

Random Variables

A *random variable* arises when we assign a numeric value to each elementary event. For example, if each elementary event is the result of a series of three tosses of a fair coin, then $X = \text{“the number of Heads”}$ is a random variable. Associated with any random variable is its *probability distribution* (sometimes called its *density function*), which indicates the likelihood that each possible value is assumed. For example, $\Pr(X=0) = 1/8$, $\Pr(X=1) = 3/8$, $\Pr(X=2) = 3/8$, and $\Pr(X=3) = 1/8$.

The *cumulative distribution function* indicates the likelihood that the random variable is less-than-or-equal-to any particular value. For example, $\Pr(X \leq x)$ is 0 for $x < 0$, $1/8$ for $0 \leq x < 1$, $1/2$ for $1 \leq x < 2$, $7/8$ for $2 \leq x < 3$, and 1 for all $x \geq 3$.

Two random variables X and Y are *independent* if all events of the form “ $X \leq x$ ” and “ $Y \leq y$ ” are independent events.

The *expected value* of X is the average value of X , weighted by the likelihood of its various possible values. Symbolically,

$$E[X] = \sum_x x \cdot \Pr(X = x)$$

where the sum is over all values taken by X with positive probability. Multiplying a random variable by any constant simply multiplies the expectation by the same constant, and adding a constant just shifts the expectation:

$$E[kX+c] = k \cdot E[X] + c .$$

For any event A , the *conditional expectation* of X given A is defined as

$$E[X|A] = \sum_x x \cdot \Pr(X=x | A) .$$

A useful way to break down some calculations is via $E[X] = E[X|A] \cdot \Pr(A) + E[X|A^c] \cdot \Pr(A^c)$.

The expected value of the sum of several random variables is equal to the sum of their expectations, e.g.,

$$E[X+Y] = E[X] + E[Y] .$$

On the other hand, the expected value of the product of two random variables is not necessarily the product of the expected values. For example, if they tend to be “large” at the same time, and “small” at the same time, $E[XY] > E[X] \cdot E[Y]$, while if one tends to be large when the other is small, $E[XY] < E[X] \cdot E[Y]$. However, in the special case in which X and Y are independent, equality *does* hold: $E[XY] = E[X] \cdot E[Y]$.

Probabilities vs. Expectations

World Airways runs daily flights from Chicago to Tokyo. They face a fixed cost of \$80,000 for each flight (basically independent of the actual number of passengers), and there are 300 seats available on their planes. One-way tickets generate revenues (after agent commissions and the like are deducted) of \$500 apiece when used, but are fully refundable if not used. The Thursday flight typically has between 5 and 95 no-shows; all intermediate values are equally likely. WA is allowed (by law) to overbook its flights, but must give compensation of \$250 to all ticketed passengers not allowed to board, and must provide those passengers with alternative transportation (the cost of providing the alternative transportation just wipes out the \$500 revenue). How many tickets should they be willing to sell for the Thursday flight?

last passenger ticketed	probability seated	probability bumped	expected profit from sale
305	100.0%	0.0%	\$500.00
315	89.0%	11.0%	\$417.58
325	78.0%	22.0%	\$335.16
335	67.0%	33.0%	\$252.75
345	56.0%	44.0%	\$170.33 = 56% · \$500 - 44% · \$250
350	50.5%	49.5%	\$129.12
355	45.1%	54.9%	\$87.91
360	39.6%	60.4%	\$46.70
364	35.2%	64.8%	\$13.74
365	34.1%	65.9%	\$5.49
366	33.0%	67.0%	(\$2.75)
370	28.6%	71.4%	(\$35.71)
375	23.1%	76.9%	(\$76.92)
380	17.6%	82.4%	(\$118.13)
385	12.1%	87.9%	(\$159.34)
390	6.6%	93.4%	(\$200.55)
395	1.1%	98.9%	(\$241.76)

