Decision Making When Outcomes Are Stochastic or Uncertain

Topics: Expected Values and Expected Utility
Applications: theory of the firm, consumer choice, criminal choices, economic regulation.

I. Introduction: Uncertainty and Time in Decision Making

A. In economics there are many choice settings in which decision makers confront problems that involve either or both uncertainties and long time horizons. These settings differ from those we've previously explored where all relevant details about the choice setting were known, there was no uncertainty about the result, and the timing of consequences or choices were neglected. Examples include:

i. A firm's choice of output levels and production methods when input prices or output prices are uncertain, or regulations are imperfectly enforced.

ii. Consumer choices to invest, take a vacation, or engage in other risky activities such as gambling, rock climbing, smoking, crime, etc..

iii. In the fields of public, environmental, and regulatory economics, there is the determination of the proper level of regulation, fine schedules, and enforcement efforts when the effects of the regulation depend on uncertain variables like the weather or future enforcement efforts.

Indeed, one can argue that most choices involve uncertainties of various kinds. Not all "facts" and "relationships" are known or can be known with certainty.

a. Many causal relationships are partly random. For example, many production costs depend on the weather (wind and temperatures).

b. Other uncertainties arise because information is incomplete or imperfect. Various macroeconomic, political, and natural crises may be surprise events. The great recession of 2008-10, the invasion of Ukraine, and the Covid Pandemic were all unanticipated by most decision makers.

c. Together, such phenomena imply that the costs and benefits of alternatives are often uncertain or probabilistic in both the short and long term.

iv. Moreover, many choice settings involve of benefits and costs that take place over relatively long periods of time, rather than simple one-shot consumer experiences.

a. The benefits and costs of many capital purchases occur over many years

b. The same tends to be true of investments in human capital, such as your investment in various college degrees at WVU. The benefits and costs of career choices span decades.

c. Regulatory and constitutional choices often involve long term commitments with both uncertainties and long time horizons. For example, the issues associated with global warming span centuries or millennia.

B. Neither uncertainty nor long time horizons make rational calculations or choices impossible, butthey do make such choices more difficult and more prone to errors of various kinds.

- Chapter 6 provides an overview of the most widely used tools (mathematical models) that economists (and most other policy analysts) employ to analyze decisionmaking under uncertainty. Chapter 7 provides an overview of the most widely used tools (mathematical models) for intertemporal decisionmaking (sometimes called dynamic choice settings).

- Together these two sets of tools allow one to model and/or assess the costs and benefits of alternative economic policies in settings where those policy decisions have long term consequences that are, at least partly, uncertain.

II. Decision Making under Uncertainty

A. The economics of risk and uncertainty was first examined thoroughly in a book by Frank Knight (1921) with the title *Risk Uncertainty and Profit.*

i. In that book he makes a distinction between risk (stochastic events in which a probability function may be accurately calculated) and uncertainty (stochastic events in which a probability function is not or cannot be accurately calculated).

ii. He also examines how risk and uncertainty (in his senses) differ in their effects on markets.

iii. Risks can be estimated and priced, and it will be redistributed among persons and firms so that the least risk averse bear larger risks than the most risk averse (because the sale and purchases of various forms of insurance).

iv. Uncertainty, in contrast, cannot be truly estimated or priced, and thus it is the primary source of extra ordinary profits and losses in the economy.

v. This distinction continues to play a role in theory, although it is not greatly emphasized in most discussions of decision making in stochastic circumstances. Unless Knight is cited, the terms risk and uncertainty are usually used interchangeably and assumed to have the same meaning.

vi. Nevertheless, Knight's distinction is a useful one to keep in the back of one's mind.

B. As a rule, economists assume that either objective probabilities can be estimated and used as an aid for decisionmaking or that intuitive approximations are used in more or less the same way to think about uncertain possibilities.

C. In many areas of choice, the benefits and costs of particular choices (or policies) are at least partly the consequence of chance.

vii. The most common models of decisionmaking in settings of uncertainty are the expected utility and expected net benefit maximizing models.

viii. **Expected value**, itself, is an idea taken from statistics and means the average result that would be expected from a series of "draws" from a stable random process of some kind.

