 2.1 as a special case of a more general formula, which they established by induction.

What is the cause that could have determined the motion of the planets and the satellites? Was it peculiar to these stars, or rather did it exert influence on all those that revolve around the Sun? The first of these questions seems to me very difficult to resolve and I confess that after long reflection and after studying carefully all hypotheses put forward to date to explain this phenomenon, I have found nothing satisfactory. As for the second question, one can answer it easily. To do so it is sufficient: 1. to calculate the mean inclination of the orbits of all observed comets and to observe the deviation from 45° in supposing the comets to have been orbited by chance, then this angle has the same likelihood of being above 45° as it has of being below 45° ; and 2. to know the ratio of the number of direct comets to the number of retrograde ones and to observe its deviation from one, for it is as likely to be greater than one as it is to be less than one. These calculations have been made by Mr. du Séjour in his excellent work on comets. This learned author found that the mean inclination of the 63 comets observed to the present was $46^\circ 16'$, which deviates but little from 45° , and that the ratio of the number of direct comets to the number of retrograde ones was $5/4$, which deviates but little from one. From this observation he concluded, with reason, that for comets there exists no single cause which forces them to move in one direction rather than in another and nearly in the same plane, and thus that the cause which determines the motion of the planets is entirely independent of the general system of the universe.

This interesting observation of Mr. du Séjour gave me the idea of submitting to analysis the probabilities that the mean inclination of the comets and the ratio of the number of direct comets to the number of retrograde ones will fall within given limits, under the assumption that they have been orbited by chance. [2, pp. 280–281]

..... Laplace, quoted in pp 31 & 32 of *The Life and Times*
 of the Central Limit Theorem, by William J. Adams.

From Laplace in the 1770s...

DES ORBITES DES COMÈTES. 301

extrémités A et B; on doit donc avoir

$$(\mu) \begin{cases} r^n - \frac{n+1}{1} (r-1)^n + \frac{(n+1)n}{1.2} (r-2)^n - \dots \\ = (n-r+1)^n - \frac{n+1}{1} (n-r)^n + \dots \end{cases}$$

en continuant l'un et l'autre membre de cette équation, jusqu'à ce qu'on arrive à un terme qui soit nul. On peut s'assurer d'ailleurs de la vérité de cette équation, en observant que l'on a

$$r^n - (n+1)(r-1)^n + \dots \mp (n+1)(r-n)^n \pm (r-n-1)^n = \Delta^{n+1}(r-n-1)^n,$$

le signe + ayant lieu si n est impair, et le signe - s'il est pair; or $\Delta^{n+1}r^n = 0$, d'où il est facile de conclure l'équation (μ).

IX.

Pour appliquer la théorie précédente à la Nature, il faudrait supposer $n = 63$, parce qu'il existe présentement soixante-trois comètes dont on a calculé les orbites; mais ce calcul serait pénible à cause de sa longueur; ainsi, l'abandonnant à ceux qui désireront de l'entreprendre, je me contenterai de supposer ici $n = 12$; j'imagine donc la droite AB, partagée en douze parties égales, dont chacune soit conséquemment de $7^{\circ}\frac{1}{2}$; on trouvera, par l'article précédent, que la probabilité que l'inclinaison moyenne des douze orbites sera comprise entre $45^{\circ} - 7^{\circ}\frac{1}{2}$ et 45° , ou entre $45^{\circ} + 7^{\circ}\frac{1}{2}$ et 45° , est égale à

$$\frac{6^{12}}{7^{12}} \left[1 - 13 \left(\frac{5}{6}\right)^{12} + \frac{13.12}{1.2} \left(\frac{4}{6}\right)^{12} - \frac{13.12.11}{1.2.3} \left(\frac{3}{6}\right)^{12} + \frac{13.12.11.10}{1.2.3.4} \left(\frac{2}{6}\right)^{12} - \frac{13.12.11.10.9}{1.2.3.4.5} \left(\frac{1}{6}\right)^{12} \right].$$

Or, en faisant le calcul, je trouve cette quantité égale à 0,339; d'où il suit : 1^o qu'il y a 839 à parier contre 161, que l'inclinaison moyenne de douze orbites sera au-dessus de $37^{\circ}\frac{1}{2}$; 2^o qu'il y a autant à parier

302 SUR L'INCLINAISON MOYENNE

qu'elle sera au-dessous de $52^{\circ}\frac{1}{2}$; 3^o qu'il y a 678 à parier contre 322, qu'elle sera entre les deux limites $37^{\circ}\frac{1}{2}$ et $52^{\circ}\frac{1}{2}$.

