# Probability Playground: Exploring Probability Distributions Through Interaction

Adam Cunningham*

**Abstract**

Interactive data visualization is widely used by data scientists, but few websites exploring probability distributions use such techniques, and none that we are aware of illustrate the relationships between distributions.

Probability Playground (http://www.probabilityplayground.com) was designed from the ground up to be a highly interactive website with a design philosophy focusing on developing intuition through discovery. It provides a consistent interface for exploring 29 probability distributions and their relationships. Unique features include:

- Animated loading of examples.
- Direct interactive editing of distribution means and variances
- Manual and automatic axis scaling
- Interactive graphing of distribution relationships
- Visualization of the processes generating each distribution

Over 150 proofs are also provided for distribution means, variances, and relationships. Students and educators looking for interactive materials to use in their teaching will find Probability Playground a novel contribution for both its scope and degree of interactivity. It provides a valuable addition to the field of statistics and data science education.

**Key Words:** interactive data visualization, data science education, educational website, interactive website, probability theory, probability distributions

## 1. Introduction

Probability Playground began as a personal project for the author while a student at the University at Buffalo taking the MA program in Biostatistics. Classes introduced a range of probability distributions, with exercises focused on the derivation and algebraic manipulation of formulas. However, the development of an intuitive visual understanding of these distributions and their relationships was mostly absent. The relationship between the visual appearance of a distribution and its parameters was never illustrated, nor was the manner in which the graph of a pdf or pmf converged to a limit, nor were the underlying causal mechanisms which generated the various distributions.

Although many websites exist for teaching probability theory and statistics, the majority of these are based on static text and graphics. The few that do take advantage of interactive web technologies are often restricted in terms of their editing functionality, and none that the author is aware of provide interactive capabilities for illustrating the relationships between distributions.

The author, who has a background as both a software engineer and math educator, was ideally placed to address this gap. Over the course of several years Probability Playground evolved into a robust and extensive resource for visually illustrating probability distributions and their relationships. The website has been tested to be compatible with all major operating systems, web browsers, and devices, providing an ”off-the-shelf” solution for educators looking to incorporate interactive and exploratory learning into their classrooms.

*University at Buffalo, 12 Capen Hall, Buffalo, NY 14260-1660, USA

# 2. Design Philosophy

Probability Playground was created from the ground up to be a highly interactive website with a design philosophy focusing on developing intuition through exploration and discovery. The website was developed with the following five core design principles in mind.

## 2.1 Everything on One Page

An additional burden is placed on the user whenever they are required to navigate between screens or to scroll to find relevant information. Visual similarities and differences between graphs are also lost if they cannot be seen together on one page.

The decision was therefore made early on that the web pages for each distribution should fit on a single page. This had the additional benefit of requiring a minimal design, with priority given regarding which information was to be displayed, the placement of page elements, and their spatial relationship to each other.

## 2.2 Interactive and Dynamic

A fundamental principle of interactive design is that there should be an instant reaction to every user action. In the context of Probability Playground, this requires that whenever any variable is changed, the change is immediately propagated to all dependent variables, graphs, and animations. The information displayed on the website at any instant is therefore always internally consistent. The D3 JavaScript library commonly used for producing dynamic, interactive data visualizations gives the website a responsive feel as the visual information displayed smoothly transitions between states in response to user input.

## 2.3 An Intuitive Interface

Probability Playground provides multiple editing options using common interface elements such as pulldown menus, buttons, text boxes and sliders. The Bootstrap framework that is popular in many commercial websites was used to provide a consistent look-and-feel that web users are already familiar with, as well as a responsive user interface that works across multiple devices.

## 2.4 Ease of Navigation

Options for navigation should be immediately available, and should be accessible using as few user actions as possible. For example, if the user has a choice between a small number of options, a group of buttons displaying the options is used instead of a pulldown menu. The available options are then immediately visible, and each option can be selected with a single user action.

Accessibility for users who have trouble using a mouse (or do not have access to one) was also a priority. Keyboard control is therefore available for all editing and navigation functions. This is also often useful for all users, as it can provide a finer degree of control than using a mouse alone.

## 2.5 A Single Unified Framework

The final principle is that there should be a single unified framework that encompasses all univariate probability distributions. This has implications for usability, maintainability, and extensibility. From the perspective of a user, a single framework means that there is a consistent set of editing and navigation interactions that is the same for all distributions. As the developer of the website, creating a general framework that encompasses all distributions simplified code development and allowed the functionality developed for one specific distribution to be immediately generalized to all. Although this required more analysis upfront, the benefits downstream as distributions were added more than outweighed the additional initial work involved.

## 3. The Map Page

The map serves as the landing page for the website, a visual overview of the relationships between distributions, and a way to navigate between them. This page is shown in figure 1 below.

