

6. Decision Theory

Not actually covered in Varian or most undergrad micro text; lots of good advanced books, but none at the right level. In order of decreasing accessibility, consider reading:

An Introduction to Bayesian Inference and Decision, by Robert L. Winkler (Holt, Reinhart and Winston, 1972).

Optimization in Economic Theory, Second Edition, by A.K. Dixit (Oxford, 1990).

The Analytics of Uncertainty and Information, by Hirshleifer and Riley (Cambridge University Press, 1992).

Optimal Statistical Decisions, by Morris DeGroot (McGraw-Hill, 1970).

I. Overview

A. Static optimization used so far neglects crucial complications such as:

1. You might have better information later, but delay is costly.
2. Your decision now may affect your opportunities (and information) later.

B. Decision theory and practice go back centuries and are still maturing, but saw major advances in the late 20th century.

C. Example applications:

1. Valuing real options, e.g., when to invest, or shut down or expand operations.
2. Valuing information and learning, e.g., when are consultants worth their fees?
3. Valuing financial options, and related questions central to modern finance.

II. Bayes Theorem

A. How should you revise your beliefs when new information arrives?

B. First breakthrough in Decision Theory, goes back to late 1700s (Rev. Bayes, Laplace).

C. Here's a basic guide, beginning with notation.

1. $m = a, b, c, \dots \in M$ are the possible **messages**, or new pieces of information, that may arrive.
 2. $s = A, B, C \dots \in S$ are the possible true **states** of Nature.
 3. $p(m, s) \geq 0$ are the joint probabilities that m is observed and s is the true state. Note that $\sum_{s \in S} \sum_{m \in M} p(m, s) = 1$.
 4. $p(m) = \sum_{s \in S} p(m, s) > 0$ is the probability of message m .
 5. $p(s) = \sum_{m \in M} p(m, s) > 0$ is the probability of state s . Bayesians call it the **prior probability** of state s .
 6. $p(m|s) = p(m, s)/p(s)$ is the conditional probability that m will be observed when the true state is s . Bayesians call it the **likelihood** of m given s .
 7. $p(s|m) = p(m, s)/p(m)$ is the conditional probability that the true state is s when m is observed. Bayesians call it the **posterior probability** of s given m .
- D. The goal is to properly compute the posterior probability after you observe a particular message m , when your prior beliefs were summarized by the $p(s)$'s. Bayes theorem tells you how to do that when you know the message likelihoods.
1. For example, suppose that you know that in the US, one person in a million with flu-like symptoms has zika, and there is a quick test that is 99.0% accurate in detecting zika when present (i.e., probability is 1% for both type I and type II errors).
 2. Suppose your friend comes in with flu-like symptoms, and then gets a positive test result. Given that bad news, what is the probability that she actually has zika?
- E. Here are three equivalent versions Bayes Theorem.

$$p(s|m) = \frac{p(m|s)}{\sum_{t \in S} p(m|t)p(t)} p(s) \quad (1)$$

$$\frac{p(s|m)}{p(t|m)} = \left[\frac{p(m|s)}{p(m|t)} \right] \left[\frac{p(s)}{p(t)} \right] \quad (2)$$

$$\ln[\text{posterior odds}] = \ln[\text{likelihood ratio}] + \ln[\text{prior odds}] \quad (3)$$

F. In words,

- the first equation says that the posterior state probability is proportional to the prior probability (with the proportion being the likelihood of the message given that state s , relative to all alternative states t).
- The second equation says that the posterior odds for any two states are equal to the prior odds times the likelihood ratio.
- The third equation just takes logs of the second equation.

G. The proof is simple, just unwinding definitions.

1. Begin with definition 7 above for $p(s|m)$, and insert def. 4 in the denominator.
2. Then rewrite the numerator $p(m, s) = p(m|s)p(s)$, using def 6.
3. Finally, again use def 6 for all summed terms in the denominator, $p(m, t) = p(m|t)p(t)$.
4. The final expression is exactly the first equation.
5. To get the second equation, simply divide the version of (1) for s by the version of (1) for another particular state t . The denominators in (1) cancel.
6. As already noted, (3) follows from (2) by taking logs.

