



COLUMBIA
BUSINESS
SCHOOL

The Normal Distribution

Fall 2001
B6014: Managerial Statistics

Professor Paul Glasserman
403 Uris Hall

1. The normal distribution (the familiar bell-shaped curve) is without question the most important distribution in all of statistics. A broad range of problems, particularly those involving large amounts of data, can be solved using the normal distribution. Although the normal distribution is continuous, it is often used to approximate discrete distributions. For example, we might say that the scores on an exam are (approximately) normally distributed, even though the scores are discrete.
2. There are actually many different normal distributions. To fix a particular normal, we must specify the mean μ and the variance σ^2 . If X has this normal distribution, we write $X \sim N(\mu, \sigma^2)$. This notation says “ X is normally distributed with mean μ and variance σ^2 .” We call μ and σ^2 the **parameters** of the normal distribution.
3. Once we specify the parameters μ and σ^2 , we have completely specified the distribution of a normal random variable. This is not true for arbitrary random variables: ordinarily, the mean and the variance do not completely determine the distribution, but for a normal random variable they do.
4. Increasing μ shifts the normal density to the right without changing its shape. Increasing σ^2 flattens the density without shifting it. See Figure 1.
5. Some examples:
 - A machine that fills bags of potato chips cannot put exactly the same weight of chips into every bag. Suppose the quantity poured into an 8 ounce bag is normally distributed with a mean of 8.3 ounces and a standard deviation of 0.2 ounces. We might ask what proportion of bags contain less than 8 ounces of chips.
 - An airplane manufacturer wants to build a smaller carrier with the same seating capacity. If the height of men is normally distributed with a mean of 5 feet 9 inches and a standard deviation of 2 inches, we might ask how low the ceiling can be so that at most 2% of men will have to duck while walking down the aisle.

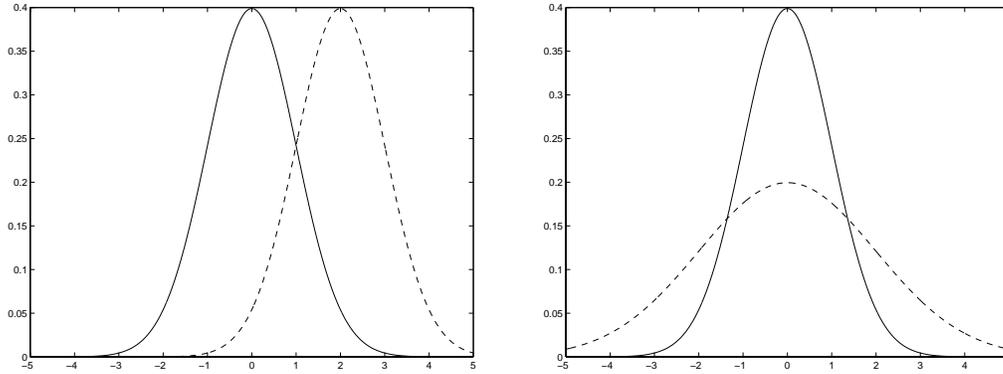


Figure 1: Left panel shows two normal distributions with different means ($\mu = 0$ for solid, $\mu = 2$ for dashed) and same standard deviation ($\sigma = 1$). Right panel shows two normal distributions with the same mean ($\mu = 0$) and different standard deviations ($\sigma = 1$ for solid, $\sigma = 2$ for dashed).

