TN1910-1: Split Pill Problem Solved

George Rebane

12 October 2019

No doubt all *RR* readers will remember my invitation way back when (30jul19) to help solve the Split Pill Problem (SPP). The anticipated crickets were not much help, so I dove in, and after a bit of noodling and writing some code, voila! kicking and screaming, the SPP solution gave itself up. Knowing that you're all eager to see it in all its glory, I have decided to post it for your enjoyment and edification. Your wait is over. (Actually, this technical note is an intro tutorial to understanding a much more relevant problem affecting public policy that I'll cover later, so pay attention.)

The actual SPP was introduced here, but there's also a little back story that actually gave rise to it. You see, our puppy Puna (who's now 10 years old - the name means red in Estonian since she's a red Dobie) who was introduced to you years ago (here and here) - has had major surgeries since last Christmas to rebuild both of her knees. During the healing process she was prescribed several types of pills, one of them being an analgesic that also reduces swelling. These pills were administered daily in half pill doses, thereby requiring me to split a pill now and then – i.e. give her a half and put back the other half when a whole one pops out of the pill bottle. Splitting the pills was a bit of a bother that made me ask the question as to when would I expect to be done splitting before the remainder of the bottle was emptied, or stated otherwise – what is the expected count of half pills when only these remain?

The more I thought about what kind of solutions might exist for such a problem, and how to go about solving it, the SPP took on a life of its own and I was hooked. It turns out that that problems such as the SPP appear in a number of areas involving systems analysis and design, albeit in different forms and versions, and I don't recall seeing one explicitly posed and solved. So I dove in.

The first insight was that the SPP has no analytical solution. This means that you can't derive a bunch of equations that have the answer embedded in them, and then solve the equations to yield the answer. The alternative is to examine in detail what happens in the pill bottle as the days go by, with each day our puppy getting a half pill. Well, on any day after you take out the first whole pill, split it, and return the unused half back into the bottle, the bottle will contain a definite number of whole and half pills. The whole pill total, call it W, and the split pill total, call it S, can be written as W/S and denotes what systems people call the 'state' of the system (i.e. the pill bottle contents).

With pencil and paper you can quickly write down a little graph of W/S nodes connected with arrows, and see how the system state can vary as the days progress. Say you start with a $N = 10$ pill bottle in the initial state 10/0. This will require 20 days of administering half pills. On day one (D1) you take out the first whole pill, split it, and pop the remaining half pill back into the bottle – the state now becomes 9/1. On D2 you know that there's a small chance you may draw the half pill, but a much bigger chance that you'll get one of the 9 whole pills. So on D2 the state can have (or transition to) either 9/0 or 8/2. On D3 the 9/0 state can only transition to 8/1, since there are no split pills, but 8/2 can transition to either 8/1 or 7/3. Writing these transitions down carefully in an array, you will soon see the possible number of states – i.e. the 'state space' - bloom a maximum of (*N*/2+1) or 6, before starting to contract again to the terminal state 0/0 when the bottle is empty. You don't have to actually do it, because I did it for you in Figure1 below.