H. Now let's solve the zika problem, and reassure your friend.

1. Let $s = A$ be the state of Nature where she actually has zika, and $s = B$ the state where she does not. So $S = \{A, B\}$.
2. Let $m = a$ be the positive test result, and $m = b$ the negative test result. So $M = \{a, b\}$.
3. Prior probabilities $p(A) = 0.000001$ and $p(B) = 0.999999$ are given, as are likelihoods $p(a|A) = 0.99$ and $p(b|B) = 0.99$. The error probabilities are $p(b|A) = p(a|B) = 0.01$.
4. Equation (1) says

$$p(A|a) = \frac{p(a|A)}{p(a|A)p(A) + p(a|B)p(B)}p(A) = \frac{.99}{(.99)(.000001) + (.01)(.999999)}(.000001) = 0.0000989.$$

5. Equation (2) says that the posterior odds are

$$\frac{p(A|a)}{p(B|a)} = \left[\frac{0.99}{0.01} \right] \left[\frac{0.000001}{.999999} \right] \approx [100][0.000001] = 0.0001.$$

6. Using base 10 logs (any base will do!), equation 3 says the posterior log odds are about $2 + (-6) = -4$.

7. Bayes Theorem tells us that the very strong prior (one in a million) dwarfs the pretty good message (about 100 to 1), so after the test your friend still has only about 1 chance in 10,000 of having zika.

8. Here's some intuition, in terms of frequencies. Out of a million folks with flu-like symptoms, on average only one actually has the disease but about 1% or 10,000 get a positive test result. Hence the posterior odds are about 1 in 10,000.

I. A moderately complicated finite problem is probably best done on a spreadsheet. Here's a good layout.

- The first column lists the true states $A, B, C, \dots \in S$.
- The second column lists the given prior probabilities.
- The next $|M|$ columns list the given likelihoods $p(m|s)$ of each message $a, b, c, \dots \in M$, one column for each message m .
- The rest of the spreadsheet calculates all other probabilities, using the formulas already introduced.
- The next set of $|M|$ columns lists the joint probabilities of each message-state pair, $p(m, s) = p(m|s)p(s)$, computed from preceding columns.
- Each column sum is the message probability $p(m) = \sum_{s \in S} p(m, s)$.
- The final set of $|M|$ columns lists the posterior probabilities $p(s|m) = p(m, s)/p(m)$, computed using a preceding column (and its sum).
- An example in excel can be found on the class website.

J. More on Bayes Theorem

1. If the message is independent of the state, then

- $p(m, s) = p(m)p(s)$ and
 - Bayes Theorem tells us that the posterior is the same as the prior, i.e., $p(s|m) = p(s)$ for all $m \in M$.
 - Such messages are uninformative, i.e., useless.
2. If you have continuous message spaces and/or continuous state spaces, then
- the formulas apply to the density functions $f(m, s)$ etc replacing the discrete probabilities $p(m, s)$ etc.
 - You just have to work with states and messages where the relevant densities are continuous and positive.
 - Mathematically inclined readers who know epsilon-delta arguments should have no problem proving this by taking the limit of discrete approximations to f .
3. What if you get several different messages, one after another? Or run several different diagnostic tests?
- Then just reapply Bayes theorem, where the new prior is the old posterior.
 - If messages are conditionally independent from each other (hopefully not independent of the true state!) then you can get a factor in (2) for each message, and
 - (3) says that you sum all the message log likelihoods (plus the log prior odds) to get the log posterior odds.
4. Does (3) remind you of maximum likelihood estimation (MLE)?
- There is a connection. Think of each observation as a separate message, and of all possible parameter values as S , the states of Nature.
 - In standard MLE, the likelihoods are functions of the unknown parameters, and one can search parameter space to find parameter values that maximize the sum of log likelihoods.
 - In standard MLE, in effect the prior odds are set to 1, so the last term ($\ln 1 = 0$) in (3) disappears.

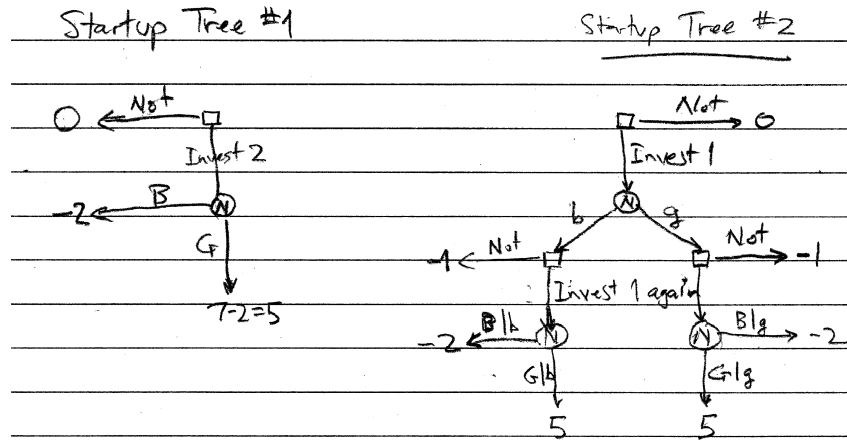


Figure 1: Your decision nodes are marked with small boxes, and random event nodes (“Nature” moves) by circles marked N.

