CHAPTER 1. ECONOMIC ANALYSIS AND ECONOMETRICS

1.1. INTRODUCTION

The study of resource allocation by the discipline of Economics is both a *pure science*, concerned with developing and validating theories of behavior of individuals, organizations, and institutions, and a *policy science*, concerned with the design of institutions and the prediction and social 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.

1.2. CAB FRANC'S RATIONAL DECISION

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

Cab faces a dilemma. He 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 Cab harvests early, he avoids the risk, but also loses the opportunity for the maximum profit. Table 1 gives Cab's profit for each possible action he can take, and each possible Event of Nature:

Table 1. Profit from Selling Grapes

		Action	
Event of Nature	Frequency	Harvest Early	Harvest Late
Wet	0.4	\$30,000	\$10,000
Dry	0.6	\$30,000	\$40,000
Expected Profit		\$30,000	\$28,000

Cab wants to maximize <u>expected</u> <u>profit</u>. In other words, he wants to make the probability-weighted average of the possible profit outcomes as large as possible. Cab is not adverse to risk; he figures that risks will average out over the years. From historical records, he knows that the frequency of early rain is 0.4. To calculate expected profit in this case from a specified action, Cab multiplies the profit he will receive in each event of Nature by the probability of this event, and sums. If he harvests early, the expected profit is ($30,000)·(0.4) + ($30,000)·(0.6) = $30,000. If he harvests late, the expected profit is ($10,000)·(0.4) + ($40,000)·(0.6) = $28,000. Then, in the absence of any further information, Cab will choose to harvest early and earn an expected profit of $30,000.

There is a specialized weather service, Blue Sky Forecasting, that sells long-run precipitation forecasts for the Napa Valley. Cab 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 of Forecasts and Outcomes

	Blue Sky Forecasts		TOTAL
Event of Nature	Early	Late	
Wet	0.3	0.1	0.4
Dry	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 rain 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 Events of Nature. Thus, Blue Sky forecasts early rain half the time, and the frequency of actual early rain is 0.4. One can also form conditional probabilities from Table 2. For example, the conditional probability of dry, given the event that late rain are forecast, equals 0.4/(0.1+0.4) = 0.8.

If Cab does <u>not</u> subscribe to the Blue Sky forecast service, then he is in the situation already analyzed, where he will choose to harvest early and earn an expected profit of $30,000. Now suppose Cab <u>does</u> subscribe to Blue Sky, and has their forecast available. In this case, he can do his expected profit calculation conditioned on the forecast. To analyze his options, Cab first calculates the conditional probabilities of early rain, given the forecast:

$$\text{Prob(Wet}|\text{Forecast Early}) = 0.3/(0.3+0.2) = 0.6$$
$$\text{Prob(Dry}|\text{Forecast Late}) = 0.1/(0.1+0.4) = 0.2\,.$$

The expected profit from harvesting early is again \$30,000, no matter what the forecast. Now consider the expected profit from harvesting late. If the forecast is for early rains, the expected profit is given by weighting the outcomes by their conditional probabilities given the forecast, or

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

This is less than \$30,000, so Cab will definitely harvest early in response to a forecast of early rain. Next suppose the forecast is for late rain. Again, the expected profit is given by weighting the outcomes by their conditional probabilities given this information,

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

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

Is subscribing to Blue Sky worth while? If Cab does not, then he will always harvest early and his expected profit is \$30,000. If Cab does subscribe, then his 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, Cab's overall expected profit if he subscribes is

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

This is \$2000 more than the expected profit if Cab 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 Cab will choose to subscribe and will earn an overall expected profit, net of the subscription cost, of \$31,000.

Cab Franc’s decision problem is a typical one for an economic agent facing uncertainty. He 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 his 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 Cab's choice. Further, the solution to the econmetrician's inference problem will be similar to Cab's solution.

A textbook definition of econometrics is *the application of the principles of statistical inference to economic data and hypotheses*. However, Cab Franc's problem suggests a deeper connection between econometric analysis and economic behavior. The decision problems faced by rational economic agents in a world with imperfect information 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 provide at least prescriptive models for 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 key concepts needed to explain economic behavior.

