

Expectations: Rational and Otherwise

The expectations of economic agents concerning the uncertain outcomes of future event often play important roles in determining market outcomes, both in the present and in the future. For example, the current demand and supplies of oil, bonds, and other commodities depend on the expectations of their future prices on the part of buyers and sellers in those markets. Thus, the current prices of these commodities depend on expectations of their future prices. Similarly, expectations of future economy-wide inflation can drive up the demand for goods and credit today, causing inflation and interest rates to rise today, and expectations of recession can lead to a fall in demand, creating a self fulfilling prophesy.

Rational Expectations

There are a variety of definitions of *rational expectations* in the economics literature. Intuitively, a rational expectation should be a best possible forecast based on all available information. We can operationalize this by saying that the rational expectation should be a mathematical expectation (expected value). Call $X_{t+j}^{e(t)}$ the expectation at time t of the future random variable X_{t+j} . Then if expectations are rational,

$$X_{t+j}^{e(t)} = \mathbf{E}_t X_{t+j}$$

where \mathbf{E}_t denotes the *conditional expectation* based on information available at time t . The specific information available at time t will define different probabilities p_i of events than the unconditional probabilities (i.e., long run frequencies) which correspond to the unconditional expectation $\mathbf{E}X_{t+j}$.

An example might be useful here. Consider the outcome of a sequence of two coin tosses. Value a head at \$1 and a tail at \$0. The probability of either a head or a tail in any toss is .5. Call the random variable Z the cumulated return to the two tosses. The unconditional expectation of the cumulated return is $\mathbf{E}(Z) = \$0.5 + \$0.5 = \$1$. However, if inbetween the two tosses (call this time t_o), we are told that the first toss was a head, then the conditional expectation of Z is now $\mathbf{E}_{t_o} Z = \$1 + \$0.5 = \$1.5$. If we had been told that the first toss had been a tail, then the conditional expectations would have been $\mathbf{E}_{t_o} Z = \$0 + \$0.5 = \$0.5$.

Now this is an easy example, since we can argue (as long as we are willing to avoid deep questions about the nature of randomness) that there are objective probabilities attached to a flip of a coin.¹ However, the macro economy is much more complicated than a controlled game like a coin toss, and this has lead many economists to question the relevance of this approach. For example, John Maynard Keynes (1935) argued that the expected returns on investment projects could not be calculated as mathematical expectations because there was no way of establishing good probabilities for the success of new products and processes or for the myriad of events that might take place in the macro economy in the future. As further support for this claim, he noted that macroeconomic events which we might try to forecast (for example the likelihood of a recession) depend themselves on market participants' expectations of them; thus my expectations must depend on my beliefs about others' expectations. Thus, Keynes believed that what is now called rational expectations (the term was coined much later, in the 1960s) is not an appropriate theory of expectations formation in the macro economy.² If it is not clear to agents in the economy how their current information

¹ In reality, the flip of the coin is deterministic. However, since we typically don't have the computational ability to forecast the result, we pragmatically assume that the outcome is truly random. Similarly, computer random number generators are deterministic (there is an algorithm that produces the numbers, but if you don't know the algorithm it's pretty hard to predict the next number). It is now standard to model quantum events as being intrinsically random, but again this is a mater of some debate (does God play dice?).

² Keynes believed that mathematical expectations *were* useful in describing expectations formation in certain situations, for example in insurance markets in which there are relatively stable long term frequencies of events (an 18 year old driving a loaded sports car will on average cost the insurance company more than a 35 year old economics professor driving a family mini-van) and good data collection.

translates into probabilities of various outcomes occurring (what is the likelihood that inflation will be at 8% five years from now?), then expectations may be indeterminate or “irrational” (the meaning of which is unclear).

