

Extremely free ruminations/ramblings on rational/disciplined/what-have-you problem solving

This all began with the following “puzzle” : *You have 243 bottles of wine, one of which is poisoned. You also have 5 rats, and if a rat ingests even the tiniest sip of poison, it will die within 1 hour (though you can't guarantee when). How can you use the rats to find the poisoned bottle in 2 hours?*

How can we think about puzzle-questions like this?

I take the view that a puzzle or problem is a gap to be filled. There's some information we have, and there's something we'd like to create or answer, and the question is: how do we get from one to the other? I like this view because it's more generally a question of connecting concepts or groups of concepts, **A** to **B**. In the case of a puzzle, **A** is information we have but it doesn't have to be that. Maybe we are inventing **A** as we go. Before hammers existed, they had to be invented to solve a problem. Or maybe **A** is a budget we have to get approved. And maybe we don't know **B** either, maybe we're just exploring where we can go from **A**.

When we view problem solving more abstractly like this, it chips away at the anthropomorphic, ego-centric notion that we are lowly humans stuck in our ignorance (**A**), needing to slave away to work at the difficult task of climbing the mountain in order to achieve our goal (**B**). It's just about finding relationships between two things. Those relationships can go in either direction, and we are free to identify with either, both, or neither.

So, that's the first step in rational problem solving. It's a little more clear now what we mean by 'rational', too. It's not about being “right”, “better”, “smarter”, or anything like that. It's just a contrast with 'intuitive' or 'emotional' problem solving. With emotions, we don't sense **A** and **B** or anything like that: we just sense the problem. We feel it as a tension, and we work instinctively to remove the tension, maybe by physically moving away, or applying knowledge. (I just burned myself while taking a pork tenderloin out of the oven and jerked my hand away. Intuitive problem solving!)

There is nothing wrong with intuitive problem solving; in fact, we probably use it a lot while we are rationally problem solving. But since it only applies to what we can more immediately see and feel, it's not very effective on its own at finding solutions that need to be planned or structured.

On the other hand, rational problem solving is about stepping away from the immediate, emotional reaction, and looking at things more soberly. We've seen that already by turning the antagonistic image of a problem being something “against” us that we have to combat,

into a more neutral picture of drawing connections between two things. So it's about stepping back, and also maybe ironically about looking closer. Because when we're not focused on our immediate reactions to a problem, we're free to look at (and perhaps uncover) things we might have missed otherwise: subtle details that can allow for more subtle and powerful connections to be drawn. In this way, we see that rational exploration isn't just part of science or other traditional "reasoning" domains, but can be part of art, poetry, music, love, etc.

Another analogy to contrast rational and intuitive thought might be speed. Intuition is about speed (maybe biologically about survival), while rational thought is about being prepared to go as slow as possible. Intuition may save your life if you're being chased by a tiger, while rational thought is needed if you're trapped in a room and only one tile can be pushed to allow for escape. You have to work thoughtfully, slowly, carefully, methodically. Clearly, both kinds of thought and problem solving are necessary, and in fact they complement and balance each other.

I think intuitive and rational problem solving work in very similar ways. They both draw connections between things, though in the intuitive approach on the one hand we're pretty exclusively working with our relevant knowledge to the situation at hand. We're drawing intuitive connections, which means we're probably reacting and responding to the most readily apparent features of whatever we're dealing with. (Which isn't to say we aren't capable of seeing very subtle things, even at a subconscious level!)

Back to **A** and **B**. So we already have this model of approaching a problem. Applying to the problem at hand, **A** is our 243 bottles of wine (1 poisoned), 5 rats, and 2 hours, while **B** is the poisoned bottle. Following the above discussion, we might take some time here and step back. Because no sooner than I've described **A** and **B**, I can't help but realized that I've phrased everything in terms of objects. Maybe I can find other ways to describe the situation. After all, while intuitive thought is often pre-verbal, rational thought has all the power of language at its disposal, and there's no reason to be satisfied with my first, gut effort.

Focusing on **B**, which I described as "the poisoned bottle", I'm slightly dissatisfied because even though that's what "I" the character in the problem would end up finding, I as the problem solver am not going to wind up with a bottle. What I need is a method, a set of rules, an algorithm, a process, for giving wine to the rats on some timed schedule, the outcome of which will allow me to determine the poisoned bottle. Specifying that process is what I need to do.

So what I've just started to do is clarify **A** and **B**. That's an important part of rational problem solving, to try to have a keen sense of what's on the table at any given point. To try to strip away aspects of **A** and **B** that may not be relevant, to uncover the heart of the matter. Precision and generalization are important here.

