1 Introduction

We begin this chapter by assuming that traders are value maximizers. In this case competitive equilibrium predicts that asset prices follow a random walk, similar to the prediction we derived for a sequential auction, and that as the liquidation date of an asset approaches the variance in its current price increases. A value maximizer cares only about the mean return of his portfolio, not its higher order moments, whereas a risk averse investor, and also a risk loving investor, has preferences over the portfolio variance as well.

When there is differential information amongst traders about product quality, the competitive equilibrium price not only equilibrates supply with demand. It also aggregates the information of the more informed traders. How much trading takes place in competitive equilibrium compared to the full information analogue depends on the precise nature of the informational asymmetries. Two extreme cases are that the competitive equilibrium is not affected at all, and that there is massive market failure with no trading. Section 3 analyzes these two extremes, as well as an intermediate case, where traders form rational expectations about the product’s value to them given the competitive price that demanders pay and suppliers receive.

In contrast to most other goods and services, which are often differentiated and tailored to individual customers and customer segments, financial assets provide services that are defined so that one unit of a given stock is a perfect substitute for every other. Furthermore the differences between financial assets are in terms of the risk characteristics, which are potentially observable over time. Finally, a large fraction of an economy’s wealth is traded in financial markets. Since studying financial markets is such a fruitful area for applying the tools we have developed to trading strategies in experimental markets, many of the examples we discuss below are concerned with stocks and currency exchange.

The final section of this chapter applies competitive equilibrium to financial markets. The last three applications in this section assume traders are risk averse expected utility maximizers. We investigate how risky assets are discounted in competitive equilibrium, relative to bonds and also relative to each other.

2 How Competitive Equilibrium Reflects Information

All our examples in this chapter have assumed that each trader is fully informed about demand and supply conditions. However a competitive equilibrium can be defined in markets where the traders are not fully informed. A feature of competitive equilibrium is that it reflects the knowledge demanders and suppliers have about tastes and technologies.

We demonstrate this in a model where traders maximize their expected value, an assumption that can plausibly applied to firms in order to explore competitive equilibrium predictions in asset markets where all traders are symmetrically informed, and then intrading games with incomplete information. In the latter parts of this chapter we shall assume that consumer exhibit risk aversion, and this leads them to take out insurance at actuarially unfair rates, and manage portfolios of financial assets with a view to looking beyond the expected return.

2.1 Present value maximization under uncertainty

Consider an economy where there is uncertainty, but all the players are risk neutral, seeking to hold assets that maximize their discounted present value. Consider for example a model in which a succession of traders enter a market to trade in an asset. For convenience we label the private valuation of each trader by the time he enters the market. Thus v_t is the valuation of the trader who enters the market at time t. Once the last trader has entered, there are T traders with valuations $\{v_t\}_{t=1}^{T}$ who trade to reach an equilibrium price of $p(v_1, v_2, \ldots, v_T)$ solved following the methods in the first section of this chapter. Given wealth maximization the price in the second last period, $T-1$, is p_{T-1} defined as

$$p_{T-1} \equiv E[p(v_1, v_2, \ldots, v_T)|v_1, v_2, \ldots, v_{T-1}]$$

and by an induction argument, the competitive equilibrium price in period t is

$$p_t \equiv E[p(v_1, v_2, \ldots, v_T)|v_1, v_2, \ldots, v_t]$$

Comparing the competitive equilibrium price in period t and $t+1$, we obtain

$$p_{t+1} - p_t = p_{t+1} - E[p(v_1, v_2, \ldots, v_T)|v_1, v_2, \ldots, v_t]$$

appealing to the law of iterated expectations as of time t yields

$$\begin{aligned} & p_{t+1} - E\{E[p(v_1, v_2, \ldots, v_T)|v_1, v_2, \ldots, v_{t+1}]|v_1, v_2, \ldots, v_t\} \\ &= p_{t+1} - E\{p_{t+1}|v_1, v_2, \ldots, v_t\} \end{aligned}$$

Once again the efficient market hypothesis reveals that the competitive equilibrium price sequence must follow a random walk. We remark that no restrictions have been placed on the probability distributions driving the sequence of random variables of $\{v_t\}_{t=1}^{T}$.

