As the name implies, the Multivariate Normal (MNVN) distribution generalizes the Normal distribution to higher dimensions.

Def: A vector of n random variables $\mathbf{X} = (X_1, \dots, X_n)$ is said to have *Multivariate Normal* distribution if every linear combination of the X_j has a Normal distribution, that is,

$$t_1 X_1 + \cdots + t_n X_n$$

is Normal for every choice of constants $t_1, \dots, t_n$. Moreover, if the linear combination is constant, we still consider the result a Normal distribution with variance 0.

In the specific case when $n = 2$, we say that $\mathbf{X}$ is *bivariate Normal.*

One consequence of this definition is the following:

Thm: If $\mathbf{X} = (X_1, \dots, X_n)$ is MVN, and $\mathbf{A}$ is the matrix for a linear transformation, then $\mathbf{AX}$ is MVN also.

Proof. Let $\mathbf{Y} = \mathbf{AX}$, with $\mathbf{Y} = (Y_1, \dots, Y_k)$. By definition of matrix multiplication, each Y_j is a linear combination of the X_j. And therefore, any linear combination of the Y_j will be a linear combination of a linear combination of the X_j, which is a linear combination of the X_j. By the MVN property of $\mathbf{X}$, this linear combination will be Normal. So $\mathbf{Y}$ is MVN as well. □

It turns out that essentially every MVN random variable arises in this way! Before we get there, we introduce some notation. Recall from the text that every MVN variable is completely specified by two collections of parameters: a vector of mean $\boldsymbol{\mu} = ([E[X_1], \dots, E[X_n]])$, and a matrix of covariances $\boldsymbol{\Sigma}$, where $\Sigma_{ij} = \operatorname{Cov}(X_i, X_j)$.

Thm: Let $\mathbf{X}$ be an n-dimensional MVN with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$, and let $\mathbf{Z}$ be a standard MVN; that is $Z_1, \dots, Z_n$ are iid $N(0, 1)$. Then there exists a matrix $\mathbf{A}$ so that $\mathbf{AZ} + \boldsymbol{\mu}$ has the same distribution as $\mathbf{X}$.

Proof. A proof of the general case is available on request after class. A proof of the bivariate case is given in Example 7.5.10. Specifically, suppose $X \sim N(\mu_x, \sigma_x^2)$ and $Y \sim N(\mu_y, \sigma_y^2)$, and $\operatorname{Corr}(X, Y) = \rho$. Let Z_1, Z_2 be iid $N(0, 1)$. Then the MVN (V, W) given by

$$V = \sigma_x Z_1 + \mu_x \qquad W = \sigma_y(\rho Z_1 + \sqrt{1 - \rho^2} Z_2) + \mu_y$$

has the same distribution as (X, Y). To complete the proof, compute the means, variances, and covaraince of V, W. □

Thus far, we've avoided discussing that actual joint PDF of the MVN. But now:

Thm: If $\mathbf{X}$ is MVN, the joint PDF of $\mathbf{X}$ is

$$f(\mathbf{x}) = \frac{1}{(2\pi)^{k/2} \sqrt{\det \boldsymbol{\Sigma}}} e^{-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu})}$$

where $\boldsymbol{\Sigma}$ is the matrix $\Sigma_{ij} = \operatorname{Cov}(X_i, X_j)$, and where $\boldsymbol{\mu} = ([E[X_1], \dots, E[X_n]])$.

Proof. Use the multivariate change-of-variables formula. □

If (X, Y) is bivariate Normal and each of X and Y are marginally $N(0, 1)$, and if $\operatorname{Cov}(X, Y) = \operatorname{Corr}(X, Y) = \rho$, then the joint PDF can be written succinctly as

$$f(x, y) = \frac{1}{2\pi\sqrt{1 - \rho^2}} \exp\left(-\frac{x^2 - 2\rho xy + y^2}{2(1 - \rho^2)} \right).$$