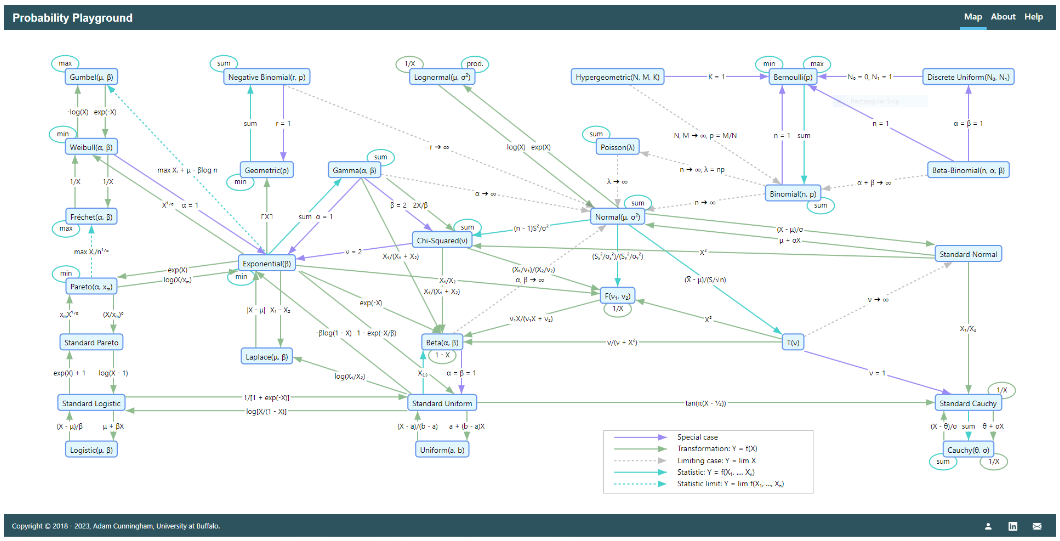


**Figure 1**: The Map page. Nodes represent distributions, while edges represent relationships. Distributions and relationships can be loaded by selecting them on this page.

Distributions are represented as nodes and relationships as edges in a directed graph. Relationships are classified into five types, with the edges color-coded by type. Edge labels describe the relationship, such as the parameter(s) held constant for special cases, the function used in a transformation of variables, or the parameter limits for limiting distributions. This is described in more detail in section 5 below.

Probability Playground has a responsive screen design that can accommodate all screen sizes. For screen sizes less than 768 pixels wide such as many current mobile phones, the map page is replaced by a vertical menu of distributions to select from.

## 4. The Distribution Page

The distribution page on a laptop or PC consists of three columns which respectively describe the distribution, graph the distribution, and graph a related distribution. In keeping with the principles of responsive web design, these are rearranged into two columns on a tablet, while on a phone screen these columns are rearranged sequentially into a single column. The same screen layouts are used throughout the website, resulting in a consistent interface and set of interactions across all univariate probability distributions.

The description given here relates to a screen with a width of 1200 pixels or greater. This page is shown in figure 2 below, using the beta distribution as an example.

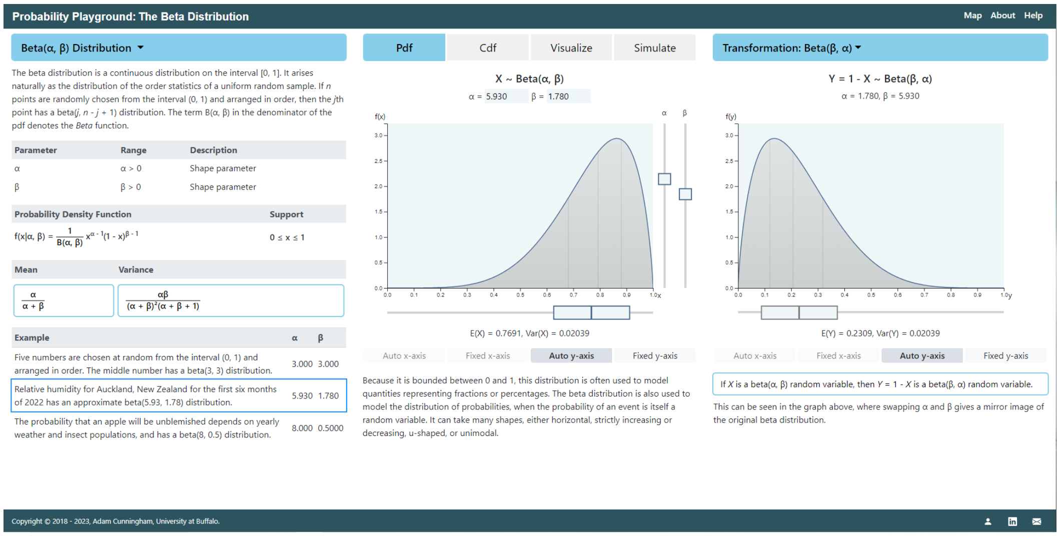


**Figure 2**: The distribution page for the beta distribution. This page provides a wide range of interactive editing capabilities, including text input for parameters, slider control over parameters, and direct editing of distribution means and variances.

In the left-hand column is a description of the distribution, its parameters, pdf/pmf, mean, variance, and some examples. The parameters for an example can be loaded by selecting the example, resulting in an animated update of the graph and the related graph. Proofs for the mean and variance can be viewed by selecting their formulas on the screen. This is described in more detail in section 4.3 below.

In the center column is a graph of the distribution. The distribution parameters are shown both as text above the graph, and through the position of the sliders to the right. The box below the graph shows the mean plus and minus one standard deviation, while the horizontal line through this box shows the interval containing the support of the distribution. This box is not displayed if the mean or variance is undefined (as, for example, with the Cauchy distribution).

Editing of distribution parameters can done in several ways. Parameters can be entered in the text boxes above the graph, which provides precise control over the display. Parameters can also be edited by using the sliders to the right of the graph, resulting in an interactive update of the graph as sliders are moved up or down. Finally, the mean and variance of the graph can be directly edited by interacting with the graph. The mean can be changed either by dragging directly on the graph, or by using the left/right arrow keys on the keyboard while either the graph or the box below it is in focus. The variance can be changed by either zooming directly on the graph, or by using the up/down arrow keys on the keyboard while either the graph or the box below it is in focus. This is described in more detail in subsection 4.1 below. Probability Playground therefore provides a rich set of interaction capabilities to support the development of users' intuition regarding the range of shapes that a distribution may take.