Maintenant, si l'on ajoute ensemble les indications des douze dernières comètes observées dont voici le Tableau :

Comètes des années.	Inclinaison des orbites.
1774.....	82.48. 0
1773.....	61.25.21
1772.....	18.59.40
1771.....	11.15.29
1771.....	31.25.55
1770.....	1.44.30
1769.....	40.42.30
1766.....	8.20. 0
1766.....	50.50.20
1764.....	53.54.19
1763.....	73.39.29
1762.....	85. 3. 2

on trouvera que leur inclinaison moyenne est de $42^{\circ}31'$. Pour soupçonner dans ces comètes une cause qui tende à les faire mouvoir dans le plan de l'écliptique, il faudrait qu'il y eût un très grand nombre à parier contre l'unité que, si elles étaient lancées au hasard, leur inclinaison moyenne surpasserait $42^{\circ}30'$; or nous venons de trouver qu'il y a 839 contre 161, ce qui ne fait pas six contre un à parier qu'elle sera au-dessus de $37^{\circ}\frac{1}{2}$, et il y a considérablement moins à parier qu'elle sera au-dessus de $42^{\circ}30'$.

X.

Sur la figure de la Terre.

Lorsque Newton voulut déterminer la figure de la Terre, il considéra

0.22 Dice for Statistical Investigations

Stigler's 1991 article Stochastic Simulation in the Nineteenth Century gives a very interesting account of the era before we had tables of, or machine generate-able', random numbers.³³

One of these methods is due to Francis Galton, published in 1890. You can find his description here: <http://galton.org/essays/1890-1899/galton-1890-dice.pdf> . This website <http://galton.org> has copies of all of his work.

- i. What if you were to replace his dice with pentakisdodecahedrons, i.e., polyhedra with 60 faces, 90 edges, and 32 vertices? What values would you write in the edges of Pentakisdodecahedron I and Pentakisdodecahedron II?
- ii. In Monte Carlo work, 'back in the days when' we had to program our own and didn't like to have to program the Inverse Gaussian CDF, many statistical experimenters used the sum of 12 Uniform(0,1) independent random variables, minus 6, as a reasonable way of generating N(0,1) observations. (Box-Muller's method was another, provided the trig functions were built in).

Why do you think a sum of 12 , rather than say 10, or 30, was used?

- iii. Can you suggest any time-savings or improvements when using Galton's dice?

³³Gosset (aka 'Student') preferred to simulate from 'real' data: see

Gosset's 750 samples of size n=4 on this Webpage

<https://jhanley.biostat.mcgill.ca/Student/750samplesOf4/>

on JH's website devoted to Student

<https://jhanley.biostat.mcgill.ca/Student/> .

"The measurements were written out on 3,000 pieces of cardboard, which were then very thoroughly shuffled and drawn at random. As each card was drawn its numbers were written down in a book, which thus contains the measurements of 3,000 criminals in a random order. Finally, each consecutive set of 4 was taken as a sample - 750 in all - and the mean, standard deviation, and correlation of each sample determined."

Full 1908 article: 'The Probable Error of a Mean' by 'Student'

<https://jhanley.biostat.mcgill.ca/bios601/Mean-Quantile/Student-t.pdf>

0.23 Same-sex or opposite-sex?

Details can be found in section 18.2.7 of 'the book'.

0.24 Did you fall into the same trap that JH did?

For many years, and especially when teaching epidemiology students, JH explained why 'variance' (σ^2) is a quantity more used by mathematical statisticians than by applied statisticians and by scientists, who tend to use (and report) standard deviations.

He explained their use of variance by the fact that many of the laws for the spread of functions of random variables (especially sums and differences) involve sums of variances. He usually uses the variability in the weights of an elevator-full of passengers to show that standard deviations don't add, but that variances (and covariances when appropriate) do.

But he always joked about the units would be if, instead of working with standard deviations, we work with variances in the case where (say) human fertility in a certain population had a mean of (say) 1.5 children per woman, and a standard deviation of (say) 1.2 children per woman! What are the units in the $1.2^2 = 1.44$?

He also justified the focus on σ by saying that whereas one could, in principle, calculate $E[|Y - \mu_Y|]$ or its sample counterpart $(1/n) \times \sum |y - \bar{y}|$ as a measure of spread, it would take two passes through the data to arrive at it. First, one would have to calculate \bar{y} and then one would have to calculate the mean of the deviations from it. As JH was well aware from the days of having to use a mechanical calculator, one could calculate both \bar{y} and s in a single pass, as long as one had two 'registers' that could cumulate the sum of the y 's and the sum of their squares. This exploits the fact that $E[(Y - \mu_Y)^2] = E(Y^2) - \mu_Y^2$.

So it was easier to work with **squared** deviations and to then at the end take the square root of the average squared deviation.

So, now that students don't have to manually cumulate y 's and y^2 's, we might ask: couldn't we just skip the squares, and take the average of the deviations from \bar{y} , only to realize that the average would be zero. So, students accept the fact that we can avoid this by squaring first, and then taking the square root (of the mean square) at the end.