At a pragmatic level, the assumption of rational expectations is often operationalized by asserting that if agents are making the best forecasts that they can given the available information, then certain conditions should hold. Particularly, consider the expectational error ϵ_{t+j} defined as

$$\epsilon_{t+j} = X_{t+j} - X_{t+j}^{e(t)}$$

This is just how far off your expectation turns out to be. If expectations are rational (i.e., people are making the best forecasts that they can), then the following conditions should hold.

1. The expectational error ought to be zero on average. I.e., the unconditional expectation of the forecast error ought to be zero:

$$\mathbf{E}\epsilon_{t+j} = 0$$

This says that expectations are not *systematically* incorrect. We do not, for example, *continuously* under- or over- estimate future rates of inflation or returns to a security over long periods of time.

2. The expectational error ought to be uncorrelated with all variables whose values are known at time t :

$$\mathbf{Cov}(\epsilon_{t+j}, Y) = 0$$

where Y is any variable whose value is known at time t . This says that all known correlations between other variables and the variable that we are trying to forecast have been exploited in our forecast. If this were not true, then we could improve our forecasts by taking these correlations into consideration. For example, if current raw materials prices are correlated with the general price level three months in the future, then information on this “leading indicator” will be exploited in the formation of inflationary expectations. Similarly, if stock prices systematically behave differently in January than in other months, this typical behavior will be anticipated by market participants.

These conditions are very powerful and are the foundation of rational expectation econometrics.

To summarize, rational expectations is a formalization that can be used in building and solving economic models and can be used as an assumption to econometrically identify economic relationships involving expectations. Whether it is more valuable as an ideal benchmark for modeling or as an empirical approximation to the real world is a matter of some debate. The Nobel prize in economics was shared in 2013 by Eugene Fama and Robert Shiller who come out on opposite sides of that debate. Fama tends to see the gyrations of prices in financial markets as reflecting efficient markets in which expectations are responding rationally to new information about economic fundamentals, while Shiller sees the same gyrations as reflecting excess volatility and bubbles driven by swings in irrational exuberance.

Below we consider a famous empirical test of Shiller’s for the presence of excess volatility in the data.

Are Stock Prices Excessively Volatile?

Rational expectations is one of the foundations of efficient market theory (EMT) in financial market theory.

According to the EMT, stock prices are based on market fundamentals rather than speculative behavior not related to market fundamentals. Consider a stock with uncertain (random) future dividends d_{t+i} . Suppose for simplicity that there is a constant interest rate R on other securities in this security’s risk class. Then the EMT predicts that the price of this stock at time t should be the rational expectation of the present value of future dividends paid on the stock:

$$\begin{aligned} (1) \quad P_t &= \mathbf{E}_t PV_t \\ &= \mathbf{E}_t \sum_{i=1}^{\infty} \left(\frac{1}{1+R} \right)^i d_{t+i} \end{aligned}$$

$$= \sum_{i=1}^{\infty} \left(\frac{1}{1+R}\right)^i \mathbf{E}_t d_{t+i}$$

Call this the *fundamental value* of the stock at time t .

We will return to the derivation of this later, to ask whether *rational* speculative bubbles might (in theory) make the price deviate from its fundamental value.

For now, however, consider Shiller's (1981) empirical test for whether stock prices deviate from their fundamental values in practice. He tests to see if actual stock prices are more volatile than should be the case if they follow expected future dividends as above.

To see how he does this, note that we can express the *actual* present value PV_t of future dividends at time t (i.e., the present value of the actual future realizations of the dividends) as our expectation of this present value, plus our forecast error ϵ_t (i.e., how wrong our expectation turns out to be):

$$\begin{aligned} PV_t &= \mathbf{E}_t PV_t + [PV_t - \mathbf{E}_t PV_t] \\ PV_t &= \mathbf{E}_t PV_t + \epsilon_t \end{aligned}$$

Then by (1), we have

$$PV_t = P_t + \epsilon_t$$

Now to look at volatility, take the variance of both sides:

$$\mathbf{Var}(PV_t) = \mathbf{Var}(P_t) + \mathbf{Var}(\epsilon_t) + 2\mathbf{Cov}(P_t, \epsilon_t)$$

