

[SQUEAKING]

[RUSTLING]

[CLICKING]

PETER KEMPTHORNE: All right. Well, welcome to the first math lecture. The topic today is linear algebra. And what I expect is that, for most of you, the topics will be very familiar. At least we'll be covering basic fundamentals of vector algebra, then dealing with special matrices, matrices that are used in applications, eigenvectors and eigenvalue, decompositions, singular value decompositions, and then, finally, the Perron-Frobenius theorem. And what's really neat about quantitative finance is that linear algebra can really be very useful in understanding different computations we want to do. So it makes our life easier.

Well, let's just begin with a discussion of basic linear algebra concepts. I'll use a notation where often I'll have a bold vector notation-- or bold letter for noting a vector. So this bold v is a vector in m -dimensional space. And it just consists of an ordered list of numbers. And we can think of it as a column vector typically. Sometimes, we'll think of it as a row vector. But special cases of vectors will be the vector of all 0's and the vector of all 1's.

And we can think of a graphical representation of a vector, where we can think of it as perhaps just a single point in, in this case, two-dimensional space. Or with an origin here, we might think of the vector representing the directed line segment from the origin to that point, where this is v_1 and this is v_2 . Depending on our application, both of these can be useful.

And a basic example for us could be the closing prices on a given day for all stocks in the S&P 500 index. So we can think of p boldface being a vector of 500 values corresponding to the closing price of stocks that are in the S&P 500 index. Importantly, the prices are positive. So that's where the R plus comes into play. And if we consider a portfolio holding shares of the S&P 500 stocks, we can represent that with a vector q . And the value of such a portfolio at the end of the day is simply the sum of the quantities, the shares times the price per share.

And later on, we'll see that $q \cdot p$, the dot product of those two vectors, gives us the value of the portfolio. What we're going to find useful is to actually add to our 500 stocks the value of a cash holding in the portfolio. So we can think of p_0 of t just being a fixed \$1. And then q_0 tells us how many dollars we hold.

And so our portfolio value, v_t , will be given by this number. And we can think about rebalancing portfolios, computing portfolio and loss on a daily basis. And we can characterize rebalancing the portfolio as shifting from shares held at the end of day t to q_j of t plus 1 to be held through day t plus 1 by adjusting the quantity for each stock by a Δ_j .

And if we constrain the sum of the change in shares across the j stocks and cash and the p_j , multiplying by the p_j of t , we can constrain this sum to be equal to 0. So this would correspond to a rebalancing of the portfolio and basically not taking any of distributions from the account.

Now, with the net gain, let's see, this terminology, PnL, it's, in a sense, a misnomer, I think, because it's really profit and loss. But it's referred to as PnL, portfolio net gain. We can just look at the change in value from day t to day $t + 1$. And this will have this simple form for the change in portfolio value-- basically, the shares that are held from the end of day t to $t + 1$ times the change in price from day t to day $t + 1$.

And so this is the portfolio net gain. And here, we're talking about portfolio management maybe on a daily basis, where, at the end of the day, we make a decision as how to rebalance our portfolio. And there's a net gain during that period. One could adjust the period of rebalancing to be intraday on a regular or irregular basis. And we'd have to have extra notation to do that.

Now, with vector concepts, if we multiply vector v by a scalar c , then we just multiply through each of the components of v by that scalar. And if we add two vectors, we simply comprise the sum of components element by element. And the dot product between two vectors is simply the sum of the product of the two.

So this allows us to work with vector algebra associated with portfolios. If we have two portfolios, we can consider, let's see-- we can think of a q_t and, in this notation, a w_t as two different portfolios. And with each of these portfolios, we can determine what the profit and loss is of each portfolio. So for q , it's the value at the end of day minus the value at the beginning of the day and similarly for portfolio w . And it's useful to include in the notation here this BOD and EOD to represent beginning of day, end of day.

And what's relevant actually is that the q_t portfolio composition, while it's indexed by t , it's actually determined at the end of day $t - 1$. And so one needs to be careful in doing computations with portfolio value measurement that you're always using portfolios that are defined using only information up to a given time point. So if we're going to be rebalancing at the end of the day for investment the next day, we can't know what will happen on the next day.

All right. Well, with these portfolios that are just simple vectors, we can look at their difference. And so if we look at d_t equal to q_t minus w_t , this defines another portfolio. We often constrain portfolios to be long only, meaning the q 's and w 's are positive or 0. But there's no reason why we can't allow component weights in portfolios to be negative. And in this case, it's negative shares. And so one can have a long short portfolio that is given by d . And we can calculate the profit and loss of that d portfolio by subtracting the profit and loss of the q and w .