Day	Doses Left	Pills Left											
0	20	10	**10/0**				**The Split Pill Problem**						
			1.000	1.000			GJRebane - 1aug2019						
1	19	9.5	0.100	**9/1**									
				1.000	0.900								
2	18	9	**9/0**	0.200	**8/2**				P(No split)	**W/S**			
			0.100	1.000	**0.900**	0.800				**P(W/S)**	P(Split)		
3	17	8.5	0.111	**8/1**	0.300	**7/3**							
				0.280	0.889	**0.720**	0.700						
4	16	8	**8/0**	0.222	**7/2**	0.400	**6/4**						
			0.031	1.000	**0.465**	0.778	**0.504**	0.600					
5	15	7.5	0.125	**7/1**	0.333	**6/3**	0.500	**5/5**					
				0.134	0.875	**0.563**	0.667	**0.302**	0.500				
6	14	7	**7/0**	0.25	**6/2**	0.444	**5/4**	0.600	**4/6**				
			0.017	1.000	**0.305**	0.750	**0.527**	0.556	**0.151**	0.400			
7	13	6.5	0.143	**6/1**	0.375	**5/3**	0.556	**4/5**	0.700	**3/7**			
				0.093	0.857	**0.463**	0.625	**0.383**	0.444	**0.060**	0.300		
8	12	6	**6/0**	0.286	**5/2**	0.5	**4/4**	0.667	**3/6**	0.800	**2/8**		
			0.013	1.000	**0.253**	0.714	**0.502**	0.500	**0.213**	0.333	**0.018**	0.200	
9	11	5.5	0.167	**5/1**	0.429	**4/3**	0.625	**3/5**	0.778	**2/7**	0.900	**1/9**	
				0.086	0.833	**0.432**	0.571	**0.393**	0.375	**0.085**	0.222	**0.004**	0.100
10	10	5	**5/0**	0.333	**4/2**	0.571	**3/4**	0.75	**2/6**	0.889	**1/8**	1.000	**0/10**
			0.014	1.000	**0.257**	0.667	**0.493**	0.429	**0.214**	0.250	**0.022**	0.111	**0.000**
11	9	4.5	0.200	**4/1**	0.500	**3/3**	0.714	**2/5**	0.875	**1/7**	1.000	**0/9**	
				0.100	0.800	**0.453**	0.500	**0.371**	0.286	**0.073**	0.125	**0.003**	
12	8	4	**4/0**	0.400	**3/2**	0.667	**2/4**	0.857	**1/6**	1.000	**0/8**		
			0.020	1.000	**0.306**	0.600	**0.492**	0.333	**0.170**	0.143	**0.012**		
13	7	3.5	0.250	**3/1**	0.600	**2/3**	0.833	**1/5**	1.000	**0/7**			
				0.142	0.750	**0.511**	0.400	**0.310**	0.167	**0.036**			
14	6	3	**3/0**	0.500	**2/2**	0.800	**1/4**	1.000	**0/6**				
			0.036	1.000	**0.414**	0.500	**0.463**	0.200	**0.088**				
15	5	2.5	0.333	**2/1**	0.750	**1/3**	1.000	**0/5**					
				0.242	0.667	**0.577**	0.250	**0.180**					
16	4	2	**2/0**	0.667	**1/2**	1.000	**0/4**						
			0.081	1.000	**0.594**	0.333	**0.325**						
17	3	1.5	0.500	**1/1**	1.000	**0/3**							
				0.477	0.500	**0.523**							
18	2	1	**1/0**	1.000	**0/2**								
			0.239	1.000	**0.761**								
19	1	0.5	1.000	**0/1**									
				1.000									
20	0	0	**0/0**										

Figure1

Note that this figure has a whole bunch of other numbers on it that we'll get to directly. But for now, just concentrate on the light tan colored *W/S* cells which denote the states of the system. The tan cell below the *W/S* cell gives the probability of that state occurring on that given day (see leftmost green column of day numbers). As an example, find the state 4/5 on D7. Below it you will see that on D7 the probability of this state, *P*(*W/S*) = *P*(4/5) = 0.383. You'll also note that 4/5 can only occur on D7, on which day we may also witness the other possible unique states of 6/1, 5/3, and 3/7, each with their own probabilities of coming true, that for any given day always add up to one (certainty that the state will transition through one of the allowable states for that day).

Notice the little key to the inset in the upper right quadrant of the figure. It describes what the surrounding numbers are for each state and its possible transitions to the next state. The cell to the left, *P*(No split), indicates the probability of having drawn a whole pill from that state. Recall that every whole and half pill has an equal chance of being drawn out of the bottle. Therefore, *P*(No split) = *S*/(*W*+*S*), which for 4/5 becomes 5/(4+5) = 0.556. To the right of the *P*(*W/S*) cell is the probability *P*(Split) of drawing a whole which does require splitting. Therefore *P*(Split) = *W*/(*W*+*S*) = 4/(5+4) = 0.444. Note that these two probabilities must always add to one, since at every state position you must either draw a W or a S.

