

Linear Algebra

Jesse Perla, Thomas J. Sargent and John Stachurski

December 4, 2020

1 Contents

- Overview [2](#)
- Vectors [3](#)
- Matrices [4](#)
- Solving Systems of Equations [5](#)
- Eigenvalues and Eigenvectors [6](#)
- Further Topics [7](#)
- Exercises [8](#)
- Solutions [9](#)

2 Overview

Linear algebra is one of the most useful branches of applied mathematics for economists to invest in.

For example, many applied problems in economics and finance require the solution of a linear system of equations, such as

$$\begin{aligned}y_1 &= ax_1 + bx_2 \\ y_2 &= cx_1 + dx_2\end{aligned}$$

or, more generally,

$$\begin{aligned}y_1 &= a_{11}x_1 + a_{12}x_2 + \cdots + a_{1k}x_k \\ &\quad \vdots \\ y_n &= a_{n1}x_1 + a_{n2}x_2 + \cdots + a_{nk}x_k\end{aligned}\tag{1}$$

The objective here is to solve for the “unknowns” x_1, \dots, x_k given a_{11}, \dots, a_{nk} and y_1, \dots, y_n .

When considering such problems, it is essential that we first consider at least some of the following questions

- Does a solution actually exist?
- Are there in fact many solutions, and if so how should we interpret them?
- If no solution exists, is there a best “approximate” solution?

- If a solution exists, how should we compute it?

These are the kinds of topics addressed by linear algebra.

In this lecture we will cover the basics of linear and matrix algebra, treating both theory and computation.

We admit some overlap with [this lecture](#), where operations on Julia arrays were first explained.

Note that this lecture is more theoretical than most, and contains background material that will be used in applications as we go along.

3 Vectors

A *vector* is an element of a vector space.

Vectors can be added together and scaled (multiplied) by scalars.

Vectors can be written as $x = [x_1, \dots, x_n]$.

The set of all n -vectors is denoted by \mathbb{R}^n .

For example, \mathbb{R}^2 is the plane, and a vector in \mathbb{R}^2 is just a point in the plane.

Traditionally, vectors are represented visually as arrows from the origin to the point.

The following figure represents three vectors in this manner.

3.1 Setup


```
In [1]: using InstantiateFromURL
        # optionally add arguments to force installation: instantiate = true,
        ↪precompile = true
        github_project("QuantEcon/quantecon-notebooks-julia", version = "0.8.0")
```

```
In [2]: using LinearAlgebra, Statistics, Plots
        gr(fmt=:png);
```

```
In [3]: x_vals = [0 0 0 ; 2 -3 -4]
        y_vals = [0 0 0 ; 4 3 -3.5]

        plot(x_vals, y_vals, arrow = true, color = :blue,
             legend = :none, xlims = (-5, 5), ylims = (-5, 5),
             annotations = [(2.2, 4.4, "[2, 4]"),
                           (-3.3, 3.3, "[-3, 3]"),
                           (-4.4, -3.85, "[-4, -3.5]")],
             xticks = -5:1:5, yticks = -5:1:5,
             framestyle = :origin)
```

```
Out[3]:
```



3.2 Vector Operations

The two most common operators for vectors are addition and scalar multiplication, which we now describe.

As a matter of definition, when we add two vectors, we add them element by element

$$x + y = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} + \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix} := \begin{bmatrix} x_1 + y_1 \\ x_2 + y_2 \\ \vdots \\ x_n + y_n \end{bmatrix}$$

Scalar multiplication is an operation that takes a number γ and a vector x and produces

$$\gamma x := \begin{bmatrix} \gamma x_1 \\ \gamma x_2 \\ \vdots \\ \gamma x_n \end{bmatrix}$$


Scalar multiplication is illustrated in the next figure

In [4]: # illustrate scalar multiplication

```
x = [2]
scalars = [-2 1 2]
vals = [0 0 0; x * scalars]
labels = [(-3.6, -4.2, "-2x"), (2.4, 1.8, "x"), (4.4, 3.8, "2x")]

plot(vals, vals, arrow = true, color = [:red :red :blue],
      legend = :none, xlims = (-5, 5), ylims = (-5, 5),
      annotations = labels, xticks = -5:1:5, yticks = -5:1:5,
      framestyle = :origin)
```

Out[4]:



In Julia, a vector can be represented as a one dimensional Array.