Now, let's see, with short selling in markets, let's see-- who in the class might offer a description of what it means to sell short a stock? So suppose that we have $d_{t,2}$ is equal to, say, negative 200. What would that mean? Or can we actually realize a portfolio where we're short one of the stocks by 200 shares? I'm sure someone-- OK. What's your name?

STUDENT: Alexander.

PETER OK, thank you, Alexander.

KEMPTHORNE:

STUDENT: I believe it is when you sell someone else's stock, and then you buy it back [INAUDIBLE].

PETER Yes. You basically sell first, buy back later. And you sell by-- you don't have the stock. So you borrow it from some
KEMPTHORNE: seller. And brokerage accounts that have margin allow you to engage in this kind of transaction. And so the policy of the brokerage is that if you want to sell short a stock, they need to locate some shares for you to borrow.

And what's rather neat about that selling short is that your brokerage account gets credited with the sale of those shares. So your brokerage account will have a positive cash flow due to the short selling. And with such a positive contribution to the brokerage account, you can actually invest long in more than, say, 100% of your wealth that's in the-- or of your bank account or brokerage account. So that's quite advantageous and allows for very interesting trading opportunities.

All right. Well, an important special case of a long short portfolio is a zero-cost portfolio. And this would be one where the sum of basically the values of the long and short positions balance each other perfectly. And so we have a portfolio, d in this case, where the d components are not all 0. But the value, the net value of the portfolio is, in fact, 0.

And we can hope for the possibility of an arbitrage portfolio, which could be a portfolio that has zero cost, but has a positive profit and loss over a period. And this kind of opportunity, generally-- well, it doesn't exist in efficient markets. But markets are never completely efficient. And so it can be feasible to identify portfolios that may be close to an arbitrage portfolio.

What, in fact, happens is that this profit and loss of a zero-cost portfolio will be random. And depending upon the market model that applies to the market we're trading in, it may be that there are opportunities that yield a positive profit all the time, in which case we would say that's an arbitrage opportunity.

What we'll see later today is the question of, under what conditions is there no arbitrage in a given market? And how do we characterize prices of assets in a market with no arbitrage? And ultimately, it turns out that option pricing theory, that was mentioned last time of Black, Scholes, and Merton, is all based upon arbitrage-free markets and the implications of that. And we'll see that, in arbitrage-free markets, there does not exist any arbitrage portfolios. And what will be neat is to understand why that's the case.

All right. So we have random prices at the end of day t . And so the profit and loss is going to be random. I guess it's relevant to note that there are strategies of statistical arbitrage. And those are strategies where we have portfolio specifications d^* for which this random profit and loss is a distribution that tends to generate positive returns. And so there's a statistical likelihood of making money. But it's often not a pure arbitrage, but what we might call a statistical arbitrage.

All right. Well, let's review just some other concepts of vectors. We have the norm of a vector being the length of a vector. So if we have a vector v here, and the norm of v is basically the length of this vector. And if we consider what the dot product is between two vectors, we have this cosine formula for what the dot product is. It's actually the length of v times the length of w times the cosine of the angle between the two.

So if we have another vector w here, then the dot product corresponds to that formula. And what's interesting about that formula is that it actually corresponds to-- the dot product corresponds to the length of the segment along w that is closest to v .

And so with this dot product notation, we actually can talk about orthogonality between vectors in m -dimensional space. And those can be very useful. We'll actually see this orthogonality applied when we talk about regression and least squares regression. So you may be familiar with some of these topics from that discussion.

All right. Well, with vectors, a really important topic is whether vectors are linearly independent of each other. And so if we have two vectors, v and w , they are linearly independent if the only linear combination of the two which gives us the 0 vector is the linear combination with $c_1 = c_2 = 0$. So linear independence is essentially the case that we have a v and w that aren't aligned with each other in two dimensions, at least vectors emanating from the origin having that property.

And the utility of linear independence is that we can work with a vector space, capital S here. And the definition of a vector space is that any vector can be rescaled by an arbitrary constant c . And it's still in the space. And with any two vectors, v and w , in the space, their sum is in the space.

With a vector space, what's very useful is a basis for the space, which corresponds to linearly independent vectors, v_1 to v_p , such that the only linear combination giving the 0 vector is the linear combination where all the coefficients are equal to 0 .