1.3. 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, the New York Stock Exchange (NYSE) value-weighted index of the prices of all the stocks listed with this exchange. The gross return to be made by taking a dollar out of a "risk-free" channel (defined here to be 90-day U.S. Treasury Bills), buying a dollar's worth of the "market" stock portfolio, and selling it one day later is $g_t = \log(M_t/M_{t-1})$, where M_t equals the NYSE index on day t, and the log gives the one-day exponential growth rate in share price. It is necessary in general to account for distributions (dividends) paid by the stocks during the time they are held; this is done automatically when g_t is reported. An arbitrager on day t-1 knows the history of the market and economic variables up through that day; let $\boldsymbol{H}_{t-1}$ denote this history. In particular, $\boldsymbol{H}_{t-1}$ includes the level M_{t-1} of the index, the pattern of historical market changes, and the overnight interest rate i_{t-1} on 90-day Treasury Bills, representing the opportunity cost of not putting or leaving a dollar in a T-Bill account. The difference $R_t = g_t - i_{t-1}$ is the profit an arbitrager makes by buying one dollar of the NYSE portfolio on t-1, and is called the *excess return* to the market on day t. (If the arbitrager sells rather than buys a dollar of the exchange index in t-1, then her profit is $-R_t$). On day t-1, conditioned on the history $\boldsymbol{H}_{t-1}$, the excess return R_t is a random variable, and the probability that it is less than any constant r is given by a distribution function $F(r|\boldsymbol{H}_{t-1})$. Then, the expected profit is $\int_{-\infty}^{+\infty} r \cdot F'(r|\boldsymbol{H}_{t-1})dr$. The argument is that if this expected profit is positive, then arbitragers will buy and drive the price M_{t-1} up until the opportunity for positive expected profit is eliminated. Conversely, if the expected profit is negative, arbitragers will sell and drive the price M_{t-1} down until the expected profit opportunity is eliminated. Then, no matter what the history, arbitrage should make expected profit zero. This argument does not take account of the possibility that there may be some trading cost, a combination of transactions charges, risk preference, the cost of acquiring information, and the opportunity cost of the time the arbitrager spends trading. These trading costs could be sufficient to allow a small positive expected excess return to persist; this is in fact observed, and is called the *equity premium*. However, even if there is an equity premium, the arbitrage argument implies that the expected excess return should be independent of history.

From the CRSP financial database and from the Federal Reserve Bank, we take observations on g_t and i_{t-1} for all days the market is open between January 2, 1968 and December 31, 1998, a total of 7806 observations. We then calculate the excess return $R_t = g_t - i_{t-1}$. The next table gives some statistics on these quantities:

Variable	Sample Average	Sample Standard Deviation
g_t	0.0514%	0.919%
i_{t-1}	0.0183%	0.007%
R_t	0.0331%	0.919%

These statistics imply that the annual return in the NYSE index over this period was 18.76 percent, and the annual rate of interest on 90-day Treasury Bills was 6.68 percent.

Now consider the efficient markets hypothesis, which we have argued should lead to excess returns that do not on average differ from one previous history to another. The table below shows the sample average excess return in the NYSE index under each of two possible conditions, a positive or a negative excess return on the previous day.

Condition	Frequency	Sample Average	Standard Error
$R_{t-1} > 0$	4082 (52.3%)	0.129%	0.013%
$R_{t-1} < 0$	3723 (47.7%)	-0.072%	0.016%

We conclude that excess return on a day following a positive excess return is on average positive, and on a day following a negative excess return is on average negative. The standard errors measure the precision of the sample averages. These averages are sufficiently precise so that we can say that the difference in sample averages did not arise by chance, and the expected excess returns under the two conditions are different.[1] We conclude that over the time period covered by our sample, the efficient markets hypothesis fails. There was a persistence in the market in which good days tended to be followed by better-than-average days, and bad days by worse-than-average days. There appears to have been a potentially profitable arbitrage strategy, to buy on up days and sell on down days, that was not fully exploited and eliminated by the arbitragers.

