Decisions Under Uncertainty

Geoffrey Churchill and Thomas Whalen

Robinson College of Business
Georgia State University
Atlanta, Georgia 30303-3083 USA

Ecclesiastes 9:11 "The race is not to the swift or the battle to the strong, nor does food come to the wise or wealth to the brilliant or favor to the learned; but time and chance happen to them all."

I. Elementary Decision Analysis

It's an inescapable fact of the human condition that what we receive, for good or ill, depends in part on what actions we freely choose to take and in part on circumstances beyond our control. "Analysis" means to distinguish the parts of something inordet to better understand the whole. Decision Analysis is based on making a clear conceptual distinction between actions we freely choose to take, and circumstances that are entirely outside of our control.

In other words, if something is partly controllable, decision analysis needs to know what part is controllable and what part is not. Thus, every Decision Analysis model must clearly specify three parts:

1. A set of **Alternative Actions**: we may choose whichever one of these we please
2. A set of possible **States of Nature**: one of these will turn out to be the real one, but we have no control as to which one
3. A set of **Outcomes**, one for each combination of an Alternative Action and a State of Nature, and a **Value**, monetary or otherwise, for each Outcome.

Decision Analysis specialists often divide decision making situations into 3 categories:

- **Certainty**: only one possible State of Nature
- **Ignorance**: several possible States of Nature
- **Risk**: several possible states of nature with an estimate of how probable each one is.

which is a useful oversimplification. In reality, though, there is more to it than that. We also restrict our discussion of Decision Analysis techniques cases where the decision maker has a countable set of alternative choices. The decision maker also faces a countable set of so-called **States of Nature**. Think of a State of Nature as one possible set of circumstances that the decision maker will discover, after committing to an Alternative Action, which determines the outcome of the decision. The decision maker *always* selects the **Alternative**. The decision maker *never* selects the State of Nature. Nature does that.

Thus in a problem that we might approach through Decision Analysis, the decision maker faces a finite set of choices, usually labeled as Alternatives. After the decision maker has selected an Alternative, the State of Nature reveals itself and the decision maker receives the corresponding Outcome. Each outcome leads to a particular Value, which we may treat as a dollar figure or which we may express in terms of Utility. Decision Analysis provides criteria and techniques for selecting alternatives.

What is the difference among Certainty, Ignorance, and Risk? They are distinctions among possible states of knowledge on the part of the decision maker with respect to what the "true" State of Nature is, and therefore what outcome will result from selecting a given alternative. In all 3 cases, we assume that the decision maker can enumerate all possible States of Nature. What the decision maker cannot know is what State of Nature will reveal itself following a decision. We assume that the decision maker always knows what set of States of Nature could occur.

Certainty In Decision Making Under Certainty (**DMUC**), we assume that the decision maker always knows with absolute certainty what the State of Nature will be, and knows it in time to make the best choice. Knowing is **not** the same as controlling. We use DMUC as a baseline, in much the same way as we used the naive forecast as a baseline.

Ignorance: Decision Making Under Ignorance (**DMUI**) assumes the exact opposite state of knowledge from DMUC. Here the decision maker can enumerate the possible States of Nature, but is not committed to any estimate of how probable or improbable any state is. Some authors call this "uncertainty" rather than "ignorance," but this is confusing since "uncertainty" also means "anything other than certainty."

Risk Decision Making Under Risk (**DMUR**) involves an in-between state of knowledge. The decision maker does not know what State of Nature will occur, but can assign a probability to each State of Nature. These probabilities may be *objective* probabilities rigorously derived from past data. They also could be *subjective* probabilities based only on the decision maker's cumulative experience with and understanding of such situations.

1.1 Two Kinds of "Easy" Decisions

Certainty: consider only one state of nature to be possible; the payoff matrix has only one column of payoffs, and the best decision is the one with the best payoff. Linear programming is an example of a technique for making decisions under the assumption of certainty.

Stochastic Dominance: If action A has a better payoff than action B under each individual state of nature, then we say that action B is stochastically dominated by action A. If the payoff matrix truly represents every thing the decision maker hopes (or fears) to receive from the decision in question, then no rational decision maker will ever choose to perform action B.

1.2 What is a Decision Rule?

In the general case, we must decide among two or more alternative actions that are not stochastically dominated, and more than one state of nature is considered possible. The payoff matrix assigns each alternative action a row of possible payoffs.

A decision rule converts this row of payoffs into a single number that somehow represents the whole row. Once this is done, it is easy to pick the alternative action whose representative number is better than the other alternative actions' representative numbers.

We look first at DMUI, and then quickly go on to DMUR. Other than our use of it as a baseline, DMUC is a no-brainer in the Decision Analysis context. Of course some other decision aids, e.g. linear programming, are forms of DMUC, but that's a horse of another cowboy.

1.3 Decision Rules: Greed and Fear Under Ignorance

Under ignorance, we can identify the possible States of Nature, but do not know their relative likelihoods. There are four commonly considered criteria for Decision Making Under Ignorance. Each makes a different assumption about the decision maker's state of mind.