In the right-hand column is a related distribution and its graph, with the parameters of the distribution shown above. The relationship is described both as a formula above the graph and as a text description below. A proof of the relationship can be viewed by selecting the text description shown outlined in blue. The pull-down menu above the related graph is used to see and select all related distributions.

In the example shown in figure 2 above it can be seen that if $X$ has a beta$(\alpha, \beta)$ distribution, then $Y = 1 - X$ has a beta$(\beta, \alpha)$ distribution. The parameters of the related distribution $Y$ are calculated from those of $X$, and are updated dynamically as the parameters of $X$ change. If further parameters are needed to specify the related distribution (such as sample size when $Y$ represents a statistic of a sample drawn from $X$), these additional parameters are shown and edited using sliders to the right of the related graph.

### 4.1 Mean-Variance Editing

For many distributions the mean and variance are controlled by separate parameters. These include location-scale families such as the Cauchy, Laplace, logistic, normal, and uniform distributions. In these cases, controlling the mean and variance of the distribution is done simply by changing the single relevant parameter.

When the mean and/or variance are determined by more than one parameter, there is still educational value in being able to hold the mean constant while changing the variance, and vice-versa. An example is where the beta($\alpha, \beta$) distribution converges to the normal as $\alpha$ and $\beta$ approach infinity. This could be illustrated by separately increasing $\alpha$ and $\beta$ in turn. However, since the mean of the beta depends on both parameters $\alpha$ and $\beta$, this results in the mean repeatedly shifting back and forth across the x-axis as this is done. A more elegant illustration involves holding the mean constant while reducing the variance by increasing both $\alpha$ and $\beta$ together at the same time. This is illustrated in figure 3 below.

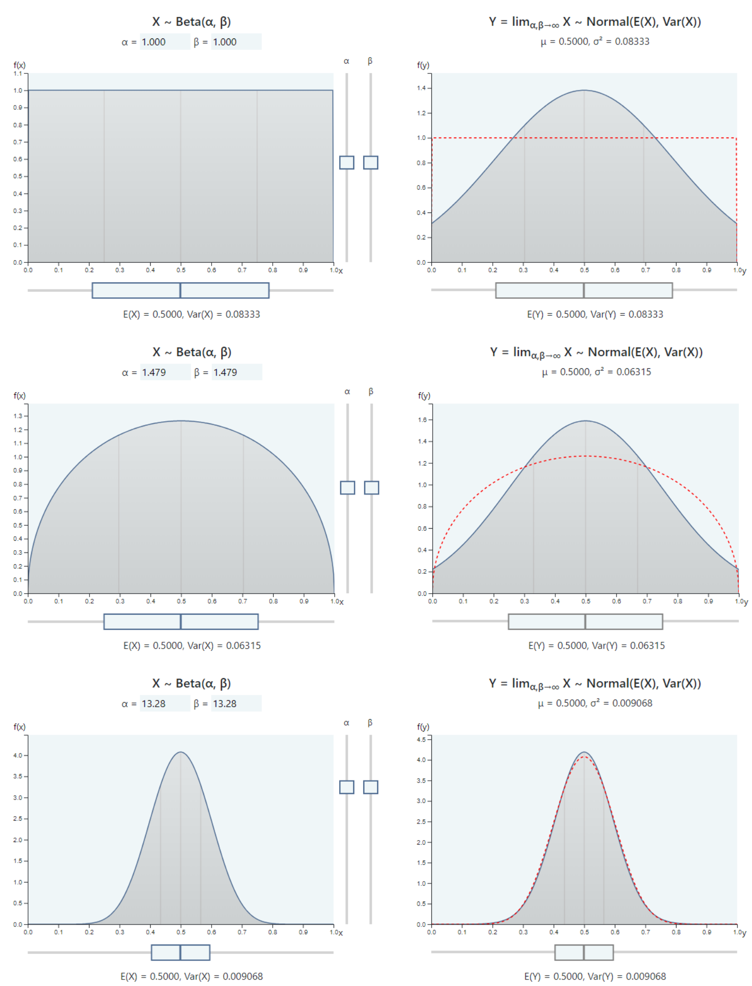


**Figure 3**: Illustration of the convergence of the beta($\alpha, \beta$) to the normal distribution as $\alpha$ and $\beta$ increase while holding the mean constant.

Accomplishing this involves changing the variance of the beta and solving for parameters $\alpha$ and $\beta$ through an inverse transformation while keeping the mean constant. Since the mean of the beta is given by the formula $\alpha/(\alpha + \beta)$, the mean can be held constant by keeping the ratio $\alpha/\beta$ constant. Solving for $\alpha$ and $\beta$ is then performed by substituting into the formula for the variance, while ensuring that neither $\alpha$ nor $\beta$ exceed the range of supported parameter values.

Means and variances can be changed directly by interacting with the graph of the distribution. The mean is changed by dragging the graph left or right, and the variance by zooming in and out. On a keyboard this is accomplished using the left/right or up/down arrow keys respectively. To support increased usability, means and variances can also be changed by interacting with the mean-variance box below the graph.

