Chapter 1. Economic Analysis and Econometrics

1. Introduction

The discipline of Economics is both a *pure science*, concerned with developing and validating theories of behavior of individuals and institutions, and a *policy science*, concerned with the forecasting and engineering of behavior. In both arenas, concepts from probability and statistics, and methods for analyzing and understanding economic data, play an important role. In this chapter, we give three introductory examples that illustrate the intertwining of economic behavior, statistical inference, and econometric forecasting. These examples contain some probability and statistics concepts that are explained in later chapters. What is important on first reading are the general connections between economic reasoning, probability, statistics, and economic data; details can be postponed.

2. Petit Verdot's Rational Decision

Petit Verdot is a typical professional economist, witty, wise, and unboundedly rational. Petit works at a university in California. Petit's life is filled with fun and excitement, the high point of course being the class she teaches in econometrics. To supplement her modest university salary, Petit operates a small vineyard in the Napa Valley, and sells her fruit to nearby wineries.

Petit faces a dilemma. She has to make a decision on whether to harvest early or late in the season. If the Fall is dry, then late-harvested fruit is improved by additional "hang time", and will fetch a premium price. On the other hand, if rains come early, then much of the late-harvested fruit will be spoiled. If Petit harvests early, she avoids the risk, but also loses the opportunity for the maximum profit.

Table 1 gives Petit's profit for each possible action she can take, and each possible Event of Nature:

Table 1.	Profit	Actions: Harvest Early	Actions: Harvest Late
Event of Nature	Early Rains	$30,000	$10,000
	Late Rains	$30,000	$40,000

There is a specialized weather service, Blue Sky Forecasting, that sells long-run precipitation forecasts for the Napa Valley. Petit has to choose whether to subscribe to this service, at a cost of $1000 for the year. Table 2 gives the historical record over the past 100 years on the joint frequency of various forecasts and outcomes.

Table 2.	Frequency	Blue Sky Forecast: Early Rains	Blue Sky Forecast: Late Rains	No Forecast TOTAL
States of Nature	Early Rains	0.3	0.1	0.4
	Late Rains	0.2	0.4	0.6
	TOTAL	0.5	0.5	1.0

The interpretation of the number 0.3 is that in 30 percent of the past 100 years, Blue Sky has forecast early rains and they do in fact occur. The column totals give the frequencies of the different Blue Sky forecasts. The row totals give the frequencies of the different States of Nature. Thus, Blue Sky forecasts early rains half the time, and the frequency of actual early rains is 0.4. One can also form conditional probabilities from Table 2. For example, the conditional probability of late rains, conditioned on the event that late rains are forecast, equals 0.4/(0.1+0.4) = 0.8.

Petit wants to maximize expected profit. In other words, she wants to make the probability-weighted average of the possible profit outcomes as large as possible. Petit is not adverse to risk; she figures that risks will average out over the years.

If Petit does not subscribe to the Blue Sky forecast service, then all she knows when she has to make her harvesting decision is the marginal frequency of early rains, 0.4, given in the last column in Table 2. To calculate expected profit in this case from a specified action, Petit multiplies the profit she will receive in each state of Nature by the probability of this state, and sums. If she harvests early, the expected profit is

$$(\$30{,}000)\cdot(0.4) + (\$30{,}000)\cdot(0.6) = \$30{,}000.$$

If she harvests late, the expected profit is

$$(\$10{,}000)\cdot(0.4) + (\$40{,}000)\cdot(0.6) = \$28{,}000.$$

Then, if Petit does not subscribe to Blue Sky, she will always choose to harvest early.