Define a *positive (negative) run of length* n to be a sequence of n successive days on which the excess return is positive (negative), with a day on each end on which the sign of excess return is reversed. The figure below plots the average excess returns conditioned on the length of an up (+) or down (-) run over the immediate past days. The horizontal lines above and below each point give

[1] The mean difference is 0.202% with a standard error of 0.010%, and the T-statistic for the hypothesis that the two conditional means is the same is 19.5.

an indication of the accuracy with which it is estimated; these are 95 percent confidence bounds. The figure suggests that expected return rises with the length of an up run, and falls with the length of a down run. A straight line is fitted through these points, taking into account the accuracy with which they are estimated, by the least squares regression method. This is also plotted in the figure, and shows that the average returns have a positive trend. Then, this figure supports our previous conclusion that the efficient markets hypothesis fails.

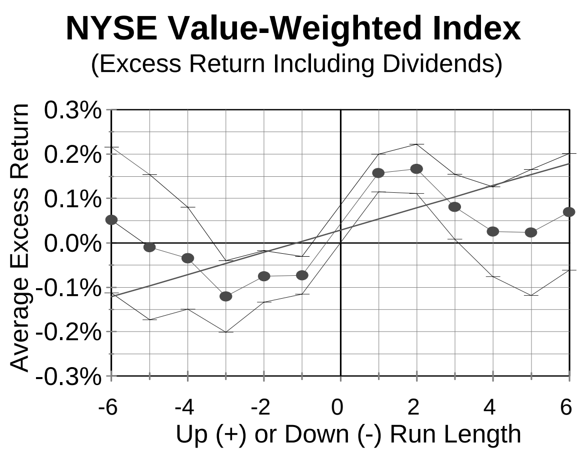

The preceding evidence suggests that arbitragers may underestimate the persistence of up and down runs, and fail to make profitable bets on persistence. The table below gives the observed counts of numbers of up and down runs of various lengths. It also includes a prediction of the numbers of runs of various lengths that you would expect to see if runs are the result of independent "coin tosses", with the probability of an "up" outcome equal to the 52.4% frequency with which the excess return was positive in our sample.[2] What one sees is that there are many fewer runs of length one and more runs of longer lengths than the "coin tossing" model predicts. There is the possibility that the differences in this table are the result of chance, but a statistical analysis using what is called a likelihood ratio test shows that the pattern we see is very unlikely to be due to chance. Then, up and down runs are indeed more persistent than one would predict if one assumed that the probability of a up or down day was independent of previous history.

[2]If P is the probability of an up day, then $P^{n-1}(1-P)$ is the probability of an up run of length n, and this multiplied by the total number of positive runs is the expected number of length n. An analogous formula applies for negative runs.

Run Length	Observed Positive	Observed Negative	Expected Positive	Expected Negative
1	1772	1772	1940.6	1938.0
2	1037	954	1015.1	924.3
3	602	495	531.0	440.8
4	313	245	277.7	210.2
5	159	120	145.3	100.3
6 or more	186	119	159.3	91.4
Total	4069	3705	4069	3705

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 an assumption 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.

1.4. 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, heterogeneity in information, or preferences that take into account risk features other than mean and variance, 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-1} 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

$$R_{kt} = \frac{P_{kt} - P_{k,t-1}}{P_{k,t-1}} + d_{k,t-1} - i_{t-1},$$

where $d_{k,t-1}$ is the announced dividend rate 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.

Consider an investor's portfolio of value A^* at the beginning of a month, and suppose A dollars are invested in stocks and A^*-A dollars are held in the risk-free account. Many investors will have $0 \le A \le A^*$. However, it is possible to have $A > A^*$, so that the investor has borrowed money (at the risk-free rate) and put this into stocks. In this case, the investor is said to have purchased stocks *on margin*. For the all-stock component of the portfolio, a fraction θ_k of each dollar in A is allocated to shares of stock k, for k = 1,...,K. The excess return to this portfolio is then AR_{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 $(\theta_1,...,\theta_K)$. The fractions θ_k are restricted to be non-negative. However, the list of "stocks" may also include *financial derivatives*, which can be interpreted as lottery tickets that pay off in dollars or stocks under specified conditions. For example, an investor may *short* a stock, which means that she in effect sells an IOU promising to deliver a share of the stock at the end of the month. She is then obligated to deliver a share of the stock at the end of the month, if necessary by buying a share of this stock to deliver in order to complete the transaction. Other elementary examples of financial derivatives are *futures*, which are contracts to deliver stocks at some future date, and *mutual funds*, which are institutions that sell shares and use the proceeds to buy portfolios of stocks. There are also more complex financial derivatives that require delivery under specified conditions, such as an increase in a stock market index of more than a specified percentage. 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 stock 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 $\sigma_{kp} \equiv \text{cov}(R_{kt},R_{pt}) = \sum_{j=1}^{K} \theta_j\sigma_{kj}$. Define the *beta* of stock k (with respect to the stock portfolio p) by the formula $\beta_k = \sigma_{kp}/\sigma_p^2$, and note that