How many no-shows do you expect?	50
How many tickets will you sell?	365
How many people do you expect to bump?	15

$$\begin{aligned}
 & E[\text{cost of uncertainty, on a day the plane could be filled}] \\
 &= \$250 \cdot E[\text{number bumped} \mid \text{bumping day}] \cdot \Pr(\text{bumping day}) \\
 &+ \$500 \cdot E[\text{number of empty seats} \mid \text{empty day}] \cdot \Pr(\text{empty day}) \\
 &+ \$0 \cdot \Pr(\text{perfect day}) \\
 &= \$250 \cdot 30.5 \cdot 60/91 + \$500 \cdot 15.5 \cdot 30/91 + \$0 = \$5027.47 + \$2554.95 + \$0 = \$7582.42
 \end{aligned}$$

Lesson: The optimal decision against the expected situation is typically *not* the same as the decision which maximizes expected payoff; the latter decision is the better one.

Marginal Cost-Benefit Analysis

Frequently, when there is a single policy variable q under a manager's control, conditions which must be satisfied by an optimal policy can be determined by (a) assuming that an optimal policy is under consideration, (b) considering a one-unit change in that policy, and (c) equating the marginal costs and benefits of that change.

For those who know calculus: Profit (q) = Benefit (q) - Cost (q). Optimally, Profit'(q) = 0 = Benefit'(q) - Cost'(q), i.e., marginal benefit equals marginal cost when q is set optimally.

For those who don't know calculus: If a one-unit increase in q yields greater benefit than cost, the original policy could not have been optimal. If a one-unit increase yields greater cost than benefit, then a one-unit decrease will yield greater benefit than cost, and again the original policy could not have been optimal.

Example 1: The single-period inventory problem. Assume that a perishable commodity is being ordered in the face of uncertain demand. Let D represent the random demand, and let C , P , and S be, respectively, the purchase cost, selling price, and salvage value of a unit. (That is, the cost of failing to meet a unit of available demand is $c_{\text{under}} = P - C$, and the cost of providing for a unit of supply, and finding no matching demand, is $c_{\text{over}} = C - S$.) Assume that Q is the optimal order quantity, and consider increasing this order to $Q + 1$ units.

$$\text{Marginal cost} = \Pr(D \leq Q) \cdot c_{\text{over}}$$

$$\text{Marginal benefit} = \Pr(D > Q) \cdot c_{\text{under}}$$

Equating marginal cost with marginal benefit, $\Pr(D \leq Q) = c_{\text{under}} / (c_{\text{under}} + c_{\text{over}})$. (This last term is called the *critical fractile*.)

Example 2: Sequencing. You have n tasks to complete, which will require t_1, \dots, t_n units of time. Until task j is completed, you accrue costs of c_j per unit time. In what order should you carry out the tasks?

Consider any ordering which puts task i immediately in front of task j . Consider switching the order of these two tasks, while keeping all others in place.

$$\begin{aligned} \text{Marginal cost of switch} &= c_i t_j && \text{(task } i \text{ gets completed } t_j \text{ units of time later)} \\ \text{Marginal benefit} &= c_j t_i && \text{(task } j \text{ gets completed } t_i \text{ units of time earlier)} \end{aligned}$$

Switching is *bad* if $c_j t_i < c_i t_j$, i.e., if $t_i / c_i < t_j / c_j$. To minimize total delay cost, compute the processing-time / delay-cost ratio for each task, and always give priority to the task with the lowest ratio.

Example 3: Meeting due dates. You have n tasks to complete, each of which has a fixed due-date. Can you get them all done on time?

Consider any ordering which puts task i immediately in front of task j , and switch the order of these two tasks, while keeping all others in place. If task j is due before task i , and if they were

both being completed on time in the original order, then they'll still both be done on time in the new order (since task j is now finished earlier, and task i, while finished later, is still finished at the same time task j originally was, and task j, due before task i, was getting completed on time in the original order).

Continuing this switching process, you'll eventually sort the tasks into earliest-due-to-latest-due order. Therefore, if there's *any* way to get all of the tasks done on time, doing them in order of their due dates will succeed.