Now suppose Petit does subscribe to Blue Sky, and has their forecast available. In this case, she can do her expected profit calculation conditioned on the forecast. In any case, the expected profit from harvesting early is \$30,000, as before. To analyze her options when she gets the forecast, Petit first calculates the conditional probabilities of early rain, given the forecast:

$$\text{Prob(Early Rains} \mid \text{Forecast Early Rains)} = \frac{0.3}{0.3+0.2} = 0.6$$

$$\text{Prob(Early Rains} \mid \text{Forecast Late Rains)} = \frac{0.1}{0.1+0.4} = 0.2 \ .$$

If the forecast is for early rains, the expected profit from harvesting late is

$$(\$10{,}000)\cdot(0.6) + (\$40{,}000)\cdot(0.4) = \$22{,}000.$$

This is less than \$30,000, so Petit will definitely harvest early if early rain is forecast. If the forecast is for late rain, then the expected profit from harvesting late is

$$(\$10{,}000)\cdot(0.2) + (\$40{,}000)\cdot(0.8) = \$34{,}000.$$

This is greater than \$30,000, so Petit will harvest late if the forecast is for late rain.

Is subscribing to Blue Sky worth while? If Petit does not, then she will always harvest early and her expected profit is \$30,000. If Petit does subscribe, then her expected profit in the event of an early rain forecast is \$30,000 and in the event of a late rain forecast is \$34,000. Since the frequency of an early rain forecast is 0.5, Petit's overall expected profit if she subscribes is

$$(\$30{,}000)\cdot(0.5) + (\$34{,}000)\cdot(0.5) = \$32{,}000.$$

This is \$2000 more than the expected profit if Petit does not subscribe, so that the <u>value</u> of the <u>information</u> provided by the subscription is \$2000. This is more than the \$1000 cost of the information, so Petit will choose to subscribe and earn an overall expected profit, net of the subscription cost, of \$31,000.

Petit Verdot's decision problem is a typical one for an economic agent facing uncertainty. She has a criterion (expected profit) to be optimized, a "model" of the probabilities of various outcomes, the possibility of collecting data (the forecast) to refine her probability model, and actions that will be based on the data collected. An econometrician facing the problem of statistical inference is in a similar situation: There is a "model" or "hypothesis" for an economic phenomenon, data that provides information that can be used to refine the model, and a criterion to be used in determining an action in response to this information. The actions of the econometrician, to declare a hypothesis "true" or "false", or to make a forecast, are similar in spirit to Petit's choice. Further, the solution to the econmetrician's inference problem will be similar to Petit's solution.

A textbook definition of econometrics is *the application of the principles of statistical inference to economic data and hypotheses*. However, Petit's problem suggests a deeper connection between econometric analysis and economic behavior. The decision problems faced by rational economic agents require statistical inference, and thus are "econometric" in nature. The solutions to these problems require the same logic and techniques that must be brought to bear more formally in scientific inference. Thus, all rational economic agents are informal working econometricians, and the study of formal econometrics can itself provide models of economic behavior. Turned around, econometrics is simply a codification of the "folk" techniques used by economic agents to solve their decision problems. Thus, the study of econometrics provides not only the body of tools needed in empirical and applied economics for data analysis, forecasting, and inference, but also the concepts needed to explain economic behavior.

3. Stock Market Efficiency

The hypothesis is often advanced that the stock market is *efficient*. Among the possible meanings of this term is the idea that arbitragers are sufficiently active and pervasive so that potential windows of opportunity for excess profit are immediately closed. Consider a broad-based stock market index, such as the S&P 500 Market Index. The profit to be made by taking a dollar out of a "risk-free" channel (usually defined to be 30-day U.S. Treasury Bills), buying a dollar's worth of a Market Index Fund, and selling it one day later is

$$R_{mt} = [M_t + d_{t-1}M_{t-1} - (1+i_{t-1})M_{t-1}]/M_{t-1},$$

where M_t equals the market index, and is by definition the price of one share of the market index fund, d_{t-1} is the dividend rate, and i_{t-1} is the interest rate, and both d_{t-1} and i_{t-1} are assumed to be known in period t-1. If the arbitrager faces a higher interest rate than the risk-free rate, then i_{t-1} should be reinterpreted as the relevant rate. The argument is that the expectation of R_{mt}, as of period t-1, must be zero, since if it were positive, the arbitrager would buy the fund until the price M_{t-1} is bid up to eliminate the profit opportunity. Similarly, if the expectation of R_{mt} in period t-1 were negative, then the arbitrager would sell the fund and bid down the price M_t until the profit opportunity is again eliminated.