- In Bayesian statistics, one explicitly estimates the posterior distribution given some prior distribution and given the likelihood functions for the observed data.
5. Behavioral economists find that in many situations, people tend to underweight prior information, i.e., over-respond to news. Did you fall into this trap on the zika example? The frequency intuition helps many people get it right.

III. Dynamic Decision Trees

A. Here is a pair of example problems to get us going.

Ex # 1: You decide now whether to invest 2 (\$ million) in a startup company.

- With probability $p(A) = .25$ it will succeed and be bought out at a price whose present value is 7.
- With probability $p(B) = .75$ it will fail and have salvage value 0.
- See Startup Tree #1.

Ex # 2: You decide now whether to invest 1 (\$ million) to allow the startup to release

a beta version of its product. After seeing whether the beta version is good or bad, you can decide whether to invest the remaining 1 (\$ million).

- See Startup Tree #2.
- The likelihoods of a good beta version for a product ultimately fated for success and for failure are $p(a|A) = 0.9$ and $p(a|B) = 0.3$.
- Using Bayes Theorem etc, the spreadsheet calculates beta version probabilities $p(a) = 0.45$ for good, and $p(b) = 0.55$ for bad; and posterior probabilities $p(A|a) = 0.5, p(B|a) = 0.5, p(A|b) = \frac{1}{21}, p(B|b) = \frac{20}{21}$.

B. Now let's solve the two example problems.

- In Tree #1, if we invest then we get expected payoff $E\pi = \sum_{s \in S} p(s)\pi(s) = p(A)\pi(A) + p(B)\pi(B) = .25(7 - 2) + .75(-2) = -0.25$. This is worse than the 0 payoff to Not investing.
- Hence we choose the Not invest branch and get payoff 0. See solved Tree #1.
- To solve Tree #2, we begin by figuring out what we would get if we invested following the beta release.
- If the release is good, then we would get $E[\pi|a, inv] = \sum_{s \in S} p(s|a)\pi(s) = p(A|a)\pi(A) + p(B|a)\pi(B) = .5(7 - 2) + .5(-2) = 1.5$.
- If the release is bad, then we would get $E[\pi|b, inv] = \sum_{s \in S} p(s|b)\pi(s) = p(A|b)\pi(A) + p(B|b)\pi(B) = \frac{1}{21}(7 - 2) + \frac{20}{21}(-2) = -\frac{35}{21} = -\frac{5}{3}$.
- Comparing $E[\pi|a, inv] = 1.5$ to $E[\pi|a, Not] = -1$, we decide to invest if the beta release is good, resulting in expected payoff $E[\pi|a] = 1.5$.
- Comparing $E[\pi|b, inv] = -\frac{5}{3}$ to $E[\pi|b, Not] = -1$, we decide Not to invest if the beta release is bad, resulting in expected payoff $E[\pi|b] = -1$.
- Given this contingency plan, we see that the expected payoff to sinking an initial investment of 1 is $E[\pi|inv] = \sum_{m \in M} p(m)E[\pi|m] = p(a)E[\pi|a] + p(b)E[\pi|b] = .45(1.5) + .55(-1) = .125$.
- Since this expected payoff exceeds the 0 payoff for Not sinking the initial

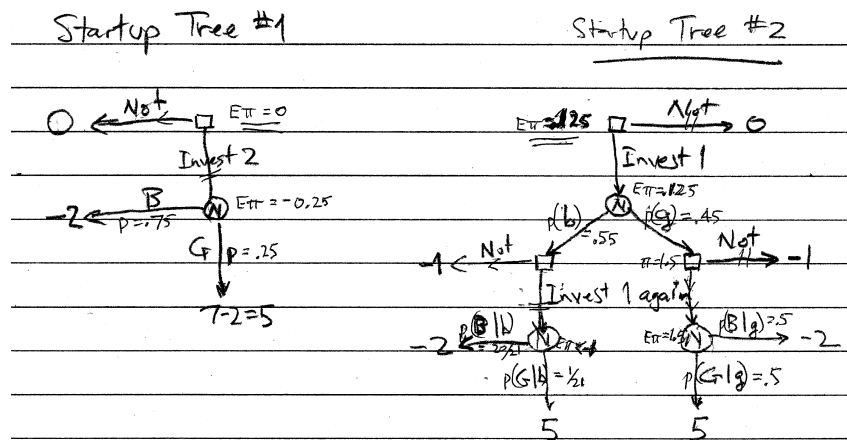


Figure 2: Solution takes $E\pi$ at N-nodes and maximizes at your decision nodes. Double slashes denote eliminated branches at each of your nodes.