D. DEF: The **mathematical expected value** of a set of possible outcomes, 1, 2, ... N with values V_1, V_2, ... V_N and probabilities of occurrence P_1, P_2 , ... P_N is

$$E(V) = \sum_{i=1}^{N} P_i V_i$$

i. Every **probability function** assigns probabilities to events (here events 1, 2, … N) such that the sum of the probabilities is 1.0. (The probability that something will actually happen is 1, is completely certain.)

- Every probability distribution has the property: $\Sigma Pi = Pi \geq 0$ for all i
- All possibilities, i, have positive probabilities of occurrence $1 \geq Pi > 0$.
- All impossibilities, j, have a zero probability of occurring and so $Pj = 0$.
- Every possibility is assigned a probability.

i. The **mathematical expected value is the sum of the values of those possibilities (here V_1, V_2 … V_N) times their probabilities of occurrence (here P_1, P_2, … P_N).** It represents the long-term *average* value of the distribution of values.

ii. The *expected* utility associated with a probabilistic setting is calculated in a similar manner:

$$E(V) = \sum_{i=1}^{N} P_i U(v_i)$$

where the *N* "value possibilities" are now measured in benefit terms associated with the affected individuals.

E. To use this formulae for expected value calculations, one has to assume that the outcomes of the "uncertain" events are finite, can be counted, can be listed, and probabilities assigned to them.

i. This is not an unreasonable assumption in many circumstances and is a reasonable first approximation of many others.

ii. Here, it bears noting that the probabilities assigned may be the result of careful empirical work (frequentist) or (Bayesian) intuitions about the likelihood of particular events that are updated as more evidence is gathered.

- (Frank Knight's *Risk, Uncertainty and Profit* argued that "risk" occurs when one can assign realistic probabilities to outcomes and "uncertainty" occurs when one cannot. He argued that true economic profits in perfectly competitive markets can only arise from "uncertainty" in perfectly competitive markets.)

iii. Most economists and most economic models are quite willing to assume that all the possible outcomes are known and that probabilities can be assigned to them.

(In most policy areas, however, the probabilities are themselves estimates that are updated as research, policies, or persuasive campaigns take place.)

F. Illustration of an expected value calculation: **the roll of a die**

i. Suppose that a single die is to be rolled. The face that turns up on top is a random event.

ii. Suppose that you will be paid a dollar amount equal to the number on the face that winds up on top.

iii. Since the probability of a particular face winding up on top is 1/6 and the value of the outcomes are 1, 2, 3, 4, 5, 6, arithmetic implies that the expected value of this game is $3.50 = (1)(1/6) + (2)(3.5) + (3)(1/6) + ...(6)(1/6) .

a. If you played the game dozens of time, your average payoff per roll would be approximately$3.50.

b. Note that the expected value of a single roll of a die is 3.5, a number that actually is impossible, rather than "expected" in the usual sense in ordinary English.

c. This is not always the case, but this example illustrates that the meaning of "expected value" is a technical one: long term average result.

d. There are many probability distributions in which the average value is also the mode, so it is also the most likely value to be observed, as with a normal distribution. iv. In most theoretical work benefits are calculated in "utility" terms.

G. E**xpected utility calculations and risk aversion**

i. Expected utility differs from expected value only in that a different "value" is being expected, namely the utility associated with a particular outcome or money value associated with an outcome.

$$E(U(V)) = \sum_{i=1}^{N} P_i \, U(V_i)$$

- Utility functions that can be used to calculate expected utility values that properly rank alternative outcomes (according to expected utility) are called **Von-Neumann Morgenstern utility functions**.
- Von-Neuman Morgenstern utility functions are all bounded and continuous.
- Von-Neuman Morgenstern utility functions for particular individuals are also "unique" up to a linear transformation(and considered by some to be a form of *cardinal* utility).