The condition of a zero expectation of R_{mt} given the information available in period t-1 implies that

$$\frac{M_t - M_{t-1}}{M_{t-1}} + d_{t-1} - i_{t-1} = \varepsilon_t,$$

where ε_t is a disturbance with mean zero given period t-1 information. The probability that the market index decreases between t-1 and t then equals the probability that $\varepsilon_t < d_{t-1} - i_{t-1}$. Assume that $d_{t-1} - i_{t-1}$ is a constant λ, and that ε_t has a distribution that is stationary over time. (These assumptions are questionable, and the conclusions below may be altered if they are false.) Then the probability that the market index decreases on any day is $Q \equiv \text{Prob}(\varepsilon_t < \lambda)$, and the probability that it increases on any day is $P \equiv \text{Prob}(\varepsilon_t > \lambda)$. Note that $Q + P \leq 1$, with equality holding if the probability of no change in the index is zero.

Define a *positive run of length* n to be a sequence of n successive days on which the index increases, with a day on each end on which the index does not increase. Similarly, define a *negative run of length* n to be a sequence of n successive days on which the index decreases, bracketed by days where price does not decrease. Note that runs in this definition can be terminated by no-change days as well as by reversals in the direction of change. The table below gives the number of positive and negative runs in the S&P 500 Market Index between July 3, 1962 and December 31, 1990. Only days on which the markets are open are counted; a total of 7165 index changes are observed over this period, of which 3753 are increases, 3366 are decreases, and 46 are no change. Then, estimates from the sample frequencies are $\hat{P}$ = 3753/7165 = 0.5238 and $\hat{Q}$ = 3366/7165 = 0.4698.

RUN LENGTH	NUMBER OF OBSERVED POSITIVE RUNS	NUMBER OF OBSERVED NEGATIVE RUNS
1	642	715
2	393	423
3	243	214
4	145	112
5	66	51
6	40	35
7	26	16
8	12	6
9	7	4
10	3	2
11	3	2
12	2	1
13	0	0
14	1	0
TOTAL	1585	1581

From probability theory, if the probability P of a price increase is constant, independent of the past history of the process, then the probability of a positive run of length n is $P^{n-1}(1-P)$. Similarly, the probability of a negative run of length n is $Q^{n-1}(1-Q)$. Using the estimates $\hat{P}$ for P and $\hat{Q}$ for Q, we can then calculate the expected numbers of runs of each length. For example, the expected number of positive runs of length 3 is given by the total number of positive runs (1585) times the probability $P^2(1-P)$ of a run of length 3, evaluated at the estimate $\hat{P}$ = 0.5238.

Then, the observed and expected numbers of runs can be compared. This is done in the table below. Because the numbers of longer runs are small, we combine all the runs of eight or more days in length.

RUN LENGTH	NUMBER OF OBSERVED POSITIVE	NUMBER OF EXPECTED POSITIVE	NUMBER OF OBSERVED NEGATIVE	NUMBER OF EXPECTED NEGATIVE
1	642	754.8	715	838.3
2	395	395.4	423	393.8
3	243	207.1	214	185.0
4	145	108.5	112	86.9
5	66	56.8	51	40.8
6	40	29.8	35	19.2
7	26	15.6	16	9.0
8+	28	17.1	15	8.0

The table shows that there are many fewer runs of length one than are expected, and many more runs of three days or more than are expected. This suggests that either the efficient market hypothesis is wrong, or that facilitating assumptions we made along the way on the behavior of $d_{t-1} - i_{t-1}$ and ε_t are wrong. Of course, one cannot readily rule out by visual inspection the possibility that the hypotheses are true, and the differences in the table are the result of the statistical noise in the data. However, we can use a "goodness-of-fit" test to make a statement about the probability that the pattern of deviations between observed and expected outcomes in the table above could arise by chance. In this case, the goodness-of-fit test establishes that it is extremely unlikely that the patterns in the table arose by chance.

