Some comments and tricks

- Always define the support of your distribution functions. In ex. 4.31 b) the answer is
$$f(x,y) = \binom{n}{y} x^y (1-x)^{n-y}, \quad x \in [0,1], y = 0,1,\ldots,n,$$
and it is not enough to just write
$$f(x,y) = \binom{n}{y} x^y (1-x)^{n-y},$$
because that is wrong for $x \notin [0,1]$ or when y is not an integer between 0 and n.

  In 4.30 b), many solved the problem by showing that $f(u,v) = f(u)f(v)$, where $U = Y/X$ and $V = X$. However, then you need to include an indicator function to show that the support of X is only $[0,1]$. Thus, we don't get $f(v) = 1$, as most people wrote, but $f(v) = I_{[0,1]}(v)$.

- Get more comfortable with calculations using expectations, variances and covariances. As an example, in 4.35 we get
$$VarX = E[Var(X|Y)] + Var(E[X|Y]) = nE[P(1-P)] + n^2 Var(P).$$
Here we can just use that
$$E[P(1-P)] = E[P - P^2] = E[P] - E[P^2],$$
and that
$$E[P^2] = Var(P) + E[P]^2,$$
and then we have solved the problem. We don't need to actually start computing all the expectations and manipulate all those α's and β's.

  Knowing that the covariance is a linear operator also helps us a lot in 4.58 b). Then we can just say that
$$Cov(X, Y - E[Y|X]) = Cov(X,Y) - Cov(X, E[Y|X]) = 0,$$
since we showed in 4.58 a) that $Cov(X,Y) = Cov(X, E[Y|X])$. This linearity comes from the fact that
$$Cov(X,Y) = E\left[(X - EX)(Y - EY)\right].$$

- Try to simplify you expressions more, if possible. It is much easier to get a good intuition about your answer if it has been simplified enough. As an example, in 4.31 c), most people give the answer as

$$P(Y = y) = \binom{n}{y} \frac{\Gamma(y+1)\Gamma(n-y+1)}{\Gamma(n+2)},$$

(without mentioning that the probability is zero for $y \notin \{0, 1, \ldots, n\}$.) This is an ugly expression that makes it hard to get a good understanding of what the distribution actually is. Knowing that y is an integer, we can simplify this to

$$P(Y = y) = \frac{n!}{y!(n-y)!} \frac{y!(n-y)!}{(n+1)!} = \frac{1}{n+1}, y = 0, 1, \ldots, n.$$

Thus, we actually see that Y is uniformly distributed between 0 and n, which makes sense since $Y|X$ is binomial, and the probability X is uniformly distributed between 0 and 1, meaning that any success-probability is equally likely, which then leads to any number of successes being equally likely.

- Read through the text and answer all the questions. More than half of the assignments forgot to find the distribution of $Y|Y + X$ in 4.15 or to identify μ^2/α as the "extra-Poisson" variation in 4.35.