H. **DEF**: An individual is said to be *risk averse* if the expected utility of some gamble or risk is less than the utility generated at the expected value (mean) of the variable that determines utility (here V).

i. A **risk averse** person is one for whom the expected utility of a gamble (risky situation) is less than utility of the expected (mean) outcome, if the latter could be obtained with certainty.

a. In mathematical terms, a person is risk averse if and only if $U(x)^e < U(X^e)$ where X is a random event and X^e is its expected value.

b. That is to say, when: $U(x)^e > U(X^e)$

c. Note that this implies that **any benefit or utility function that is strictly concave with respect to income, exhibits risk aversion**.

d. (Why? Because expected utilities are convex combinations of utilities. Draw a one-dimensional illustration.)

ii. A ***risk neutral*** individual is one for whom the expected utility of a gamble (risky situation) and utility of the expected (mean) outcome are the same.

$U(x)^e = U(X^e)$

iii. A ***risk preferring*** individual is one for whom the expected utility of a gamble is greater than the utility of the expected (mean) outcome.

$U(x)^e > U(X^e)$

iv. The degree of risk aversion is often measured using the *Arrow-Pratt* measure of (absolute) risk aversion: r(Y) = - U"(Y)/U'(Y)]

I. The utility Functions that imply risk-averse behavior are all strictly concave, as illustrated below.

i. The figure below illustrates a choice setting in which an individual is risk averse and facing a risky environment in which either an outcome with the value V_1 or another outcome with the value V_2 will occur. The individual cannot influence which outcome it will be, but knows that the probability of V_1 is P, which implies that the probability that V_2 is (1-P) (Recall that that the probabilities for the only two possible events has to add up to one.)

ii. Lets refer to the individual as Al. Al's utility function is strictly concave, which means that a cord connecting any two points on it lies below the utility function (except for the two points used as end points, which by definition are not part of the cord).

iii. Note that if we consider the expected utility of this choice setting it can be written as

iv. Assume that Al confronts an uncertain environment in which V_1 occurs with probability P and V_2 occurs with probability (1-P) Al's expected utility in that case is:

- $U^e = PU(V_1) + (1-P)U(V_2)$
- As P increases from 0 to 1, the expected utilities trace out the cord between $U(V_1)$ and $U(V_2)$ and so will be below the utility function if it is strictly concave.
- (Recall the definition of strict concavity in Chapter 2 using alpha.)
- It is also the case that a sufficient condition for concavity in this case is that Al's utility function has a positive first derivative for V and a negative second derivative for V.
- In other words, Al's utility function is concave if it exhibits diminishing marginal returns from V.
- The expected value of V is $V^e = PV_1 + (1-P)V_2$
- Note that if U is strictly concave then $U(V^e) > PU(V_1) + (1-P)U(V_2)$
- Note that this looks just like the definition for concavity except that we've substitute "P" for "α".
- This is shown in the diagram below for a probability P approximately equal to 0.5, but it would be true for all probabilities $0 < P < 1$ and all strictly concave utility functions.