- Suppose that weekly fluctuations in the price of XYZ stock are well approximated by a normal distribution with a mean of .3% and a standard deviation of .4%. What is the probability that the price will drop 1% or more in one week?
6. We cannot answer these types of questions just by using a calculator because there is no formula to evaluate probabilities for the normal distribution. Instead, we use the table of **cumulative** probabilities for the **standard normal distribution**. The standard normal distribution is $N(0, 1)$; i.e., the normal distribution with mean 0 and variance 1. Probabilities for any normal distribution $N(\mu, \sigma^2)$ can be found from a table for $N(0, 1)$. (The table appears at the end of these notes.) To see this, we need a few properties of normal random variables.
 7. Converting a probability question concerning a general normal distribution $N(\mu, \sigma^2)$ into one concerning the standard normal $N(0, 1)$ is called **standardizing**. Intuitively, switching from $N(\mu, \sigma^2)$ involves thinking about the problem in terms of standard deviations away from the mean.
 8. The key to standardizing is the following property: If $X \sim N(\mu, \sigma^2)$ then Z defined by

$$Z = \frac{X - \mu}{\sigma}$$

is a standard normal random variable: $Z \sim N(0, 1)$. Notice that

$$X = \mu + Z\sigma$$

so Z measures X in standard deviations away from μ .

9. We need the following property: any **linear transformation** of a normal random variable is normal: if X is normal, then so is $a + bX$ for any constants a and b .

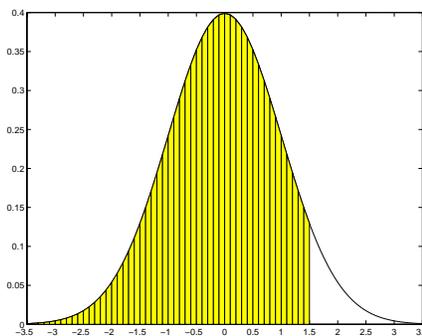


Figure 2: The normal table gives the area under the normal curve to the left of z for values of z ranging from 0 to 3.49. The figure illustrates the area to the left of $z = 1.50$, which is given in the table to be 0.9332. This is the probability $P(Z \leq 1.50)$.

10. This transformation from X to Z is called **standardizing** the random variable X . Notice that is a linear transformation with $a = (-\mu/\sigma)$ and $b = 1/\sigma$; i.e.,

$$Z = (-\mu/\sigma) + (1/\sigma)X.$$

11. The standard normal table gives $P(Z \leq z)$ for various values of z , with $Z \sim N(0, 1)$. The row and column labels determine the value of z (the column labels determine the last digit). The values inside the table give the area to the left of z under the normal density, as illustrated in Figure 2.
12. Using the table: If $X \sim N(\mu, \sigma^2)$, then to find $P(X \leq x)$, for any x , first note that

$$\begin{aligned} P(X \leq x) &= P\left(\frac{X - \mu}{\sigma} \leq \frac{x - \mu}{\sigma}\right) \\ &= P\left(Z \leq \frac{x - \mu}{\sigma}\right). \end{aligned}$$

If $(x - \mu)/\sigma$ is positive, look it up in the table under “ z ”. The corresponding entry gives the desired probability.

13. Example: Suppose $X \sim N(10, 4)$ so that $\sigma = 2$. Suppose we want to find $P(X \leq 13)$. By standardizing, we find that

$$P(X \leq 13) = P\left(\frac{X - 10}{2} \leq \frac{13 - 10}{2}\right) = P(Z \leq 1.5).$$

Now look up 1.50 in the table to get 0.9332, as illustrated in Figure 2. We conclude that $P(X \leq 13) = 0.9332$. Interpretation: 13 is 1.5 standard deviations above the mean for X , so $P(X \leq 13)$ is the same as $P(Z \leq 1.5)$.

14. In spreadsheets, the function `NORMSDIST` returns values from the cumulative normal distribution. For example, in the calculation just illustrated, `NORMSDIST(1.5)` returns the value 0.9332 we found in the table.

15. Keep in mind that the normal distribution is continuous, so $P(X \leq x) = P(X < x)$ for all x . In other words, $P(X = x) = 0$ for all x .
16. Notice that the standard normal table only gives probabilities $P(Z \leq z)$ for positive values of z . To find $P(Z \leq -z)$ for negative values $-z$, we use the symmetry of the normal distribution. If $z > 0$, then

$$P(Z \leq -z) = P(Z \geq z) = 1 - P(Z \leq z).$$

Thus, to find $P(Z \leq -z)$, look up $P(Z \leq z)$ in the table and subtract the tabulated value from 1.