So how do we calculate the probability *P(W/S)* of any state being realized. We simply look at the transitions from its previous states, and do a little probability arithmetic. Each state can have at most two ways to get to it. You can get to state *W/S* from either 'parent state' *W*+1/*S*-1 or *W*/*S*+1. Near the state space boundaries there will be only one way of getting to *W/S*. So the required calculation is to multiply the probability of the parent state times the transition probability to the state in question. If there are two parents, then do that for each and add them together.

For our example state *W/S* = 4/5 we have parents 5/4 and 4/6 with *P*(5/4) = 0.527 and *P*(4/6) = 0.151. For 5/4 its *P*(split) = 5/9 = 0.556, and for 4/6 its *P*(No split) = 6/10 = 0.600. So that lets us compute *P*(4/5) = 0.527*0.556 + 0.151*0.600 = 0.383. And you do that for all the possible states starting with *P*(10/0) = 1.000 giving *P*(9/1) = 1.000*1.000 = 1.000. After that, the possible states start bifurcating as described above and shown in Figure1 which actually explains the whole 20-day rigamarole – i.e. all the possible state transition paths - for a *N* = 10 pill bottle.

The Most Probable Path. We now can determine which is the most probable path through state space over the 20-day period. This is done by starting at D1 with state 9/1 and finding the highest probability 'offspring state' or 'child state' of 9/1 on D2. From Figure1 that D2 state is 8/2 having probability 0.900 (the other child 9/0 has probability 0.100). This process is repeated for each succeeding day's state transition. Figure2 below shows the highest probability path through this dynamic state space by following the green arrows.

You can see that the most probable path is in the shape of an arc shown by the translucent blue-gray arrows. What is important for us to see here is on what day and which state does our most probable path encounter the All Splits Boundary which defines the states of the pill bottle when it only contains split pills. This tells us the answer to our primary question of 'what is the expected count of half pills when only these remain?' From Figure2 the answer is three, since the most probable path hits the boundary at state 0/3. From there it can only go to the end state 0/0 in the one way shown by the arrows.