To be more concrete, we consider an asset that pays off dividends of d_t at fixed or random intervals through time until some date T when the asset is liquidated for some value r_T. In addition to these common component there is a private value v_n the n^{th} player receives at the end of the game for each unit of the asset held. The expected present value at date s is then the sum of the discounted dividends plus the scrap value at the end of the game, namely:

$$\pi_s = E_s\left[\sum_{t=s}^{T} \beta_t d_t\right] + E_s[\beta_T r_T]$$

If everyone was risk neutral, and sequentially maximized the expected value of their terminal wealth, the value of their portfolio would be

$$E_s\left[\sum_{t=s}^{T} \beta_t d_t\right] + E_s[\beta_T r_T] + v_n$$

From the discussion in the previous section it is to compute the competitive equilibrium price at time s. Let p_0 denote the equilibrium price when there are neither dividends nor scrap value, that is $\pi_s = 0$ for all s and $r_T \equiv 0$. Clearly p_0 is determined using the methods we developed above. Then the equilibrium price at time s is $p_0 + \pi_s$. The change in price from one period to the next is then

$$\pi_{s+1} - \pi_s = E_s\left[\sum_{t=s}^{T} \beta_t d_t\right] + E_s[\beta_T r_T] - E_s\left[\sum_{t=s}^{T} \beta_t d_t\right] + E_s[\beta_T r_T]$$

$$=$$

Taking expectations we obtain the efficient markets hypothesis that prices behave like a random walk

$$E_s[\pi_{s+1}] - \pi_s = 0$$

The validity of the random walk hypothesis depends on three basic assumptions, that dividends received at different times are not discounted by different amounts, that traders are risk neutral, and that information about value is fully revealing. Aside from these conditions, the efficient markets hypothesis is a prominent feature of competitive equilibrium in a wide variety of trading environments.

For example consider a simple model of foreign currency exchange markets. Suppose U.S. export companies sporadically receive Euro and yuan injections from demanders for their goods in the E.U. between dates 0 and T. Similarly European (and Chinese) exporters sporadically receive injections of dollars and yuan (euros) for their sales in the U.S. and China (Europe). Export firms in each country also purchase domestic currency on the foreign exchange market between date 0 and T. At date T all export companies are liquidated, foreign currency earnings that have not been repatriated through the exchange market are expropriated, so no further value is placed on holding foreign currency. We also assume the U.S. dollar is a dominant currency, meaning all currency prices are quoted in dollars. We assume each export firm maximizes its expected dividend payments paid in domestic currency before the liquidation date T. The liquidation value is unknown at all dates t $<$ T, but as new information arrives about foreign trade throughout the trading phase, the traders become more informed about the value of foreign currency.

In competitive equilibrium the price of each exporter is the expected value of its dividend flow plus its liquidation value. Thus the efficient markets hypothesis implies in this model implies that the exchange rates follow a random walk. For supposing the dollar price of yuan is lower in date t than its expected price in date s $>$ t, Chinese exporters buying yuan at date t are not maximizing their value, because the expected value of their companies would be higher if they postponed yen purchases until date s. A symmetric argument applied to U.S. exporters explains why the the dollar price of yuan is not higher in date t than its expected price in date s $>$ t, and similar arguments apply to the dollar euro exchange market.

Exercise *Consider a sequence of entrants.*

1. *Derive and map the sequence of competitive equilibrium prices.*

2. *Plot the sequence of trading prices*

3. *Regress the sequence of trading prices on the sequence of competitive equilibrium prices. Is the regression coefficient significantly different from zero?*

4. *Square the residuals from the sequence of trading prices and the competitive equilibrium price sequence, and plot the results. Regress them on time. Can you detect evidence of learning about the competitive equilibrium price?*

2.2 Liquidity

In the previous example we may assume without loss of generality that there is continuous trading in the asset up until a common liquidation date T when the capitalized value of all the firms are recognized. How would prices behave if consumers had limited opportunities to enter and exit the market, effectively segmenting the market into different time markets? We consider the following variation on the foreign exchnage application to show the consequnces of illiquid markets.