17. Examples: $P(Z > -1.5)$ is the shaded area in Figure 3. By symmetry, this is the same as the shaded area in Figure 2. So, to find $P(Z > -1.5)$ we look up $P(Z < 1.5)$ which is 0.9332. Now suppose we want $P(Z < -1.5)$. This is the unshaded area in Figure 3. By symmetry, this is the same as the unshaded area in Figure 2. Since the total area under the curve is 1, this unshaded area is given by $1 - 0.9332 = 0.0668$.

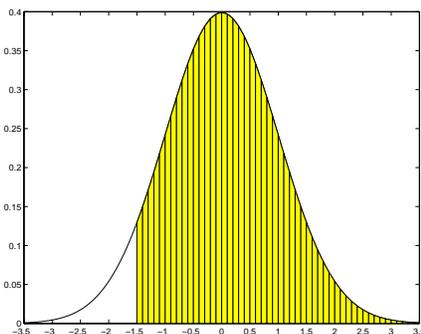


Figure 3: By symmetry of the normal distribution, $P(Z > -1.5)$, the shaded area, is the same as $P(Z < 1.5)$, the shaded area in Figure 2. The unshaded areas in the two figures are also the same and are equal to $P(Z < -1.5)$ and $P(Z > 1.5)$.

18. How do we find $P(-1.5 < Z < 1.5)$ the shaded area in Figure 4? Write it as $P(Z < 1.5) - P(Z < -1.5)$: it's the area to the left of 1.5 but with the area to left of -1.5 taken away. Using the values we found previously, this is $0.9332 - 0.0668 = 0.8664$.
19. Example: Consider the potato chip problem above. We want to find $P(X < 8)$ if $X \sim N(8.3, (0.2)^2)$. By standardizing, we find that

$$P(X < 8) = P\left(\frac{X - 8.3}{0.2} < \frac{8 - 8.3}{0.2}\right) = P(Z < -1.5).$$

We now look up the value 1.5 in the table to get .9332. Now subtract from 1 to get $1 - .9332 = .0668$. We conclude that $P(X < 8) = P(Z < -1.5) = 1 - P(Z \leq 1.5) = .0668$.

20. By inverting the steps that take us from X to Z , we get, once again,

$$X = \mu + Z\sigma.$$

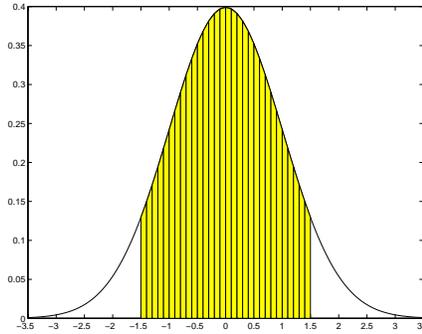


Figure 4: The shaded area is $P(-1.5 < Z < 1.5)$. We find it from the normal table by expressing it as $P(Z < 1.5) - P(Z < -1.5)$.

As noted above, the standardized r.v. Z tells us how many standard deviations away from its mean the random variable X lands. In the potato-chip example, we should think of the 8-ounce minimum required as 1.5 standard deviations below the mean of 8.3 ounces.

21. Sometimes, we use the table to go in the opposite direction. Example: Suppose the potato-chip dispenser can be adjusted to any mean while leaving the standard deviation unchanged. To what mean μ should we set the machine so that only 5% of bags contain less than 8 ounces?