Such interactive functionality is available in Probability Playground for the beta, beta-binomial, gamma, Gumbel, lognormal, Pareto, and Weibull distributions, all of which have mean and/or variance formulas that depend on more than one parameter. The same editing actions also works for all distributions where the mean and variance can each be mapped back to a single parameter, such as location-scale families. Hence, Probability Playground provides the same interactive editing capabilities for a wide range of distributions within a single framework.

### 4.2 Automatic Axis Scaling

A significant amount of analysis has been performed with regard to automatic axis scaling. Given that the support of many distributions covers an infinite range, displaying the full graph of the distribution is not possible in these cases. Automatic axis scaling involves choosing which range of values to display without input from the user. In line with the design principles of maximizing interactivity and usability, we would like to provide a "reasonable" display of any probability mass or density function over the range of supported parameters. Although there are several websites which display these functions for a variety of distributions, we believe that our treatment of automatic axis scaling is the most sophisticated to date, allowing the smooth automatic updating of graphs while users are interacting with the website.

The intervals to display are defined as functions of the distribution parameters. For the x-axis this interval is given by $(X_{\min}, X_{\max})$ and for the y-axis by the interval $(Y_{\min}, Y_{\max})$. To provide visually intuitive and appealing graphing, the functions defining the displayed x- and y-intervals need to satisfy the following four constraints.

1. Nearly all of the probability mass should always be visible on the graph. The graph should not appear to be truncated on the left or right, and the value of the pdf/pmf at the limits of the x-axis relative to the maximum pdf/pmf value should be close to zero.

2. The graph of each distribution should be clearly visible across the whole range of parameters. For example, the distribution should never appear to be so horizontally compressed that the distribution shape is difficult to see.

3. X-axis limits need to be consistent across related distributions. For example, since the geometric($p$) is a special case of the negative binomial($r, p$) (setting $r$ = 1), then the interval used to display the geometric needs to be identical to that for the negative binomial when $r$ equals 1.

4. The relative position of the distribution mean $\mu$ on the x-axis should be an increasing function of $\mu$.

This last constraint relates to the previous section 4.1 on mean-variance editing. Probability Playground provides the ability to "drag" the probability mass up and down the x-axis by interacting directly with the graph. This is interpreted as changing the location of the mean $\mu$ of the distribution relative to the displayed interval of x values. For this to function correctly, increasing $\mu$ should always result in the relative location of the mean on the x-axis also increasing (and vice-versa).

This relative location is defined as the position of $\mu$ relative to the displayed interval $(X_{\min}, X_{\max})$:

$$\mu_{rel} = \frac{\mu - X_{\min}}{X_{\max} - X_{\min}} \tag{1}$$

Stated formally, for $\mu_{rel}$ to be a strictly increasing function of $\mu$, we must have:

$$\frac{d\mu_{rel}}{d\mu} > 0 \tag{2}$$

The definition of $X_{\min}$ and $X_{\max}$ that satisfies these constraints is dependent on the specifics of the distribution. However, as we show in the following sections, distributions can be categorized into four main classes such that the form of the functions defining $X_{\min}$ and $X_{\max}$ is similar for distributions in the same class.

Automatic scaling of the y-axis presents a much simpler problem. To maximize the area of the graph covered by a probability distribution, the maximum value $Y_{\max}$ of the y-axis is usually set to be 1.1 times the greatest value of the pdf/pmf over the range of values displayed. $Y_{\min}$ is naturally set to zero in all circumstances, since the value of a pdf/pmf is bounded below by zero. When displaying the cdf (which has a maximum value of 1 for all distributions), $Y_{\max}$ is usually set directly to the value 1.1.

#### *4.2.1 Bounded Distributions*

For bounded distributions i.e. those for which the support lies in a finite interval, a natural range of x values to display is the support itself. Such distributions include the Bernoulli, beta, beta-binomial, binomial, discrete uniform, hypergeometric, and uniform distributions.

For the cases where $X_{\min}$ and $X_{\max}$ are constant, we can see from equation (1) that:

$$\frac{d\mu_{rel}}{d\mu} = \frac{d}{d\mu}\frac{\mu - X_{\min}}{X_{\max} - X_{\min}} = \frac{1}{X_{\max} - X_{\min}} > 0 \tag{3}$$

where the inequality follows since $X_{\max} > X_{\min}$ in all circumstances, and hence $X_{\max} - X_{\min} > 0$. The constraint that $\mu_{rel}$ is an increasing function of $\mu$ is therefore satisfied, so dragging the mean up or down the x-axis works as required.

Mean editing is dealt with in an intuitive fashion for cases such as the binomial$(n, p)$ distribution that involve counting the numbers of successes. Although the mean $np$ depends on both $n$ and $p$, the choice is made to only change $p$ when the mean is edited, and to keep $n$ constant.

A slight adjustment is made for the hypergeometric$(N, M, K)$ distribution, where $K$ elements are selected from a population of size $N$ containing $M$ elements have some given property. Although $K$ elements are selected, the support of the random variable $X$ representing the number of items selected having the given property may not fully span the set $\{0, 1, \ldots, K\}$. However, displaying this full interval was found to be more visually intuitive and enables the restriction of the support to a subinterval of $(0, K)$ to be clearly seen through the horizontal line under the graph as shown in figure 4 below.