Suppose exporters face the threat of their foreign earnings being confiscated, or there are incomplete markets that limit savings and borrowing opportunities in domestic markets. Then exporters have an incentive to immediately capitalize their foreign earnings by converting them to domestic dividends and distributing them as dividends.

In this case successive prices in the foreign exchange market would exhibit mean reversion. At the other extreme to the random walk observed in perfectly liquid market, prices in disconnected markets are independently distributed, and in a stationary environment, have the same conditional mean. Thus the hypothesis that asset prices follow a random walk might be interpreted as a test of liquidity.

2.3 Variance bounds

It has been argued that the the random walk hypothesis is quite weak, and that there are other implications from rational behavior on price processes in liquid markets, restrictions for higher order moments. As rational consumers process information, this affects the variances of prices. More specifically, variances of discounted summed dividend streams that condition on more information are larger than those which condition on less.

Let:

$$p_s^* = \sum_{t=s}^{T} d_t$$

denote the summed value of future dividends from time s until time T, the liquidation date. Also define:

$$p_s = E_s\left[\sum_{t=s}^{T} d_t\right]$$

as the competitive equilibrium asset price at time s, which is conditional on past dividend performance.

The competitive equilibrium price of the stock, p_s, is based on less information than p_s^*. From the definitions p_s^* incorporates more information than the competitive equilibrium price p_s about the value of future dividends and other events affecting the liquidation value of the asset.We now prove that:

$$var(p_s) < var(p_s^*)$$

Intuitively, impounding information into prices creates variation that reflects updating in the asset's value. To prove the inequality, note that since:

$$p_s^* \equiv p_s + (p_s^* - p_s)$$

it follows that:

$$var(p_s^*) = var(p_s) + var(p_s^* - p_s) + 2cov(p_s^* - p_s, p_s)$$

But

$$var(p_s^* - p_s) > 0$$

and:

$$\begin{aligned} cov(p_s^* - p_s, p_s) &= E[(p_s^* - p_s)p_s] - E[p_s^* - p_s]E[, p_s] \\ &= E[(p_s^* - p_s)p_s] \\ &= E\left\{\left(\sum_{t=s}^{T} d_t - E_s\left[\sum_{t=s}^{T} d_t\right]\right)E_s\left[\sum_{t=s}^{T} d_t\right]\right\} \\ &= 0 \end{aligned}$$

the last line following from the law of iterated expectations. The inequality bounding the variances of prices now follows directly.

Experiment 20.8

Random walk and variance bounds. The tests we have devised above exploit the knowledge that since the structure of the model is known to the designer of the experiment, the competitive equilibrium can be computed and then compared with the experimental outcome. It is, however, unnecessary to derive the equilibrium price sequence in order to check the efficient markets hypothesis. One weak test is to confirm that prices follow a random walk.

3 Aggregating Information in Competitive Equilibrium

In the trading games described above, all traders have the same information, and trade occurs only because of differences in stochastic endowments and preferences. We now investigate how differences in information between traders affect competitive equilibrium. We preface this discussion with the claim that if people have differential information about an asset they would all value the same way were they fully informed, then no trading will take place in competitive equilibrium. The proof is that if one trader party strictly benefits from the trading then the other party must lose. Since all traders anticipate this, it follows that no buyer and seller can mutually benefit from trade. So to motivate trade in markets where there is incomplete information, we remark that trade can only occur in competitive equilibrium, or for that matter any other solution to a (voluntary) trading mechanism, if there are differences in endowments and preferences, rather than differences in information alone.

3.1 Fully revealing competitive equilibrium prices

Competitive equilibrium economizes on the amount of information traders need to optimize their portfolios. Indeed a peculiar feature of competitive equilibrium is that in some situations it fully reveals private information to those who are less informed about market conditions.

