Expected Value of Information
Brian Tomasik
12/31/2007

The "expected value of information" is colloquially defined as the expected value after learning a given piece of information minus the expected value before learning that piece of information. This wording is a little imprecise. For instance, suppose we currently think our best action has expected value +10 and then we learn new information that tells us that the new best action has expected value +5. Clearly, the value of the information wasn't -5. Below, I give a simple, explicit formula for the value of information.

For simplicity, suppose there are finitely many possible worlds, indexed by j, to which we assign nonzero probability. These possible worlds are complete descriptions of all the facts in a particular universe, so for instance, a world in which fish are sentient is a different world from one in which fish are not sentient. Let P(j) denote our subjective probability of being in possible world j. Let i be an index over our (assumed finite) set of possible actions, and let U(i, j) denote the utility of taking action i given that we're in world j. (Note that this is not an expected utility, because the possible world we're in gives a complete description of all facts, i.e., there's no uncertainty.) The expected utility of action i is

EU(i) := sum_j P(j) U(i,j).

Utilitarians choose the action

[http://en.wikipedia.org/wiki/Arg_max argmax]_i EU(i).

Suppose we're considering whether to expend resources to find out more information. If we knew for certain the outcome k of our information gathering, we could update our probability distribution over possible worlds. Let P(j,k) denote the probability of possible world j given information k. With this information, we compute new expected utilities:

EU(i,k) := sum_j P(j,k) U(i,j),

and we now choose action

argmax_i EU(i,k).

What have we gained from our information? Instead of taking action argmax_i EU(i), we take action argmax_i EU(i,k). (Of course, these may turn out to be the same action, in which case our information didn't provide any value). The expected value of this change in action is

EU(argmax_i EU(i,k),k) – EU(argmax_i EU(i),k),

i.e., we reevaluate the expected value of our old action argmax_i EU(i) based on our new information. This expression can be rewritten as

max_i EU(i,k) – EU(argmax_i EU(i),k).

Now, of course, we don't actually know what information we'll discover; the best we have is a subjective probability distribution over what that information will turn out to be (assume the number of possibilities is finite). Let Q_k be the probability of discovering information k. Then the expected value of acquiring new information is

sum_k Q_k [max_i EU(i,k) – EU(argmax_i EU(i),k)].