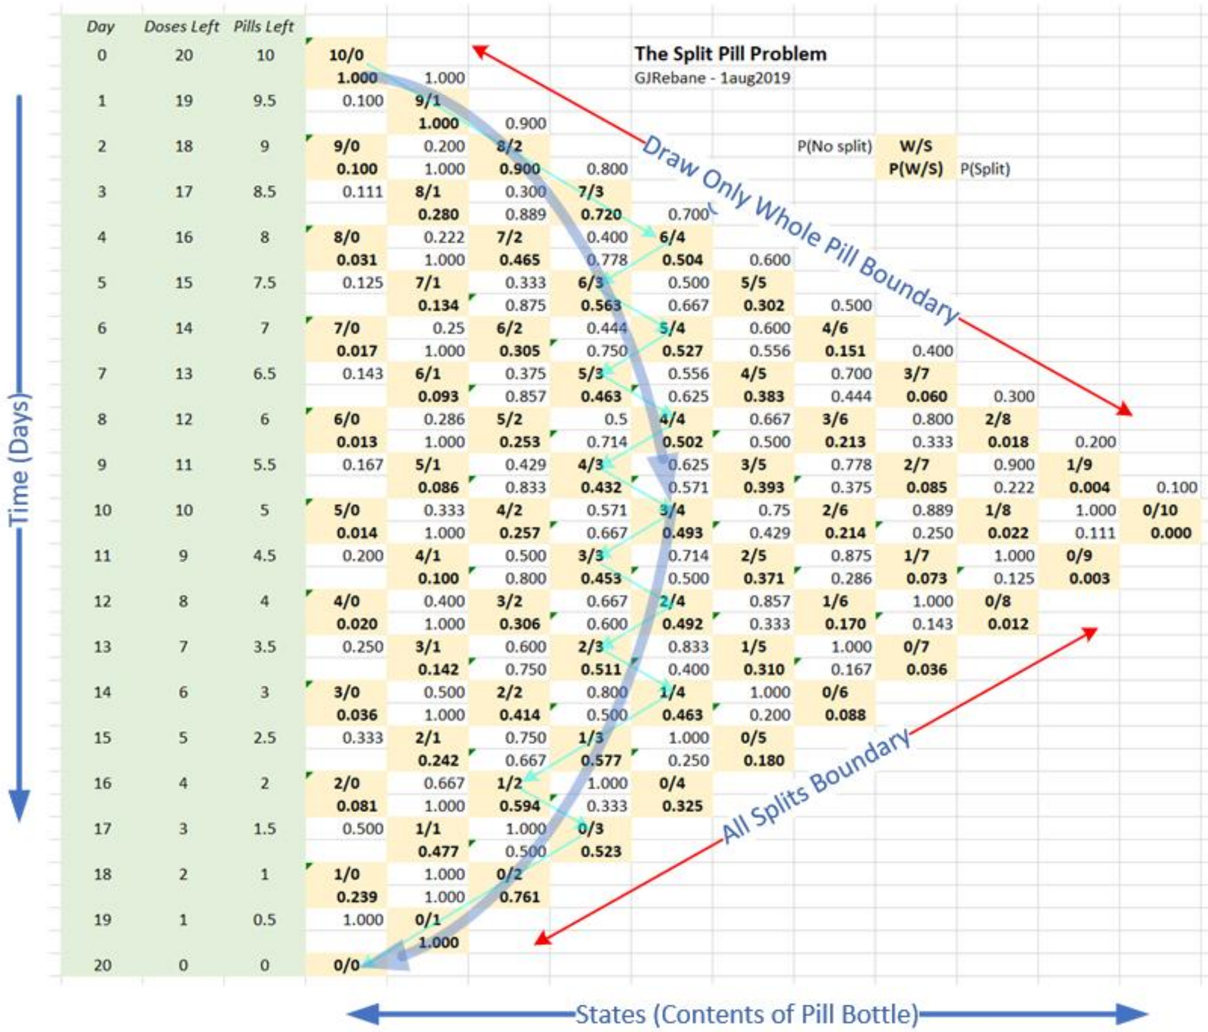

Figure2

The numerate reader will by now have noted that there can exist no simple formula for computing that answer given a pill bottle with ten whole pills. You have set up a state-space model and calculate it out day by day until you hit state 0/*X* where *X* is then the answer. You can also see that the state space for each day forms a discrete probability density function (pdf) in which the state probabilities sum to one or certainty, since the state must needs transition through one of the allowable states for any given day. Note that for a given starting number of pills, each succeeding day has only certain allowable states. Starting with N pills, the maximum number of allowable states is $N/2 + 1$ and occurs on day *N* or D*N*.

For any given pill dispensing experiment, we cannot guarantee that the path through state space will follow the most probable path. After all we are dealing with what is called a 'stochastic process' (a process dependent on more than one simple random process like flipping a coin or rolling a die), and each experiment will most likely yield a different path that we expect will follow the general arc shown in Figure2. Two such stochastic paths were generated, and are shown in Figure3.

Day				Two Random Paths Through State Space							
0	2	0	0	0	0	0	0	0	0	0	0
1	0	2	0	0	0	0	0	0	0	0	0
2	0	0	2	0	0	0	0	0	0	0	0
3	0	1	0	1	0	0	0	0	0	0	0
4	0	0	1	0	1	0	0	0	0	0	0
5	0	0	0	1	0	1	0	0	0	0	0
6	0	0	1	0	0	0	1	0	0	0	0
7	0	0	0	1	0	1	0	0	0	0	0
8	0	0	0	0	1	0	1	0	0	0	0
9	0	0	0	0	0	2	0	0	0	0	0
10	0	0	0	0	0	0	2	0	0	0	0
11	0	0	0	0	0	2	0	0	0	0	0
12	0	0	0	0	1	0	1	0	0	0	0
13	0	0	0	1	0	0	0	1	0	0	0
14	0	0	1	0	0	0	1	0	0	0	0
15	0	0	0	1	0	1	0	0	0	0	0
16	0	0	1	0	1	0	0	0	0	0	0
17	0	1	0	1	0	0	0	0	0	0	0
18	0	0	2	0	0	0	0	0	0	0	0
19	0	2	0	0	0	0	0	0	0	0	0
20	2	0	0	0	0	0	0	0	0	0	0