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

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 $\theta_1,\theta_2,...,\theta_K$. Assume that the investor cares <u>only</u> 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. The investor prefers to be as far to the northwest in this figure as possible, where mean is high and standard deviation is low, and will have indifference curves that specify her tradeoffs between mean and standard deviation. (Note that an investor with low risk aversion will have indifference curves that are almost horizontal, while an extremely risk-averse investor will have indifference curves that are almost vertical.) If preferences between mean and standard deviation are derived from a utility function that is concave in consumption, then the indifference curve will be convex. A specified mix $(\theta_1^1,...,\theta_K^1)$ of stocks determines a one-dollar all-equity portfolio with particular values for standard deviation and mean, and an all-stock portfolio of value A^* invested in this mixture will have a mean and standard deviation that are A^* times those of the one-dollar portfolio. This point is indicated in the diagram as Portfolio 1. By holding A of the all-equity portfolio and $A^* - A$ of the riskless asset in some combination, the investor can attain mean and standard deviation combinations anywhere on a straight line through the origin and the all-equity Portfolio 1 point. (Points between zero and the all-equity portfolio correspond to investing only part of the investor’s assets in stocks, while points on the line to the right of the Portfolio 1 point correspond to borrowing money to take a margin position in stocks.) Another mix $(\theta_1^2,...,\theta_K^2)$ gives the point in the diagram indicated as Portfolio 2. Again, by holding the riskless asset and Portfolio 2 in some combination, the consumer could attain any point on the straight line connecting the origin and Portfolio 2. There will be a frontier envelope on the north-west boundary of all the mean and standard deviation combinations that can be attained with all-equity portfolios with value A^*; this envelope is drawn as a heavy curve in the figure, and represents *efficient* portfolios in terms of tradeoffs between mean and standard deviation in all-equity portfolios.