This example shows how an economic hypothesis can be formulated as a condition on a probability model of the data generation process, and how statistical tools can be used to judge whether the economic hypothesis is true. In this case, the evidence is that either the efficient markets hypothesis does not hold, or that there is a problem with one of the assumptions that we made along the way to facilitate the analysis. A more careful study of the time-series of stock market prices tends to support one aspect of the efficient markets hypothesis, that expected profit at time t-1 from an arbitrage to be completed the following period is zero. Thus, the elementary economic idea that arbitragers discipline the market is supported. However, there do appear to be longer-run time dependencies in the market, as well as heterogeneity, that are inconsistent with some stronger versions of the efficient markets hypothesis.

4. The Capital Asset Pricing Model

The return that an investor can earn from a stock is a random variable, depending on events that impinge on the firm and on the economy. By selecting the stocks that they hold, investors can trade off between average return and risk. A basic, and influential, theory of rational portfolio selection is the Capital Asset Pricing (CAP) model. This theory concludes that if investors are concerned only with the mean and variance of the return from their portfolio, then there is a single portfolio of stocks that is optimal, and that every rational investor will hold this portfolio in some mix with a riskless asset to achieve the desired balance of mean and variance. Since every investor, no matter what her attitudes to risk, holds the same stock portfolio, this portfolio will simply be a share of the total market; that is, all investors simply hold a market index fund. This is a powerful conclusion, and one that appears to be easily refutable by examining the portfolios of individuals. This suggests that other factors, such as irrationality, transactions costs, or heterogeneity in information, are influencing behavior. Nevertheless, the

CAP model is often useful as a normative guide to optimal investment behavior, and as a tool for understanding the benefits of diversification.

To explain the CAP model, consider a market with K stocks, indexed k = 1,2,...,K. Let P_{kt} be the price of stock k at the end of month t; this is a *random variable* when considered before the end of month t, and after that it is a number that is a *realization* of this random variable. Suppose an investor can withdraw or deposit funds in an account that holds U.S. Treasury 30-Day Bills that pay an interest rate I_t during month t. (The investor is assumed to be able to borrow money at this rate if necessary.) Conventionally, the T-Bill interest rate is assumed to be risk-free and known to the investor in advance. The profit, or *excess return*, that the investor can make from withdrawing a dollar from her T-Bill account and buying a dollar's worth of stock k is given by

$$R_{kt} = \frac{P_{kt} + D_{kt} - P_{k,t-1}}{P_{k,t-1}} - I_t,$$

where D_{kt} is the dividend paid by the stock at the end of month t. The excess return R_{kt} is again a random variable. Let r_k denote the *mean* of R_{kt}. Let σ_k^2 denote its variance, and let σ_{kj} denote the covariance of R_{kt} and R_{jt}. Note that σ_{kk} and σ_k^2 are two different notations for the same variance. The square root of the variance, σ_k, is called the standard deviation.

Now consider a portfolio in which *A* dollars are invested in stocks, with a fraction θ_k of each dollar allocated to shares of stock k, for k = 1,...,K. The excess return to this portfolio is then $A R_{pt}$, where