Answer: Let X be the weight of chips in a bag when the machine is set to a mean of μ ; $X \sim N(\mu, (0.2)^2)$. We want to set μ so that $P(X < 8) = .05$; i.e.,

$$\begin{aligned} .05 &= P(X < 8) \\ &= P\left(\frac{X - \mu}{0.2} < \frac{8 - \mu}{0.2}\right) \\ &= P\left(Z < \frac{8 - \mu}{0.2}\right). \end{aligned}$$

Clearly, μ will have to be bigger than 8, so $(8 - \mu)/0.2$ is negative. We should therefore write

$$.05 = P\left(Z < \frac{8 - \mu}{0.2}\right) = 1 - P\left(Z < \frac{\mu - 8}{0.2}\right);$$

i.e.,

$$P\left(Z < \frac{\mu - 8}{0.2}\right) = .95.$$

Now we look for .95 in the body of the table, since this is our target probability. We find that this corresponds to $z = 1.65$. Thus, we must choose μ so that

$$\frac{\mu - 8}{0.2} = 1.65.$$

In other words,

$$\mu = 8 + 1.65(0.2) = 8.33.$$

This tells us that we must set the machine 1.65 standard deviations above the threshold of 8 to make sure that only 5% of bags fall below the threshold.

22. Correctly carrying out the standardization is the key to solving virtually all problems involving the normal distribution.
23. Important fact: Sums of normally distributed random variables are normally distributed. The usual rules for means and variances apply: if $X_1 \sim N(\mu_1, \sigma_1^2)$ and $X_2 \sim N(\mu_2, \sigma_2^2)$, then $X_1 + X_2$ is normal with a mean of $\mu_1 + \mu_2$. The variance of the sum depends on the correlation between X_1 and X_2 . If they are independent, then $Var[X_1 + X_2] = \sigma_1^2 + \sigma_2^2$.

The Central Limit Theorem

1. The importance of the normal distribution is in large part explained by the **central limit theorem**. This result states that if X_1, X_2, \dots, X_n are random variables
 - all having the same distribution, with mean μ and variance σ^2 ; and
 - all independent,

then the **sample mean**

$$\bar{X} = (X_1 + X_2 + \dots + X_n)/n$$

is approximately normally distributed with mean μ and variance σ^2/n .

2. Notice that we have not assumed that the X_i 's are normally distributed. (If they were, we would not need to say "approximately.") Thus, no matter what the distribution of the underlying data, as we take averages we get something that looks like a normal distribution.
3. Make sure you understand the two hypotheses for the central limit theorem: identically distributed and independent random variables are required.
4. Another way to state the central limit theorem is that

$$Z = \frac{\bar{X} - \mu}{\sigma/\sqrt{n}}$$

tends towards the standard normal distribution.

5. The quality of the normal approximation in the central limit theorem improves as n increases. Generally, $n \geq 30$ is recommended.
6. The central limit theorem will be the justification for using the normal distribution when we talk about confidence intervals and hypothesis tests.

Cumulative Probabilities for the Standard Normal Distribution.
 This table gives probabilities to the left of given z values for the
 standard normal distribution