To demonstrate this remarkable property of competitive equilibrium we now consider a market where consumer valuations for the good are private knowledge, and aggregate supply is fixed. In pariticular suppose no demander wants to consume more than one unit of the good, and each demander draws an identically and independently distributed random variable that determines their valuation for the first unit. The competitive equilibrium price does not depend on whether each trader observes the valuations of the others. Hence every trader acts the same way as he would if he were fully informed about aggregate demand. Therefore the price of the good does not depend on whether each trader knows the demand of all the others or not.

More generally, suppose there are N traders, and the nth trader receives a signal sn about the the state of the economy, where n = 1,2, . . ., N. We may suppose that in general, the competitive equilibrium price vector p depends on all the signals the traders receive. If, however, p is an invertible mapping of all the relevant information s available to traders, then every trader acts the same way as he would if he were fully informed. In this case p(s) has an inverse which we call f(p). Each trader realizes that seeing p is as good as seeing s. In the example above s is aggregate demand, and p is monotone increasing in s.

A more interesting example of how competitive equilibrium is when there is differential information about product quality. Accordingly suppose a component of each demander valuation is common, and traders have differential information about that component. The more favorable the signal to the informed traders, the greater is their demand, and hence the higher is the market clearing price. As in the previous example, uninformed traders compute their demands, deducing that if the market clearing price is p, then the common component is f(p). Thus informed traders cannot benefit from their superior information in competitive equilibrium.

Implications for trading mechanisms

Teh quality of as before we assume consumers are indexed by v and have utility function $u(v,w)$, where w is drawn from distribution function $G(w)$. If the productivity shock w was observed then all demanders with v satisfying the inequality $u(v,w) \geq p$

would purchase the product. Setting $u(v_0, w) = p$ and integrating over all $v \geq v_0$ we obtain the aggregate demand for the product $1 - F(v_0)$. Given an industry marginal curve that depends on quantity, in competitive equilibrium

$$c(q) = u(v_0, w)$$

In competitive equilibrium higher prices reveals the product's higher quality. Since p^e is a strictly increasing function of w it follows that consumers could use it as a sufficient statistice for w. Rather than even consulting about productivity they simply form a demand function

$$u(v, w(p)) > p$$

By construction the amount demanded by each consumer and the equilibrium allocations are the same. But supple every unit of the good is identical, then the rational expectations competitive equilibrium price fully reveals its quality. This concept is called a fully revealing rational expectations equilibrium, and implies that consumers. Price not only clears the market, but is a function of the aggregate shock that hits demanders.

But what if some are informed, and they increase the aggregate demand?

The argument that concumers should can infer quality from the price is debatable. Another definition of competitive equilibrium would set the realization of the productivity aftere the the market has cleared. Thus prices would account for the distribution of the signal but not its realization. Thus the threshold consumer solves

$$\int_{\underline{w}}^{\vec{w}} u(v_0, w) dG(w) = p^e$$

and realizes an expost utility of

$$u(v_0, w^*) - p^e$$

from the productivity shock w^*.

Our experiments compare a full information equilibrium with the competitive equilibrium and a rational expectations equilibrium. How would we construct an "uninformed competitive equilibrium" if there are in fact informed traders? An argument for using the concept of equilbrium that conditions on the information. And how many treaders does it take?

We compare rational expections with a competitive equilibrium where traders only now the distribution of the shock, and do not use price to infer its realization.

Some economists have used this theoretical result to argue that markets are good at aggregating the information that traders have about the preferences of demanders and the technologies of suppliers. Other economists have argued there is limited investment in acquiring new information relevant to suppliers and demanders, because those who use up resources to become better informed cannot recoup the benefit from their private information. Both arguments implicitly assume that a competitive equilibrium accurately predicts price and resource allocations from trading.

3.2 Insider trading in competitive equilibirum

When do competitive equilibrium prices hide information from uninformed traders? There are two reasons when competitive equilibrium prices are not fully revealing . First, the mapping from signals to prices p(s) might not be invertible. That is, two or more values of a signal, s1 and s2, would yield the same fully revealing competitive equilibrium price if everyone observed the signal's value, meaning pfr(s1) = pfr(s2). The second reason is that different units of the product might not be identical, although they are traded on the same market, and these differences are observed by some but not all the traders.