In the same vein (or: put another way), here is where it's important to embark on a separation of concerns, or disentanglement. We want to pull apart, if we can, the various features of **A** and **B**. This will firstly allow us to focus on one thing at a time (which is important because the focus needed for rational thought is pretty all-consuming (unlike with intuitive thought, which is largely subconscious)), and also allows us to see which pieces of information or features we've not managed to connect or use. We might even be able to tell that a problem is unsolvable without using a certain piece of information, in which case it's crucially important to know whether we've used it. Et cetera.

Back to the problem at hand. We're looking for a method. What can we say about that method? Let's turn to **A**. We're giving wine to the rats; they'll die or they won't; then we draw a conclusion. Okay, make it more precise. We have two hours, but it takes up to an hour to decide whether a rat was poisoned. That means each rat can only be tested twice. But there are 5 rats and 243 bottles of wine. If we only use one bottle per test, there's no way to solve the problem. But this line of reason now solves the problem if there are only 11 bottles of wine (or 49 hours) (or 121 rats).

(Note that in the preceding paragraph, the decisions of where to turn, and how to compare information, was led largely by intuition. Even if the decision to pick up a hammer is made rationally, the actual picking up is guided by intuition — we don't micromanage our muscles all the way down.)

These conclusions we are drawing are a result of separation of concerns / disentanglement / generalization: we're noticing that if we make certain budgets to the givens or allowances, we could solve the problem. That might seem pointless, but it's not. Remember that in many situations, **A** is not fixed, but maybe is negotiable, like a budget. If we could ask for more rats or hours or fewer bottles, we can solve the problem. Also, when we approach problem solving in this way, we might come up with a new "sub"-problem that is easier to solve.

But assuming we can't change the numbers of bottles, rats, or hours, we have to make a budget somewhere else. We'll have to allow ourselves to mix samples from multiple bottles in the tests.

All this might seem obvious, but that is the hallmark of rational thought: staying close to the most obvious and essential aspects of **A** and **B**. That would be classified under simplicity. To summarize what we've clarified so far. We're looking for a testing method which runs as follows: We'll give some wine mixtures to the five rats. Then we'll see which rats die. (That means if all the rats die, we have to be able to find the poisoned bottle after the first hour!) If there are still rats alive, we'll give some wine mixtures to the remaining rats. We'll see which rats die, and from those two outcomes, we have to be able to find the poisoned bottle.

What is left to be specified? The mixtures. How do we make the specification? With the constraint that the death pattern can tell us about where the poisoned bottle is.

Now I notice (intuitively) that 243 hasn't come into play in any of the above. Or 5 for that matter. One avenue to explore next would be the mathematical properties of these numbers, to draw some inspiration there.

But I think we might be able to say more. Above I wrote "the death pattern [has to tell us] where the poisoned bottle is". Okay, so what can we say about that? If we give a wine mixture to a rat, and it dies, then the poisoned bottle is among that mixture. So one instinctive path towards a solution (simple!) might be to divide the 243 bottles into disjoint groups (49, 49, 49, 49, 47?), and test those. That would narrow it down from 243 to at worst 49 after one hour! But then what? 49 bottles and 4 rats left. We could split them 13, 13, 13, 10. Then 13 bottles and 3 rats. Split them 5, 5, 3. Now 5 bottles and 2 rats. Split 3, 2. One rat and 3 bottles. Can't be done. But I probably could fudge the above numbers and make it work, but this process takes 5 hours. (Yes, you can do it, I realized, because at each step you can leave out one bottle — if all rats survive then the one you left out is poisoned. Doing it this way works in 5 hours.) The previous budging only worked in 49 hours, so this method is already a huge improvement.

But! We've learned something, namely that we can't come close to solving the problem unless we let the wine mixtures overlap. (Meaning that the same bottle might be a part of different mixtures.) This of course makes the solution much more subtle — nearly impossible to stumble upon. There are probably millions of ways to guess at how to create these overlapping mixtures of wine. If there isn't a methodical / mathematical way to solve the problem, it would require the brute force power of a computer.

So, rather than guessing at the mixtures anymore (instinct), let's try to describe the structure of the outcomes of the testing (rational). Whatever mixtures we use, each rat will either die after the first hour, or after the second, or not at all. So one rat yields one of three possible outcomes. Now, if each of the five rats can yield one of three outcomes, then there are $3 \times 3 \times 3 \times 3 \times 3$ conceivable outcomes of the testing. And this number happens to be 243. Cool! This answers a lot of general questions about the problem. If we hadn't been told the number 243, we could have concluded that this is the largest number of bottles we could possibly test in this scenario. We still don't know that we can actually do it, but if the problem had been stated with 244 bottles, we could confidently answer that the problem is impossible to solve as stated.