z	0.00	0.01	0.02	0.03	0.04	0.05	0.06	0.07	0.08	0.09
0	0.5000	0.5040	0.5080	0.5120	0.5160	0.5199	0.5239	0.5279	0.5319	0.5359
0.1	0.5398	0.5438	0.5478	0.5517	0.5557	0.5596	0.5636	0.5675	0.5714	0.5753
0.2	0.5793	0.5832	0.5871	0.5910	0.5948	0.5987	0.6026	0.6064	0.6103	0.6141
0.3	0.6179	0.6217	0.6255	0.6293	0.6331	0.6368	0.6406	0.6443	0.6480	0.6517
0.4	0.6554	0.6591	0.6628	0.6664	0.6700	0.6736	0.6772	0.6808	0.6844	0.6879
0.5	0.6915	0.6950	0.6985	0.7019	0.7054	0.7088	0.7123	0.7157	0.7190	0.7224
0.6	0.7257	0.7291	0.7324	0.7357	0.7389	0.7422	0.7454	0.7486	0.7517	0.7549
0.7	0.7580	0.7611	0.7642	0.7673	0.7704	0.7734	0.7764	0.7794	0.7823	0.7852
0.8	0.7881	0.7910	0.7939	0.7967	0.7995	0.8023	0.8051	0.8078	0.8106	0.8133
0.9	0.8159	0.8186	0.8212	0.8238	0.8264	0.8289	0.8315	0.8340	0.8365	0.8389
1	0.8413	0.8438	0.8461	0.8485	0.8508	0.8531	0.8554	0.8577	0.8599	0.8621
1.1	0.8643	0.8665	0.8686	0.8708	0.8729	0.8749	0.8770	0.8790	0.8810	0.8830
1.2	0.8849	0.8869	0.8888	0.8907	0.8925	0.8944	0.8962	0.8980	0.8997	0.9015
1.3	0.9032	0.9049	0.9066	0.9082	0.9099	0.9115	0.9131	0.9147	0.9162	0.9177
1.4	0.9192	0.9207	0.9222	0.9236	0.9251	0.9265	0.9279	0.9292	0.9306	0.9319
1.5	0.9332	0.9345	0.9357	0.9370	0.9382	0.9394	0.9406	0.9418	0.9429	0.9441
1.6	0.9452	0.9463	0.9474	0.9484	0.9495	0.9505	0.9515	0.9525	0.9535	0.9545
1.7	0.9554	0.9564	0.9573	0.9582	0.9591	0.9599	0.9608	0.9616	0.9625	0.9633
1.8	0.9641	0.9649	0.9656	0.9664	0.9671	0.9678	0.9686	0.9693	0.9699	0.9706
1.9	0.9713	0.9719	0.9726	0.9732	0.9738	0.9744	0.9750	0.9756	0.9761	0.9767
2	0.9772	0.9778	0.9783	0.9788	0.9793	0.9798	0.9803	0.9808	0.9812	0.9817
2.1	0.9821	0.9826	0.9830	0.9834	0.9838	0.9842	0.9846	0.9850	0.9854	0.9857
2.2	0.9861	0.9864	0.9868	0.9871	0.9875	0.9878	0.9881	0.9884	0.9887	0.9890
2.3	0.9893	0.9896	0.9898	0.9901	0.9904	0.9906	0.9909	0.9911	0.9913	0.9916
2.4	0.9918	0.9920	0.9922	0.9925	0.9927	0.9929	0.9931	0.9932	0.9934	0.9936
2.5	0.9938	0.9940	0.9941	0.9943	0.9945	0.9946	0.9948	0.9949	0.9951	0.9952
2.6	0.9953	0.9955	0.9956	0.9957	0.9959	0.9960	0.9961	0.9962	0.9963	0.9964
2.7	0.9965	0.9966	0.9967	0.9968	0.9969	0.9970	0.9971	0.9972	0.9973	0.9974
2.8	0.9974	0.9975	0.9976	0.9977	0.9977	0.9978	0.9979	0.9979	0.9980	0.9981
2.9	0.9981	0.9982	0.9982	0.9983	0.9984	0.9984	0.9985	0.9985	0.9986	0.9986
3	0.9987	0.9987	0.9987	0.9988	0.9988	0.9989	0.9989	0.9989	0.9990	0.9990
3.1	0.9990	0.9991	0.9991	0.9991	0.9992	0.9992	0.9992	0.9992	0.9993	0.9993
3.2	0.9993	0.9993	0.9994	0.9994	0.9994	0.9994	0.9994	0.9995	0.9995	0.9995
3.3	0.9995	0.9995	0.9995	0.9996	0.9996	0.9996	0.9996	0.9996	0.9996	0.9997
3.4	0.9997	0.9997	0.9997	0.9997	0.9997	0.9997	0.9997	0.9997	0.9997	0.9998