Pure Greed: The Maximax Rule "Go for the gold!" Select the *alternative* that, if things turn out for the *best*, provides the highest payoff. In other words, **Maxi**mize the **max**imum payoff. Either the decision maker is extremely optimistic, or he or she is in such dire straits that any but the best payoff is of no value. Remember that you *assume* that you will get the best State of Nature given your chosen alternative, but you don't really get to choose it and there is ***no*** guarantee that it will happen.

Pure Fear: The Maximin Rule "Cover your assets" Select the *alternative* that, if things turn out for the *worst*, provides the highest payoff. **Maxi**mize the **min**imum payoff. This is a very conservative and pessimistic criterion. Nature is out to get me, and I must protect myself. Even we paranoids have enemies.

Combining Greed and Fear: The Hurwicz Rule "Seek a balance" This is a compromise between the maximax and maximin rules; simple pick an "optimism" coefficient α, and represent each alternative by α times its best payoff plus $1-\alpha$ times its worst payoff.

Fear of Gult: The Minimax Regret Rule "I'd hate myself if it went wrong" Select the alternative that will **Mini**mize the **max**imum regret. Here, we define "regret" as being the opportunity loss of having selected a choice other than the one that turns out to be the best. This decision maker can accept some risk, but doesn't want to look bad after the fact. Perhaps someone is Monday morning quarterbacking him or her. In effect, regret focuses attention where our actions have the most effect.

Equal Averaging: The Laplace-Bayes role, or " Principle of Insufficient Reason " Named for 2 great early theoreticians, this is sometimes called the Principle of Equal Likelihood. If you have no idea that any State of Nature is more probable than any other, this criterion holds that this is equivalent to saying that all are equally probable. In that case, you just compute the average payoff across all States of Nature for each alternative, and select the *alternative* with the best average! Note that this rule, unlike the previous ones, assumes that there is no connection between how likely something is and how likable it is.

Example: An Auto Race Promoter's Problem

Roger Pinky, a talented and wealthy businessman, has committed to promote an IndyCar race at Road Alpharetta next March. Roger would have preferred a date later in the Spring, but this was the best date available considering Road Alpharetta's and the IndyCar Series' schedules. He estimates that it will cost him $2,000,000 to put on the race, plus an average variable cost per spectator of $10. On a warm, dry day, he estimates that he will draw 62,500 spectators the first year. Of course, if it is cold and wet, he won't do as well; he figures he might get 25,000 hard core fans. Cold and dry would improve on that by 10%. Since rain races can be very dramatic, if it is wet but warm he can probably draw 30% more fans than on a cold wet day. Including tickets and his cut of concessions and souvenirs, he figures he will bring in $75 from the average spectator. Parking at Road Alpharetta is plentiful and free, so that won't bring in any revenue.

Of course, not even Roger Pinky can control the weather. Next March any of the 4 states of nature might happen. There *are* some things Roger might do to alter the scenario.

MARTA (Metropolitan Alpharetta Random Transit Authority) has offered him a transportation deal that is hard to refuse. For a mere $500,000 MARTA would provide free (to the rider) 2 way transportation between the track and essentially any point served by MARTA, all weekend. Roger figures, since folks like to drink and raise hell at the races, this might draw a lot of people who would rather not have a DUI on their license. On a dry day, he estimates that it would boost attendance by 10%. On a wet day, when people risk getting their cars stuck in the infield mud, it's probably worth a 40% boost in attendance. That would really help cut the risk from rain.

IndyCars have never really pulled big crowds in the South; this is NASCAR country. NASCAR has offered him the possibility of a Busch Grand National taxicab race for a total cost to him of $500,000. Roger is tempted. It might be a way of educating some NASCAR fans about IndyCars, and he thinks that a BGN support race might boost attendance 20%. It is worth considering.

So Roger's alternatives are to put the race on with or without MARTA and with or without a BGN support race. He sees the possible attendance outcomes like this:

ATTENDANCE	Cold	Cold	Warm	Warm
	Wet	Dry	Wet	Dry
No Busch, no Marta	25,000	27,500	32,500	62,500
Busch, no Marta	30,000	33,000	39,000	75,000
Marta, no Busch	35,000	30,250	45,500	68,750
Marta and Busch	42,000	36,300	54,600	82,500

That would give Roger this revenue picture:

REVENUES	Cold	Cold	Warm	Warm
@ $75 per ticket	Wet	Dry	Wet	Dry
No Busch, no Marta	$1,875,000	$2,062,500	$2,437,500	$4,687,500
Busch, no Marta	$2,250,000	$2,475,000	$2,925,000	$5,625,000
Marta, no Busch	$2,625,000	$2,268,750	$3,412,500	$5,156,250
Marta and Busch	$3,150,000	$2,722,500	$4,095,000	$6,187,500

This looks nice until we realize that there are, after all, some expenses that Roger will have to meet. He has to cover the basic $2 million. The two million is a fixed cost that we could have omitted without any effect on the decision, but it is useful to include it to show the distinction between maximizing gains and cutting losses. He must also pony up another half million bucks each for MARTA and the BGN support race if he decides to cut them in on the deal, and the approximate $10 cost for each spectator. That adds up to a considerable bundle:

EXPENSES	Cold	Cold	Warm	Warm
	Wet	Dry	Wet	Dry
No Busch, no Marta	$2,250,000	$2,275,000	$2,325,000	$2,625,000
Busch, no Marta	$2,800,000	$2,830,000	$2,890,000	$3,250,000
Marta, no Busch	$2,850,000	$2,802,500	$2,955,000	$3,187,500
Marta and Busch	$3,420,000	$3,363,000	$3,546,000	$3,825,000

When you deduct the expenses from the revenues, you get this picture, known as a ***Payoff Table***:

NET PAYOFF	Cold	Cold	Warm	Warm
	Wet	Dry	Wet	Dry
No Busch, no Marta	($375,000)	($212,500)	$112,500	$2,062,500
Busch, no Marta	($550,000)	($355,000)	$35,000	$2,375,000
Marta, no Busch	($225,000)	($533,750)	$457,500	$1,968,750
Marta and Busch	($270,000)	($640,500)	$549,000	$2,362,500

It's pretty clear that if it is a cold day, Roger takes a cold bath. Roger's a big boy; he can take it. He didn't get rich by running from risks, but he didn't do it by making careless decisions, either. Let's see if Decision Making Under Ignorance can throw some light on the situation for Roger.

Using a Decision Rule to Simplify a Payoff Matrix

Each row of the table represents the possible outcomes to the decision maker of one Alternative Action he or she might choose. Each of the criteria below begins by reducing each row to one number that is, in that criterion's sense, representative of the whole row. Then we select the row whose number is "best" according to the criterion. The choice or decision corresponding to that row is then the one the decision maker should select if he or she wants to follow the dictates of the particular criterion.

In Roger's payoff table, each alternative corresponds to a row of 4 different outcome values, one for each possible State of Nature. The decision is a hard one because each alternative is better than all the others under some States of Nature, but worse under other States of Nature. The goal of each criterion is to reduce each row of possible values into a single number which somehow stands for the whole row. If Roger could be satisfied that the single number for each alternative was really representative of the whole row of possibilities, then it would be a simple matter to pick the alternative whose number was best.

LaPlace-Bayes

By far the most familiar and common way to turn a collection of numbers (such as the possible payoffs from selecting an alternative) into a single representative value is to take the simple average. For example, the first alternative, No Busch, No MARTA, has possible payoffs of -$375,000, -$212,500, $112,500, and $2052,500. The average of these is $396,875. The decision criterion that compares alternatives on the basis of the simple arithmetic mean is called the **LaPlace-Bayes** criterion.

Roger is a rationalist. The argument for the **LaPlace-Bayes** criterion is very appealing. If he really doesn't have an idea what the weather will be like, then that tells him that he does not think that warm is more or less probable than cold, nor does he think that dry is more or less probable than wet. That says to him that they must be equally probable. (In fact, that's probably what he had in mind, explicitly or implicitly, when he set up the four categories of weather; "cold" means, roughly, the coldest 50% of possible weather conditions for the date in question, and similarly for "Warm," "Wet," and "Dry.') Technically, this calls for a probability weighted average (an expected value) where all weights (probabilities) are equal. Averaging across all the possible states of nature, the LaPlace-Bayes criterion tells him to go for broke. By this criterion, he should put up the extra million bucks and get both the support race and the free spectator rides.

NET PAYOFF	Cold	Cold	Warm	Warm	
	Wet	Dry	Wet	Dry	Mean
No Busch, no Marta	($375,000)	($212,500)	$112,500	$2,062,500	$396,875
Busch, no Marta	($550,000)	($355,000)	$35,000	$2,375,000	$376,250
Marta, no Busch	($225,000)	($533,750)	$457,500	$1,968,750	$416,875
Marta and Busch	**($270,000)**	**($640,500)**	**$549,000**	**$2,362,500**	**$500,250**

Maximax

If Roger is in a really "up" optimistic mood, he might try the **Maximax** criterion. In the best of all possible worlds, Roger's IndyCar race will be on a warm dry day. In that case, the Busch Grand National Race helps his profits, but MARTA doesn't. He maximizes his maximum gain with the choice of a Busch Grand National support race, but no help from MARTA.

NET PAYOFF	Cold	Cold	Warm	Warm	
	Wet	Dry	Wet	Dry	Maximum
No Busch, no Marta	($375,000)	($212,500)	$112,500	$2,062,500	$2,062,500
Busch, no Marta	**($550,000)**	**($355,000)**	**$35,000**	**$2,375,000**	**$2,375,000**
Marta, no Busch	($225,000)	($533,750)	$457,500	$1,968,750	$1,968,750
Marta and Busch	($270,000)	($640,500)	$549,000	$2,362,500	$2,362,500

Maximin

If Roger thinks the world is out to get him, he might take a different approach. Or if he can't take a very big financial hit, that could make him conservative. Either way, this might incline him to use the **Maximin** criterion. In that case, he would examine each possible decision alternative in light of the worst thing that could happen to him if he selected that alternative. The worst that can happen if he chooses to go with neither the BGN race nor MARTA is that he loses $375 grand. That's not as "worst" as the worst that can happen with the other 3 possible choices.