Julia Arrays allow us to express scalar multiplication and addition with a very natural syntax

In [5]: `x = ones(3)`

Out[5]: 3-element Array{Float64,1}:
1.0
1.0
1.0

In [6]: `y = [2, 4, 6]`

Out[6]: 3-element Array{Int64,1}:
2
4
6

In [7]: `x + y`

Out[7]: 3-element Array{Float64,1}:
3.0
5.0
7.0

In [8]: `4x` # equivalent to `4 * x` and `4 .* x`

Out[8]: 3-element Array{Float64,1}:
4.0
4.0
4.0

3.3 Inner Product and Norm

The *inner product* of vectors $x, y \in \mathbb{R}^n$ is defined as

$$x'y := \sum_{i=1}^n x_i y_i$$

Two vectors are called *orthogonal* if their inner product is zero.

The *norm* of a vector x represents its “length” (i.e., its distance from the zero vector) and is defined as

$$\|x\| := \sqrt{x'x} := \left(\sum_{i=1}^n x_i^2 \right)^{1/2}$$

The expression $\|x - y\|$ is thought of as the distance between x and y .

Continuing on from the previous example, the inner product and norm can be computed as follows

```
In [9]: using LinearAlgebra
```

```
In [10]: dot(x, y)           # Inner product of x and y
```

```
Out[10]: 12.0
```

```
In [11]: sum(prod, zip(x, y)) # Gives the same result
```

```
Out[11]: 12.0
```

```
In [12]: norm(x)           # Norm of x
```

```
Out[12]: 1.7320508075688772
```

```
In [13]: sqrt(sum(abs2, x)) # Gives the same result
```

```
Out[13]: 1.7320508075688772
```

3.4 Span

Given a set of vectors $A := \{a_1, \dots, a_k\}$ in \mathbb{R}^n , it's natural to think about the new vectors we can create by performing linear operations.

New vectors created in this manner are called *linear combinations* of A .

In particular, $y \in \mathbb{R}^n$ is a linear combination of $A := \{a_1, \dots, a_k\}$ if

$$y = \beta_1 a_1 + \dots + \beta_k a_k \text{ for some scalars } \beta_1, \dots, \beta_k$$

In this context, the values β_1, \dots, β_k are called the *coefficients* of the linear combination.

The set of linear combinations of A is called the *span* of A .

The next figure shows the span of $A = \{a_1, a_2\}$ in \mathbb{R}^3 .

The span is a 2 dimensional plane passing through these two points and the origin.

```


In [14]: # fixed linear function, to generate a plane
         f(x, y) = 0.2x + 0.1y

         # lines to vectors
         x_vec = [0 0; 3 3]
         y_vec = [0 0; 4 -4]
         z_vec = [0 0; f(3, 4) f(3, -4)]

         # draw the plane
         n = 20
         grid = range(-5, 5, length = n)
         z2 = [ f(grid[row], grid[col]) for row in 1:n, col in 1:n ]
         wireframe(grid, grid, z2, fill = :blues, gridalpha =1 )
         plot!(x_vec, y_vec, z_vec, color = [:blue :green], linewidth = 3, labels=
↳= "",
           colorbar = false)

```

Out[14]:



3.4.1 Examples

If A contains only one vector $a_1 \in \mathbb{R}^2$, then its span is just the scalar multiples of a_1 , which is the unique line passing through both a_1 and the origin.

If $A = \{e_1, e_2, e_3\}$ consists of the *canonical basis vectors* of \mathbb{R}^3 , that is

$$e_1 := \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \quad e_2 := \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \quad e_3 := \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

then the span of A is all of \mathbb{R}^3 , because, for any $x = (x_1, x_2, x_3) \in \mathbb{R}^3$, we can write

$$x = x_1e_1 + x_2e_2 + x_3e_3$$

Now consider $A_0 = \{e_1, e_2, e_1 + e_2\}$.

If $y = (y_1, y_2, y_3)$ is any linear combination of these vectors, then $y_3 = 0$ (check it).

Hence A_0 fails to span all of \mathbb{R}^3 .

3.5 Linear Independence

As we'll see, it's often desirable to find families of vectors with relatively large span, so that many vectors can be described by linear operators on a few vectors.

The condition we need for a set of vectors to have a large span is what's called linear independence.

In particular, a collection of vectors $A := \{a_1, \dots, a_k\}$ in \mathbb{R}^n is said to be

- *linearly dependent* if some strict subset of A has the same span as A
- *linearly independent* if it is not linearly dependent

Put differently, a set of vectors is linearly independent if no vector is redundant to the span, and linearly dependent otherwise.

To illustrate the idea, recall [the figure](#) that showed the span of vectors $\{a_1, a_2\}$ in \mathbb{R}^3 as a plane through the origin.