The covariance is zero under rational expectations; since today's price is observed today, the forecast error should not be correlated with it. Thus,

$$\mathbf{Var}(PV_t) = \mathbf{Var}(P_t) + \mathbf{Var}(\epsilon_t) \geq \mathbf{Var}(P_t)$$

(variances can not be negative). Thus, we have the prediction, from the EMT, that the variance of stock prices should *not* be greater than the variance of the present value of actual future dividends on the stock. Intuitively, the price of a stock should be based on the *expected* values of future dividends, which are less volatile (over time for example) than the *actual* realizations of future dividends.

This is a testable hypothesis, since we have historical data on stock prices and subsequent dividends. Since $\left(\frac{1}{1+R}\right)^i$ becomes small as i becomes large, we can get a decent measure of PV_t with a finite number of subsequent periods. Shiller constructed time series for stock prices and (partial) present values of subsequent dividends for a variety of stocks, and finds that the sample variances of price tends to be *greater than* the sample variances of dividends, violating the prediction of the EMT.

Fundamental Equilibrium and Rational Bubbles in Asset Markets

Note that Shiller's test is a test of the joint hypothesis that the price of a stock at any time is equal to the expected present value at that time of future dividends and that expectations are rational. The first part of the hypothesis (equation (1) above) can be derived from a simple arbitrage condition stating that competition should drive the expected return on the stock to be equal to the expected return on any other asset that is perceived to be a close substitute investment.

It turns out that equation (1), however, is only one of a large set of possible solutions to this arbitrage condition. This solution is often called the *fundamental solution*, as it says that the price of the asset is based solely on its economic fundamentals. The other solutions can be thought of as rational bubble solutions.

Let's derive first equation (1) and then the full set of rational expectations equilibria for the price of the stock based on a simple arbitrage condition.

The expected one period holding return to a stock from t to $t + 1$ is the expected dividend plus the expected capital gain as a percentage of the purchase price.

$$h_t = \frac{\mathbf{E}_t d_{t+1} + (\mathbf{E}_t P_{t+1} - P_t)}{P_t}$$

If other securities in a similar risk class yield an expected rate of return R , then arbitrage would require that $h_t = R$. Rewriting,

$$(2) \quad P_t = \frac{1}{1+R} \mathbf{E}_t d_{t+1} + \frac{1}{1+R} \mathbf{E}_t P_{t+1}$$

Equation (2) is the basic arbitrage condition that any rational expectations equilibrium path of the stock price must satisfy.

The Fundamental Solution

We can solve equation (2) as follows to get equation (1). We consider a representative investor who understands that the stock price will be determined by equation (2) today and in each future period t and also believes that all other investors also understand this and are rational.

Knowing that the price level at *any time* t follows (2), the investor can form her forecasts of future prices *based on* (2). From (2) she knows that the price level tomorrow will actually be

$$P_{t+1} = \frac{1}{1+R} \mathbf{E}_{t+1} d_{t+2} + \frac{1}{1+R} \mathbf{E}_{t+1} P_{t+2}$$

Thus, her rational expectation of this price level must be

$$\mathbf{E}_t P_{t+1} = \frac{1}{1+R} \mathbf{E}_t [\mathbf{E}_{t+1} d_{t+2}] + \frac{1}{1+R} \mathbf{E}_t [\mathbf{E}_{t+1} P_{t+2}]$$

which is³

$$(3) \quad \mathbf{E}_t P_{t+1} = \frac{1}{1+R} \mathbf{E}_t d_{t+2} + \frac{1}{1+R} \mathbf{E}_t P_{t+2}$$