In this subsection we consider adding dimensions to the sources of uncertainty. In the previous example the uninformed segment of the population could infer the true state because a mapping exists from the competitive equilibrium price to the shock defining the product quality. Now we introduce a second shock. Those traders who observe one shock can infer the other from the competitive equilibrium price. Those who observe neither can only form estimates of what both shocks are from the competitive equilibrium price.

In the next example both the supply of teh commidity and its quality are random variables. Product quality is only known by some of the demanders, and the aggregate quantity supplied is not directly observed by anyone. In this case an uninformed demander cannot infer product quality from the competitive equilibrium price, because a high price could indicate high demand from informed traders or low supply. Appealing to the arguments we developed in the previous subsection informed demanders can however deduce aggreagte supply from the competitive equilibrium price. The reason why the uninformed segment of the population can infer the true state of affairs in the example above is because there is a mapping between the shock and the resulting competitive equilibrium price. They benefit from the fact that demand by uninformed traders is less than it would be if they were fully informed when product quality is high, depressing the price for high quality goods, and vice versa.

In our last examples on this topic we introduce a second shock. Those people who know one of the shocks can infer the other from the competitive equilibrium price. Those who know neither can only form estimates of what both shocks are from the competitive equilibrium price. Consider the following example. There are two groups of players, respectively called the suppliers and demanders of an asset. All demanders have the same valuation of the asset, either Suppliers are endowed with a random amount of the asset, but the proportion with each is random too, so total supply is unknown.

Compare a framework in which there are all private values with one in which there is a common value, and noisy signals to everyone, another where one person is differentially informed but has an additive valuation in his common and private component. The value of insider trading

We have just shown that there is no value from having private information about the prospects of a firm in competitive equilibrium if there are no other sources of uncertainty in the market. This argument implies insiders will not use up resources acquiring information about a firm's prospects if there is only one source of uncertainty. Only by recognizing multiple sources of uncertainty is there value from investing resources in market intelligence to gain inside information about the common value of an asset.

3.3 Differential information about heterogeneity across units

Suppose the quality of the individual units varies, and that traders are differentially informed it. What would competitive equilibrium theory predict about the price and quantity traded? We mentioned above the second reason why a competitive equilibrium is not fully revealing: the characteristics of the good, such as its quality, vary across units, but these differences are only observed by some of the traders, and for that reason, they are traded at the same price on a single market instead of several. Since traders condition their individual demand and supply on their information, more informed traders gain at the expense of the less informed. Two implications follow immediately. First there is no sorting between those consumers who have differential valutations for quality. Second, there there are fewer high quality units traded on the market than would be in a market where characteristics are fully observed, because uninformed potential buyers discount their willingness their reservation price to refelct the fact that sellers might trade them a low quality product. The prospect of being exploited by a well informed trader discourages a poorly informed player from trading. If every unit has a different quality, inefficiencies may result both in terms of the amount provided and their allocation to different demanders.

For example, consider a used car market. Suppose there are less cars than commuter traders, and no one demands more than one car. The valuation of a trader for owning one car is identically and independently distributed across the population. The quality of each car is independently and identically distributed across the population. Each owner, but no one else, knows the quality of his car. The amenity value from car travel is the product of the commuter's valuation and the quality of the car he owns.

In some industries where there is product differentiation it is hard to evaluate product quality before consumption. For example consider the market for wine from small vintners. Now supposing that the reputations of the various suppliers has not been established, we might suppose that all brands sell for the same uniform price. Entry into the industry then equates the expected value of the lowest demander with

$$\int_{\underline{w}}^{\vec{w}} u(v_0, w) dG(w) = c = u(\underline{v}, \underline{w})$$

Therefore $v_0 < \underline{v}$ and hence $w_0 < \underline{w}$. Clearly wine producers with quality greater than $\underline{w}$ are worse off than when there is labelling, while the lower quality producers with $w_0 \in (w_0, \underline{w})$ are better off. So high quality wine growers have an incentive to label the quality of their wine, and will do so if it is not costly. Similarly consumers with lower reservation prices for wine are better off. We know this for the new consumers. What about the high end of demand?