If we take a third vector a_3 and form the set $\{a_1, a_2, a_3\}$, this set will be

- linearly dependent if a_3 lies in the plane
- linearly independent otherwise

As another illustration of the concept, since \mathbb{R}^n can be spanned by n vectors (see the discussion of canonical basis vectors above), any collection of $m > n$ vectors in \mathbb{R}^n must be linearly dependent.

The following statements are equivalent to linear independence of $A := \{a_1, \dots, a_k\} \subset \mathbb{R}^n$.

1. No vector in A can be formed as a linear combination of the other elements.
2. If $\beta_1 a_1 + \dots + \beta_k a_k = 0$ for scalars β_1, \dots, β_k , then $\beta_1 = \dots = \beta_k = 0$.

(The zero in the first expression is the origin of \mathbb{R}^n)

3.6 Unique Representations

Another nice thing about sets of linearly independent vectors is that each element in the span has a unique representation as a linear combination of these vectors.

In other words, if $A := \{a_1, \dots, a_k\} \subset \mathbb{R}^n$ is linearly independent and

$$y = \beta_1 a_1 + \dots + \beta_k a_k$$

then no other coefficient sequence $\gamma_1, \dots, \gamma_k$ will produce the same vector y .

Indeed, if we also have $y = \gamma_1 a_1 + \dots + \gamma_k a_k$, then

$$(\beta_1 - \gamma_1)a_1 + \dots + (\beta_k - \gamma_k)a_k = 0$$

Linear independence now implies $\gamma_i = \beta_i$ for all i .

4 Matrices

Matrices are a neat way of organizing data for use in linear operations.

An $n \times k$ matrix is a rectangular array A of numbers with n rows and k columns:

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1k} \\ a_{21} & a_{22} & \cdots & a_{2k} \\ \vdots & \vdots & & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nk} \end{bmatrix}$$

Often, the numbers in the matrix represent coefficients in a system of linear equations, as discussed at the start of this lecture.

For obvious reasons, the matrix A is also called a vector if either $n = 1$ or $k = 1$.

In the former case, A is called a *row vector*, while in the latter it is called a *column vector*.

If $n = k$, then A is called *square*.

The matrix formed by replacing a_{ij} by a_{ji} for every i and j is called the *transpose* of A , and denoted A' or A^\top .

If $A = A'$, then A is called *symmetric*.

For a square matrix A , the i elements of the form a_{ii} for $i = 1, \dots, n$ are called the *principal diagonal*.

A is called *diagonal* if the only nonzero entries are on the principal diagonal.

If, in addition to being diagonal, each element along the principal diagonal is equal to 1, then A is called the *identity matrix*, and denoted by I .

4.1 Matrix Operations

Just as was the case for vectors, a number of algebraic operations are defined for matrices.

Scalar multiplication and addition are immediate generalizations of the vector case:

$$\gamma A = \gamma \begin{bmatrix} a_{11} & \cdots & a_{1k} \\ \vdots & \vdots & \vdots \\ a_{n1} & \cdots & a_{nk} \end{bmatrix} := \begin{bmatrix} \gamma a_{11} & \cdots & \gamma a_{1k} \\ \vdots & \vdots & \vdots \\ \gamma a_{n1} & \cdots & \gamma a_{nk} \end{bmatrix}$$

and

$$A + B = \begin{bmatrix} a_{11} & \cdots & a_{1k} \\ \vdots & \vdots & \vdots \\ a_{n1} & \cdots & a_{nk} \end{bmatrix} + \begin{bmatrix} b_{11} & \cdots & b_{1k} \\ \vdots & \vdots & \vdots \\ b_{n1} & \cdots & b_{nk} \end{bmatrix} := \begin{bmatrix} a_{11} + b_{11} & \cdots & a_{1k} + b_{1k} \\ \vdots & \vdots & \vdots \\ a_{n1} + b_{n1} & \cdots & a_{nk} + b_{nk} \end{bmatrix}$$

In the latter case, the matrices must have the same shape in order for the definition to make sense.

We also have a convention for *multiplying* two matrices.

The rule for matrix multiplication generalizes the idea of inner products discussed above, and is designed to make multiplication play well with basic linear operations.

If A and B are two matrices, then their product AB is formed by taking as its i, j -th element the inner product of the i -th row of A and the j -th column of B .

There are many tutorials to help you visualize this operation, such as [this one](#), or the discussion on the [Wikipedia page](#).

If A is $n \times k$ and B is $j \times m$, then to multiply A and B we require $k = j$, and the resulting matrix AB is $n \times m$.

As perhaps the most important special case, consider multiplying $n \times k$ matrix A and $k \times 1$ column vector x .