Similarly, $\mathbf{E}_t P_{t+2}$ can be based on what is known about the actual behavior of P_{t+2} according to (2)

$$P_{t+2} = \frac{1}{1+R} \mathbf{E}_{t+2} d_{t+3} + \frac{1}{1+R} \mathbf{E}_{t+2} P_{t+3}$$

so

$$(4) \quad \mathbf{E}_t P_{t+2} = \frac{1}{1+R} \mathbf{E}_t d_{t+3} + \frac{1}{1+R} \mathbf{E}_t P_{t+3}$$

Now if we continue in this vein for $\mathbf{E}_t P_{t+3} \dots$ etc., and substitute (3) into (2) and then (4) into the resulting equation, and so forth, we get

$$P_t = \frac{1}{1+R} \mathbf{E}_t d_{t+1} + \left(\frac{1}{1+R}\right)^2 \mathbf{E}_t d_{t+2} + \left(\frac{1}{1+R}\right)^3 \mathbf{E}_t d_{t+3} + \dots$$

³ My best guess today about what my expectation will be tomorrow about the following period's price (or dividend) must just be my best guess today about that price (or dividend), so that e.g., $\mathbf{E}_t [\mathbf{E}_{t+1} P_{t+2}] = \mathbf{E}_t P_{t+2}$. If you are not convinced of this, you might think about the coin toss earlier in the handout. Before the first coin is tossed, what is your best guess as to what your best guess of Z will be after you discover the value of the first toss? Since you don't yet know what the value of the first toss is, it is just your current best guess of Z , which is \$1.0.

which can be written

$$P_t = \sum_{i=1}^{\infty} \left(\frac{1}{1+R}\right)^i \mathbf{E}_t d_{t+i}$$

which is just (1).

Thus, according to this solution, the rational expectations of future price levels resolve into rational expectations of future dividends. Consequently, so does the current equilibrium price level. By this theory, a stock's price is determined purely by rational expectations about the stock's fundamentals.

Example

Suppose that investors correctly believe that the dividend on the stock in each period t is determined by $d_t = \bar{d} + \varepsilon_t$, where \bar{d} is a constant and ε is a random variable with zero mean and finite variance. Then today's rational expectation of the dividend in any future period is just \bar{d} .⁴

Plugging this into (1), we get

$$\begin{aligned} P_t &= \sum_{i=1}^{\infty} \left(\frac{1}{1+R}\right)^i \mathbf{E}_t d_{t+i} \\ &= \sum_{i=1}^{\infty} \left(\frac{1}{1+R}\right)^i \bar{d} \\ &= \frac{\bar{d}}{R} \end{aligned}$$

So if, for example, $\bar{d} = \$10$ and $R = 0.05$, then the stock price will be \$200 in every period t . In this rational expectations equilibrium, the price should never vary from \$200, unless \bar{d} or R changes, even though dividends are fluctuating over time. In this example *any* price volatility not tied to rational expectations of future changes in R or \bar{d} would be excess volatility, not predicted by the EMT, of the sort that Shiller found in his data.

Bubble Solutions

However, (1) is only one solution to (2). To see this, let's consider (2) under a very simple case. Specifically, simplify things by assuming perfect foresight and setting the actual and expected dividend stream constant over time. Then, rewriting (2) so that P_{t+1} is on the left hand side, we have

$$(5) \quad P_{t+1} = -\bar{d} + (1+R)P_t$$

The steady state equilibrium is $P^* = \frac{\bar{d}}{R}$, the fundamental solution. Notice that you can solve for it here by simply setting $P_t = P_{t+1}$.

However, since the coefficient on P_t is $(1+R)$ which is greater than one, the steady state solution is unstable. If the price starts at $\frac{\bar{d}}{R}$ it will stay there forever. If it starts above this value, however, the price level will follow a rational speculative bubble diverging from P^* over time at rate R . One way to see this is to simply plug a starting value for the price above the fundamental value into equation (12) above, and see that the price will be even higher in the next period. If we iterate forward, the price will continue to grow away from the fundamental value over time at growth rate R .