All right. Well, let's just go through all of matrix algebra, pretty much. We can extend from vectors to matrices where we have m rows and n columns in a matrix A , a rectangular matrix. And one very useful representation of a matrix is as a collection of column vectors that are stacked one after the other.

And so we can denote $\mathbf{a}_1$ through $\mathbf{a}_n$ as the column vectors. And with matrix algebra, let's see, well, we can multiply a matrix, capital A , by a constant. And we end up multiplying each of the components, elements of the table of A values by the same constant c .

With a transpose operation, a transpose is simply to flip rows and columns. And so if we have a a_{ij} being the i -th row j -th column element of A , then A_{ji} corresponds to the j -th row and i -th column. Now, if we think of the transpose of each of the column vectors in A , that's a row vector. And our transpose of the matrix A can be thought of as stacked row vectors corresponding to the same vectors, $\mathbf{a}_1$ through $\mathbf{a}_n$.

Now, one of the reasons for thinking of a matrix as being a composition of different column vectors is that, when we multiply a matrix A by a vector v , then this corresponds to the linear combination of the columns of A with coefficients corresponding to v_1 through v_n .

So we have a linear combination of A 's columns. And this representation of matrix vector multiplication turns out to be very, very useful. Perhaps less useful is that, if we look at the matrix vector product, it's also a vector of dot products of A 's rows with v . And so being comfortable with dot products is potentially useful. But for the moment, thinking of linear combinations of columns of A is very useful. So I guess, if we think of the columns of A corresponding to potentially a basis of vectors, then Av corresponds to another vector in the space spanned by those columns.

All right. Now, when we multiply two matrices together, we actually have a matrix A and a matrix B , each with their respective columns, vectors. And we can basically expand the matrix B into b_1 through b_p , and then do the matrix vector product for each term.

This is equivalent, actually, to looking at element C_{ij} , which are just these double sums. When I learned linear algebra decades ago, I used these double sums, which were not very intuitive to interpret, although one with a simple algebra notation can prove everything. But things are much easier to prove with this other notation.

Now, what's important when we multiply two matrices together is that the matrices have to be conformal, meaning we need to be able to multiply A times the column vectors of B . So the m by n matrix of A has to be able to multiply the m by p matrix of B . Actually, I think there's a-- all right, let's see. Yeah, sorry. And there are n columns in A and n rows in B . And so the result is an m by p matrix.

All right. Well, with different matrix operations, you learn in linear algebra all of these rules, these laws associated with matrix algebra computations. Everything is pretty obvious, except maybe when you transpose a product, you reverse the order of the product of transposes.

And if you consider reversing the order of matrix multiplications, if one can reverse the order, if A and B have the same number of rows and columns, it turns out that the commutative law for multiplication doesn't apply to matrices. So you want to be careful with those distinctions.

All right. So special matrices, when we have a symmetric matrix, then its transpose equals itself. So the matrix has basically-- well, let's see. The off-diagonal entries are matched one by one with A_{ij} equaling A_{ji} . The identity matrix is a matrix with 1's along the diagonal. And it can be convenient to represent this as the matrix of columns where E_j is a simple matrix or a simple vector, which is all 0's, but then a single 1, where this is the j -th row. And different computations can be made easy with that sometimes.

Now, there's a notation of using the J matrix to be the matrix of all 1's. And the matrix J turns out to be a vector 1 times its transpose. And we have diagonal matrices where the matrix is all 0's except for the diagonal entries, which may or may not be 1, and are given by the list D_1 through D_n . All right.

OK. Well now, we can turn to some applications that use matrices. And the first one we want to discuss is stochastic matrices. So these are matrices which will consider a stochastic matrix that's square. And the columns of A sum to 1. So we'll have an A with A_1, A_2 up to A_m or A_n . And basically, the sum of A_{ij} , i equaling 1 to m , is going to equal to 1.

This kind of matrix is associated with Markov chains, where we think of there being m possible states. So we basically have m rows and m columns. And we have this notation A_{ij} , or lowercase a_{ij} , is equal to the probability that the state, the next state, equals i , given the previous state equals j is given by this A_{ij} .

Let's see. For each column j , we have, if the previous state at time t was j , then there's a transition to a state at time t plus 1. And these transition probabilities, in some cases, can be assumed to be stable or stationary over t . And with a Markov chain, we can think of there basically being different states which are indexed from 1 up to m .