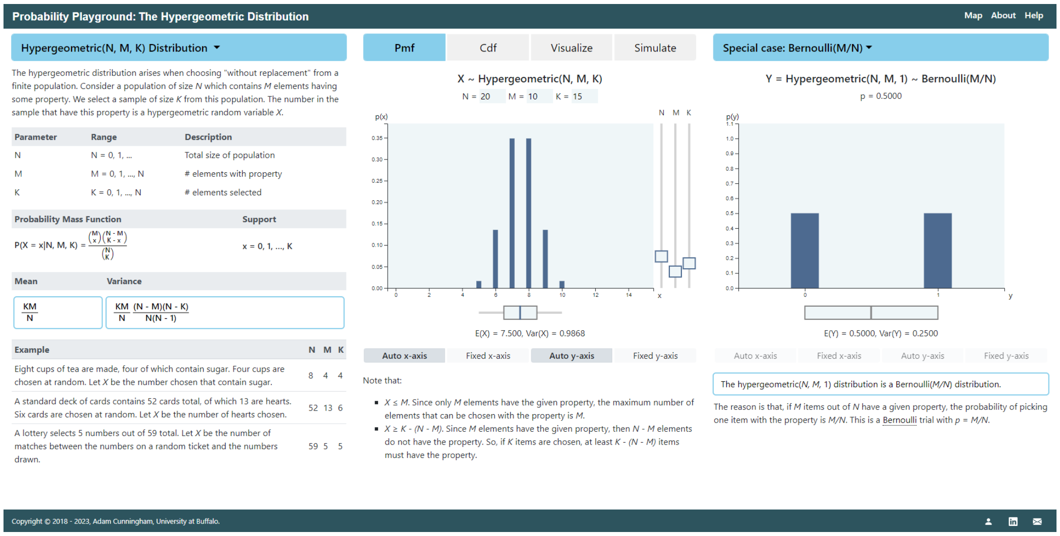


**Figure 4**: The distribution page for the hypergeometric distribution. The displayed x interval is set to $(0, K)$. The horizontal line below the graph in the center column shows that the support of the distribution is restricted in this case to a subinterval of $(0, K)$.

### *4.2.2 Partially Bounded Distributions*

A second class of distributions contains those that are bounded below but not above. Such distributions include the chi-squared, exponential, F, Fréchet, gamma, geometric, lognormal, negative binomial, Pareto, Poisson, and Weibull.

The support of the majority of these distributions is bounded below by zero. The exceptions are the Pareto (bounded below by $x_m$), the geometric (bounded below by 1), and the negative binomial (bounded below by $r$). However, for simplicity and to clearly show the lower limit of the support, we set $X_{\text{min}}$ to be zero in all cases. This simplifies equation (1) to:

$$\mu_{rel} = \frac{\mu}{X_{\text{max}}} \tag{4}$$

For the upper limit $X_{\text{max}}$, we require a function that satisfies inequality (2) i.e. the relative location $\mu_{rel}$ should be a strictly increasing function of the population mean $\mu$. A simple solution is to consider $\mu_{rel}$ as a power of $\mu$:

$$\mu_{rel} = A\mu^B \tag{5}$$

where $A$ and $B$ are both constants with $A > 0$. To satisfy inequality (2), we must therefore have:

$$\frac{d\mu_{rel}}{d\mu} = \frac{d}{d\mu}A\mu^B = AB\mu^{B-1} > 0 \tag{6}$$

Note that, since these distributions are bounded below by zero, we must have $\mu > 0$, and hence $B > 0$ for inequality (2) to be satisfied. Substituting equation (6) into equation (4) and rearranging yields:

$$X_{\text{max}} = \frac{\mu}{\mu_{rel}} = \frac{\mu}{A\mu^B} = \frac{\mu^{1-B}}{A} \tag{7}$$

The power $B$ can take a range of possible values. At one extreme, with $B$ = 1, then $X_{\text{max}}$ has the constant value $1/A$ independent of $\mu$. While this both satisfies inequality

(2) and is simple to understand, it was found to be unsatisfactory in practice for many distributions since it overly restricts the range of parameters that can be graphed while still satisfying constraints (1) and (2). At the other extreme, with $B = 0$, the relative location $\mu_{rel}$ is directly proportional to $\mu$, with the result that the distribution mean is always displayed at the same relative location on the x-axis. Direct editing of the mean by dragging on the graph is then not possible, so this option is ruled invalid.

A better option is to choose the power $B$ such that $0 < B < 1$. Experimentation showed that the value $B$ = 0.3 provided visually pleasing results, displaying distribution pdf/pmfs well across a wide range of parameters. The constant $A$ is then chosen to be as small as possible while still satisfying constraint (1), such that nearly all of the probability mass is still visible even when $\mu$ has its largest possible value within the supported range of parameters.

This functional form for $X_{\max}$ was found to work well for nearly all distributions for which the mean $\mu$ is always defined, and is used for the chi-squared, exponential, gamma, geometric, negative binomial, and Weibull distributions.

Automatic X-axis scaling proved more difficult to generalize for heavy-tailed distributions such as the Fréchet, lognormal, and Pareto. In each case the formula for $X_{\max}$ was customized to the specifics of the distribution, while still satisfying the constraints described earlier in this section. For example, the formula for $X_{\max}$ used for the lognormal is given by $8e^{0.9\mu}$, where $\mu$ is the mean of the related normal distribution. Although the formulas developed for $X_{\max}$ for such distributions work adequately, there is still scope for some further optimization and generalization.