The general solution to (5) is

$$(6) \quad P_t = P^* + (1+R)^t (P_0 - P^*)$$

What is the intuition for these rational bubble solutions? Let's continue with our example above in which $\bar{d} = 10$ and $R = 0.05$, so that $P^* = 200$. Then if you believe that the price will stay constant at

⁴ $\mathbf{E}_t d_{t+i} = \mathbf{E}_t(\bar{d} + \varepsilon_{t+i}) = \bar{d} + 0 = \bar{d}$.

\$200, then the expected capital gain on the stock is zero, and the expected return is $\bar{d}/200$ which is exactly 0.05. So, in that case, you should be willing to pay up to \$200 for the stock, as that is the price that leaves you indifferent between the stock and other securities. This is indeed the fundamental value of the stock at which the EMT predicts that it will be priced.

Now, under what condition would you pay \$300 for this stock? If you expected the price to rise by \$5 during the period, then this expected capital gain of \$5 plus the \$10 dividend would yield the required expected holding return of 5%. Therefore, $P_t = \$300$ is a rational expectations equilibrium if $\mathbf{E}_t P_{t+1} = \$305$.

And under what condition will the price actually be \$305 in $t + 1$? For that to be the case, we need $\mathbf{E}_t P_{t+2} = \$310.25$. I.e., if we expect the price to continue to grow away from \$200 at 5% per year, every year in the future, then this expectation will be self fulfilling, and the price level will follow exactly that path.

Similarly, suppose that you thought that the price may rise next period or the bubble might burst. Suppose for example that you believe that with probability 0.2 the price will fall back to \$200 next period (the bubble will burst), and with probability 0.8 the price will rise next period (the bubble will continue). Then, it is rational to pay \$300 for the stock today, if you believe that, *if* the bubble does not burst tomorrow, the price tomorrow will be \$331.25. In that case, you believe that there is an 80% chance of getting a \$31.25 capital gain, and a 20% chance of a \$100 capital loss, so that your expected capital gain is \$5). Adding the expected dividend of \$10, your expected rate of return on the stock is exactly 5%.

So ‘rational’ bubbles are theoretically possible. Whether they are plausible or empirically relevant is a matter of some controversy.

Among possible objections to these bubble solutions are the following. First, how will a large number of independent economic actors coordinate on a particular initial value of the price level other than $P = P^*$. Of course the same criticism can be levied at the fundamental solution $P = P^*$. Second, why should anyone expect that the price of the stock can rise indefinitely in the absence of rising dividends? Once a bubble starts, above, it must continue on forever. If the public does not believe that the price can grow forever, then the only solution to (2) left is the fundamental solution.

However, an alternative to believing that a bubble *will* last forever is to expect any bubble to burst eventually, but not to know with certainty when that will happen. The public might, for example, collectively coordinate the timing of bubbles with the outcomes of otherwise irrelevant events. Such equilibria are called *sunspot equilibria*. Again, skeptics would argue that these rational bubbles are unreasonable in that they require the public to believe that the bubble *could* continue for a long time (and to act on this belief).

Some skeptics of *rational* bubbles therefore prefer the fundamental rational expectations solution, while others argue that bubbles *are* indeed relevant, but that they are not grounded in rational behavior, rather in irrational or boundedly rational behavior. For example, some investors may buy into a bubble because they are blindly following the trend, or they may believe that the bubble will burst but are optimistic that they will be among the ones to get out first leaving others to hold the bag. As noted before, Shiller falls into the category of economists who believe both that there is widespread excess volatility and that this volatility is driven by less than fully rational behavior.

Bounded Rationality and Learning

While rational expectations is a very strong modeling assumption, it is considered to be an important modeling benchmark (how does the model behave if agents in it have rational expectations).

There may also be a back-door rationale for using rational expectations in macroeconomic and financial modeling. Suppose that in reality, many economic agents have fairly naive expectations, but learn over time from their mistakes. Then this learning may cause expectations to converge to rational expectations over time.