The assumption of constant costs is critical in this analysis. We now suppose that costs are increasing in quatiliy. Depending on potential for supply, In this case might be no competitive equilibrium except at the lowest possible quality of wine. If in addition

$$c'(w) < u_2(F^{-1}[G(w)], w)$$

then the same equilibrium exsits for the fully informed case.

4 Risk Sharing in Competitive Equilibrium

Models of competitive equilibrium have been extensively applied in portfolio management in financial markets, and the final section in this chapter turns to portfolio management and asset pricing. We investigate the role risk aversion plays in portfolio choice, as it applies to the relationship between financial returns of different assets and, within dynamic settings, the consumption profiles of investors.

Our point of departure is that if traders only cared about the mean return of an asset, it is hard to justify why assets could have different mean returns. Because there is abundant evidence that assets have different mean returns, traders apparently care about other moments of the probability distribution apart from the first. For example, supposing traders are risk averse, rather than risk neutral, they would hold diversified portfolios.

4.1 Risk aversion in financial markets

Modern asset pricing literature is concerned about how traders finance their consumption schemes, and how disruptive volatility in asset pricing is in affecting their plans. The less revisions these plans require and the higher consumption they permit, the better of they are. Asset portfolios that support such plans are judged superior. We remark straightaway that volatility in individual asset returns is neither good nor bad necessarily; the key is whether such volatility helps to offset other disturbances to financing consumption plans or exacerbates them. with the trade off between expected return and its variance. The motivating issue for an investor ultimately depends on the type of lifestyle

Discounting Risky Stocks

A corollary of the fact that the volatility declines with the term of the asset is that if traders are risk averse then the asset value converges from below to its liquidation value. The martingale property is violated in a specific way.

Experiment 20.9

Review the evidence above. Is there an upward trend in the prices of securities as the liquidation date approaches?

4.2 Portfolio choices

Quadratic preferences

In the simplest version of CAPM, we do not constrain traders to hold positive amounts of q_k. This effectively allows investors to make short sales, selling assets they do not have. However they remain liable for those assets, and thus reap the negative of the return on the asset, to account for the fact that to rebalance their position at the end of the game, they must purchase the same number of units they shorted on and absorb the return. This version of the model implies that at the end of the game they may become insolvent, because they cannot pay.

To simplify the analysis, we shall focus on an interior first order condition in which it is optimal to hold strictly nonzero amount of each security. This assumption can be rationalized in two ways. We could assume that is it optimal for the investor to hold strictly positive amounts of each security. Whether that is justified or not depends on our assumptions about the means and variances of the securities. Alternatively, we could relax the assumption that investors are required to be solvent at the end of trading, in other words relax the and permit short sale A third approach is to simply focus on those assets which are purchased in positive quantities and ignore those assets are not held.

Recall from Chapter 4 that an investor with quadratic preferences

For all interior $k \in \{1, \ldots, K\}$

$$E[c(\pi_k - \pi_0)] = 0$$

To derive the equation that characterizes CAPM, we define the return on the market portfolio as

$$\pi_m \sum_{n=1}^{N} c_0 = \sum_{n=1}^{N} c_n$$

It now follows that

$$E[\pi_m(\pi_k - \pi_0)] = 0$$

or

$$E[\pi_m \pi_k] = E[\pi_0]E[\pi_m]$$

Subtracting $E[\pi_m]E[\pi_k]$ from both sides of this equation yields

$$E[\pi_m \pi_k] - E[\pi_m]E[\pi_k] = -E[\pi_k]E[\pi_m] + E[\pi_0]E[\pi_m]$$

Applying the definition of a covariance to the left side we see that

$$cov(\pi_m, \pi_k) = -E[\pi_k - \pi_0]E[\pi_m]$$

This condition is not only true for each of the individual securities, but also hold for the market index itself, a fact that can be derived directly, or by including the market index as one of the K risky securities in the original definition of the problem. Consequently

$$var(\pi_m) = -E[\pi_m - \pi_0]E[\pi_m]$$

Define the coefficient β_k as the ratio of the covariance of the stock return with the market return to the variance of the market return

