Theorem The maximum of n mutually independent Bernoulli random variables is a Bernoulli random variable.

Proof The Bernoulli distribution has probability mass function

$$f(x) = P(X = x) = \begin{cases} 1 - p & x = 0 \\ p & x = 1 \\ 0 & \text{otherwise.} \end{cases}$$

Let $X_i \sim \text{Bernoulli}(p_i)$ for $i = 1, 2, \ldots, n$. Then,

$$Y = \max\{X_1, X_2, \ldots, X_n\} \sim \text{Bernoulli}\left(1 - \prod_{i=1}^{n}(1 - p_i)\right)$$

because the support of Y is $\{0, 1\}$ and $Y = 0$ occurs only when $X_1 = X_2 = \cdots = X_n = 0$. Thus because $X_1, X_2, \ldots, X_n$ are assumed to be mutually independent, $P(Y = 0) = (1 - p_1)(1 - p_2) \ldots (1 - p_n)$ and $P(Y = 1) = 1 - P(Y = 0)$.

APPL verification: The APPL statements

```
MaximumIID(BernoulliRV(p),n);
simplify(op(%[1])[2]);
```

yield expressions that are consistent with the result for the special case where $X_1, X_2, \ldots, X_n$ are assumed to be mutually independent and identically distributed.