To see this, let’s consider condition (2) again, but drop the assumption of rational expectations. Let’s also simplify by assuming that the dividend stream is constant at \bar{d} , and that the investor knows this. In

this case, (2) is:

$$(2') \quad P_t = \frac{1}{1+R} \bar{d} + \frac{1}{1+R} P_{t+1}^{e(t)}$$

Now let's assume that the investor has a naive expectation that that the price level in the future will be constant at some level \tilde{P} . Then, by (2'), the actual price level today will be

$$(7) \quad P_t = \frac{1}{1+R} \bar{d} + \frac{1}{1+R} \tilde{P}$$

which can be written in terms of P^* as follows

$$(8) \quad P_t = P^* + \left(\frac{\tilde{P} - P^*}{1+R} \right)$$

We can see by (8) that P_t lies between P^* and \tilde{P} . So given the expectation that the price in the next period will be \tilde{P} , arbitrage between the stock and the bond makes the actual equilibrium price lie between this expectation and the fundamental rational expectations equilibrium P^* . If investors stick to their fixed expectations, then they will see over time that this expectation is systematically wrong every period. If they eventually *learn* to recognize their mistake, they will adjust their expectation toward the actual price, which is also in the direction of the rational expectations equilibrium price P^* . Thus, over time, learning will drive both the expected price \tilde{P} and the actual equilibrium price P toward P^* .

For example, if $\bar{d} = \$10$ and $R = 0.05$, so that $P^* = 200$, suppose that the average investor believed that the price tomorrow would be \$250. Then by arbitrage (2') the price today will be $P_t = \frac{260}{1.05} = \247.62 . Now suppose that after seeing that the price today was \$247.62, investors update their expectations in the following period so that $\tilde{P}_{t+1} = \$247.62$. Then the actual price in period $t+1$ will be $P_{t+1} = \frac{257.62}{1.05} = \245.35 . So as expectations adapt to the actual history of prices, both expected and actual price gravitate in the direction of $P^* = 200$. Once \tilde{P} reaches 200, then P and \tilde{P} will both stay at \$200.

So we see that in this very simple case, learning should eventually drive expectations to be rational. Indeed, in this example, learning selects the *fundamental* rational expectations equilibrium (not rational bubble solutions to the model). So at least in this very simple example, learning supports the EMT as a longer run outcome, rather than as a condition that is fulfilled in every period.

Before we get too excited about this result however, note that for adaptive expectations to respond rationally in a changing environment, the learning process would have to be substantially more sophisticated. Some economists have argued that we should expect this to be the case, while others (including myself) remain skeptical. See Georges (2008) for a discussion of how learning can fail to converge to rational expectations and Georges and Pereira (2020) for examples with machine learning algorithmic traders.

References

- Georges, Christophe “Bounded Memory, Overparameterized Forecast Rules, and Instability,” *Economics Letters*, **98**, 2008, 129–135.
- Georges, Christophe, and Javier Pereira. “Market Stability with Machine Learning Traders.” Manuscript, Department of Economics, Hamilton College. July 2018, revised January 2020.
- Keynes, John Maynard, *The General Theory of Employment, Interest, and Money*, 1935.
- Shiller, Robert, “Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?,” *American Economic Review*, **71**, 1981, 421–436.

Appendix (optional)

Digression on Solutions to Difference Equations

Consider a first order linear difference equation in a variable x with constant coefficients a and b :

$$(9) \quad x_{t+1} = a + b x_t$$

Here are some useful results concerning the behavior of x under (9). The steady state value of x is $x^* = \frac{a}{1-b}$. If $0 < b < 1$, then x converges monotonically to x^* from any starting value. If $b > 1$, then x diverges from x^* monotonically at rate $b - 1$, if it starts away from x^* .⁵ To solve for the path of x , either plug in a known initial value of x on the right side of (9), and solve forward iteratively, or use the general solution which can be calculated as follows. Calculate the value of x^* . Subtracting x^* from both sides of (9) gives

$$(x_{t+1} - x^*) = b(x_t - x^*)$$


Iterating this forward from some time zero implies that

$$(x_t - x^*) = b^t(x_0 - x^*)$$

or

$$x_t = x^* + b^t(x_0 - x^*)$$

This formula gives the time t value of x for *any* initial value x_0 and any value of b . Again, we can see that if $|b| < 1$, x converges to x^* over time, since b^t gets smaller over time, and if $b > 1$, the gap between x and x^* grows at constant growth rate $(b - 1)$.