But this seemed like just a trick, and so JH would instead offer a way that seemed more natural: why not just drop the sign, and take the average (mean) of the absolute deviations? So he would calculate the mean age in the class, and then ask each student to call out their deviation from this, but without

the sign. He would ask students to listen to all of the (absolute) deviations being called out, and to make an informal rough summary, either a mean or a median.

And he used to say (without much evidence or testing) that the average absolute deviation or (more accurately) the *mean* average deviation (MAD) (which would be easier to explain to people who no longer knew to calculate a square root) would be quite close to the ‘standard’ deviation obtained as a root mean square.

All this changed last year, when working 1 on 1 with a (bios601) student from the Math Stat department!

It started when he mentioned how easy it was to get some students to predict that a sample variance based on $n = 5$ would be smaller than one based on (say) $n = 10$ and some to predict the opposite. [We will test this ourselves in Q 0.6]

But then, JH went on to say why he thought the RMS deviation or SD was not a natural measure, and how the MAD made more sense to him. And he did a few simulations, only to find out that the ‘MA’ deviation was not at all close to the ‘standard’ deviation. Indeed, when sampling from the $N(0, \sigma)$ distribution, the MAD was only about 80% of σ !

Exercise³⁴

- i. Would you have expected this?
- ii. Check, by simulation, that JH’s simulations were done correctly. Hint: you can just evaluate `mean(abs(rnorm(100000)))` a few times! or use brute force in R: `integrate(function(x) abs(x) * dnorm(x), -Inf, Inf)`
- iii. Presumably, the correct answer is some function of e or π or some such. Any ideas as to what it is?
- iv. Can you establish the ratio formally from mathematical-statistics theory? If not, maybe Google for articles/posts that do so.*
* Relationship Between MAD and Standard Deviation for a Normally Distributed Random Variable
- v. Do you think MAD is smaller than σ for all possible distributions? Explain your answer.

³⁴To keep it simpler, assume we know that $\mu = 0$.

0.25 Air New Zealand to weigh passengers before they board the airplane



CNN May 2023:

Remove your shoes. Take your keys out of your pocket. Step on the scale?

That's right: New Zealand's Civil Aviation Authority is asking that its national airline weigh passengers departing on international flights from Auckland International Airport through July 2, 2023.

The program, which Air New Zealand calls a passenger weight survey, is a way to gather data on the weight load and distribution for planes, the airline said.

"We weigh everything that goes on the aircraft – from the cargo to the meals onboard, to the luggage in the hold," Alastair James, the airline's load control improvement specialist said in a statement. **"For customers, crew and cabin bags, we use average weights, which we get from doing this survey."**

Still, weight is a personal thing that not everyone wishes to disclose. In order to protect individuals' privacy, the airline says it has made the data anonymous.

Travelers will be asked to stand on a digital scale when they check in for their flight. The information about their weight is then submitted to the survey but will not be viewable on the agent's screen.

They will also place their luggage on another identical scale for separate weighing. "We know stepping on the scales can be daunting. We want to reassure our customers there is no visible display anywhere. No one can see your weight, not even us," James said. This isn't the first time that Air NZ has asked passengers to step on the scales before boarding their flights. Domestic passengers took part in a survey in 2021, but the one for international travelers was delayed due to the pandemic. Among the people who may be asked to take part in the survey are those traveling on the direct flight from Auckland to New York City's JFK Airport.

The 17-hour flagship route was launched last fall as a lynchpin of Air NZ's post-pandemic strategy. It's also one of the longest flights in the world.

Exercise

Q10, assigned in previous years, was something that JH just thought up, but this application is real.

From what JH can find online, the Auckland to New York City direct flight will use a 787-9+Dreamliner. It weights 115,300 kg when empty; the maximum cargo weight is 60,395 kg, the maximum takeoff weight is 252,650 kg and it burns about 5,500 litres of fuel per hour. (The maximum fuel capacity is about 125,000 litres, which would allow it to fly about 22.5 hours but they have to leave a safety margin!) It can seat 296 passengers (two-classes), and has a quoted range of 14,010 km.

For now, leave aside the cargo, crew and luggage and focus on the possible (combined) weight of a full complement of 296 passengers.

And for now assume that their survey of $n = 10,000$ international passengers yields a mean (M) of 178 pounds and a standard error of the mean (SEM) of 0.4 pounds. [JH came up with these numbers based on the weights of US adults in the NHANES surveys from 2014 to 2018. See 16.3.3 Example 2 of the [online book](#) he started 2 months before COVID-19 and hopes to get much further with in the next few years.]

Think ahead to post-survey flights, where the airline will have to rely on the average from the survey.

