Estimating the parameter from infinitely many IID Bernoulli distributions

written by DRMacIver on Functor Network
original link: https://functor.network/user/3130/entry/1260

These are some notes working out a particular unbiased estimator that I needed for making some improvements to the AWRS algorithm in genlm-control. You can see where this gets used in practice in the AWRS implementation.

Suppose you've got an infinite sequence of IID Bernoulli(p) distributions, X_i. You want an unbiased estimator for p. How can you efficiently read the values of the X_i to get one?

We're particularly interested in the case where sampling the X_i is "expensive" in some sense, so you want to use as few values as possible. I won't make this precise here, it's just the motivating intuition.

The first, and most obvious, strategy, is that for some n you unconditionally return $\frac{1}{n}\sum_{i=1}^{n} X_i$.

This approach is fine, but you always "spend" n samples. One of the cases that we're interested in is that when p is close to 1 we would like this process to be cheap, which means that spending n samples each time isn't ideal.

A natural idea is to sample until you see your first success. Unfortunately, this idea doesn't work, for the following reason:

Theorem: Let $X \sim \text{Geom}(p)$. The only unbiased estimator for p from a single sample of X is to return 1 if $X = 0$ and 0 if $X > 0$.

Proof:

Suppose we have some unbiased estimator f.

Then $p = \text{E}(f(X)) = \sum_{i \geq 0} f(i)p(1-p)^i$

So $1 = \sum_{i \geq 0} f(i)(1-p)^i$.

The right hand side is a power series in $1 - p$, so by the uniqueness of power series, we must have that $f(0) = 1$ and $f(i) = 0$ for $i > 0$ as desired. $\square$

However, you can get a good estimator for p by sampling until you see $s > 1$ successes and record the number of failures, n. This gives you a negative binomial $\text{NB}(s, p)$ distribution, and this has a standard minimum variance unbiased estimator for p of $\frac{s-1}{s+n-1}$.

This works just fine.

However, another consideration that we might care about is that we'd like our estimator to not be too expensive when p is small. This has expected value $\frac{s(1-p)}{p}$, so when p is small we will have to take many samples.

The solution is to "round towards zero" at a certain point. The reason we need unboundedly many draws is that we need to distinguish between very small values of p. If we're happy to return an estimator of 0 for some small non-zero values of p, we can bound the cost of sampling using the following theorem:

Fix integers $r, s > 0$. Sample from the X_i until we've seen either r failures or s successes, and let R be the number of failures we saw and and S the number of successes.

These variables have the joint law:

- $P(R = i, S = s) = \binom{i+s-1}{s-1}(1-p)^i p^s$
- $P(R = r, S = j) = \binom{r+j-1}{r-1}(1-p)^r p^j$

Consider the case $S = s$. This result arises precisely when the first $i + s$ draws contain exactly s successes, of which the last one is a success.

Each such sequence of assignments occurs with probability $p^s(1-p)^i$, so to calculate the probability of this we just need to count sequences.

There are $i + s$ such sequences, distinguished only by where the successes are. One success must be at the end, so there are $s - 1$ successes to distribute among $i + s - 1$ positions. This can be done in $\binom{i+s-1}{s-1}$ ways, giving the desired result.

The other case is proved identically. $\square$

Now, suppose $r, s > 1$.

Define Q as follows: If $R = r$ (i.e. we stopped because we hit the maximum number of failures), $Q = \frac{S}{R+S+1}$. Otherwise (i.e. we stopped because we hit the maximum number of successes), let $Q = \frac{S-1}{R+S+1}$.

Theorem: Q is an unbiased estimator for p.

Proof:

This will follow from the Rao-Blackwell theorem.

First, note that (R, S) is a sufficient statistic for p, as all dependence on p is via its value.

Now, consider the trivial unbiased estimator X_1. i.e. we estimate 1 if X_1 succeeds and 0 otherwise.

Let us consider the conditional expectation $E(X_1|R, S)$.

First, suppose $S = s$.

Then

$$
\begin{aligned}
E(X_1|R=i,S=s) &= P(X_1|R=i,S=s) \\
&= \frac{P(X_1 \wedge R=i \wedge S=s)}{P(R=i,S=s)} \\
&= \frac{P(X_1)P(R=i \wedge S=s|X_1)}{P(R=i,S=s)} \\
&= \frac{p\binom{i+s-2}{s-2}(1-p)^i p^{s-1}}{\binom{i+s-1}{s-1}(1-p)^i p^s} \\
&= \frac{\binom{i+s-12}{s-2}}{\binom{i+s-1}{s-1}} \\
&= \frac{(i+s-2)!(s-1)!i!}{(i+s-1)!(s-2)!i!} \\
&= \frac{s-1}{i+s-1}
\end{aligned}
$$

Where we used that $P(R=i \wedge S=s|X_1) = \binom{i+s-2}{s-2}(1-p)^i p^{s-1}$ because this is just exactly the same process but with one fewer success required to complete (because it was already provided by the X_1.

We could do the algebra for the other direction, but instead we exploit the symmetry of the process. By swapping success and failure, we get the same process but with r and s swapped and $p \to 1-p$. Thus $\frac{r-1}{r+j-1}$ is an unbiased estimator for $1-p$, and:

$$
\begin{aligned}
1 - \frac{r-1}{r+j-1} &= \frac{r+j-1-(r-1)}{r+j-1} \\
&= \frac{j}{r+j-1}
\end{aligned}
$$

as desired.

These values are precisely Q as we previously defined it, completing the proof. $\square$