According to the preceding rule, this gives us an $n \times 1$ column vector

$$Ax = \begin{bmatrix} a_{11} & \cdots & a_{1k} \\ \vdots & \vdots & \vdots \\ a_{n1} & \cdots & a_{nk} \end{bmatrix} \begin{bmatrix} x_1 \\ \vdots \\ x_k \end{bmatrix} := \begin{bmatrix} a_{11}x_1 + \cdots + a_{1k}x_k \\ \vdots \\ a_{n1}x_1 + \cdots + a_{nk}x_k \end{bmatrix} \quad (2)$$

Note

AB and BA are not generally the same thing.

Another important special case is the identity matrix.

You should check that if A is $n \times k$ and I is the $k \times k$ identity matrix, then $AI = A$.

If I is the $n \times n$ identity matrix, then $IA = A$.

4.2 Matrices in Julia

Julia arrays are also used as matrices, and have fast, efficient functions and methods for all the standard matrix operations.

You can create them as follows

```
In [15]: A = [1 2
              3 4]
```

```
Out[15]: 2×2 Array{Int64,2}:
 1  2
 3  4
```

```
In [16]: typeof(A)
```

```
Out[16]: Array{Int64,2}
```

```
In [17]: size(A)
```

```
Out[17]: (2, 2)
```

The `size` function returns a tuple giving the number of rows and columns.

To get the transpose of `A`, use `transpose(A)` or, more simply, `A'`.

There are many convenient functions for creating common matrices (matrices of zeros, ones, etc.) — see [here](#).

Since operations are performed elementwise by default, scalar multiplication and addition have very natural syntax

```
In [18]: A = ones(3, 3)
```

```
Out[18]: 3×3 Array{Float64,2}:
 1.0  1.0  1.0
 1.0  1.0  1.0
 1.0  1.0  1.0
```

```
In [19]: 2I
```

```
Out[19]: UniformScaling{Int64}
 2*I
```

```
In [20]: A + I
```

```
Out[20]: 3×3 Array{Float64,2}:
 2.0  1.0  1.0
 1.0  2.0  1.0
 1.0  1.0  2.0
```

To multiply matrices we use the `*` operator.

In particular, `A * B` is matrix multiplication, whereas `A .* B` is element by element multiplication.

4.3 Matrices as Maps

Each $n \times k$ matrix A can be identified with a function $f(x) = Ax$ that maps $x \in \mathbb{R}^k$ into $y = Ax \in \mathbb{R}^n$.

These kinds of functions have a special property: they are *linear*.

A function $f: \mathbb{R}^k \rightarrow \mathbb{R}^n$ is called *linear* if, for all $x, y \in \mathbb{R}^k$ and all scalars α, β , we have

$$f(\alpha x + \beta y) = \alpha f(x) + \beta f(y)$$

You can check that this holds for the function $f(x) = Ax + b$ when b is the zero vector, and fails when b is nonzero.

In fact, it's [known](#) that f is linear if and *only if* there exists a matrix A such that $f(x) = Ax$ for all x .

5 Solving Systems of Equations

Recall again the system of equations (1).

If we compare (1) and (2), we see that (1) can now be written more conveniently as

$$y = Ax \tag{3}$$

The problem we face is to determine a vector $x \in \mathbb{R}^k$ that solves (3), taking y and A as given.

This is a special case of a more general problem: Find an x such that $y = f(x)$.