$$R_{pt} = \sum_{k=1}^{K} \theta_k R_{kt}$$

is the excess return to the one dollar stock portfolio characterized by the shares $(\theta_1,\ldots,\theta_K)$. The excess return R_{pt} is again a random variable, with mean $r_p = \sum_{k=1}^{K} \theta_k r_{kt}$ and variance $\sigma_p^2 = \sum_{k=1}^{K} \sum_{j=1}^{K} \theta_k\theta_j\sigma_{kj}$. This implies that the portfolio with *A* dollars invested has an excess return with mean Ar_p and variance $A^2\sigma_p^2$ (or standard deviation $A\sigma_p$). The covariance of R_{kt} and R_{pt} is given by $\sigma_{kp} \equiv \mathrm{cov}(R_{kt},R_{pt}) = \sum_{j=1}^{K} \theta_j\sigma_{kj}$. Define the *beta* of stock k (with respect to the portfolio p) by the formula $\beta_k = \sigma_{kp}/\sigma_p^2$, and note that

$$\sigma_p^2 = \sum_{k=1}^{K} \theta_k \left(\sum_{j=1}^{K} \theta_j\sigma_{kj} \right) \equiv \sum_{k=1}^{K} \theta_k\sigma_{kp} \equiv \sum_{k=1}^{K} \theta_k\beta_k\sigma_p^2,$$

and hence that $\sum_{k=1}^{K} \theta_k\beta_k = 1$.

Now consider the rational investor's problem of choosing the level of investment *A* and the portfolio mix $\theta_1,\theta_2,\ldots,\theta_K$. Assume that the investor cares only about the mean and standard deviation of the excess return from her portfolio, and always prefers a higher mean and a lower standard deviation. The investor's tastes for risk will determine how mean and standard deviation are traded off. The figure below shows the alternatives available to the investor. For a specified mix $(\theta_1^1,\ldots,\theta_K^1)$, the locus of portfolios obtainable by varying *A* is a straight line through the origin. Another mix $(\theta_1^2,\ldots,\theta_K^2)$ will give another straight line through the origin. The location of these lines reflects the operation of diversification to reduce risk; i.e., by holding a mix of stocks, some of which are likely to go up when others are going down, one may be able to reduce standard deviation for a given level of the mean. There will be an *optimal* mix $(\theta_1^*,\ldots,\theta_K^*)$ that gives a line that is rotated as far to the northwest as possible in the figure. Different investors will locate at different points along this optimal line, by picking different *A* levels, to maximize their various preferences for mean versus standard deviation. However, all investors will choose exactly the same mix $(\theta_1^*,\ldots,\theta_K^*)$ for the stocks that they hold. But if every investor holds stocks in the same proportions, then these must also be the proportions that prevail in the market as a whole. Then $(\theta_1^*,\ldots,\theta_K^*)$ will be the shares by value of all the stocks in the market. Such a portfolio is called a *market index fund*. The CAP model then concludes that rational investors who care only about the mean and standard deviation of excess return will hold only the market index fund, with the levels of investment reflecting their heterogeneous tastes for risk.

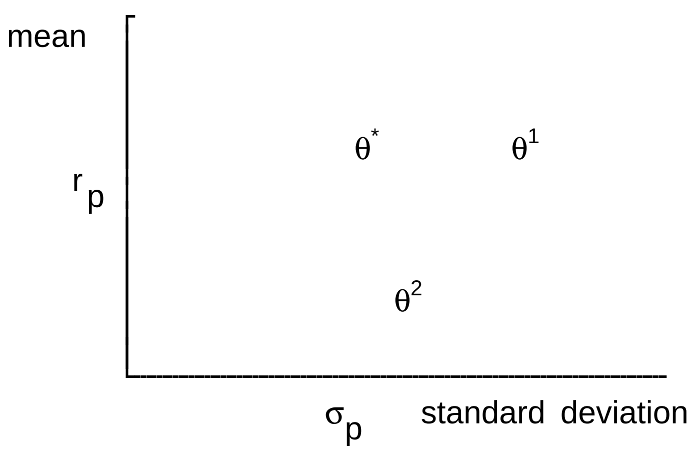

The problem of determining the optimal portfolio mix $(\theta_1^*,\ldots,\theta_K^*)$ is most easily solved by considering a closely related problem. An investor’s portfolio is characterized by *A* and $(\theta_1,\ldots,\theta_K)$. Given a choice among all the portfolios that achieve a specified level of mean return, the investor would want to choose the one that minimizes variance. In the figure, this corresponds to getting as far to the left as possible when constrained to the feasible portfolios on a specified horizontal line. From the previous discussion, the solution to this problem will be a portfolio with the optimal mix of stocks, and the only difference between this problem and the one of maximizing preferences will be in determining the overall investment level *A*. The problem of minimizing variance for a given mean is one of constrained minimization:

$$\text{Choose } A,\theta_1,\ldots,\theta_K \geq 0 \text{ to minimize } A^2 \sum_{k=1}^{K} \sum_{j=1}^{K} \theta_k\theta_j\sigma_{kj},$$

$$\text{subject to } A \sum_{k=1}^{K} \theta_k r_k = c \text{ (a constant) and } \sum_{k=1}^{K} \theta_k = 1.$$

The first-order (Kuhn-Tucker) conditions for this problem are

(1) $$2A^2 \sum_{j=1}^{K} \theta_j^* \sigma_{kj} \geq \lambda A r_k + \mu, \text{ with equality unless } \theta_k^* = 0, \text{ for } k = 1,\ldots,K$$

(2) $$2A \sum_{k=1}^{K} \sum_{j=1}^{K} \theta_k^*\theta_j^*\sigma_{kj} = \lambda \sum_{k=1}^{K} \theta_k^* r_k,$$

where the scalars λ and μ are Lagrange multipliers. Multiply (1) by θ_k^* and sum over k to obtain the result

$$A\left(2A \sum_{k=1}^{K} \sum_{j=1}^{K} \theta_k^*\theta_j^*\sigma_{kj} - \lambda \sum_{k=1}^{K} \theta_k^* r_k\right) = \mu \sum_{k=1}^{K} \theta_k^*.$$

Using (2), this implies $\mu = 0$. Then, (1) implies that the optimal θ_k^* satisfy

(3) $\sum_{j=1}^{K} \theta_j^* \sigma_{kj} \geq \gamma r_k$, with equality unless $\theta_k^* = 0$, for k = 1,...,K,

where γ is a scalar defined so that the θ_k^* sum to one. Since by the earlier comments this mix is simply the mix in the total market, equality will hold for all stocks that are in the market and have positive value.

Assume now that $r_p = \sum_{k=1}^{K} \theta_k^* r_k$ and $\sigma_p^2 = \sum_{k=1}^{K} \sum_{j=1}^{K} \theta_k^* \theta_j^* \sigma_{kj}$ refer to the optimal portfolio. The left-hand-side of (3) equals $\sigma_{kp} \equiv \beta_k \sigma_p^2$, so this condition can be rewritten

(4) $\beta_k \sigma_p^2 \geq \gamma r_k$, with equality unless $\theta_k^* = 0$, for k = 1,...,K.

Multiplying both sides of this inequality by θ_k^* and summing yields the condition

$$\gamma r_p \equiv \gamma \sum_{k=1}^{K} \theta_k^* r_k = \sigma_p^2 \sum_{k=1}^{K} \theta_k \beta_k = \sigma_p^2,$$

or $r_p = \sigma_p^2/\gamma$. Substituting this into (4) gives us the final form of a main result of the CAP model,

$r_k \leq \beta_k r_p$, with equality if the stock is held, for k = 1,...,K

The mean returns are not observed directly, but the realizations of monthly returns on individual stocks and the market are observed. Write an observed return as the sum of its mean and a deviation from the mean, $R_{kt} = r_k + \varepsilon_{kt}$ and $R_{pt} = r_p +$