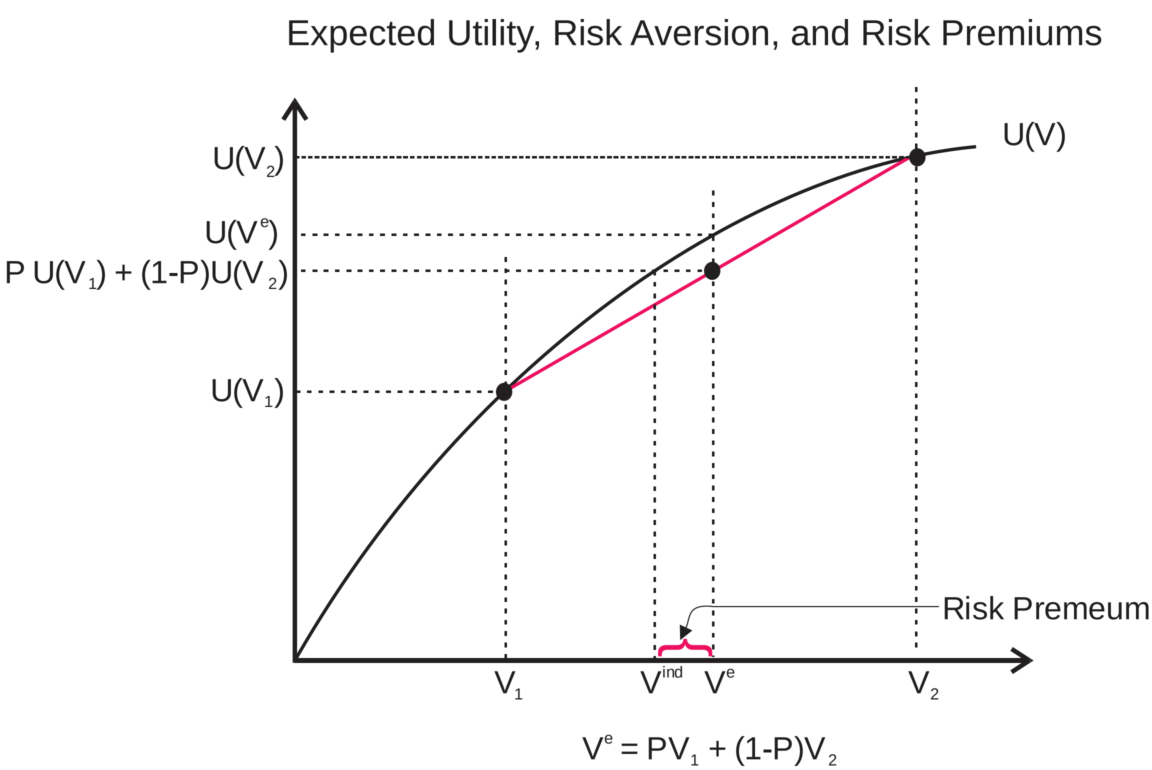

-
- This diagram can also be used to determine how much an individual would be willing to pay to have a certain payoff rather than face a risky or uncertain future.
- This is done by looking at the certain outcome that a person would be equivalent in their mind to the risky event.

- If we go to the left from the expected utility associated the two probabilistic outcomes over to the utility function and then down to the horizontal axis we find the value (labeled V^{ind}) that Al would find equivalent to the risky one faced. (V^{ind} is the certain outcome that generates the same expected utility as the risky one faced.)
- The difference in values, $V^e - V^{ind}$, (assuming that the values are in money terms) is the highest price that Al would pay to avoid the risk.
- It is also the lowest value that Al would accept to bear the risky environment shown rather than have outcome V^{ind} with certainty.
- That difference is called Al's **risk premium** for this "gamble." Al would accept the gamble (risky environment) rather than V^{ind} only if the expected value of the payoff is at least that much greater than V^{ind}.

III. Applications: Expected Benefits Maximization and Uncertain Product Quality

i. Suppose that Al is considering purchasing some produce from a farm and knows that some of the produce will be of high quality (H) and some will be of low quality (L) but either has to choose blindly (by ordering a particular number of units from the farm or pulling them out of a box) or simply can't tell the difference between the two types or produce at the time of purchase, as is true of many types of produce (corn, potatoes, tomatoes, squash, etc)..

- Suppose that the probability of high quality is F and that price per unit is simply P
- Suppose that the benefits of high-quality units is B(Q, H) and the benefits from quality units is B(Q, L) where B(Q,H) > B(Q,L) for every Q.
- How many units will Al purchase? .

ii. Al's expected net benefit from purchasing produce is expected benefits less expected costs:

- $N^e = FB(Q,H) + (1\text{-}F)B(Q,L) - PQ$
- To find Q*, differentiate Ne with respect to Q and set the result equal to zero.
- $F\ (dB^H/dQ + (1\text{-}F)(dB^L/dQ) - P = 0$
- The first two terms of the expected marginal benefit of the produce and the last is its marginal cost.