investment, we see that our optimal course of action is: *sink the initial 1, and invest another 1 if the beta release is good, otherwise to abandon the startup.*

- See solved tree #2 in Figure 2.

IV. These examples illustrate the general rule for solving dynamic decision problems. Here is the cookbook.

A. First, construct the decision tree as follows.

1. Consider the logical sequence of events and the important contingencies, described as messages, states and choices. What will you know at the time you make each choice, and what are the possibilities after?
2. Sketch the resulting tree, carefully distinguishing your choice nodes from Nature nodes (where messages or states are revealed).
3. Fill in the probabilities of branches emanating from each N-node, using Bayes Theorem as necessary.
4. Fill in the final payoffs at terminal nodes (the most remote branch tips), contingent on all prior choices and events.

B. Then solve the decision tree, using Backward Induction as follows.

1. Solve each node just prior to terminal nodes.
 - If it is a N-node, then take the expected payoff over all branches.
 - If it is a C-node, then take the maximum payoff over all branches.
2. Replace that node and all its branches by a terminal node whose payoff is the solution just computed.
3. Iterate the previous steps on the truncated tree, keeping track of where each node solution came from. This is called induction.
4. When you get to the initial node you are done, and you know the expected payoff for the optimal plan. Write out that plan, i.e., the node-by-node solution to the decision tree.

C. Writing and solving trees fluently takes practice. Try a few practice problems and try your own examples.

D. So far we have represented the final payoffs as profits, and we will continue to do so for the sake of simplicity. But this implicitly assumes risk neutrality, and the off-

cial theory says that the payoffs really are the utility of the final outcomes. [DF's personal opinion: In most practical applications, the extra layer of complication is not worth the trouble.]

E. The economic value of information is the expected payoff in the solved tree when the information is available minus the expected payoff in the solved tree for which the information is not available.

- You can think of Startup tree #2 as informed, and #1 as uninformed.
- In this case, the value of information is $V = 0.125 - 0.0 = 0.125$, or \$125,000.

F. Notice that information only has economic value to the extent that it leads you to alter a decision.

- In the Startup example, you would make the second stage investment following a good beta version, but without information you would not invest.
- In a different example where no message would alter your optimal plan, the economic value of information would be zero. (Of course, it might have some psychological value.)

V. Extensions of Decision Theory

A. The basic recipe is always the same: specify the sequence of information events and decisions to create a well-posed optimization problem, and solve the problem by backward induction.

B. Many decision problems involve trading off immediate payoffs against future payoffs.

1. Backward induction here leads to the following general insight.
2. When solving a decision node, consider the “continuation value” obtained from downstream (already solved) nodes.
3. The optimal choice at each node maximizes the sum [current payoff plus continuation value].

C. Bellman's equation is a widely used tool in discrete time optimization.

1. It is a recursion equation that basically says the value at time t is the maximum among current choices of [current payoff plus continuation value].
2. It can be written for infinite horizon (i.e., the last period is not known from the outset) problems where backward induction cannot be done directly.
3. A few of numerous examples include
 - a. job search: at each stage, compare the payoff to accepting the best job currently on offer to the continuation value of continuing the search
 - b. deferral options: compare the payoff to making an irreversible investment now to the continuation value, knowing that better information will arrive in the meantime (this has the same flavor as Startup tree #2)
 - c. Financial call options: at each stage, compare the value of exercising the option (to purchase the underlying financial asset at the given strike price) to the continuation value of waiting (possibly enabling the underlying price to go higher and thus provide a higher net payoff).

D. A standard technique to evaluate numerically the value of financial options is to build a large but finite binomial event tree

- Each N-node is has 2 branches: the underlying asset price goes up or down by a small percentage (say .01) with known probability.
- Each C-node also has 2 branches: exercise the option or let it run.
- You solve the tree exactly as above.
- Example: MS thesis of Cagney and Oh, about 20 years ago.