- Calculate (and give the associated probabilities as to) how far off the total weight of 298 passengers on a flight could be from the estimate based on the average in the survey.
- And state your assumptions, and ways they might be violated – and why you might need to be a bit more conservative with your 'margins'. [See hint in section 16.4 of the online book, or [this item](#)]

0.26 Characterizing CLT-friendly and CLT-less-friendly distributions

Here is some R code³⁵ that uses convolutions to compute the probability mass function of a sum or mean of n i.i.d. integer-valued random variables [in the code, y has probability mass pmf].

```
PMF = function(y, pmf, n.max) {

  # y : a vector of consecutive integers (possibly shifted by a constant)
  # pmf: probability mass function [no need to scale so they add to 1]
  # n.max : sample sizes 1 to n.max

  pmf = pmf/sum(pmf) # in case not already scaled

  n.values = length(y)
  SUM = PMF = MEAN = vector("list",n.max)
  E = SKEWNESS = KURTOSIS = vector("list",n.max)

  PMF[[1]] = pmf; SUM[[1]] = y ; MEAN[[1]] = y/1
  E[[1]] = sum(y*pmf); DEV = MEAN[[1]]-E[[1]]
  V = sum( PMF[[1]] *DEV^2 ) ; SD = sqrt(V )
  SKEWNESS[[1]] = sum( PMF[[1]] * DEV^3 )
  KURTOSIS[[1]] = sum( PMF[[1]] * DEV^4 )

  if(n.max > 1){
    for(i in 2:n.max){
      P = outer(pmf, PMF[[i-1]] )
      #if( any( dim(P) ) ==1 ) new = P
      #if( dim(P)[1]>1 & dim(P)[2]>1 )
      new = sapply( split(P,
        col(P)+row(P)) , sum )
      #print(length(new))
      PMF[[i]] = as.vector(new)
      SUM[[i]] = (i*min(y)):(i*max(y))
      MEAN[[i]] = SUM[[i]]/i
      E[[i]] = sum( MEAN[[i]] * PMF[[i]] )
      DEV = MEAN[[i]]-E[[i]]
      V = sum( PMF[[i]] *DEV^2 )
      DEV = DEV/sqrt(V)
      SKEWNESS[[i]] = sum( PMF[[i]] * DEV^3 )
      KURTOSIS[[i]] = sum( PMF[[i]] * DEV^4 )
    } # i
  } # n > 1
  return(list(PMF=PMF,SUM=SUM,MEAN=MEAN,E=E,
    SKEWNESS=SKEWNESS,
    KURTOSIS=KURTOSIS) )
}
```

Your **exercise** is to try it out on a variety of ‘sampled from’ distributions, and to discover which shapes are user friendly, and which are less so – i.e., for what shapes does it take a larger/smaller n for the sampling distribution of the sum or mean to come close to Normality?

Obviously, if you start with $y \sim N$, then sample means or sample sums are exactly N even at $n = 1$. [To test the code, you could ‘mimic’ a N distribution with the probability mass limited to 61 integers, and look at $n = 2$:]

```
y= -30:30 ; pmf = dnorm(y,0,10) ; result = PMF(y,pmf,n.max=2)
```

What if the concern is with the mean or total number of points showing when you toss n fair dice (1:6 points each, with equal probability)? Try it out:

```
result = PMF(y = 1:6-mean(1:6), # centered at 0, w.l.o.g.
  pmf=rep(1/6,6), # Uniform
  n.max=15)
```

```
par(mfrow=c(3,5), mar=c(2,3,1,1))
for(n in 1:15 ){
  XLIM=range(result$MEAN[[n]])
  YLIM=c(0,max(result$PMF[[n]]))
  plot(result$MEAN[[n]],
    result$PMF[[n]],
    type="h", lend=1, xlim=XLIM,
    ylim=c(0,max(result$PMF[[n]])),
    main=paste0("n = ",n))
  abline(v=result$E[[n]],
    col="blue", lty="dotted", lwd=1)
  text(XLIM[2],0.96*YLIM[2],
    paste0("S: ",
      format(round(
        result$SKEWNESS[[n]],2),nsmall=2),
      "\nK:"),
    format(
      round(
        result$KURTOSIS[[n]],2),nsmall=2)
    ),
    adj=c(1,1),cex=0.5
  )
}
```

EXERCISE Try various shapes [minimal support(just 0-1)*, symmetric/not, uni/multi-modal, long/short tails, etc.] and come up with you ranking, from most-to-least CLT-friendly.

You may use ones from Notes/Exercises, but also go outside that box. E.g., the no. of cars per household was a 3-point distribution; could limit the variable to just 0/1*, (e.g have internet) or use say 0:10 for no. of persons in household, with realistic pmf; or no. of devices (0:? pmf?). Or how much McGill students pay monthly in rent (in \$100 units) or phone plans (\$10 units) or how many times a day (units?) they/you look at their/your phone.

³⁵Also available [here](#).