iii. To find a specific value we would need to use concrete functional forms for the two benefit functions, as with $B^H = HQ^{.5}$. and $B^L = LQ^{.5}$, with H>L, in which case our first order condition is:

- $.5FH/Q^{.5} + .5(1\text{-}F)L/Q^{.5} = P$
- Multiplying both sides by $2Q^{.5}$ yields $HF + L(1\text{-}F) = 2PQ^{.5}$
- Which implies that $Q^* = [HF + L(1\text{-}F)]^2/4P^2$
- **The quantity Al purchases** rises with F (the probability of the high-quality type) and with H (an indication of the quality of the high-quality type) and falls as L increases or price increases.

IV. Applications: Expected Values and the Logic of Crime and Punishment

i. The economic analysis of crime derives from a classic paper written by Gary Becker, who subsequently won a Nobel prize in economics. In that paper, and in many others published since then, a criminal is modeled as a rational agent interested in maximizing his EXPECTED income or utility, given some probability of punishment.

ii. This type of model can be used to model theft and violations of other laws.

- In the real world, criminal laws are only imperfectly enforced, and both criminals and ordinary persons who occasionally think about violating a law or two know this.
- For example, a net income maximizing criminal would maximize an expected function like $\Pi^e = PQ - cQ^2 - p(Q)F$ where Q is the number of crimes (thefts), price is the average price received by "fencing" the stolen goods, p(Q) is a probability function describing the way that that the probability of being caught and convicted varies with the number of crimes and F is the financial penalty assessed (or if jail time is spent, the opportunity cost of the time spent in jail and any subsequent losses in earnings).
- The rational theft chooses Q* such that $\Pi^e_Q = 0$, which in this case requires Q* to satisfy $P-2cQ - p_QF = 0$ or $P = 2cQ + p_QF$ (set the marginal revenue from theft equal to its **expected marginal cost**, which is not known with certainty).
- Let's give the probability function a concrete form as with: $p = aQ_2$ then $p_Q = 2aQ$ and the above first order condition becomes $P-2cQ - 2aQF = 0$ or $P= 2cQ + 2aQF$, which can be solved for Q.
- $P = Q(2c +2aF) \rightarrow$ **$Q^* = P/(2c + 2aF)$**
- Note that this implies that the rational criminal responds to incentives, his or her crime rate falls as the probability of being caught and convicted rises (e.g. with 2a), as the fine increases, and as the marginal cost of theft increases.
- Note also that there are tradeoffs between the size of the fine and the probability that a criminal is caught in terms of their overall effect on the criminal.
- (This model provides a short form of Gary Becker's classic 1968 paper on crime and punishment.)

iii. Many other examples from law and economics can also similarly modelled. One does not have to be a more or less professional criminal for this logic to apply.

- One can think of choices to drive faster than the speed limit on a highway or to park without putting money in a parking meter, or to trespass on a neighbor's property, fail to report some income on one's taxes, and so on in much the same manner.

V. Applications: Expected Values and the Effects of Regulation

i. One can also use this type of model to model the effects of economic regulation.

- For example, in the area of environmental regulations, firms will take account of their overall net benefits from pollution including both cost savings and anticipated regulatory fines when choosing their production methods.
- In the absence of fines or fees for pollution and in the absence of enforcement of fines greater than 0, firms will choose their production methods to minimize their production costs—as in the models developed in the first part of the course (prior to the midterm).
- (This does not necessarily mean that firms will pay no attention to air or water pollution, but they will do so only insofar as it affects the firm's expected profit through productivity and cost effects. Air or water quality that *affects the productivity of the firm's workforce* will be taken account of, but not spillovers on others outside the firm.)

ii. In the real world, regulations are only imperfectly enforced, and firms know this.

- Consequently, it is not simply the magnitude of the fine or penalty schedule that affects a firm's decision to "pollute illegally or not," but also the probability that a person that violates the law will be caught, convicted and punished.
- Analyzing regulatory law and its enforcement on a firm's choice of production method and output level requires taking account of both the "expected cost" and "expected marginal cost" of any fines or penalties that might be associated with its production and output decisions.
- (In addition, firms might face a loss of reputation and therefore reduced demand for their products if they are found guilty of violating regulatory law, but that effect will be ignored or assumed to be part of the fine.)

iii. Consider a case in which production methods are fixed and output is regulated—which is the easiest case to model.