Example 4: The classical EOQ formula. Consider a long-term policy of ordering Q units per cycle. If such a policy is optimal, then the marginal cost of shifting to an order size of $Q + 1$ in the first cycle, and then reverting to an order quantity of Q in all subsequent cycles, should be equal to the marginal benefit of such a shift.

Marginal cost = $c_H Q/D$ (the $(Q+1)$ -st unit must be carried for Q/D units of time)

Marginal benefit = $[c_S D/Q + c_H Q/2]/D$ (the cost of ordering and carrying the unit in a later cycle is avoided)

Equating marginal cost and marginal benefit, $Q = \sqrt{(2c_S D/c_H)}$.

Example 5: One-period discounts. Assume that a per-unit price discount of r is available for the current order cycle, after which per-unit price will rise to and remain at p . Let Q^* be the economic order quantity at a price of p . Assume that Q is an optimal order quantity for the one-time discounted price, and consider increasing this order to $Q + 1$ units. Let F be the relevant inventory fraction.

Marginal cost = $(p-r)F \cdot Q/D$

Marginal benefit = $r + p F \cdot Q^*/D$ (a purchase-price savings of r , together with the cost of ordering and carrying the marginal unit in a later cycle; $p F \cdot Q^*$ is the annual inventory-related cost of ordering Q^* per period)

Equating marginal cost with marginal benefit, $Q = rD/(p-r)F + (p/(p-r))Q^*$.

Example 6: Imputed stockout cost. Assume that the principal cost of stocking out is a cost of c_B per unit of demand during the stockout period. Let SS be the optimal amount of safety stock to keep, and consider increasing this to $SS + 1$.

Marginal cost (per year) = c_H (the extra unit is carried throughout most of the year)

Marginal benefit (per year) = $c_B \cdot \text{Pr}(\text{stockout in a cycle}) \cdot D/Q^*$
(each time a stockout occurs, one less unit of demand goes unsatisfied)

Equating marginal cost with marginal benefit, $\text{Pr}(\text{stockout}) = c_H Q^*/(c_B D)$.

The Akerlof “Lemons” Problem

An owner of a used car is negotiating with a prospective buyer. The quality of the car is known only to the seller; expressed in terms of the car's value to the seller, the buyer believes it equally likely to be worth any amount between \$0 and \$500. The buyer, who would utilize the car to a greater extent, would derive 50% more value from its ownership. At what price might a sale take place?

This example was discussed in class in order to introduce the general notion of *adverse selection*. To reprise the key points:

You are subject to adverse selection whenever

1. You offer to engage in a transaction with another party, and that party can either accept or refuse your offer.
2. The other party holds information not yet available to you concerning the value to you of the transaction.
3. The other party is most likely to accept the offer when the information is “bad news” to you.

In the example, (1) you made me an offer, (2) I knew how well the car had been maintained (and you didn't), and this was of relevance in determining the value of the car to you, and (3) I would accept your offer if the car was worth less than that to me, i.e., relatively poorly maintained, and reject your offer if the car was worth more to me., i.e., was relatively well-maintained.

When you are subject to adverse selection, your expected benefit from making any particular offer is

$$E[\text{benefit} \mid \text{offer is rejected}] \cdot \text{Prob}(\text{offer is rejected}) \\ + E[\text{benefit} \mid \text{offer is accepted}] \cdot \text{Prob}(\text{offer is accepted}),$$

and you must take into account that the second conditional expectation is typically less than

$$E[\text{benefit} \mid \text{other party is forced to accept the offer}].$$

The acceptance of your offer conveys information to you. In the used-car example, your offer of \$x would yield you an expected ultimate benefit of

$$(500-x)/500 \cdot 0 + (x/500) \cdot (E[\text{value of car to you} \mid \text{I say “yes”}] - x).$$

The critical observation was that, if I'd only accept your offer if it was at least as great as the value of the car to me, then the expected value of the car to you given that I say “yes” is only $1.5 \cdot (x/2)$, and therefore your expected benefit if the offer is accepted is negative for any positive offer. The spreadsheet on the class webpage shows how the optimal offer can be determined in less-extreme cases.

(Auditing, and the use of state-contingent contracts (warranties), can also facilitate a deal here.)