Again, it may seem silly to talk about all these general conclusions, but in fact, reinforcing and understanding the connections between our givens starts to hem in the space of possibilities in which we can find further connections (and a solution). Once again clearing away some of the legwork and rephrasing our understanding of the problem: We are

going to give some necessarily overlapping wine mixtures to the rats as described above. Since there are 5 rats and 3 outcomes for each rat, our experiment yields $3^5 = 243$ distinct outcomes, and since any one of the 243 bottles may be poisoned, we have to ensure that the mixtures are created in such a way that each outcome corresponds to the correct bottle.

It's worth saying at this point that instinct might have led someone to factor the suspicious-looking number 243 from the start, and back-construct the relevance of 3 and 5. Furthermore, if they know a little more math (technically something taught in grade school), a reasonable-seeming solution might immediately present itself. Something to be said for the connection between top-down, rational investigation and bottom-up, instinct-driven, information-fueled investigation.

Okay, so back to the mixtures. Let's see what we can conclude from the experiment with one particular rat. In the first hour, we give it some mixture of bottles, call it M1. If the rat dies, we know the poison bottle is among M1. If it doesn't, we give it another mixture, call it M2. Note that we could include any of the bottles from M1 in the M2 mixture, as we know they are safe for this rat. But as we learn nothing new from retesting these bottles, it seems sweetly reasonable to say that all the M2 bottles should be new.

[A word about this term: 'sweetly reasonable' is to **B** as 'opportunistic' is to **A**. If we see an opportunity to fruitfully use our information, that helps us connect from **A** towards **B**. If we are looking at the goal, **B** (in this case our mixture strategy), doing something sweetly reasonable is a way to reasonably constrain the types of solutions we are looking for, in order to make one easier to find.]

Moreover, we must include some new bottles in M2, otherwise the rat will not die and we learn nothing new. It is sweetly reasonable to assume that because we have exactly enough rats and hours to cover 243 bottles, we shouldn't waste any opportunity to conclude information.

If the rat dies, we know that the poison bottle is among M2. Otherwise, the poison bottle is among the untested bottles, which we could call M0.

[Another digression. Notice as we are working we are introducing some ad-hoc terminology: M0, M1, M2. The issue of naming is a crucial part of rational problem-solving, and whole essays have been written on the topic. What to name, when to name, naming part-whole relationships, naming 'opposite' concepts (like 'heavy' instead of 'non-light'); and then in mathematical domains, the issue arises whether the name or notation is just intended to be a shorthand, or whether it is suitable for doing calculation (uninterpreted manipulation of named symbols and formulas). Naming is important to rational exploration because it allows you to talk and think about concepts more clearly.]

My instinct at this point, given the ‘tightness’ of 243 in this problem, is that M0, M1, and M2 should all have the same number of bottles, namely $243 \div 3 = 81$. Also, when considering the symmetry of the problem, all the bottles are indistinguishable from one another. There is no way to ensure which of M0, M1, M2 will have the poisoned bottle, and we don’t want one outcome to leave us with more bottles to test than in other outcomes. So, that would be a sweetly reasonable constraint, and of course the symmetry of the problem once again would tell us that, at least when considering just a single rat, it doesn’t matter one whit which bottles we put in M0, M1, M2, as long as they each have 81 distinct bottles. It would be sweetly reasonable yet again to assume that each mixture for each rat will be comprised of 81 distinct bottles (although the rats presumably need different mixtures — otherwise there’s no point in having different rats).

I’ll continue exploring this path a little later. But right now it strikes me that there is another sweetly reasonable option: link the 243 potential outcomes with the 243 bottles. Why? Because the numbers are the same. So for example, one of the outcomes is “Rats 1, 3, and 4 die in the first hour, rat 2 dies in the second hour, and rat 5 doesn’t die”. So we assign that outcome to one of the bottles, call it Bottle X. Each bottle gets assigned exactly one of the 243 outcomes of the testing.

Now, does this pay off? That is, does it tell us how to create the mixtures? Let’s see. If Bottle X is poisoned, we need to be able to identify it at the end of the testing. How could we do that? Well, the only way to identify Bottle X is by the outcome associated with it, which claims that rats 1, 3, and 4 die in the first hour, rat 2 dies in the second hour, and rat 5 doesn’t die. So, if we want to identify Bottle X, we had better follow what it says, and give it to rats 1, 3, and 4 in the first hour, rat 2 in the second hour, and not at all to rat 5. The other bottles are safe, so this outcome will indeed come to pass no matter how we distribute the wine in the other bottles.

But of course, we don’t know which bottle has the poison, so we had better follow the outcome associated with every bottle. Therefore we have hit on a solution: Assign the 243 outcomes arbitrarily to the 243 bottles, and then create the mixtures to follow the predicted outcomes on the bottles. The outcome on the poisoned bottle will come to pass, while all other outcomes (being different) will be incorrect.