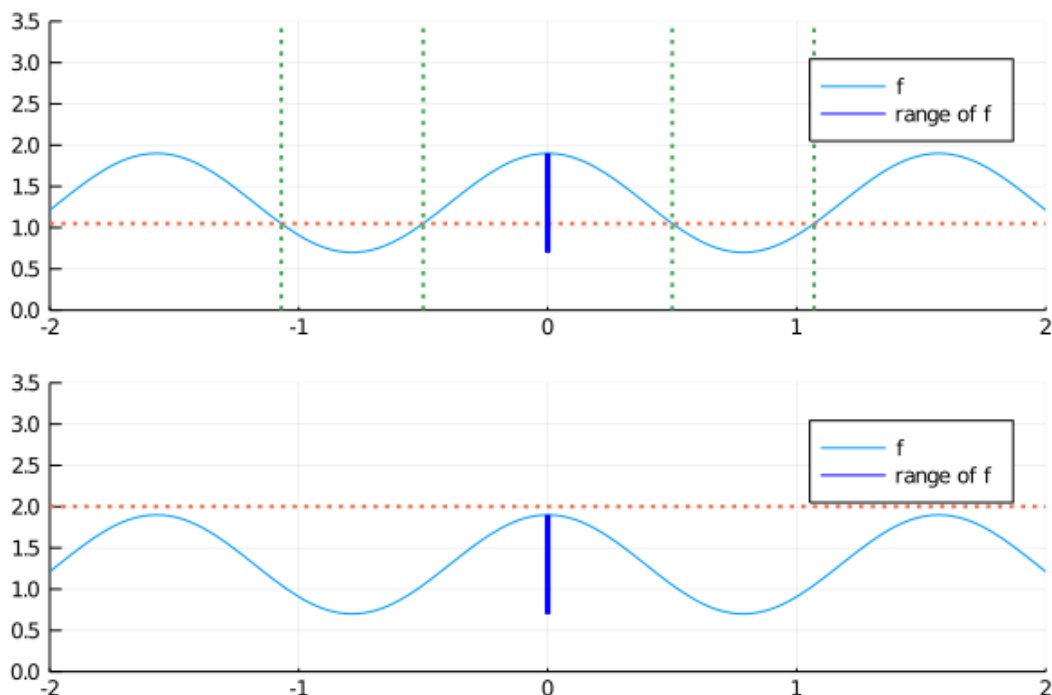
Given an arbitrary function f and a y , is there always an x such that $y = f(x)$?

If so, is it always unique?

The answer to both these questions is negative, as the next figure shows

```
In [21]: f(x) = 0.6*cos(4x) + 1.3
grid = range(-2, 2, length = 100)
y_min, y_max = extrema( f(x) for x in grid )
plt1 = plot(f, xlim = (-2, 2), label = "f")
hline!(plt1, [f(0.5)], linestyle = :dot, linewidth = 2, label = "")
vline!(plt1, [-1.07, -0.5, 0.5, 1.07], linestyle = :dot, linewidth = 2,
label = "")
plot!(plt1, fill(0, 2), [y_min y_min; y_max y_max], lw = 3, color = :blue,
label = ["range of f" ""])
plt2 = plot(f, xlim = (-2, 2), label = "f")
hline!(plt2, [2], linestyle = :dot, linewidth = 2, label = "")
plot!(plt2, fill(0, 2), [y_min y_min; y_max y_max], lw = 3, color = :blue,
label = ["range of f" ""])
plot(plt1, plt2, layout = (2, 1), ylim = (0, 3.5))
```

Out[21]:



In the first plot there are multiple solutions, as the function is not one-to-one, while in the second there are no solutions, since y lies outside the range of f .

Can we impose conditions on A in (3) that rule out these problems?

In this context, the most important thing to recognize about the expression Ax is that it corresponds to a linear combination of the columns of A .

In particular, if a_1, \dots, a_k are the columns of A , then

$$Ax = x_1a_1 + \dots + x_ka_k$$

Hence the range of $f(x) = Ax$ is exactly the span of the columns of A .

We want the range to be large, so that it contains arbitrary y .

As you might recall, the condition that we want for the span to be large is **linear independence**.

A happy fact is that linear independence of the columns of A also gives us uniqueness.

Indeed, it follows from our **earlier discussion** that if $\{a_1, \dots, a_k\}$ are linearly independent and $y = Ax = x_1a_1 + \dots + x_ka_k$, then no $z \neq x$ satisfies $y = Az$.

5.1 The $n \times n$ Case

Let's discuss some more details, starting with the case where A is $n \times n$.

This is the familiar case where the number of unknowns equals the number of equations.

For arbitrary $y \in \mathbb{R}^n$, we hope to find a unique $x \in \mathbb{R}^n$ such that $y = Ax$.

In view of the observations immediately above, if the columns of A are linearly independent, then their span, and hence the range of $f(x) = Ax$, is all of \mathbb{R}^n .

Hence there always exists an x such that $y = Ax$.

Moreover, the solution is unique.

In particular, the following are equivalent

1. The columns of A are linearly independent
2. For any $y \in \mathbb{R}^n$, the equation $y = Ax$ has a unique solution

The property of having linearly independent columns is sometimes expressed as having *full column rank*.

5.1.1 Inverse Matrices

Can we give some sort of expression for the solution?

If y and A are scalar with $A \neq 0$, then the solution is $x = A^{-1}y$.

A similar expression is available in the matrix case.

In particular, if square matrix A has full column rank, then it possesses a multiplicative *inverse matrix* A^{-1} , with the property that $AA^{-1} = A^{-1}A = I$.

As a consequence, if we pre-multiply both sides of $y = Ax$ by A^{-1} , we get $x = A^{-1}y$.