- **In a regulatory environment with fines, a pragmatic firm's expected profits equal its total revenues less its production costs less its expected fines**: $\Pi = R - C - F^e$ where $F^e = PF$
- Suppose that Acme's output is sold in a competitive market, its cost function is $C=cQ^2wr$ and that its expected fine is the probability of being caught and convicted, which increases with output in excess of the regulatory limit, $p(Q-Q^R)$ and a fine schedule that increases with the extent of the violation $f(Q-Q^R)$ for $Q> Q^R$.
- $\Pi^e = PQ - cQ^2wr - p(Q-Q^R)f(Q-Q^R)$
- To make the functional form a bit more concrete, let us assume that $P(Q-Q^R) = a(Q-Q^R)$ and $f(Q-Q^R) = b(Q-Q^R)$. In this case, Acme's expected profits are:
- $\Pi^e = PQ - cQ^2wr - a(Q^R-Q)^2\, b(Q^R-Q) = PQ - cQ^2wr - ab(Q-Q^R)^2$
- Assume that the regulatory constraint is binding on Acme, and so it will take the expected fine schedule into account when making its output decision. Its expected profit maximizing output can be characterized by differentiating the above function with respect to Q, which is a bit more complex than usual because of the "$Q-Q^R$" terms.
- $\Pi^e_Q = P - 2cQwr - 2ab(Q-Q^R) = 0$

iv. This can be solved for Q* . First, shift the Q terms to the left side of the equal sign:

- $P = 2cQwr + 2ab(Q-Q^R) = Q(2cwr + 2ab) - 2abQ^R$
- Adding $2abQ^R$ and dividing yields:
- $\mathbf{Q^* = (P+2abQ^R)/(2cwr + 2ab)}$
- **This is Acme's supply function in the regulatory environment modeled.**

v. Note that its output now varies with the regulatory standard (Q^R) its input costs (w and r) and parameters of the probability of being fined and fine schedules (a and b).

- Acme's output declines as input prices and the expected fines increase (w, r, a, or b increase) and increases as the regulatory threshold (Q^R) increases.
-
- (Another possible output is simply Q^R, but this cannot be modeled with calculus because of a discontinuity in the expected cost function at that quantity. See below.)