This approach doesn't just apply to Roger. One important question in ***any*** business decision process should usually be "If worst comes to worst, how badly could this choice turn out?". This question forms the basis of the **Maximin** criterion, which we also, not too surprisingly, call **Worst Case**

Analysis. In this model, we represent the collection of possible payoffs for any action by the *worst* payoff in the set.

In Roger's problem, No Busch, No MARTA, has possible payoffs of -$375,000, -$212,500, $112,500, and $2052,500, so the representative payoff under **Maximin** is the worst one, -$375,000. Note that the worst possible payoff for some alternatives corresponds to the State of Nature Cold, Wet while the worst case for the other alternatives arises from the State of Nature Cold, Dry. We represent each alternative with the payoff (and, implicitly, the State of Nature) that is worst for that alternative independently of any other alternative. Among the 4 alternatives available to Roger, No Busch, No MARTA has a worst case that is less awful than the worst case for any other alternative. Thus No Busch, No MARTA is the action selected by the **Maximin** criterion. It is the *best of the worst*.

In decisions where the State of Nature is actually the response of a malevolent opponent who is "out to get you", the **Maximin** criterion is often the best way to play the game. But it is also an important element of many other decision processes where it is important to ***cover your assets***.

Maximin	Cold	Cold	Warm	Warm	
	Wet	Dry	Wet	Dry	Minimum
No Busch, no Marta	**($375,000)**	($212,500)	$112,500	$2,062,500	**($375,000)**
Busch, no Marta	**($550,000)**	($355,000)	$35,000	$2,375,000	($550,000)
Marta, no Busch	($225,000)	**($533,750)**	$457,500	$1,968,750	($533,750)
Marta and Busch	($270,000)	**($640,500)**	$549,000	$2,362,500	($640,500)

Hurwicz

If Roget decides to "split the difference' between the maximin and the maximax decision rule, his indicated choice is "Busch, no MARTA."

Hurwicz	Cold	Cold	Warm	Warm	(MAX+
	Wet	Dry	Wet	Dry	MIN)/2
No Busch, no MARTA	**-$375,000**	-$212,500	$112,500	**$2,062,500**	$843,750
Busch, no MARTA	**-$550,000**	-$355,000	$35,000	**$2,375,000**	**$912,500**
MARTA, no Busch	-$225,000	**-$533,750**	$457,500	**$1,968,750**	$717,500
MARTA and Busch	-$270,000	**-$640,500**	$549,000	**$2,362,500**	$861,000

Minimax Regret

The LaPlace-Bayes criterion treats all States of Nature as being equally likely to happen, the Maximin criterion emphasizes States of Nature that are unfavorable, and the Maximax criterion emphasizes States of Nature that are favorable. Another criterion, **Minimax Regret**, strives to put the emphasis on the States of Nature where *our actions make the most difference*. In order to do this, we must define a new quantity known as ***regret***, which we also often refer to as *opportunity loss.* (Be sure you don't confuse opportunity loss with the quite different economic concept of opportunity cost.)

The first step in calculating regret is to indulge in a little wishful thinking. Suppose Roger could magically obtain an envelope containing an *absolutely infallible* weather forecast for race day. Let's assume that the envelope is available only ***after*** he has absolutely committed to holding the IndyCar race - he can't cancel out completely. His best course of action, which we represent by an imaginary alternative called ***Ideal***, would be to open the envelope and choose the best ***real*** alternative based on the forecast he finds in the envelope. If the envelope forecasts Cold, Wet weather, he chooses MARTA, No Busch and loses $225,000. If the forecast is for Cold, Dry weather, he will select No Busch, No MARTA and hold his losses down to $212,500. If his forecast is for Warm, Wet weather he will select MARTA and Busch and earn $549,000. If he is really lucky and gets a forecast of Warm, Dry weather, he will select Busch, No MARTA and hit the $2,375,000 jackpot!

The best choice for each state of nature is easy enough to identify. Just don't get confused and start thinking that you get to pick the state of nature you like. Or that you choose the best outcome for each decision (really the same thing).

PAYOFF	Cold	Cold	Warm	Warm
	Wet	Dry	Wet	Dry
No Busch, no MARTA	($375,000)	**($212,500)**	$112,500	$2,062,500
Busch, no MARTA	($550,000)	($355,000)	$35,000	**$2,375,000**
Marta, no Busch	**($225,000)**	($533,750)	$457,500	$1,968,750
Marta and Busch	($270,000)	($640,500)	**$549,000**	$2,362,500
Ideal	($225,000)	($212,500)	$549,000	$2,375,000

To get the **regrets** (we also call them opportunity losses), all we have to do is subtract each payoff from the best payoff for its state of nature and collect the results in a Regret Matrix. The **Minimax Regret** criterion then tells us to represent the row of regrets corresponding to each Alternative Action by the highest (i.e. the worst) regret in the row, and select the action whose greatest regret is smaller than the greatest regret for any other action. If Roger buys the idea of minimizing his maximum regret, then he will go with free transportation for spectators, but drop the idea of a Busch series race.