$$\beta_k = \frac{cov(\pi_m, \pi_j)}{var(\pi_m)}$$

The of CAPM equation is obtained by combining the two equations to eliminate $E[\pi_m]$:

$$E[\pi_k - \pi_0] = \beta_k E[\pi_m - \pi_0]$$

In words, CAPM predicts that the excess return on a stock over the risk free rate is proportional in expectation to the excess return of the market portfolio, where the coefficient of proportionality is β_k as defined above. That is the price of stocks at the beginning of the period to make this equation hold if players are to choose an interior solution.

4.3 Dynamic considerations

The demand for financial assets comes from individuals and households holding wealth in financial securities to defer consumption. For example parents save for the education of their children, individuals save for retirement, and the wealthy bequest future generations with their largesse. Half of the value of the stock market is held by a very small fraction of individuals. Nevertheless more than 50 percent of households hold financial securities of some form or other. Collectively, these groups, including foreign investors, create the demand for financial securities. On the supply side, the main financial securities are stocks, bonds and their derivatives issued by corporations and private enterprises to finance their operations, mortgage backed securities that bundle loans on housing stock, bonds issued by governments (local, state and federal) to help finance their public expenditures, and finally fiat money or currency, and foreign exchange issued national governments, and managed by the banking system.

Starting with the basic model of inter-temporal consumption smoothing, we derive the fundamental equation that determines how financial assets are priced in competitive equilibrium. This leads into a discussion of how theories of risk sharing based on competitive equilibrium pricing can be tested using experimental methods.

Consider the the lifetime utility of a consumer whose labor supply and wage income are determined outside the model. We focus on their preferences under uncertainty about the timing of consumption. Let:

$$\sum_{t=0}^{T-1} \beta^t u(c_t)$$

denote the lifetime household utility function, which for convenience is assumed additively separable in $c_t \geq 0$, consumption in period t, where $t \in \{0, 1, \ldots, T\}$ is the period or year, $\beta \in (0,1)$ is a subjective discount factor, and $u(c_t)$ is current utility, which we assume is concave increasing throughout its domain.

Shortselling is prohibited in this model, so households do not hold any debt, and a period by period household budget constraint limits consumption each period.

Suppose there are person has assets, which can either consumed or invested in J financial securities $(q_{t1},\ldots,q_{tJ})$: Then where q_{tj} is the amount of the j^{th} security bought at the beginning of period t, r_{tj} is the return on the j^{th} security announced at the end of period t.

The household's maximization problem can be formulated as follows. Given her choices up until period t, and anticipating her future choices from period $t+1$ onwards, the consumer chooses consumption ct and her assets $(q_{t1},\ldots,q_{tJ})$ to maximize:

$$u\left(w_t - \sum_{j=1}^{J} q_{tj}\right) + \beta E\left[u\left(\sum_{j=1}^{J} r_{tj}q_{tj}\right)\right]$$

subject to her period t budget constraint

$$c_t \le \sum_{j=1}^{J}(r_{t-1,j}q_{t-1,j} - q_{tj})$$

where is the expectations operator based on information at time t.

Noting $u(c_t)$ is strictly increasing, in other words the conumer prefers more consumption to less, regardless of his level, all wealth is consumed, all the budget constraints are met with strict equality, yielding the expression for consumption:

$$c_t = \sum_{j=1}^{J}(r_{t-1,j}q_{t-1,j} - q_{tj})$$

This implies we can express for consumption into the objective function, and reformulate the consumer investor's problem as sequentially choosing the vector of assets to maximize:

$$u\left(w_t - \sum_{j=1}^{J} q_{tj}\right) + \beta E\left[u\left(\sum_{j=1}^{J} r_{tj}q_{tj}\right)\right]$$

The fundamental equation of portfolio choice comes from the first order condition of this problem. The interior first order condition requires that for each asset held by the consumer:

$$u'\left(\sum_{j=1}^{J}(r_{t-1,j}q_{t-1,j} - q_{tj})\right) = \beta E_t\left[r_{tj}u'\left(\sum_{j=1}^{J}(r_{tj}q_{tj} - q_{t+1,j})\right)\right]$$

Substituting the definition of consumption back into the first order condition we obtain:

$$u'(c_t) = \beta E_t[r_{tj}u'(c_{t+1})]$$

or

$$E_t\left[r_{tj}\frac{\beta u'(c_{t+1})}{u'(c_t)}\right] \equiv E_t[r_{tj}m_t] = 1$$

Asset return is discounted by the expected marginal rate of substitution between current and future consumption.