And we might think of these as being nodes in a graph. And there's a probability of moving from one node to another for each time point. And these probabilities of transitions might be stationary in a stationary Markov chain model.

Well, if we have a vector, π , which is going to be basically π of t equaling π_1 of t , π_2 of t , down to π_m of t , so this m vector, π , tells us at time t what's the likelihood of being in any of the states. Then, if we want to compute the probabilities of states at time t plus 1, then it turns out to be simply the matrix vector product of A times π of t . So this formula gives us the model for how the probabilities of states occurring changes from time t to t plus 1.

Now, if we consider what is the probability vector at time t plus 2, well, we have the same matrix vector model where the stationary transition probabilities A are multiplied by the vector of probabilities at t plus 1. And we can substitute in what that product is, which is $A\pi$. And so we basically get the square of the transition probability matrix times the π vector at time t .

Now, what's interesting to ask with Markov chains is, if we start out with a π vector at time 0 and we have the model that the likelihood of different states evolves such that the t period transition probabilities from π_0 to t is given by this formula, then does this limit exist or not? And so let me just ask that question to the class. Anyone remember from stochastic processes perhaps or linear algebra whether that has a limit? Yes?

STUDENT: The stationary distribution.

PETER If it exists, it's the stationary distribution. And so basically, if π^* is the stationary distribution, it is equal to a π^* so that we're at a limiting distribution where there's no change in probabilities, continuing from t equal to infinity to infinity plus 1. Then this would be the formula. Now, are there circumstances where such a stationary distribution doesn't exist? Yes, OK. And when is that? I see some solid nods. Yes?

STUDENT: When it's periodic.

PETER When it's periodic, yep. So you could have cycles in the process. And a simplest case would be two states that just alternate even odd being one state or the other as an example. So we need to have acyclic distributions. Or the Markov chain needs to be acyclic.

And actually, all these topics are quite interesting when you study them. It turns out that you can prove that a Markov chain is acyclic if all of its components of the A matrix, if every single one, is nonzero, so that there's always a possibility of transitioning from one state to another.

That's a bit too restrictive. But it turns out that, if you consider a particular power of A , that for some power, not necessarily very big, but there's at some point of looking at powers of A , the A to that power is a matrix of all positive entries. So it's interesting just how these mathematical features can be shown.

Now, let's see. All right. So if it does exist, we have a stationary distribution. And that's nice. Let me just point out that, I guess, in the first problem set for the course that's been posted-- that's due a week from Friday. So this is a healthy amount of time away, I hope, for everybody-- we're asking you to look at some simple Markov chains and explore the stationary distribution of those cases. And the Markov chain model is used for modeling whether, say, trades at the buy or sell side-- whether successive trades in a stock will be at the buy side or sell side, depending on what the previous trade was at

All right. Well, let's move on to another class of matrices that's of interest. Actually, before doing that, I'm going to jump ahead just in terms of topics. But this formula for the stationary distribution is a formula that's very specific to eigenvalues and eigenvectors. And so if you think of the matrix A having eigenvalues and eigenvectors, what's the eigenvalue and what's the eigenvector in this case? Does anybody-- yes, go ahead.

STUDENT: Eigenvalue value of 1, and the eigenvectors would be all the π stars [INAUDIBLE].

PETER KEMPTHORNE: So this basically is a case where one of the eigenvalues is 1. Or sorry, one of the eigenvalues is 1. And this is satisfied. We'll see later that, basically, the other eigenvalues may be smaller than 1. And when you multiply the A matrix by itself to many, many powers, then it's only this one eigenvector that doesn't get shrunk to 0 in the limit.

All right. Well, let's turn to positive matrices. Positive matrices arise in economics, where we think of there being, say, prices of goods under different states of the market or costs of goods under states of the market. And those are all positive. And we can think of this arising in a single period market model, where we have n assets indexed by i and, let's see, we have two periods-- or sorry, a single period with a beginning and end, t_0 and t equals capital T .

And if we think of what is the price of assets at time 0, we may know what those are at time 0. And then what are the subsequent prices at the end of the single period? Capital T . And those prices may or may not be known. And so let me just give an example here with two assets.

And so we're going to think of the two-period model. And we're going to have n equal to 2. We're going to consider basically a bond, B , which takes on values that will be B_0 at t equals 0 and B_T at time capital T . And we'll consider the case of a bond which basically pays interest at the interest rate R_F .