REGRET	Cold	Cold	Warm	Warm	Max
	Wet	Dry	Wet	Dry	Regret
	$150,000	$0	**$436,500**	$312,500	$436,500
Busch, no MARTA	**$325,000**	$142,500	$514,000	$0	$514,000
MARTA, no Busch	$0	$321,250	$91,500	**$406,250**	**$406,250**
MARTA and Busch	$45,000	**$428,000**	$0	$12,500	$428,000

So there we have it. Decision Making Under Ignorance can give very different solutions, depending on what criterion you choose. And what criterion you should choose depends on your situation. Properly applied, any one of these criteria makes more sense than a random choice. But if you select the criterion at random, you might as well have made your choice at random.

Decision Making Under Risk

In most situations, the idea of Decision Making under Risk makes a lot more sense than Decision Making under Ignorance. There are few situations where someone has the responsibility to make a decision of great consequence with no knowledge of the relative likelihood of the outcomes involved and no way to acquire that knowledge. Without some effort, that knowledge of how probable the various states of nature may not be explicit or formalized, but it is present or available.

In DMUR, we must assign a probability to each State of Nature. In some cases, research will reveal historical relative frequency information that we might conclude reveals the underlying probabilities. In other cases, there are formal models designed for the job. The Weather Bureau uses such a model to arrive at the probability of thundershowers for tomorrow. Formal probability models for Blackjack (also known as 21) ended up completely changing the way casinos run their games.

Often, the decision maker must rely on subjective probabilities. Don't dismiss this as mere intuition. If you are familiar with the domain in which the problem arises, you can assess very useful, if imperfect probabilities. You do this informally every day. I doubt that there's a student alive who doesn't jaywalk. You, in effect, evaluate the probability of getting hit before you step off the curb. If that probability is too great, you don't step off. And every now and then, somebody gets hit. Low probability events occur, even if infrequently.

In fact, there is a sense in which all probabilities are subjective, since the only absolutely "objective" probability that something will happen is 1 if it happens, or 0 if it doesn't. Any other probability is an expression of the incomplete information possessed by a particular decision maker at a particular point in time.

Why do we need probabilities? DMUR is based on using probability weighted averages as the basis for choice. To compute those averages, we must have probabilities.

Expected Monetary Value (EMV)

The **EMV** criterion is fundamental to DMUR.

S_i The i^{th} State of Nature from our list of possible States of Nature.

A_j The j^{th} Action from our list of possible decisions.

$P(S_i)$ The probability that S_i will occur.

V_{ij} The Payoff to the decision maker if he or she chooses A_j and S_i occurs.

EMV_j Expected Monetary Value. The long term average payoff if we could repeat the same decision many times under the same circumstances.

$$EMV_j = \Sigma_{i=1}^{n} P(S_i) \times V_{i,j}$$

According to the EMV criterion, we should choose the decision alternative which has the highest EMV associated with it. We'll call this EMV the Expected Value Under Initial Information (EVUII) to distinguish it from what the EMV might become if we later get more information. ***Do not*** make the common student error of believing that the EMV is the payoff that the decision maker will get. The actual payoff will be the $V_{i,j}$ for that alternative (j) and for the State of Nature (i) that actually occurs.

Expected Monetary Value is a ***long run*** criterion. The theory holds that a decision maker who consistently applies the EMV criterion will do better in the long run than one who does not. This is sound theory, *provided* that the amount of money at risk is small relative to the decision maker's total wealth. (Note that in this context, the decision maker could be a large corporation, represented by the individual who actually makes the decision.) If the amount at risk is substantial, then we need to consider the possibility that a few bad outcomes could drive total wealth to zero. When total wealth is zero, you no longer have access to good outcomes. This concept is sometimes called the case of Gambler's Ruin.

We can evaluate the risk associated with an Alternative Action by calculating the variance of the payoffs. Depending on your willingness to accept risk, an Alternative Action with only a moderate EMV and a small variance may be superior to a choice that has a large EMV and also a large variance. The variance of the payoffs for an Alternative Action is defined as

$$Variance_j = \Sigma_{i=1}^{n} P(S_i) \times (EMV_j - V_{i,j})^2$$

Given Action j, we square the difference between each payoff and the EMV, weight it with the probability of the State of Nature that leads to that payoff, and add up those products. Most of the time, we wand to make EMV as large as possible and variance as small as possible, Unfortunately, the maximum-EMV alternative and the minimum-variance alternative are usually not the same, so that in the end it boils down to an educated judgment call. Since, in the context of payoffs in the $millions, variances can be humongous numbers, we very commonly end up scaling them down by dividing all of them by the same large number.

Expected Value of Perfect Information

This is where we make use of the idea of Decision Making under Certainty. Suppose that although you couldn't *control* the future, you could *foresee* the future with perfect accuracy. That's what Certainty is all about. If you could foresee the future perfectly, then you would have Perfect Information. Perfect Information is certainly the very best kind of information. Imperfect Information can't be as valuable. It is certainly too bad that perfect information doesn't really exist.

