The following problems are difficult, but can be done with the techniques we've discussed. Take a look at them if you're interested; ignore them if you aren't.

1. One of the first probability problems ever considered seriously was the "problem of the points": if two gamblers are playing a game where the first to reach a certain number of points wins and they are forced to abandon the game in the middle, what is the probability that the first gambler would have won? Formally, suppose that the game consists of a series of independent Bernoulli trials with parameter p and that the first gambler gets a point on success while the second gambler gets a point on failure. What is the probability that the first gambler wins the game if the first gambler needs n more points to win and the second gambler needs m more points to win?

2. A related (but much more difficult problem) is the following: Suppose that a sequence of independent Bernoulli trials with probability p of success is run. For notational convenience, let $q = 1 - p$. Show that the probability that there will be a run of n consecutive successes before a run of m consecutive failures is $\frac{p^{n-1}(1-q^m)}{p^{n-1}+q^{m-1}-p^{n-1}q^{m-1}}$. (*Hint:* Start by conditioning on whether the first trial is a success or a failure. If we let S be the event that the first trial is a success and E be the event that there is a run of n consecutive successes before a run of m consecutive failures, try to find two equations relating $P(E|S)$ and $P(E|S^C)$ and then use linear algebra to solve for these quantities.)

3. Suppose that two urns are filled with red and blue balls. The probabilities of drawing blue balls from the first and second urns are p and q respectively. Balls are drawn one at a time (with replacement) from the urns according to the following procedure: With probability α, a ball is initially drawn from the first urn and with probability $1 - \alpha$, it is chosen from the second urn. For subsequent selections, follow the rule that if the previous ball was blue then the next ball is drawn from the same urn and that if the previous ball was red then the next ball is drawn from the other urn. Show that $\lim_{n\to\infty} P(\text{the } n\text{th ball is blue}) = \dfrac{p+q-2pq}{2-p-q}$. (*Hint:* Start by computing the probability that the nth ball will be drawn from the first urn.)

4. A pair of dice is rolled repeatedly. Show that the probability that there will be a roll on which the sum of the dice is 5 before a roll on which the sum of the dice is 7 is $\frac{2}{5}$. Then, generalize: show that if A and B are disjoint events, then the probability that the event A will occur on a roll before the event B occurs on a roll is $\frac{P(A)}{P(A)+P(B)}$. (*Hint:* For the special case, let E_n be the event that no 5 or 7 appears on the first $n-1$ rolls and that a 5 appears on the nth roll. Then you want to calculate $P(\bigcup_{n=1}^{\infty} E_n)$.)

5. Given events H and E consider the probabilities: $P(H)$, $P(E)$, $P(H|E)$, $P(H|E^C)$, $P(E|H)$, and $P(E|H^C)$. Pick any three of these and find formulas for the other three in terms of the three that you picked. (This can be done for any of the $\binom{6}{3} = 20$ ways of picking three. You don't really understand conditional probability until you've done this for all 20.)

6. If you did the previous problem, this one ought to be a breeze: Researchers have developed a new test to detect a certain disease. They've found that the test gives positive results 3% of the time, gives positive results 99% of the time if the person actually has the disease, and gives negative results 98% of the time if the person doesn't have the disease. Show that about 1 person in every 97 has the disease.