Bayesian 95% confidence interval for $\pi_1 - \pi_2$ using Table 3.14

	Y=1	Y=2
X=1	21	2
X=2	15	3

and uniform "prior" probability density function $p(\pi_1, \pi_2) = 1$ for $0 < \pi_1 < 1$ & $0 < \pi_2 < 1$ else 0:

$$p(a,b) = \lim \frac{P((\pi_1, \pi_2) \in A)}{\text{area}(A)}$$

where the limit is taken over small rectangles A converging to a point $\{(a,b)\}$. Treating the rows as independent binomials,

$$\frac{1}{c_0} P(\text{data} \mid \pi_1, \pi_2) = \pi_1^{21}(1-\pi_1)^2 \, \pi_2^{15}(1-\pi_2)^3 .$$

$$\frac{1}{c_0} P(\text{data}) = \int_0^1 \left(\int_0^1 P(\text{data} \mid \pi_1, \pi_2)\, p(\pi_1, \pi_2)\, d\pi_1 \right) d\pi_2 / c_0$$

c_0 is a constant we don't need.

$$= 1/94140288 \approx 1.062 \times 10^{-8}$$

For events A, B, $P(B|A) = P(A|B)\,P(B)/P(A)$

says Bayes' Theorem. For us, $A =$ data.

If B is "$(\pi_1, \pi_2) \in C$", then in the limit where C converges to a point $\{(a,b)\}$,

$$\frac{P(B|A)}{\text{area}(C)} \longrightarrow p(a,b \mid \text{data})$$ and this equals

$$\lim \frac{P(A|B)P(B)}{P(A)\,\text{area}(C)},$$ which equals $$\frac{P(\text{data} \mid \pi_1, \pi_2)}{P(\text{data})}$$

(because of Bayes' Thm and because $p(a,b) = 1$).

$p(a,b \mid \text{data})$ is the "posterior" probability density of $(\pi_1 = a, \pi_2 = b)$. To obtain posterior probabilities, integrate:

$$P((\pi_1, \pi_2) \in C \mid \text{data}) = \iint_{(\pi_1, \pi_2) \in C} p(\pi_1, \pi_2 \mid \text{data})\, d\pi_1\, d\pi_2$$

For example, $P(\pi_1 < \pi_2 \mid \text{data}) = \int_0^1 \left(\int_0^{\pi_2} p(\pi_1, \pi_2 \mid \text{data})\, d\pi_1 \right) d\pi_2$

≈ 0.226; $P(\pi_2 < \pi_1 \mid \text{data}) = \int_0^1 \left(\int_{\pi_2}^1 p(\pi_1, \pi_2 \mid \text{data})\, d\pi_1 \right) d\pi_2$

≈ 0.774. ($P(\pi_1 = \pi_2 \mid \text{data})$, $P(\pi_1 = \pi_2) = 0$ because the line $\pi_1 = \pi_2$ has area zero.)

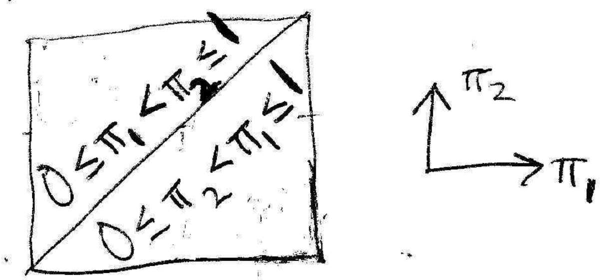

For our confidence interval, we will use the equal tails method, finding L and U such that $P(\pi_1 - \pi_2 < L \mid \text{data}) = P(\pi_1 - \pi_2 > U \mid \text{data}) = 0.025$.

Even before computing L & U, we can see that our CI will include $\pi_1 - \pi_2 = 0$ because

$0.226, 0.774 > 0.025$.

For $L \le 0$, $P(\pi_1 - \pi_2 < L \mid \text{data}) = \int_{-L}^{1} \left(\int_0^{\pi_2 + L} f(\pi_1, \pi_2)\, d\pi_1 \right) d\pi_2$

where $f(\pi_1, \pi_2) = p(\pi_1, \pi_2 \mid \text{data})$. For $U \ge 0$,

$P(\pi_1 - \pi_2 > U \mid \text{data}) = \int_0^{1-U} \left(\int_{\pi_2 + U}^{1} f(\pi_1, \pi_2)\, d\pi_1 \right) d\pi_2$

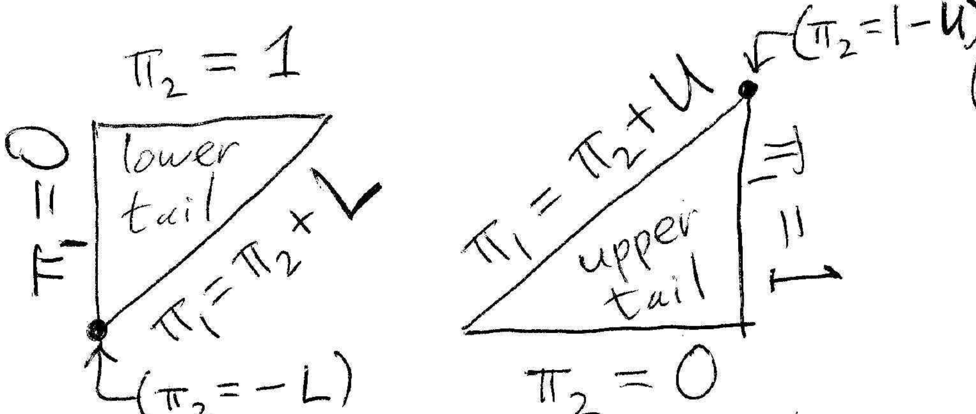

Using a computer algebra system on my calculator or computer ...

... I can get exact polynomial formulas for the above two tail probabilities. However, that doesn't tell me much about where they equal 0.025. We can estimate this using a table of values of for $P(\ldots L \ldots)$ for $-1 \le L \le 0$ and a table of $P(\ldots U \ldots)$ for $0 \le U \le 1$, or we can get fancy and use Newton's method or bisection ...

Just using two crude tables of values (up to 100 each), we can find:

L	tail prob.	U	tail prob.
⋮	⋮	⋮	⋮
-0.13	0.0231	0.30	0.0261
-0.12	0.0285	0.31	0.0216
⋮	⋮	⋮	⋮

← $P = 0.025$ somewhere in between.

So, a slightly conservative 95% CI is

$\pi_1 - \pi_2 \in [-0.13, 0.31]$.

I've attached a very inefficient but very simple Python script for accomplishing the above.
It uses the sympy and numpy packages.
(The alternative to symbolic integration is Monte Carlo...)