If perfect information ***did*** exist, it would probably be a lot like the envelope we imagined Roger to have when we discussed Minimax Regret. Only now that we are considering probabilities, we now know the probability that the envelope will perfectly predict any given State of Nature. Of course we still don't know what's in the envelope until decision time arrives. We'll return to the envelope idea when we come back to Roger's problem.

In many situations, you can at least get *bette*r information before you commit to an Alternative Action, often at a price. Before tooling up to produce a radical new product, you can do consumer research using prototype units. Before buying land for a plant, you can survey its neighbors to learn whether they will oppose your plans. Race tire manufacturers often rent a race track for a day so that teams can test to see if a new tire design is actually faster than the existing model.

The **E**xpected **V**alue of **P**erfect **I**nformation (**EVPI)** provides an absolute upper limit on the value of additional information (ignoring the value of reduced risk). It measures the amount by which you could improve on your best EMV if you had perfect information. It is the difference between the **E**xpected **V**alue **U**nder **P**erfect **I**nformation (**EVUPI**) and the EMV of the best action (EVUII).

$$EVUPI = \Sigma_{i=1}^{n} P(S_i) \times \max_{j=1}^{m} (V_{i,j})$$

$$EVUII = \max_{j=1}^{m} (EMV_j)$$

$$EVPI = EVUPI - EVUII$$

These formulas look worse than they are. The EVUPI formula simply says to go through the possible States of Nature. For each State of Nature, find the best Payoff. Multiply it by the probability of that State of Nature happening. Sum the products. EVPI is even easier. Subtract the EVUII from the EVUPI. The result is the EVPI.

Expected Value of Perfect Information measures how much better you could do on this decision, averaging over repeating the decision situation many times, if you could always know what State of Nature would occur, just in time to make the best decision for that State of Nature. Remember that it does ***not*** imply control of the States of Nature, just perfect prediction. Remember also that it is a long run average. It places an upper limit on the value of additional information.

Back to Roger's Problem

Roger Pinky is not actually in a state of ignorance. He is in a state of risk. He has a pretty good idea what the weather can be like at Road Alpharetta in March. He's not worried about snow. If it snows, he gets to put on the event the next weekend. His concern is just with rain and with cold. He had a subordinate call the Weather Bureau to get the facts. In the last 100 years, the weekend he has the event scheduled was a cold, wet one 10 times, a cold, dry one 15 times, a warm, wet one 40 times, and it was a warm, dry weekend 35 times. He is willing to assume that those figures tell him the probabilities of the events.

He'll try the EMV criterion to see what it tells him.

NET PAYOFF	Cold	Cold	Warm	Warm		Variance/
	Wet	Dry	Wet	Dry	EMV	1,000,000
No Busch, no MARTA	($375,000)	($212,500)	$112,500	$2,062,500	$697,500	1,028,259
Busch, no MARTA	($550,000)	($355,000)	$35,000	$2,375,000	$737,000	1,480,694
MARTA, no Busch	($225,000)	($533,750)	$457,500	$1,968,750	$769,500	**895,980**
MARTA and Busch	($270,000)	($640,500)	$549,000	$2,362,500	**$923,400**	1,290,211
Probability	0.1	0.15	0.4	0.35		

It turns out that EMV tells him several interesting things. The combination of MARTA and a Grand National support race has the highest Expected Monetary Value by a healthy margin. But using MARTA with no Busch race is the least risky choice by a substantial amount. He will probably go with the low variance choice. The first two choices (the "plain vanilla" race and the Busch, no MARTA option) are clearly inferior to the MARTA, no Busch option since they have lower EMV and higher variance (less average payoff and more risk).

If Roger does go with MARTA and Busch, remember that he won't make **$923,400**. He either will lose $270,000 or $640,500 ***or*** he will make $549,000 or $2,362,500. Choosing via EMV is "going with the odds". Those who "go with the odds" come out ahead in the long run, *if* they can afford to stay in the game. That's how the casinos at Lost Wages make their money.

Roger has to accept some risk, but he doesn't have to like it. If he can reduce the risk at an affordable cost, he certainly would consider it. That's where the EVPI comes in. The EVPI can tell him the most he should consider paying for new information about the weather next March. And that would be for perfect foreknowledge of the weather, something he isn't going to get in this lifetime. Clearly, better information about next March's weather is valuable, but not that valuable. And he has no motive to pay the full value of the information. He'll want to pay less so that his bottom line can improve.

Let's turn our thoughts back to the imaginary case where Roger can get a sealed envelope containing a perfect weather forecast for race day. How much better off would he be with the envelope than without it? While the envelope remains sealed, Roger doesn't know what it contains. He *does know* that the probability that it contains a Cold, Wet forecast is 0.1, that the probability that it contains a Cold, Dry forecast is 0.15, that the probability that it holds a Warm, Wet forecast is 0.4, and that the chance that it has a Warm, Dry forecast is 0.35.