A very simple example might be useful. Suppose that $a = 0$, so that the difference equation that we want to solve is

$$x_{t+1} = b x_t$$

Then the steady state equilibrium is $x^* = 0$. Suppose, however, that x starts at time 0 at $x_0 = 1$. First let $b = 2$. Then $x_1 = 2$, $x_2 = 4$, $x_3 = 8$, etc. $x - x^*$ grows at rate $b - 1 = 1$ (100% per period). Now let $b = 0.5$. Then $x_1 = 0.5$, $x_2 = 0.25$, $x_3 = 0.125$, etc. x converges on x^* over time.

In (5), since the coefficient on P_t is greater than one, P^* is an unstable steady state equilibrium. If P_0 starts either above or below P^* it will grow further away from P^* over time.

⁵ If $-1 < b < 0$, x also converges to x^* , but oscillates, and for $b < -1$ x undergoes explosive oscillations away from x^* .

Review of Statistics: Notation and Definitions

Consider two random variables X and Y defined over m distinct possible events. Event i occurs with probability p_i , in which case X and Y take on values x_i and y_i . Thus the probabilities of the various events occurring are p_1, p_2, \dots, p_m , and X and Y take on possible values x_1, x_2, \dots, x_m and y_1, y_2, \dots, y_m respectively. If we have considered all possible events, then it must be that the sum of the m probabilities is one: $\sum_{i=1}^m p_i = 1$.

For example, consider the following game. We flip a coin once. In the event that it lands heads-up, Beth and Bob each get \$1. In the event the coin lands tails-up, Beth gets \$2 and Bob gets nothing. Here there are two possible events (heads and tails), which occur with equal probability $p_1 = p_2 = 0.5$. The payoffs to Beth and Bob are each random variables which we could call X and Y .

We will use the following notation

$$\begin{aligned}\mathbf{E}(X) &= \text{The Expected Value of } X \\ \mathbf{Var}(X) &= \text{The Variance of } X \\ \mathbf{Cov}(X, Y) &= \text{The Covariance of } X \text{ and } Y \\ \mu_X &= \mathbf{E}(X) \\ \sigma_X &= \text{The Standard Deviation of } X \\ \rho_{XY} &= \text{The Correlation of } X \text{ and } Y\end{aligned}$$

These quantities are defined as follows.

$$\begin{aligned}\mathbf{E}(X) &= p_1x_1 + p_2x_2 + \dots + p_mx_m \\ &= \sum_{i=1}^m p_ix_i \\ \mathbf{Var}(X) &= p_1(x_1 - \mu_X)^2 + p_2(x_2 - \mu_X)^2 + \dots + p_m(x_m - \mu_X)^2 \\ &= \sum_{i=1}^m p_i(x_i - \mu_X)^2 \\ &= \mathbf{E}[(X - \mu_X)^2] \\ \sigma_X &= \sqrt{\mathbf{Var}(X)} \\ \mathbf{Cov}(X, Y) &= p_1(x_1 - \mu_X)(y_1 - \mu_Y) + p_2(x_2 - \mu_X)(y_2 - \mu_Y) + \dots + p_m(x_m - \mu_X)(y_m - \mu_Y) \\ &= \sum_{i=1}^m p_i(x_i - \mu_X)(y_i - \mu_Y) \\ &= \mathbf{E}[(X - \mu_X)(Y - \mu_Y)] \\ \rho_{XY} &= \frac{\mathbf{Cov}(X, Y)}{\sigma_X\sigma_Y}\end{aligned}$$

