Suppose you take a poll by asking a random set of n people a yes/no question.

Then n is the sample size, and the percentage p of the people who would answer Yes in the whole population is called the population proportion. The percentage of people who actually answered Yes in your sample is called the sample proportion, written p^ (the ^ should go directly above the p, but that's hard to do without special software).

We can make relative frequency histograms for p^, showing the percentages of samples for each value of p^. For example, suppose we ask "Do you support Heineman in the coming election?" According to recent Rasmussen polling, p = 66%. If n = 2, then p^ can be either 0% (neither of the two people we poll support Heineman), or p^ can be 50% (one supports Heineman and one doesn't) or p^ can be 100% (both people we poll support Heineman). Assuming p=66%, it turns out that we would expect that (1/3)*(1/3) = 1/9 = 11.1% of all samples of 2 people would be samples in which neither person supports Heineman, while 2*(1/3)*(2/3) = 4/9 = 44.4% of all samples of 2 people would be samples where only one of the two people support Heineman, and (2/3)*(2/3) = 4/9 = 44.4% of all samples of 2 people would be samples where both of the two people support Heineman. This gives the following histogram, where the horizontal axis gives the possible values of p^ and the vertical axis gives the probability of getting a particular value of p^ when you take poll n = 2 randomly chosen people. I.e., there's an 11% chance that in a sample of 2 randomly chosen people, neither one will support Heineman.

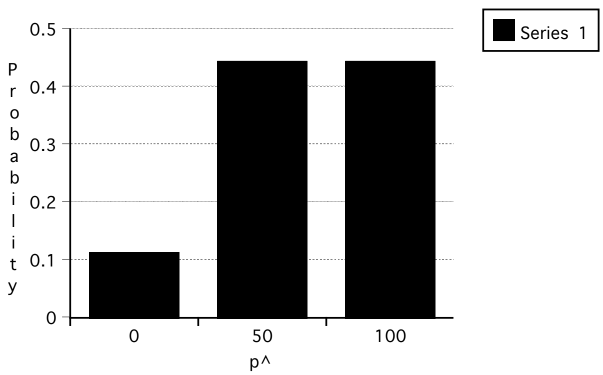

When you take a random sample of size n = 2, the probability that p^ is within 20 percentage points of the actual percentage of p = 66 is just the sum of the bars for p^'s ranging from p^ = 66 - 20 to p^ = 66 + 20. In this case, there's only one bar in that range, of height 44.4%. So this means that the chance that the measured level of support p^ for Heineman is within 20 percentage points of the actual value p is only 44.4% if you take a poll with a sample size of n = 2.

To get a higher probability of a more accurate result we need to make n bigger. Here's what happens with n = 100:

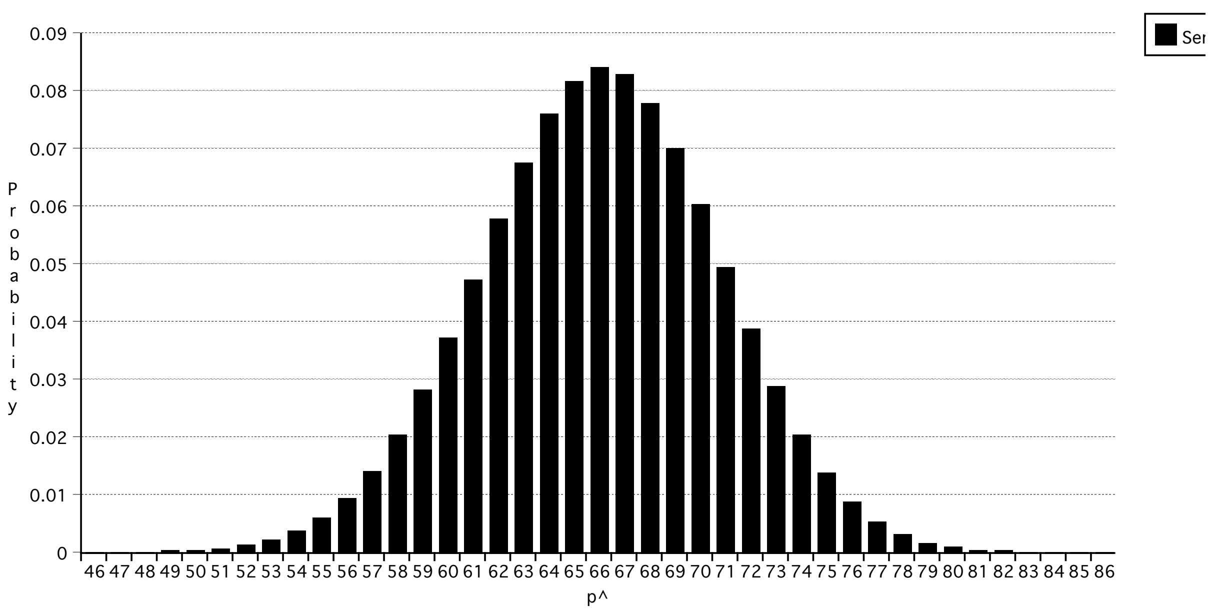

So for example, if p^ = 58%, the height of the bar is 0.02 = 2%, so in 2% of all samples of 100 people, 58 of the 100 would be Heineman supporters.

Here's the actual data used to make the bar chart:

x	Probability that p^ = x	Probability that p^ ≤ x
46	1.84077e-05	3.17051e-05
47	4.10544e-05	7.27595e-05
48	8.79952e-05	0.000160755
49	0.000181272	0.000342027
50	0.000358919	0.000700946
51	0.000683064	0.00138401
52	0.00124945	0.00263346
53	0.00219659	0.00483005
54	0.00371124	0.00854129
55	0.0060253	0.0145666
56	0.00939871	0.0239653
57	0.0140835	0.0380488
58	0.0202683	0.0583171
59	0.0280078	0.0863249
60	0.0371515	0.123476
61	0.0472903	0.170767
62	0.0577444	0.228511
63	0.0676111	0.296122
64	0.075876	0.371998
65	0.0815753	0.453573
66	0.0839746	0.537548
67	0.0827212	0.620269
68	0.0779268	0.698196
69	0.0701541	0.76835
70	0.0603089	0.828659
71	0.0494663	0.878125
72	0.0386759	0.916801
73	0.0287965	0.945598
74	0.0203956	0.965993
75	0.0137251	0.979718
76	0.00876407	0.988482
77	0.00530263	0.993785
78	0.00303522	0.99682
79	0.00164078	0.998461
80	0.000836074	0.999297
81	0.000400732	0.999698
82	0.000180243	0.999878
83	7.58785e-05	0.999954
84	2.98094e-05	0.999984
85	1.08923e-05	0.999995
86	3.68788e-06	0.999998

Polling facts:

If the sample size n is not too small and if the population size N is big compared to n, the bar graph for the percentages of samples for each value of p^ will be approximately normal. This means 95% of samples of size n will be in the range $p \pm 2\sigma$. To make use of this we need to know how to compute σ in terms of p and n.

But first, how big is big enough for n? The standard test is that n should be bigger than both 9(1-p)/p and 9p/(1-p).

The standard deviation σ is given by the formula: $\sigma = \sqrt{(p(1-p)/n)}$.

Typically we use p^ to estimate σ, so $s = \sqrt{(p\hat{}(1-p\hat{})/n)}$ is an estimate for σ; we call s the "standard error".

Note: In these formulas, p and p^ must be written as decimals, not as percentages (i.e., as 0.45 and not as 45%).

The main fact that polling is based on is: there's a 95% chance in a random sample of size n that p^ will be in the range $p \pm 2\sigma$. I.e., that p and p ^ are within 2σ of each other, and thus, using s to estimate σ, we expect there to be a 95% chance that p is in the range p^ ± 2s, no matter how big N is (as long as N is not too small). We call ±2s the "margin of error".

Thus, as long as n and N are not too small, when you take a random sample, typically even just n = 1000 people where N is huge (like 100,000,000), there's a 95% chance that the actual value p will be within the margin of error of the measured value p^. All you know for sure are the sentiments of the people in the sample, but you're pretty sure (95% sure) that the sentiments of the whole population (no matter how big N is, even if N is way bigger than n) are pretty close (i.e., within the margin of error)