NET PAYOFF	Cold	Cold	Warm	Warm		Best
	Wet	Dry	Wet	Dry	EMV	EMV
No Busch, no MARTA	($375,000)	($212,500)	$112,500	$2,062,500	$697,500	
Busch, no MARTA	($550,000)	($355,000)	$35,000	$2,375,000	$737,000	
MARTA, no Busch	($225,000)	($533,750)	$457,500	$1,968,750	$769,500	
MARTA and Busch	**($270,000)**	**($640,500)**	**$549,000**	**$2,362,500**	**$923,400**	**$923,400**
Ideal	*($225,000)*	*($212,500)*	*$549,000*	*$2,375,000*	*$996,475*	***$996,475***
Probability	0.1	0.15	0.4	0.35	Difference	$73,075
EVUPI=	$996,475	EVUII =	$923,400		EVPI=	$73,075

If Roger could get the magic envelope, he would be, on average, $73,075 better off than he is in the real world. This is because his expected return without the magic envelope is $923,400 while with it his average return is $996,475, or $923,400 + $73,075.

Roger's Expected Value Under Perfect Information is $996,475. Exactly what does that mean? **IF** it could really happen this way, it means that if it is going to turn out to be:

a cold, wet weekend then he will know in time to choose MARTA, no Busch and only lose $225,000.

a cold, dry weekend then he will know in time to choose no Busch, no MARTA and only lose $212,500.

a warm, wet weekend then he will know this in time to choose MARTA and Busch, and clear $549,000.

a warm, dry weekend then he will find this out in time to choose Busch, no MARTA, earning $2,362,500.

On the average, this would be $73,075 better than doing without perfect information. This is his EVPI for this decision.

Expected Regret ("Expected Opportunity Loss")

The Expected Regret criterion, also known as the ExpectedOpportunity Loss criterion, is an alternative to EMV. Opportunity loss is the same as the Regret that we looked at under Ignorance.

S_i The ith State of Nature from our list of possible States of Nature.

A_j The jth Action from our list of possible decisions.

$P(S_i)$ The probability that S_i will occur.

V_{ij} The Payoff to the decision maker if he or she chooses A_j and S_i occurs.

OL_{ij} The Opportunity Loss if the decision maker chooses A_j and S_i occurs. This is zero for the best A_j for a given S_i.

EOL_{ij} Expected Opportunity Loss. The long term average opportunity loss if we could repeat the same decision many times under the same circumstances.

$$OL_{i,j} = \max_{i=1}^{n}(V_{i,j}) \quad - \quad V_{i,j}$$

$$EOL_j = \Sigma_{i=1}^{n} P(S_i) \times OL_{i,j}$$

The EOL criterion says that we should minimize the Expected Opportunity Loss. There are 2 interesting things to note about this. If we choose the Alternative Action with the lowest EOL, we choose the same Alternative Action that we would select if we chose the one with the highest EMV! Not only that, but the EOL for that Alternative Action is the EVPI. If you think carefully about this, it actually makes sense. But it is certainly not "intuitively obvious".

Let's look at Expected Opportunity Loss for Roger's problem.

OPPORTUNITY LOSS (Regret)	Cold Wet	Cold Dry	Warm Wet	Warm Dry	EOL *expected regret)	Best EOL
No Busch, no MARTA	$150,000	$0	$436,500	$312,500	$298,975	
Busch, no MARTA	$325,000	$142,500	$514,000	$0	$259,475	
MARTA, no Busch	$0	$321,250	$91,500	$406,250	$226,975	
MARTA and Busch	$45,000	$428,000	$0	$12,500	**$73,075**	**$73,075**
Probability	0.1	0.15	0.4	0.35		

I'll be doggoned! We made the same choice as before, and the minimum EOL is the same as the EVPI. Just like I said. The two criteria are equivalent. It's often a good idea to work a problem both ways as a check on your math; if the two answers don't agree, there's an error in there somewhere! Another useful check is that the EMV plus the EOL for each Alternative Action should add up to the same quantity, which equals the EMV of the imaginary "ideal" action.

Decision Trees

The Payoff Table and the Opportunity Loss Table are two very similar ways of looking at a Decision Analysis problem. Another way of seeing the structure of the problem is the ***Decision Tree***. For a "simple" single stage problem like Roger Pinky's problem, it is easy to do without decision trees. When we reach sequential decision problems, they will become virtually indispensable.

To draw a decision tree, you must figure out in what order the decision maker (here it's Roger) makes decisions (chooses alternatives) and has information (here, States of Nature) revealed to him or her. In Roger's case, it is pretty easy to figure that out. First Roger must decide on an alternative (No Busch, no MARTA, *or* Busch, no MARTA, *or* MARTA, no Busch, *or* MARTA and Busch). Then, when the weekend of his IndyCar race arrives, whatever weather happens, happens. It isn't always that simple.

It doesn't matter in what order things happen. A Decision Tree is driven by the order in which the decision maker has information and acts on it!