This is the solution that we're looking for.

5.1.2 Determinants

Another quick comment about square matrices is that to every such matrix we assign a unique number called the *determinant* of the matrix — you can find the expression for it [here](#).

If the determinant of A is not zero, then we say that A is *nonsingular*.

Perhaps the most important fact about determinants is that A is nonsingular if and only if A is of full column rank.

This gives us a useful one-number summary of whether or not a square matrix can be inverted.

5.2 More Rows than Columns

This is the $n \times k$ case with $n > k$.

This case is very important in many settings, not least in the setting of linear regression (where n is the number of observations, and k is the number of explanatory variables).

Given arbitrary $y \in \mathbb{R}^n$, we seek an $x \in \mathbb{R}^k$ such that $y = Ax$.

In this setting, existence of a solution is highly unlikely.

Without much loss of generality, let's go over the intuition focusing on the case where the columns of A are linearly independent.

It follows that the span of the columns of A is a k -dimensional subspace of \mathbb{R}^n .

This span is very “unlikely” to contain arbitrary $y \in \mathbb{R}^n$.

To see why, recall the [figure above](#), where $k = 2$ and $n = 3$.

Imagine an arbitrarily chosen $y \in \mathbb{R}^3$, located somewhere in that three dimensional space.

What's the likelihood that y lies in the span of $\{a_1, a_2\}$ (i.e., the two dimensional plane through these points)?

In a sense it must be very small, since this plane has zero “thickness”.

As a result, in the $n > k$ case we usually give up on existence.

However, we can still seek a best approximation, for example an x that makes the distance $\|y - Ax\|$ as small as possible.

To solve this problem, one can use either calculus or the theory of orthogonal projections.

The solution is known to be $\hat{x} = (A'A)^{-1}A'y$ — see for example chapter 3 of these notes

5.3 More Columns than Rows

This is the $n \times k$ case with $n < k$, so there are fewer equations than unknowns.

In this case there are either no solutions or infinitely many — in other words, uniqueness never holds.

For example, consider the case where $k = 3$ and $n = 2$.

Thus, the columns of A consists of 3 vectors in \mathbb{R}^2 .

This set can never be linearly independent, since it is possible to find two vectors that span \mathbb{R}^2 .

(For example, use the canonical basis vectors)

It follows that one column is a linear combination of the other two.

For example, let's say that $a_1 = \alpha a_2 + \beta a_3$.

Then if $y = Ax = x_1 a_1 + x_2 a_2 + x_3 a_3$, we can also write

$$y = x_1(\alpha a_2 + \beta a_3) + x_2 a_2 + x_3 a_3 = (x_1 \alpha + x_2) a_2 + (x_1 \beta + x_3) a_3$$

In other words, uniqueness fails.

5.4 Linear Equations with Julia

Here's an illustration of how to solve linear equations with Julia's built-in linear algebra facilities

```
In [22]: A = [1.0 2.0; 3.0 4.0];
```

```
In [23]: y = ones(2, 1); # A column vector
```

```
In [24]: det(A)
```

```
Out[24]: -2.0
```

```
In [25]: A_inv = inv(A)
```

```
Out[25]: 2x2 Array{Float64,2}:
 -2.0  1.0
  1.5 -0.5
```

```
In [26]: x = A_inv * y # solution
```

```
Out[26]: 2x1 Array{Float64,2}:
 -0.9999999999999998
  0.9999999999999999
```

```
In [27]: A * x # should equal y (a vector of ones)
```

```
Out[27]: 2x1 Array{Float64,2}:
 1.0
 1.0000000000000004
```

```
In [28]: A \ y # produces the same solution
```

```
Out[28]: 2x1 Array{Float64,2}:
 -1.0
  1.0
```

Observe how we can solve for $x = A^{-1}y$ by either via `inv(A) * y`, or using `A \ y`.

The latter method is preferred because it automatically selects the best algorithm for the problem based on the types of **A** and **y**.

If **A** is not square then `A \ y` returns the least squares solution $\hat{x} = (A'A)^{-1}A'y$.

6 Eigenvalues and Eigenvectors

Let A be an $n \times n$ square matrix.

If λ is scalar and v is a non-zero vector in \mathbb{R}^n such that

$$Av = \lambda v$$

then we say that λ is an *eigenvalue* of A , and v is an *eigenvector*.

Thus, an eigenvector of A is a vector such that when the map $f(x) = Ax$ is applied, v is merely scaled.

The next figure shows two eigenvectors (blue arrows) and their images under A (red arrows).