#### *4.2.3 Unbounded Distributions*

A third class of distributions contains those that are bounded neither above nor below. These include the Cauchy, Gumbel, Laplace, logistic, normal, and T distributions. The ability to independently edit the mean and variance of these distributions places additional constraints on $X_{\min}$ and $X_{\max}$. Specifically:

1. Changing the mean $\mu$ while holding the variance constant should not change the length $X_{range}$ of the displayed x-axis interval.
2. The ratio of the standard deviation $\sigma$ to $X_{range}$ should be a strictly increasing function of $\sigma$.

These constraints are necessary to ensure that mean and variance editing work in an intuitive way. The first constraint ensures that dragging the probability mass on the graph only changes the relative location of the mean without changing the shape of the graph. The second constraint ensures that increasing the variance results in a perceived increase in the variance on the graph (and similarly when the variance is decreased).

A set of functions which satisfy these constraints is as follows.

$$\begin{aligned} X_{range} &= c\sqrt{\sigma} && (8) \\ X_{\min} &= \mu(1 - k * X_{range}) - X_{range}/2 && (9) \\ X_{\max} &= X_{\min} + X_{range} && (10) \end{aligned}$$

where $c$ and $k$ are constants specific to the distribution. It can be seen that:

- $X_{range}$ is independent of $\mu$. Hence, constraint (1) above is satisfied.

- The ratio of the standard deviation $\sigma$ to $X_{range}$ is given by $\sigma/(c\sqrt{\sigma}) = \sqrt{\sigma}/c$, which is a strictly increasing function of $\sigma$. Hence, constraint (2) above is also satisfied.
- When $\mu = 0$, the interval $(X_{\min}, X_{\max})$ is given by $(-c\sqrt{\sigma}/2, c\sqrt{\sigma}/2)$. This is symmetric about the origin as would be intuitively expected.

One further constraint that needs to be satisfied is that the relative location $\mu_{rel}$ of the mean $\mu$ on the x-axis is a strictly increasing function of $\mu$. From equation (1) we have:

$$\mu_{rel} = \frac{\mu - X_{\min}}{X_{\max} - X_{\min}} = \frac{\mu - [\mu(1 - k * X_{range}) - X_{range}/2]}{X_{range}} = k\mu + 1/2 \qquad (11)$$

Hence, $\mu_{rel}$ is a strictly increasing function of $\mu$ as required.

The constants $c$ and $k$ are chosen to satisfy constraints (1) and (2) given at the start of section 4.2. Constant $k$ controls the relative position of the mean on the x-axis, and constant $c$ controls the scaling of the standard deviation relative to the length $X_{range}$. These constants are both maximized while still ensuring that the graph of the pdf does not appear truncated either to the left or to the right over the entire range of supported parameter values.

#### *4.2.4 Standard Distributions*

Finally, Probability Playground includes five standard distributions - the standard Cauchy, logistic, normal, Pareto, and uniform. Since the parameters of these distributions are fixed, they are not editable, and the graphs of their pdfs are always displayed within constant x-axis limits.

To ensure consistency across the website, these constant x-axis limits are chosen to be equal to the limits that would be calculated by the x-axis automatic scaling functions for the standardized version of the distribution. For example, the graph of a standard normal distribution will be displayed with the same x-axis limits as a normal distribution with parameters $\mu = 0$ and $\sigma^2 = 1$.

#### *4.2.5 Manual Fixing of Axis Limits*

Probability Playground also supports the ability to manually fix the axis limits at constant values. This can be done independently for the x- and y-axes for both the main and related graphs, and is accomplished by selecting the "Fix x-axis" and "Fix y-axis" buttons below the graph. The axes limits are then fixed at their current values until the "Auto x-axis" and "Auto y-axis" buttons are selected, at which point automatic axis scaling resumes.

### 4.3 Proofs

Courses in probability theory and statistics frequently include proving some of the main results relating to each probability distribution. Mathematical rigor is provided on the website through the inclusion of over 150 proofs of distribution means, variances, and the relations between distributions. In line with the principle of ease of navigation, such proofs are accessed simply by selecting the formula of interest. A modal dialog box showing the proof is then displayed as an overlay, with the result to be proved still visible on the web page beneath as shown in figure 5 below.

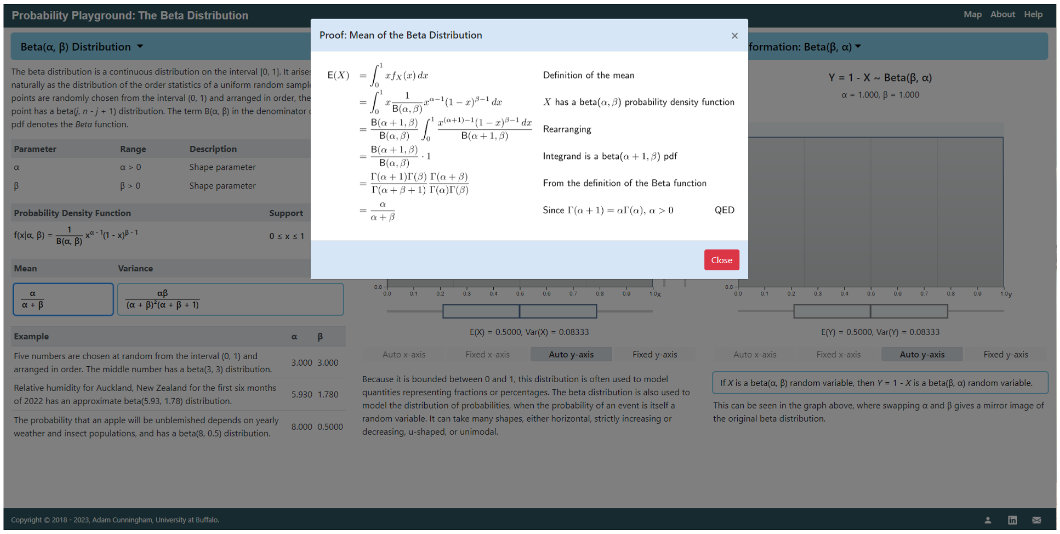

**Figure 5**: Proof of the mean of the beta distribution. Over 150 proofs are provided for distribution means, variances, and the relationships between distributions. The proof of a formula is accessed directly by selecting the formula.