In Roger's case, the tree represents

1. Roger's decision on how to support his IndyCar race
2. The weather that happens.

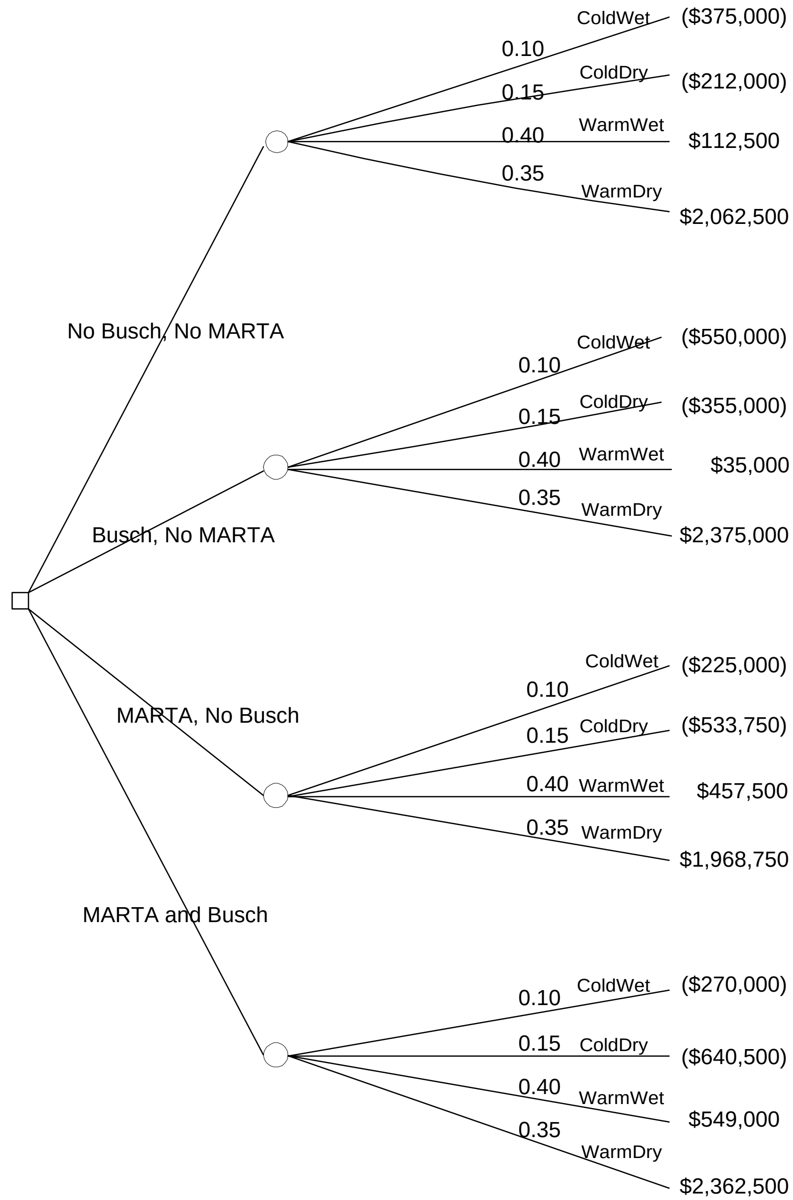

Suppose Roger could buy some sort of long range weather forecasting study. In that case the tree would need to represent

1. Roger's decision on whether to buy a long range weather forecast
2. If he buys the forecast, the results of the forecast

3. With or without the forecast, his race support decision
4. The weather that happens.

If you draw the tree in any sequence other than this one, you get stuck. You can't finish it.

Let's suppose that you do understand the order in which the decision maker has information and acts on it. There are some conventions we follow in constructing a Decision tree.

1. The tree begins at a single "node", usually a decision.
2. We show decision nodes as little boxes; branches from such nodes represent alternatives.
3. We show chance nodes as little circles. Branches from such nodes represent outcomes or States of Nature.
4. The outermost branches end at "terminal points", where we show the payoffs.
5. We label chance branches with their probabilities.
6. In fact, we label everything as clearly as possible.

The initial tree shows us the structure of Roger's problem, but it doesn't solve it. How do we solve the problem (again)?

First, we go to the outermost chance nodes, and we apply the probabilities to the payoffs to compute the EMV. Now that we have the EMV, we treat the node as a terminal point. That is to say, we will now act as though the EMV were an actual payoff for reaching that node. We will work our way from the "top" of the tree back toward the base until we reach a decision node. At each decision node, we select the one branch leaving that node which has the highest value. That highest value could be a "real" payoff, or it could be an EMV. We leave that branch intact, and "cut off" all other branches leaving that decision node. A pair of hash marks through the branch indicates that we have "cut off" the branch, along with any branches that follow it.

On the next page, you'll see Roger's finished tree. The Expected Monetary Value for the alternative MARTA and Busch is (again) $923,400, which is higher than for any other alternative. We have "pruned" the branches representing the other alternatives. The $923,400 EMV for that branch is also the EMV for the tree.

Notice that only ***decision branches*** get pruned. The only way to lose a chance branch is if it grows from a pruned decision branch. We prune ***all*** branches following the pruned decision branch with it.

The completed decision tree amounts to a ***decision rule***. In this case it is very simple decision rule. The decision rule simply says to support the IndyCar race with a Busch race and give free MARTA rides, then live with whatever weather you get.

.