Assume $X_n \sim \text{Binomial}(K_n, p)$, for $n = 1, 2, \ldots, N$. Notice that K is also subscripted by n, so as to suggest that each new random variable has a possibly different number of trials. For instance, if X_n records the number of bakers that burn their baked good in an episode of the Great British Baking Show, then K_n recognizes that each episode has a different number of contestants on it (since one contestant was kicked off on the previous episode).

1. Find the maximum likelihood estimator for p, call it $\hat{p}$.
2. Let's double check our calculus above using simulation using $p = 0.5$.
3. For $n = 1, \ldots, 1001$, randomly sample $K_n \sim \text{Uniform}(10, 20)$. Call this vector `K`.
4. Use `K` to randomly sample $X_n \sim \text{Binomial}(K_n, p)$, for some p that you pick. Note that `rbinom` allows the argument `size` (which is not the sample size) to be a vector; thank you vectorization. Call this vector `x`.
5. Simply using `x` and your calculus solution above, estimate p in `R`.
6. Since the expected value is not easy to calculate, let's use the simulation above to approximate it. If you take the mean of the `x`, you'll get an estimate of $\mathbb{E}(X)$. Try this.
7. What might you guess is the population expected value? Does this make sense given $K_n \sim \text{Uniform}(10, 20)$?