So this B_0 price of the bond at time 0 grows to B sub capital T , which is basically growing at a simple interest rate of R_F . And let's see. We can talk about R_B , the return of this, just being B_T minus B_0 . So this is the absolute return. And we can also define the percent return, which would be looking at R sub B equal to the absolute return divided by the initial value. So this would be the percentage return.

And let's see. If we subtract this here, we end up getting that this is $B_0 R_F$ times T . So this turns out to be equal to simply $R_F T$ as the percent return. And let's see. We can also talk about the simple rate of return, which would be given by-- well, let's see. Actually, maybe, here I should write R sub B and here R sub B instead of R sub F because I want to define R sub f to be equal to basically 1 over t , the average rate of return. So it's equal to this times B_T minus B_0 over B_0 .

And so this, simply, in this case, is equal to that R_B . And I guess, I'll be using a notation of characterizing an interest rate as being a risk-free interest rate, if this bond is guaranteed to have that final value. And so in quantitative finance, we often talk about risk-free rates. And nothing in the world is risk-free. But when governments can print money to pay off their liabilities, then it essentially is a risk-free rate. And we worry less about it with countries with stronger economies.

Now, this is the first asset. Let's also consider a stock, S , which has at time 0 a value of S_0 . And then at time T , it's going to be random. And so we have S sub T is going to equal different possible outcomes. And we'll consider a special case where it's S_{Tu} and S_{Td} , where we have an underlying state being an up or a down market.

So we have a random outcome for S_T . We're not yet specifying what probabilities are associated with these different states. And so there are two states at time T , ω equal to-- well, maybe I'll just write u and ω equal to d for that.

And so in thinking of a single market model with n assets, not just two, we are defining basically two assets-- I guess I'm using j for the assets instead of i . But we basically have known asset prices at time 0 and potentially random prices at time capital T . And so we can model this by associating different states or scenarios at time capital T that are indexed w_1 to w_m . And we then can have different prices in that.

So here's the framework for this right now, at least the framework with two assets. And we can also characterize portfolios of the n equals 2 assets. And we can define a π vector being simply equal to a weight in the bond and a weight in the stock, units of the stock, units of the bond, given by this vector.

And so what we can then do is deal with modeling the outcomes of this portfolio. If we consider v_0 , the value of the portfolio at time 0, this would be just $\pi_B B_0$ plus $\pi_S S_0$. And at time capital T , then our portfolio would have values $\pi_B B_T$ plus $\pi_S S_T^u$ for the up case and $\pi_B B_T$ plus $\pi_S S_T^d$ when w is equal to d , when the state is down. So this value of the portfolio will be random, depending upon the likelihood of these up and down states.

Well, with this setup, it's useful at this point to introduce you to some concepts, which is contingent claims and replicating portfolios. And so suppose that we have a contingent claim C , which has a value C_0 , is the value at t equal to 0. And C_t will be the value of this contingent claim depending upon whether we have an up or down-- well, actually, I'll reverse it, a down or an up market.

What we end up doing in these market models is being able to model the prices of contingent claims when we can replicate the contingent claim with a portfolio of the underlying assets. So we can ask, does there exist a portfolio of B and S such that $\pi_B B_T$ plus $\pi_S S_T^u$ equals C_T^u and $\pi_B B_T$ plus $\pi_S S_T^d$ equals C_T^d . I've got these notations wrong-- equals to C_T^u .

So if this were the case, then-- if so, then we could have C_0 equaling $\pi_B B_0$ plus $\pi_S S_0$. And we would replicate this contingent claim payoff at time capital T . Now, mathematically, can we solve for portfolio coefficients that satisfy this contingent claim? Any thoughts?

Well, let's see. If we look at C_T^d versus C_T^u -- and here's the origin. If we had a portfolio that was just the bond, we might have B_T here and B_T here. And here's the point corresponding to B_T , B_T . And here's the vector corresponding to the portfolio that's just the bond.

And if we look at the stock, the portfolio, that's just the stock, we might have S_T^d here, and then way up here maybe S_T^u , the value of the stock when the market is up. And so this would be our ST vector in the two-dimensional space corresponding to down market state versus up market state. Now, given this figure, if we think of just an arbitrary point in this space of contingent claims, can we construct a portfolio that matches that? Yes?

STUDENT: Sorry, can you explain what [INAUDIBLE]?