### 4.4 Accessibility

Keyboard control is a key principle in web accessibility. Users with motor disabilities or hand tremors may not have the fine motor control needed to use a mouse, and are therefore reliant on a keyboard. Users without disabilities may also find using a keyboard easier or more efficient in some circumstances.

To accommodate such users, all interactions with Probability Playground can be performed using either a mouse or a keyboard. Parameters can be changed using up and down arrow keys, or can be entered as text via a keyboard. Distribution means and variances can be modified using the left/right or up/down arrow keys respectively. Navigation between controls can be performed using the TAB key, with controls ordered left to right then top to bottom within columns for ease of navigation. Probability Playground has therefore been designed from the beginning to be fully accessible to users in need of such support.

Accessibility concerns for users with visual impairments have also been addressed. The color scheme used has been tested against different types of color blindness to ensure sufficient visual contrast for all such users, and text descriptions are provided for all controls.

## 5. Related Distributions

Probability Playground is unique for its interactive graphing of a wide selection of related distributions. Whenever a distribution $Y$ is related to a distribution $X$, the graph of $Y$ is immediately updated to reflect any changes to the graph of $X$. Although some websites do display such graphs for specific relationships, Probability Playground is the first to include this for all such relationships as a standard part of its functionality.

Graphing such relationships is intended to aid students in developing their intuition. For example, many distributions converge to a normal distribution as certain parameters approach infinity (such as the beta, binomial, gamma, negative binomial, and Poisson). The beta-binomial and hypergeometric converge to the binomial, which itself converges to the Poisson under certain circumstances. By plotting both a distribution $X$ and its limit $Y$ on the same graph, students are able to see this convergence occurring as the relevant parameters change in addition to understanding the relevant algebraic manipulations.

Probability Playground classifies the relationships between a random variable $X$ and a related random variable $Y$ using the five following categories.

**Special Cases** These indicate that the distribution of $Y$ has been obtained by setting one or more of the parameters of $X$ to constant values. Examples include setting the parameter $r$ of the Negative Binomial$(r, p)$ distribution equal to 1, which yields a Geometric(p) distribution.

**Transformations** A transformation indicates that $Y = f(X)$, where $f$ is some function. Thus, for any set A,

$$P(Y \in A) = P(f(X) \in A)$$

**Limiting Cases** A limiting distribution can be expressed as

$$Y = \lim_{\mathbf{p}_X \to \mathbf{P}} X$$

where $\mathbf{p}_X = \{p_1, p_2, \ldots\}$ are parameters of $X$ and the vector $\mathbf{P}$ defines the limiting values they each approach. Examples include the convergence of the binomial$(n, p)$ distribution to the Poisson$(\lambda)$ distribution as $n \to \infty$ and $np$ stays small, where the parameter $\lambda$ of the Poisson is defined by $np$.

**Statistics** A statistic is any quantity computed from values in a sample. This can be expressed as

$$Y = f(\mathbf{X})$$

where $\mathbf{X} = X_1, X_2, \ldots, X_n$ is an independent and identically distributed sample of size $n$ drawn from the distribution of $X$. Distributions based on statistics include those obtained as the minimum, maximum, or sum of an iid sample, as well as the chi-squared, F, and T distributions derived from a normal sample.

**Statistic Limit** A statistic limit is a limiting distribution of a statistic. Examples in Probability Playground usually relate to extreme value distributions, such as when a suitable standardized maximum of an iid sample drawn from an exponential distribution converges to a Gumbel distribution as the sample size approaches infinity.

Each of these different types of relationship is considered as belonging to a single general form as illustrated in figure 6 below.

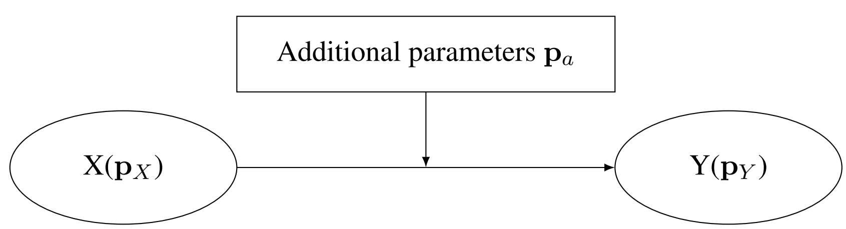


**Figure 6**: General form of the relationship between distributions. The vector $\mathbf{p}_X$ contains the parameters of $X$, $\mathbf{p}_Y$ is a vector of the parameters of $Y$, and $\mathbf{p}_a$ is a vector of any additional parameters needed to define the distribution of $Y$.