As expected, the image Av of each v is just a scaled version of the original

```
In [29]: A = [1 2
              2 1]
          evals, vecs = eigen(A)

          a1, a2 = evals
          eig_1 = [0 0; vecs[:,1]]'
          eig_2 = [0 0; vecs[:,2]]'
          x = range(-5, 5, length = 10)
          y = -x

          plot(eig_1[:, 2], a1 * eig_2[:, 2], arrow = true, color = :red,
               legend = :none, xlims = (-3, 3), ylims = (-3, 3), xticks = -3:3,
               yticks = -3:3,
               framestyle = :origin)
          plot!(a2 * eig_1[:, 2], a2 * eig_2, arrow = true, color = :red)
          plot!(eig_1, eig_2, arrow = true, color = :blue)
          plot!(x, y, color = :blue, lw = 0.4, alpha = 0.6)
          plot!(x, x, color = :blue, lw = 0.4, alpha = 0.6)
```

Out[29]:



The eigenvalue equation is equivalent to $(A - \lambda I)v = 0$, and this has a nonzero solution v only when the columns of $A - \lambda I$ are linearly dependent.

This in turn is equivalent to stating that the determinant is zero.

Hence to find all eigenvalues, we can look for λ such that the determinant of $A - \lambda I$ is zero.

This problem can be expressed as one of solving for the roots of a polynomial in λ of degree n .

This in turn implies the existence of n solutions in the complex plane, although some might be repeated.

Some nice facts about the eigenvalues of a square matrix A are as follows

1. The determinant of A equals the product of the eigenvalues.
2. The trace of A (the sum of the elements on the principal diagonal) equals the sum of the eigenvalues.
3. If A is symmetric, then all of its eigenvalues are real.
4. If A is invertible and $\lambda_1, \dots, \lambda_n$ are its eigenvalues, then the eigenvalues of A^{-1} are $1/\lambda_1, \dots, 1/\lambda_n$.

A corollary of the first statement is that a matrix is invertible if and only if all its eigenvalues are nonzero.

Using Julia, we can solve for the eigenvalues and eigenvectors of a matrix as follows

```
In [30]: A = [1.0 2.0; 2.0 1.0];
```

```
In [31]: evals, vecs = eigen(A);
```


In [32]: evals

```
Out[32]: 2-element Array{Float64,1}:  
 -1.0  
  3.0
```

In [33]: evecs

```
Out[33]: 2×2 Array{Float64,2}:  
 -0.707107  0.707107  
  0.707107  0.707107
```

Note that the *columns* of `evecs` are the eigenvectors.

Since any scalar multiple of an eigenvector is an eigenvector with the same eigenvalue (check it), the eig routine normalizes the length of each eigenvector to one.

6.1 Generalized Eigenvalues

It is sometimes useful to consider the *generalized eigenvalue problem*, which, for given matrices A and B , seeks generalized eigenvalues λ and eigenvectors v such that

$$Av = \lambda Bv$$

This can be solved in Julia via `eigen(A, B)`.

Of course, if B is square and invertible, then we can treat the generalized eigenvalue problem as an ordinary eigenvalue problem $B^{-1}Av = \lambda v$, but this is not always the case.

7 Further Topics

We round out our discussion by briefly mentioning several other important topics.

7.1 Series Expansions

Recall the usual summation formula for a geometric progression, which states that if $|a| < 1$, then $\sum_{k=0}^{\infty} a^k = (1 - a)^{-1}$.

A generalization of this idea exists in the matrix setting.

7.1.1 Matrix Norms

Let A be a square matrix, and let

$$\|A\| := \max_{\|x\|=1} \|Ax\|$$

The norms on the right-hand side are ordinary vector norms, while the norm on the left-hand side is a *matrix norm* — in this case, the so-called *spectral norm*.

For example, for a square matrix S , the condition $\|S\| < 1$ means that S is *contractive*, in the sense that it pulls all vectors towards the origin Section ??.

7.1.2 Neumann's Theorem

Let A be a square matrix and let $A^k := AA^{k-1}$ with $A^1 := A$.

In other words, A^k is the k -th power of A .

Neumann's theorem states the following: If $\|A^k\| < 1$ for some $k \in \mathbb{N}$, then $I - A$ is invertible, and

$$(I - A)^{-1} = \sum_{k=0}^{\infty} A^k \quad (4)$$

7.1.3 Spectral Radius

A result known as Gelfand's formula tells us that, for any square matrix A ,

$$\rho(A) = \lim_{k \rightarrow \infty} \|A^k\|^{1/k}$$

Here $\rho(A)$ is the *spectral radius*, defined as $\max_i |\lambda_i|$, where $\{\lambda_i\}_i$ is the set of eigenvalues of A .