The expected value of X is an average or mean value of the realizations x which occur with frequencies p . The variance of X is a measure of the average amount of variation of the realizations x around their mean value. The covariance of X and Y is a measure of the degree to which values of X which are larger than average tend to coincide with values of Y which are larger or smaller than average. For example, a negative covariance indicates that when the realization of X is greater than μ_X , the corresponding realization of Y tends, on average, to be less than μ_Y . The correlation of X and Y is the covariance normalized so that this value falls between -1 and 1 . We say that X and Y are perfectly correlated if this correlation is -1 or 1 .

Returning to our example of Beth and Bob, above, we have $\mathbf{E}(X) = \$1.50$, $\mathbf{E}(Y) = \$0.50$, $\mathbf{Var}(X) = \mathbf{Var}(Y) = 0.25$, $\sigma_X = \sigma_Y = 0.5$, $\mathbf{Cov}(X, Y) = -0.25$, and $\rho_{XY} = -1$.

Review of Statistics: Identities

Let a , b , and c be arbitrary constants. The following identities follow from the definitions above.

$$\begin{aligned}\mathbf{E}(a) &= a \\ \mathbf{E}(a + X) &= a + \mathbf{E}(X) \\ \mathbf{E}(bX) &= b\mathbf{E}(X) \\ \mathbf{E}(a + bX) &= a + b\mathbf{E}(X) \\ \mathbf{E}(X + Y) &= \mathbf{E}(X) + \mathbf{E}(Y) \\ \mathbf{E}\left(\sum_{i=1}^n X_i\right) &= \sum_{i=1}^n \mathbf{E}(X_i)\end{aligned}$$

In the last identity, X_1, X_2, \dots, X_n are n random variables.

$$\begin{aligned}\mathbf{Var}(a) &= 0 \\ \mathbf{Var}(a + X) &= \mathbf{Var}(X) \\ \mathbf{Var}(bX) &= b^2\mathbf{Var}(X) \\ \mathbf{Var}(X + Y) &= \mathbf{Var}(X) + \mathbf{Var}(Y) + 2\mathbf{Cov}(X, Y) \\ \mathbf{Var}\left(\sum_{i=1}^n X_i\right) &= \sum_{i=1}^n \mathbf{Var}(X_i) + \sum_{i=1}^n \sum_{\substack{j=1 \\ j \neq i}}^n \mathbf{Cov}(X_i, X_j)\end{aligned}$$

It follows from the last identity that **if** n random variables are mutually uncorrelated, the variance of their sum is equal to the sum of their variances.

$$\begin{aligned}\mathbf{Cov}(X, Y) &= \mathbf{Cov}(Y, X) \\ \mathbf{Cov}(X, X) &= \mathbf{Var}(X) \\ \mathbf{Cov}(a, X) &= 0 \\ \mathbf{Cov}(a + X, b + Y) &= \mathbf{Cov}(X, Y) \\ \mathbf{Cov}(bX, Y) &= b\mathbf{Cov}(X, Y) \\ \mathbf{Cov}(bX, cY) &= bc\mathbf{Cov}(X, Y) \\ \mathbf{Cov}(X, Y + Z) &= \mathbf{Cov}(X, Y) + \mathbf{Cov}(X, Z)\end{aligned}$$

These identities can be proved easily by employing the definitions given above. For example, the identity $\mathbf{Var}(a + X) = \mathbf{Var}(X)$ can be proved as follows:

$$\begin{aligned}\mathbf{Var}(a + X) &= \mathbf{E}[(a + X) - \mathbf{E}(a + X)]^2 \\ &= \mathbf{E}[(a - \mathbf{E}(a)) + (X - \mathbf{E}(X))]^2 \\ &= \mathbf{E}[(X - \mathbf{E}(X))^2] \\ &= \mathbf{Var}(X)\end{aligned}$$