Draw only Whole Pill Boundary

All Splits Boundary

Figure3

Nevertheless, if we repeat the experiment a large number of times for some *N*-bottle, then we should expect the fraction of times the trajectory went through any given state on any given day to closely approximate the theoretical pdf that are computed and shown in Figure1 and Figure2. You can see what this result looks like by first appreciating the beauty of the 3D surface made of stacking together the daily state pdfs for a 100-bottle (*N*=100). This shown in Figure4 below which is annotated with the All Splits Boundary and Whole Pill Boundary. Note which direction the days run in this figure. The 'most probable berm' forms a graceful arc through state space that defines the region through which the overwhelming number of state trajectories (pill dispensing histories) will run. (Such beautiful surfaces and figures turn up more often than not from numerical calculations of complex processes, and, correctly plotted, they often provide valuable insights in the analysis and understanding of the underlying process/es.)

The Arc of Life. The most probable berm through state space shown in Figure4 has another message embedded in it. Consider how we humans go through life. We all start out as post-partum infants within a very limited state space that offers us few opportunities beyond the hopeful one that someone will nurture and protect us. As we grow older and more capable, additional alternatives begin to present themselves and we start having more choices as we proceed on our life's path. And maturing even more toward adulthood, the available paths multiply rapidly, expanding available possibilities even more. Somewhere in the middle of life we max out in the number of choices available to us, and as we age past that, we notice that our life's path starts becoming more constrained with few opportunities – the state space of life starts contracting. And this age of ever more limits continues until we finally arrive again at a singular state from which we exit life.

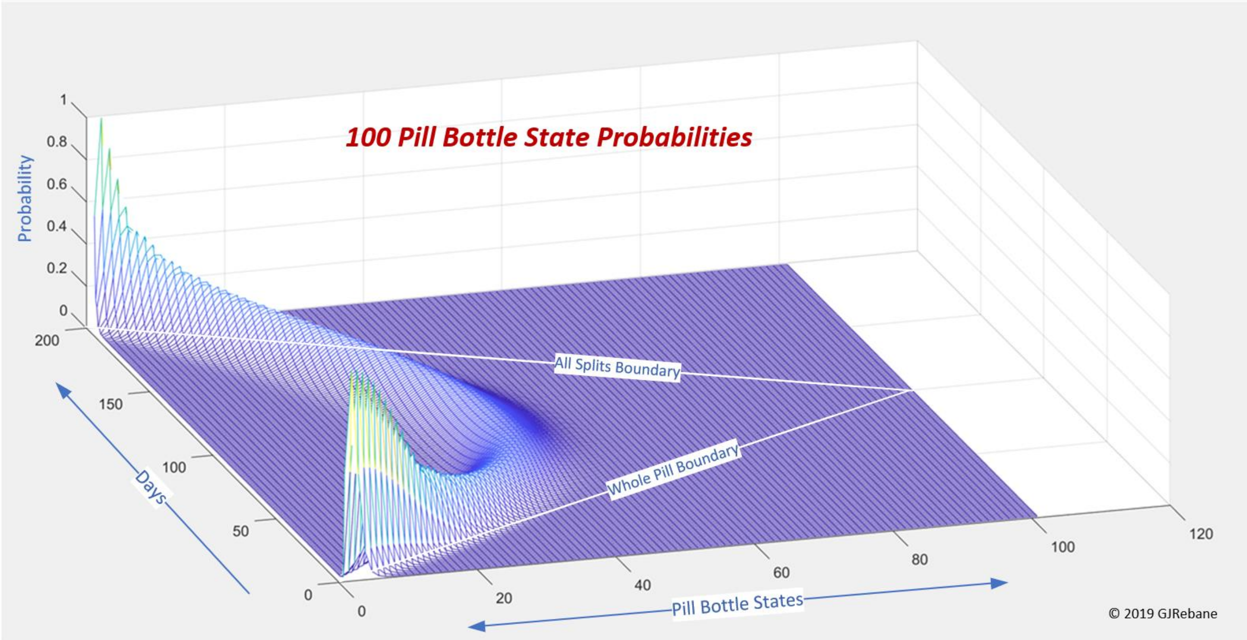

Figure4