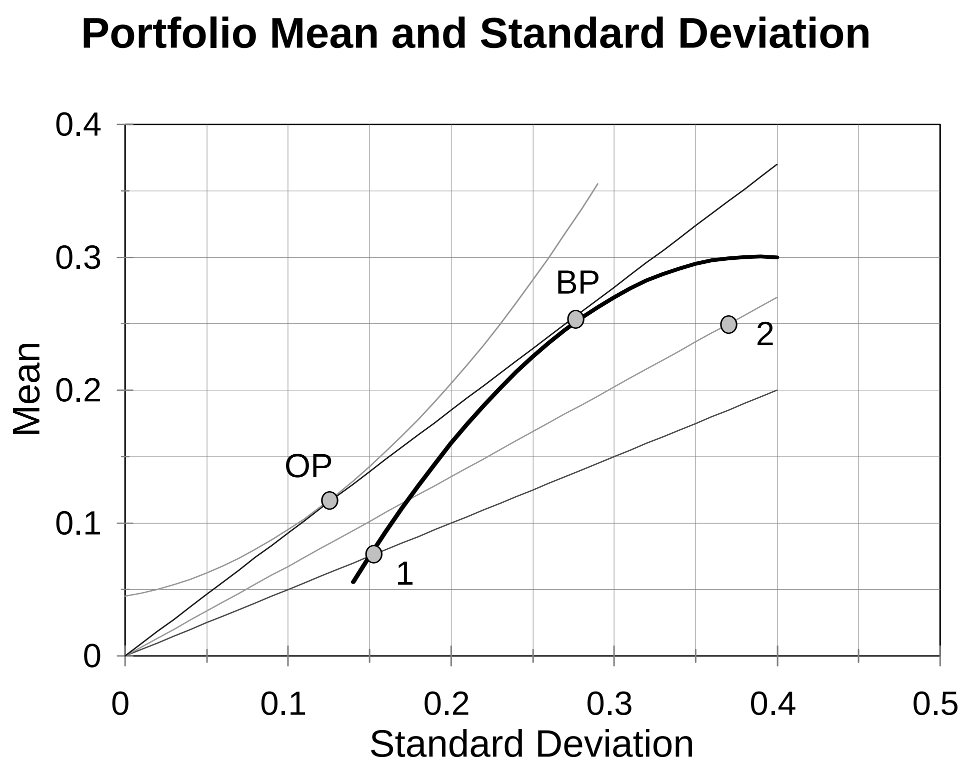

Portfolio 1 is efficient, while Portfolio 2 is not because it is southwest of the all-equity portfolio frontier. Note however that the consumer can be made better off with a portfolio that is some combination of Portfolio 2 and the riskless asset than with any combination of Portfolio 1 and the riskless asset, because the line through Portfolio 2 is always northwest of the line through Portfolio 1. Consider all the lines that connect the origin and all-equity portfolios. 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 *efficient* mix $(\theta_1^*,...,\theta_K^*)$, labeled BP (for Best Portfolio) in the figure, that gives a line that is rotated as far to the northwest as possible. No matter what the specific tastes of the investor for mean versus standard deviation, she will maximize preferences somewhere alone this tangency line, using some combination of the riskless asset and the Best Portfolio. The diagram shows for a particular indifference curve how the optimal portfolio, labeled OP in the figure, is determined. Different investors will locate at different points along the 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 BP mix $(\theta_1^*,...,\theta_K^*)$ 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 the Best Portfolio $(\theta_1^*,...,\theta_K^*)$ 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. Investors may purchase individual stocks in the BP proportions; however, there are mutual funds that do precisely this, and the investor may then simply put her stock portfolio into the market index mutual fund.

The problem of determining the optimal portfolio mix $(\theta_1^*,...,\theta_K^*)$ is most easily solved by considering a closely related problem. An investor's portfolio is characterized by A and $(\theta_1,...,\theta_K)$. 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:

Choose $A,\theta_1,...,\theta_K \geq 0$ to minimize $A^2 \sum_{k=1}^{K} \sum_{j=1}^{K} \theta_k\theta_j\sigma_{kj}$, subject to $A \sum_{k=1}^{K} \theta_k r_k = c$ and $\sum_{k=1}^{K} \theta_k = 1$, where c is a constant that can be varied parametrically. The first-order (Kuhn-Tucker) conditions for this problem are

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

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

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 best portfolio. The left-hand-side of (3) equals $\sigma_{kp} \equiv \beta_k\sigma_p^2$, so this condition can be rewritten

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

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

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

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 facilitating assumptions that make the statistical analysis tractable but whose economic plausibility is weak or undetermined.

1.5. CONCLUSION

A traditional view of econometrics is that it is the special field that deals with economic data and with statistical methods that can be employed to use these data to test economic hypotheses and make forecasts. If this were all there was to econometrics, it would still be one of the most important parts of the training of most economists, who in their professional careers deal with economic hypotheses, policy issues, and planning that in the final analysis depend on facts and the interpretation of facts. However, the examples in this chapter are intended to show that the subject of econometrics is far more deeply intertwined with economic science, providing tools for modeling core theories of the behavior of economic agents under uncertainty, and a template for rational decision-making under incomplete information. This suggests that it is useful to understand econometrics at three levels: (1) the relatively straightforward and mechanical procedures for applied econometric data analysis and inference that are needed to understand and carry out empirical economics research, (2) a deeper knowledge of the theory of statistics in the form that it is needed to develop and adapt econometric tools for the situations frequently encountered in applied work where conventional techniques may not apply, or may not make efficient use of the data, and (3) a deeper knowledge of the concepts and formal theory of probability, statistics, and decision theory that enter models of the behavior of economic agents, and show the conceptual unity of econometrics and economic theory.