ε_{pt}. Note that $R_{pt} = \sum_{k=1}^{K} \theta_k^* R_{kt}$, so then $\varepsilon_{pt} = \sum_{k=1}^{K} \theta_k^* \varepsilon_{kt}$. For all stocks held in the market, the CAP model implies $r_k = \beta_k r_p$. Use the form $R_{kt} = r_k + \varepsilon_{kt}$ to rewrite the equation $r_k = \beta_k r_p$ as $R_{kt} - \varepsilon_{kt} = \beta_k(R_{pt} - \varepsilon_{pt})$. Define $\nu_{kt} = \varepsilon_{kt} - \beta_k \varepsilon_{pt}$. Then, the equation becomes

(5) $$R_{kt} = \beta_k R_{pt} + \nu_{kt}.$$

This equation can be interpreted as a relation between *market risk*, embodied in the market excess return R_{pt}, and the risk of stock k, embodied in R_{kt}. The *disturbance* ν_{kt} in this equation is sometimes called the *specific risk* in stock k, the proportion of the total risk in this stock that is not responsive to market fluctuations. This disturbance has the following properties:

(a) $\mathbf{E}\nu_{kt} = 0$;

(b) $\mathbf{E}\nu_{kt}^2 = \mathbf{E}(\varepsilon_{kt} - \beta_k \varepsilon_{pt})^2 = \sigma_k^2 + \beta_k^2 \sigma_p^2 - 2\beta_k \sigma_{kp} \equiv \sigma_k^2 - \beta_k^2 \sigma_p^2$

(c) $\mathbf{E}\nu_{kt} R_{pt} = \mathbf{E}(\varepsilon_{kt} - \beta_k \varepsilon_{pt})\varepsilon_{pt} = \sigma_{kp} - \beta_k \sigma_p^2 = 0$.

Equation (5) is a *linear regression* formula. Properties (a)-(c) are called the *Gauss-Markov conditions*. We will see that they imply that an estimate of β_k with desirable statistical properties can be obtained by using the *method of least squares*. Then, the CAP model's assumptions on behavior imply an econometric model that can then be fitted to provide estimates of the *market betas*, key parameters in the CAP analysis.

The market beta's of individual stocks are often used by portfolio managers to assess the merits of adding or deleting stocks to their portfolios. (Of course, the

CAP model says that there is no need to consider holding portfolios different than the market index fund, and therefore no need for portfolio managers. That these things exist is itself evidence that there is some deficiency in the CAP model, perhaps due to failures of rationality, the presence of transactions cost, or the ability to mimic the excess return of the market using various subsets of all the stocks on the market because the optimal portfolio is not unique.) Further, statistical analysis of the validity of the assumptions (a)-(c) can be used to test the validity of the CAP model.

The β's in formula (5) convey information on the relationship between the excess return on an individual stock and the excess return in the market. Subtract means in (5), square both sides, and take the expectation to get the formula

$$\sigma_k^2 = \beta_k^2\sigma_p^2 + \delta_k^2,$$

where $\delta_k^2 = \mathbf{E}\nu_{kt}^2$ is the variance of the disturbance. This equation says that the risk of stock k equals the market risk, amplified by β_k^2, plus the specific risk. Stock k will have high risk if it has large specific risk, or a β_k that is large in magnitude, or both. A positive β_k implies that events that influence the market tend to influence stock k in the same way; i.e., the stock is *pro-cyclic*. A negative β_k implies that the stock tends to move in a direction opposite that of the market, so that it is *counter-cyclic*. Stocks that have small or negative β_k are *defensive*, aiding diversification and making an important contribution to reducing risk in the market portfolio.

The CAP model described in this example is a widely used tool in economics and finance. It illustrates a situation in which economic axioms on behavior lead to a widely used statistical model for the data generation process, the linear regression model with errors satisfying Gauss-Markov assumptions. An important feature of this example is that these statistical assumptions are implied by the economic axioms, not attached *ad hoc* to facilitate data analysis. It is a happy, but rare, circumstance in which an economic theory and its econometric analysis are fully integrated. To seek such harmonies, theorists need to draw out and make precise the empirical implications of their work, and econometricians need to develop models and methods that minimize the need for extra facilitating assumptions.