PETER KEMPTHORNE: Yes. OK. The contingent claim is a payoff that is contingent upon the state of the market at that time. And it's either a down market or an up market over the one period. And the payoff, if it's an up market, is C_T^u . And the payoff, if it's a down market, is C_T^d . And in this space of contingent claims, we can draw vectors corresponding to the stock and the bond as these vectors.

And so what's curious to think about is-- suppose we have a contingent claim that-- well, let's see. Suppose we consider a call option contingent claim. Well, let's see. So there would be a strike price of-- let's see-- A strike price of K to buy the stock at time T .

And in order for this to be of interest, we might have K being smaller than ST_u and greater than ST_d . So we have a call option with strike price K to buy the stock. And the contingent claim of the call would be 0, if we have the down market. And it would be ST_u minus K if it's the up market.

So we could consider 0 in the down market. And basically, maybe this would be CT minus K here. So this vector here would correspond to the payoff of this call option. So under what circumstances can we take a linear combination of these two vectors and get this vector? Yes?

STUDENT: So you've got to short one of the stocks?

PETER Yeah. I mean, basically, we can solve for this here by substituting in for this 0 and P minus K , and then just solve
KEMPTHORNE: it to get our two equations and two unknowns. So if this system of equations, two equations and two unknowns, has a solution, then we can solve it.

And what is interesting is to consider what are the space of all possible contingent claims. And let me just turn to this note. This is on Canvas. But this note that's been posted in Canvas talks about a one-period economy with two assets in section 1. That's what I've just outlined on the chalkboard here.

One can generalize that to a single period model with more than two assets. That's chapter section 2. And then one can generalize to-- not generalize, but one can consider visualizing the assets and portfolios in section 3 here. And so what I do in the last section here is looking at the vector of payoffs of the bond and the stock, colored differently.

And so this is a bond that may pay off the same amount in the down versus the up market. And the stock pays off more in the up market than the down. And then one can consider collections of portfolios that weight the bond and the stock so that the positive weights sum to 1. And so these vectors correspond to the payoff vectors of those long portfolios.

And if we consider-- well, let's see-- looking at those long portfolios, instead of representing vectors as pointed segments from the origin, one could just represent them as points in the space. And that representation is helpful when we want to consider a variety of other portfolios in the mix.

And so here's a graph showing the value of the payoff if we consider portfolios that can invest up to 100% in each of the bond or the stock or both. And so if we invest 100% in both, we basically get the payoffs that are really high. And if we don't invest much, close to 0 in both, then we get the 0 vector. But this graph shows the contingent claims that are mapped or equivalent by replicating portfolios of these types.

And so if we then allow for shorting, this extends the contingent claims to allow for negative payoffs, which are graphed here, basically with the payoff in the down market versus the up. And what's of interest here, although it's not quite easy to see, is there are portfolios where the initial value of the portfolio is 0. And so that would correspond to a long short portfolio that is implemented with zero cost.

So this is the payoff vector at T and looking at the value where we white out those portfolios that have zero cost. Now, what's useful here to see is that, with these kinds of portfolios, we can replicate different contingent claims which correspond to different points in this space.

If we were able to invest more than 100% of our portfolio, long or short, then we could actually realize any point in this space, in some cases. And so any contingent claim could be solved for with this equation here. And that's rather neat. But it also is reasonable to consider if we make certain assumptions on the market, such as, as you buy more of an asset, you don't change the price of the asset. There's basically a fixed price is suitable regardless of how much. And there's no constraint on short selling the asset, that you can do that with liberty.

So there are these kinds of properties that arise. Now, with the case of extending from the bond and the single stock to m -- sorry, little n assets, then everything generalizes. And we end up having a positive price matrix at time T corresponding to different payoffs of the individual assets at time capital T for different states ω_1 through ω_m .

And with this kind of setup, one can represent the initial value of a portfolio Q being the sum of the Q weights times the initial prices in the A_0 vector. And then the payoff of the portfolio at time capital T is given by the same weights q_j times the time capital T prices. But we have different prices depending upon what state occurs.

And so in this case, we have basically our matrix A times Q giving us the vector of state-dependent values at time capital T . And so with this model, we can contemplate there being an arbitrage portfolio where the initial value is not positive. I guess one could constrain this to be 0. But it's easier to just say less than or equal to 0.

If we have a negative cost, that means we get paid to implement this portfolio. And it will be an arbitrage portfolio if the likelihood of a negative payoff is 0 and there's some likelihood of a positive payoff. Such a portfolio would be risk-free, in terms of downside risk. We never lose money. We either make money all the time in at least one of the states. So this is an arbitrage portfolio.