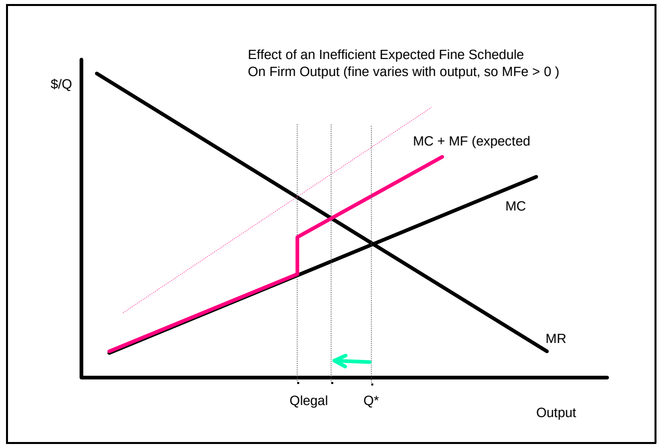

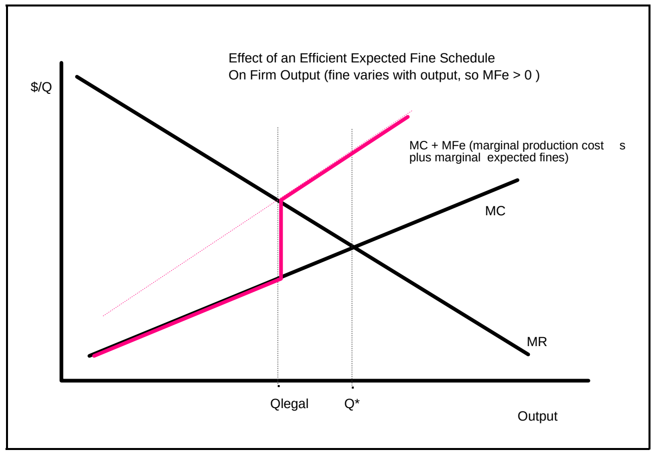

A. The diagram to the left illustrates Acme's decision in this type of setting (with somewhat simpler probability and fine schedules).

B. For students that have had public economics, note the similarities between Pigovian taxes and optimal enforcement with fines.

   If the regulation attempts solve an externality problem and achieve Pareto efficiency, Q**, then the smallest fine sufficient to induce the target Q** has the **same expected value** as a Pigovian tax at Q** (with $Q^R \leq$ Q**). The expected fine should equal the expected marginal damages done by the Q**th unit of output.

C. Note that there is always policy-tradeoff between the probability of conviction and the optimal level of punishment. [Recall that the expected fine is $F^e = PF$]

   i. The larger the fine, the smaller the probability of capture can be to generate the same effect on individuals.

   ii. The larger is the probability the smaller the fine can be and still have the same effect.

      a. The effect is determined by the expected fine, PF, in this case.

      b. The probability that an illegal activity is detected and punished varies with the resources used to enforce the law and the flagrancy of the violation, so the probability of being caught and punished tends to vary with law enforcement budgets and the size of the violation.

      c. The politics of enforcement and penalties are partly determined by error rates in detecting criminal activities--sometimes the wrong person is singled out for punishment.

- Puzzle. Given this, how would you pick the appropriate punishment for speeding? for theft? For murder? etc.
- Puzzle. How would the relative importance of the probability of detection and the expected fine be affected by the process of a jury trial and a long delay between being detected and being fined? (Some ideas for doing so are provided in the next chapter.)
- Puzzle: Write down an expected profit function for a firm facing a fine schedule that is imperfectly enforced, but where the fine increases as Q exceeds Q legal. Find the first order conditions and compare them to the above diagrams.
- Puzzle: draw examples of a perfectly enforced and imperfectly enforced "fixed fine schedule." (Such fines do **not** affect expected marginal costs.) Compare your graph with the mathematics of expected profit maximization in this case. Are such fines always irrelevant?

-

VI. Optional Appendix: expected values with continuous probability functions

A. The above is developed for cases in which the stochastic factor(s) is (are) take only discrete such as the value generated when rolling dice or yes/no types of outcomes, such as being fined or not for violating a law or regulation.

B. Continuous cases—where the domain of the probability density function is continuous rather than a discrete function (as in the case of rolls of a die) can be represented with integrals.

C. Given f(Π|Q) being the conditional probability that a particular profit level is realized when output Q is produced, expected profits can be written as an integral of the following sort.

- $\Pi^e = \int_{-\infty}^{+\infty} \Pi f(\Pi|Q)\ d\Pi$
- The optimal quantity of output that maximizes expected profits in this case is the one that sets the following expression equal to zero.
- $\Pi_Q^e = \int_{-\infty}^{+\infty} \Pi f(\Pi|Q)_Q\ d\Pi = 0$
- where subscripts indicate derivatives with respect to the variables subscripted.

- Note that the integrals are (usually) just carried forward and the initial total variables integrated are replaced with their relevant marginal values. (The integral written above is the expected (or average) marginal profit associated with output Q, which is zero at the expected profit maximizing output.)
- Note also that the domain of the integral is determined by the terms integrated, which in most expected value cases is determined by that of the probability density function. Density functions such as the normal distribution have unbounded domains, whereas the uniform distribution usually has a bounded domain.
- One uses the term "probability density function" (pdf) rather than "probability function" here, because probabilities are associated with integrals of (areas under) the density function. Thus, the total area under both a conditional and unconditional probability function is 1 (by definition).