The relationship between distributions $X$ and $Y$ is hence defined by the relationship between the parameters $\mathbf{p}_X$, $\mathbf{p}_a$, and $\mathbf{p}_Y$ through the multivariable function $g$:

$$\mathbf{p}_Y = g(\mathbf{p}_X, \mathbf{p}_a) \tag{12}$$

By defining the relationship between distributions in such a general manner, Probability Playground is able to graph any kind of relationship within a single unified framework. Note that the parameter vector $\mathbf{p}_Y$ does not necessarily depend on all the parameters $\mathbf{p}_X$. For example, when $Y$ is a special case of $X$, one or more of the parameters of $X$ are either set to a constant value or are dropped entirely.

## 6. Animations of Generating Processes

To help students visualize the causal mechanism which generates each distribution, Probability Playground includes an animated simulation for each process. This is illustrated in figure 7 below for the binomial distribution, where a set of Bernoulli trials are run and the random variable $X$ is the total number of successes.

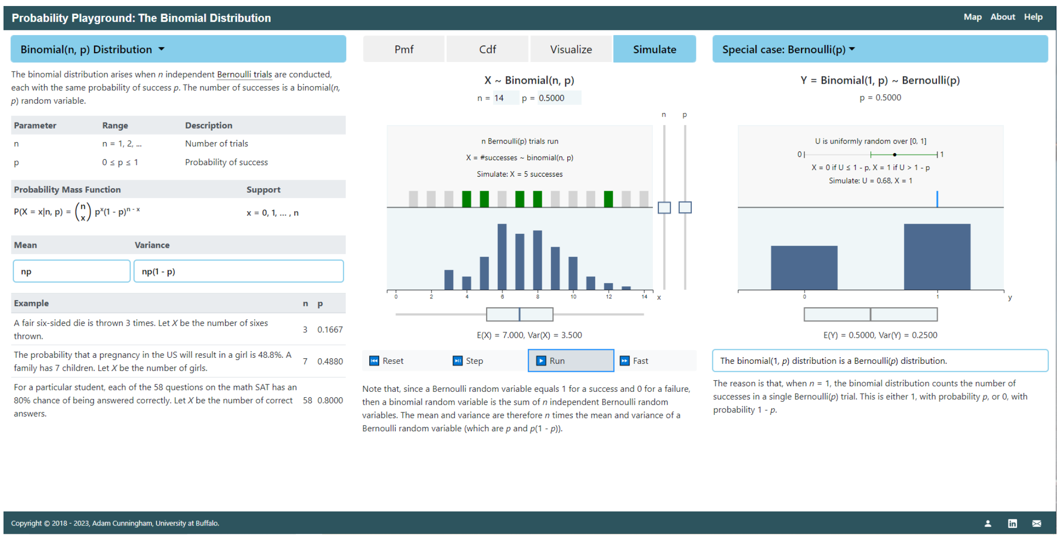


**Figure 7**: Generating process for the binomial distribution shown in the center column. Successful Bernoulli trials are shown in green and failures in gray. The simulation is controlled through a set of video-style controls. A histogram below the simulation accumulates the results.

Video-style controls allow the simulation to be stopped, run, reset, or stepped through one simulation update at a time. A histogram below the simulation accumulates the results, which shows how the shape of the distribution is generated. In keeping with the principle of full interactivity, the parameters of the distribution can be changed or examples loaded while the simulation is running. This resets the histogram and continues the simulation with the new set of parameters.

The visual metaphors used for each distribution are chosen to be consistent across the website. For example, the same visual representation of a Bernoulli trial success or failure shown in figure 7 above is used for all distributions based on the Bernoulli such as the binomial, geometric, and negative binomial distributions. Similarly, distributions based on a Poisson process such as the Poisson, exponential, and gamma all use the same visual representation of events occurring randomly and independently in time.

## 7. Conclusion

As an educational tool to develop students' visual intuition regarding probability distributions and their relationships, Probability Playground is unique for its scope, functionality, and degree of interactivity. It provides a single unified framework for exploring a wide range of univariate distributions that are commonly encountered by students of probability theory and statistics. Interactive editing of distribution parameters, means, and variances using either a mouse, touchscreen, or keyboard allows students to easily visualize the various forms that distributions can take, while animated simulations of the causal mechanisms generating each distribution help to develop their understanding of the real world situations in which they arise. Relationships between distributions are illustrated both through a directed graph on a navigation page, and by the interactive graphing of related distributions, allowing convergence to limiting distributions to be seen. Over 150 easily accessible proofs for the formulas of means, variances, and relationships provide additional mathematical rigor. Probability Playground therefore offers an "off-the-shelf" solution for educators looking to incorporate interactive and exploratory learning into their classrooms.

## REFERENCES


Beaird, J., and George, J. (2014), *The Principles of Beautiful Web Design* (3rd ed.), Melbourne, Australia: Sitepoint.

Casella, G., and Berger, R. L. (2001), *Statistical Inference* (2nd ed.), Boston, Massachusetts: Cengage Learning.

Forbes, C., Evans, M., Hastings, N., and Peacock, B. (2011), *Statistical Distributions* (4th ed.), Hoboken, New Jersey: Wiley.

Graham, L. (1998), *Principles of Interactive Design*, Clifton Park, New York: Delmar Cengage Learning.

Lett, J. (2019), *Bootstrap Reference Guide*, Sterling Heights, Michigan: Bootstrap Creative.

Murray, S. (2017), *Interactive Data Visualization for the Web* (2nd ed.), Sebastapol, California: O'Reilly Media, Inc.

Norman, D. (2013), *The Design of Everyday Things*, New York, New York: Basic Books.

Steane, J. (2022), *The Principles and Processes of Interactive Design*, New York, New York: Bloomsbury Visual Arts.