Now, with this model, there can be conditions that ensure that there is no arbitrage in the market. And so that's one issue to think about. What conditions would ensure no arbitrage? And another condition that's of interest is called market completeness, which is, are there portfolios that can realize any payoff vector?

So in looking at the contingent claims payoff, can we span this space fully with portfolios of the underlying assets in the market? And so if we can represent any contingent payoff as a portfolio, then we say that the market is complete. And it turns out that the conditions of no arbitrage and market completeness are related to whether a pricing measure exists on scenarios.

And so think about Q^* being a probability distribution over the underlying states, ω_1 through ω_m . And it turns out that if the initial price of the j -th asset is the discounted expected terminal price of that j -th asset across the different states, then Q^* is called a pricing measure. So our initial prices equal a discount factor times the expected terminal time t values.

If such a pricing measure exists, then there's no arbitrage. So it's no arbitrage if all of the q_j^* stars are positive. So if this pricing measure gives positive weight to each of the underlying states, then there's no arbitrage. And we say that the market is complete with no arbitrage if this pricing measure is unique.

So it's an interesting set of properties that actually are quite, well, useful to know about. And these models actually underlie much of option pricing theory and arbitrage relationships in quantitative finance. And in this single period model, note that we have distributed there's this reference of this book by Albanese and Campolieti, where they go through proofs of these. And they can get a bit technical. But the basic ideas, I think, are relatively accessible.

Now, let's see. Moving on to other applications of matrices, we can consider systems of linear equations and their solution, where we basically have m equations in n unknowns. And so with these m equations, we have Ax equals B , a B vector.

And to solve such a system of equations, well, one case where we can solve these is where our number of unknowns and number of equations are the same and our coefficients matrix A in the system of equations has full rank. So A being square and of full rank has a determinant that's non-zero and its inverse matrix exists. And so we can actually just take our system of equations and multiply both sides by the inverse of A and get our answer.

When the number of equations is smaller than the number of unknowns, then we have an undetermined system of equations. When m is greater than n , it's overdetermined and so forth. So these are different cases that you likely are familiar with from your study of linear algebra.

All right. The next topic in our remaining few minutes is eigenvalues and eigenvectors. When we have square matrices A , we can consider the existence of an eigenvalue λ and an eigenvector V such that the matrix A times v is equal to just the scalar λ times v . So the matrix vector product just rescales v by the factor λ .

And one can solve for eigenvalues and eigenvectors by just re-expressing the eigenvalue-eigenvector relationship as this system of equations. And so this system of equations will correspond to-- Let's see. This system of equations can be solved by saying that the determinant of $A - \lambda I$ is equal to 0. And so if λ equals 1 of the roots of this polynomial, then that λ will be an eigenvalue. And one can then, for that fixed eigenvalue, solve for the eigenvector v that satisfies that.

Now, there are properties with eigenvalues and eigenvectors that are very useful, one of which is that the maximum eigenvalue magnitude corresponds to the maximum length of Av for vectors v that are of length 1. And when we have a matrix A that has independent eigenvalues or eigenvectors, then, in this special case, we can consider defining the matrix S to be the matrix whose columns are the different eigenvectors.

And if you look at this equation here, if we multiply the original matrix A by this S matrix, then each column of the product is the corresponding eigenvalue times the corresponding eigenvector. This equation here is very straightforward from the notation above.

And if we consider the inverse of this S matrix that exists if the eigenvectors are linearly independent, then we can write A to be $S \Lambda S^{-1}$. And this representation of the matrix A also gives us the way of diagonalizing A by pre and post-multiplying by S^{-1} and S respectively.

And so with powers of a diagonalizable matrix A , they have a very nice form. We basically get cancellation of S and S inverse in the products of A times itself. And so powers of A simply are the same matrix S times λ , the diagonal matrix of eigenvalues to the k -th power times S inverse. And so these are very useful in defining state equations and Kalman filters.

All right. So let me just-- I guess I'll finish there. There's some interesting key theorems in the remainder of these notes. I think they're pretty readable. And I'll maybe point out these during our next lecture. But there we are. Also, let's see, I posted an RStudio program that computes portfolios of S&P 500 stocks and looks at the changing value over time. And so I hope people can explore those programs and try to use those in the RStudio Cloud.