[Note that we might shorten Bottle X’s outcome to 12110, where these five digits correspond to the five rats; the ‘1’ and ‘2’ stand for the hour of death, while ‘0’ stands for no death. Then we may recognize that the 243 outcomes correspond to the numbers 0 through 242, as represented in ternary or base-3 notation. Someone who is familiar with this sort of notation might intuit this solution almost instantaneously upon realizing that $243 = 3^5$ or by recognizing that the rats and hours are code for 243 combinations. And speaking of naming: this code ‘12110’ is —unlike M0, M1, M2— not just a shorthand. Base representational systems are highly flexible and powerful ways of representing number.]

Okay, but we can't always count on our hunches to work out, so I want to return to the original plan and see if we can use reason to figure out the appropriate wine mixtures.

To do this, I want to see if I can't simplify the problem first to get a grasp on what's going on. Of course, ideally I want to simplify in such a way that whatever I figure out for the simpler case will "scale up" to the problem at hand. To do this, I am not only guided by taste, but also by my understanding of the structure of the problem and the connections we have drawn already. In particular, I am thinking about the relationship between 3, 5, and 243. We have already seen that 243 is a function of 3 and 5, so let's see about changing those. My taste is to simplify 5, the number of rats, though we could also try lowering the number of hours.

What if there were only one rat? Then the number of outcomes is $3^1 = 3$, so we can handle 3 bottles of wine. Now let us recall the idea of M0, M1, M2. We test the bottles in M1 in hour 1, and those in M2 in hour 2. If the rat dies in hour 1, the poisoned bottle is in M1; if the rat dies in hour 2, the poisoned bottle is in M2; otherwise the poisoned bottle is in the untested M0. Now, once the rat dies, we can do no more testing. And once the two hours are up, we can do no more testing. This means that M0, M1, and M2 can have at most 1 bottle each — if there are two or more bottles in one of these mixtures, we can't narrow it down. But also by design, each bottle of wine is in either M0, M1, or M2. Hence the only possibility is to have each mixture consist of one of the three bottles.

That might seem like overkill, to apply so much reasoning to an obvious case. But this again is the point of rational exploration. If we'd used intuition to solve the simple case, we'd have nothing to guide our generalization when we move to a more complex case.

So, now consider two rats. The number of outcomes is $3^2 = 9$, so we can handle 9 bottles of wine. Let's say the first rat has mixtures M0, M1, M2 (not the same as the previous case — we have to figure them out again because there are more bottles). Now, a few paragraphs ago we said it would be sweetly reasonable to suppose that each of these three mixtures would have wine from an equal number of bottles, 3 in each. But I don't want to make that assumption — I just want to reason on. So the only thing I know about the M mixtures is that every bottle is in one of the mixtures.

The second rat will receive mixtures N0, N1, N2, and again every bottle is represented in one of these mixtures. Now it's time to generalize the conclusion from before:

Now, once the rat dies, we can do no more testing. And once the two hours are up, we can do no more testing. This means that M0, M1, and M2 can have at most 1 bottle each — if there are two or more bottles in one of these mixtures, we can't narrow it down.

How does this observation generalize to the case of two rats? Once both rats die, we can do no more testing. And once the two hours are up, we can do no more testing. This means

that the Ms and the Ns can have at most 1 bottle in common — if there are two or more bottles in an M/N combination, we can't narrow it down. To explain by example: Suppose the first rat dies in the first hour and the second rat doesn't die. Then the poisoned bottle is in M1 and N0. If M1 and N0 have two or more bottles in common, we can't narrow it down. So every M/N combination can have at most 1 bottle in common, and since there are 9 such combinations and every bottle appears in a combination (every bottle is tested or not by each rat), we conclude that every M/N combination has to correspond to exactly 1 bottle.

So we can arbitrarily label the bottles M0/N0, M0/N1, etc and that will tell us our groupings. (Or if that top-down method isn't obvious we could do it concretely: Which bottles will be in M0? One that overlaps from N0, one from N1, one from N2. So we pick three arbitrary bottles and put pour one in N0, one in N1, one in N2, and a little from all three in M0. Then we put those bottles aside and repeat with M1 and M2.) Note that this indeed means that each mixture contains wine from 3 bottles. Sweetly reasonable as that assumption would have been, we didn't need to make it.

The generalization to 5 (or more!) rats now seems extremely clear. Whatever the mixtures are, they must be designed in such a way that every one of the 243 combinations of mixtures has exactly 1 bottle in common. At this point we might drop the Ms and Ns and just apply labels like 01200 or 12110, the latter of which would mean for example that wine from that bottle should be put into the mixture given to the first, third, and fourth rats in hour 1, the second rat in hour 2, and not at all to the fifth rat. (And again we indeed see that each mixture contains wine from 81 bottles, but we didn't need to make that assumption.)

February 27th, 2017
NYC