Not all assets are necesarily held by all consumers, and there are side conditions for not holding certain assets. Since $u(c_t)$ is concave increasing it follows that if no units of the jth asset is held, that is:

$$q_{tj} = 0$$

then:

$$u'(c_t) \geq \beta E_t[r_{tj} u'(c_{t+1})]$$

This equation says that the distribution of returns on the jth asset are too low to justify buying any units of the asset.

In principle the market clearing conditions can be used to derive the equilibrium asset prices. For each trader a first order condition applies to each positively consumed asset; otherwise it is not held. These conditions imply there is a solution to the asset allocation of each trader as a function of his asset endowment at the beginning of the period and the joint probability distribution governing asset returns. To express the market clearing conditions, we temporarily superscript asset endowments and allocations $q_{tj}^{(n)}$ for trader n. Market clearing in competitive equilibrium means that every period the demand for each asset exactly offsets its supply:

$$\sum_{n=1}^{N} \left(r_{t-1,j} q_{t-1,j}^{(n)} - q_{tj}^{(n)} \right) = 0$$

We can derive the risk free rate that a consumer requires to be induced to save from the fundamental equation of portfolio choice. If a risk free (interest) rate called rt exists, it must satisfy the equation too:

$$r_t \equiv 1 + i_t = 1/E_t[m_t]$$

where i_t is teh one period interest rate in period t or

$$r_t E_t[m_t] = 1$$

Risk corrections. Recall from the definition of a covariance:

$$cov_t(r_{tj}, m_t) = E_t[r_{tj} m_t] - E_t[r_{tj}] E_t[m_t]$$

Dividing both sides of the equation by the expected value of the marginal rate of substitution yields

$$E_t[r_{tj}] - E_t[r_{tj} m_t]/E_t[m_t] = -cov_t(r_{tj}, m_t)/E_t[m_t]$$

Using the formulas for teh interest rate and the markginal rate of substitution in the equation above we now obtain the risk correction for the mean return on the jth asset :

$$E_t[r_{tj}] - r_t = -r_t cov_t(r_{tj}, m_t)$$

The mean-variance frontier. From the risk correction formula for the mean return on the jth asset, we can write:

$$E_t[r_{tj}] - r_t = -r_t \sigma_{jt} \sigma_{mt} \frac{cov_t(r_{tj}, m_t)}{\sigma_{jt} \sigma_{mt}} \equiv -r_t \sigma_{jt} \sigma_{mt} \rho_t$$

where ρ_t is the correlation coefficient between mt and rjt, is σ_{jt}^2 the variance of rjt, and

$$\sigma_{mt}^2 \equiv E_t[m_t^2] - E_t[m_t]^2$$

is the variance of mt. Since the absolute value of every correlation coefficient is bounded by one, that is $|\rho_t| \leq 1$, the following inequality must be satisfied in a competitive equilibrium:

$$\left| \frac{E_t[r_{tj}] - r_t}{r_t} \right| \leq \sigma_{jt}\sigma_{mt}$$

Testing the fundamental equation

In experiments we specify the utility functions of the traders and the dividend processes. The portfolio choices qjt and the returns rjt are experimental outcomes.

It follows that we can compute (by integration), by transaction, quantities such as that enter the two basic variations on the fundamental equation, namely including ,

and

Experiment 20.10

Estimate risk aversion parameter and then use it to calculate $E_t[M_t]$. In the special case where all payoffs are consumed at the same time, we require $E_t[M_t] = 1$.

5 Summary

This final chapter analyzed competitive equilibrium when there is uncertainty.