As a consequence of Gelfand's formula, if all eigenvalues are strictly less than one in modulus, there exists a k with $\|A^k\| < 1$.

In which case (4) is valid.

7.2 Positive Definite Matrices

Let A be a symmetric $n \times n$ matrix.

We say that A is

1. *positive definite* if $x'Ax > 0$ for every $x \in \mathbb{R}^n \setminus \{0\}$
2. *positive semi-definite* or *nonnegative definite* if $x'Ax \geq 0$ for every $x \in \mathbb{R}^n$

Analogous definitions exist for negative definite and negative semi-definite matrices.

It is notable that if A is positive definite, then all of its eigenvalues are strictly positive, and hence A is invertible (with positive definite inverse).

7.3 Differentiating Linear and Quadratic forms

The following formulas are useful in many economic contexts. Let

- z, x and a all be $n \times 1$ vectors
- A be an $n \times n$ matrix
- B be an $m \times n$ matrix and y be an $m \times 1$ vector

Then

1. $\frac{\partial a'x}{\partial x} = a$

2. $\frac{\partial Ax}{\partial x} = A'$
3. $\frac{\partial x'Ax}{\partial x} = (A + A')x$
4. $\frac{\partial y'Bz}{\partial y} = Bz$
5. $\frac{\partial y'Bz}{\partial B} = yz'$

Exercise 1 below asks you to apply these formulas.

7.4 Further Reading

The documentation of the linear algebra features built into Julia can be found [here](#).

Chapters 2 and 3 of the [Econometric Theory](#) contains a discussion of linear algebra along the same lines as above, with solved exercises.

If you don't mind a slightly abstract approach, a nice intermediate-level text on linear algebra is [1].

8 Exercises

8.1 Exercise 1

Let x be a given $n \times 1$ vector and consider the problem

$$v(x) = \max_{y,u} \{-y'Py - u'Qu\}$$

subject to the linear constraint

$$y = Ax + Bu$$

Here

- P is an $n \times n$ matrix and Q is an $m \times m$ matrix
- A is an $n \times n$ matrix and B is an $n \times m$ matrix
- both P and Q are symmetric and positive semidefinite

(What must the dimensions of y and u be to make this a well-posed problem?)

One way to solve the problem is to form the Lagrangian

$$\mathcal{L} = -y'Py - u'Qu + \lambda' [Ax + Bu - y]$$

where λ is an $n \times 1$ vector of Lagrange multipliers.

Try applying the formulas given above for differentiating quadratic and linear forms to obtain the first-order conditions for maximizing \mathcal{L} with respect to y, u and minimizing it with respect to λ .

Show that these conditions imply that

1. $\lambda = -2Py$
2. The optimizing choice of u satisfies $u = -(Q + B'PB)^{-1}B'PAx$
3. The function v satisfies $v(x) = -x'\tilde{P}x$ where $\tilde{P} = A'PA - A'PB(Q + B'PB)^{-1}B'PA$

As we will see, in economic contexts Lagrange multipliers often are shadow prices

Note

If we don't care about the Lagrange multipliers, we can substitute the constraint into the objective function, and then just maximize $-(Ax+Bu)'P(Ax+Bu)-u'Qu$ with respect to u . You can verify that this leads to the same maximizer.

9 Solutions

Thanks to [Willem Hekman](#) and Guanlong Ren for providing this solution.

9.1 Exercise 1

We have an optimization problem:

$$v(x) = \max_{y,u} \{-y'Py - u'Qu\}$$

s.t.

$$y = Ax + Bu$$

with primitives

- P be a symmetric and positive semidefinite $n \times n$ matrix.
- Q be a symmetric and positive semidefinite $m \times m$ matrix.
- A an $n \times n$ matrix.
- B an $n \times m$ matrix.

The associated Lagrangian is :

$$L = -y'Py - u'Qu + \lambda'[Ax + Bu - y]$$

9.1.1 1.

Differentiating Lagrangian equation w.r.t y and setting its derivative equal to zero yields

$$\frac{\partial L}{\partial y} = -(P + P')y - \lambda = -2Py - \lambda = 0,$$

since P is symmetric.

Accordingly, the first-order condition for maximizing L w.r.t. y implies

$$\lambda = -2Py.$$

9.1.2 2.

Differentiating Lagrangian equation w.r.t. u and setting its derivative equal to zero yields

$$\frac{\partial L}{\partial u} = -(Q + Q')u - B'\lambda = -2Qu + B'\lambda = 0.$$

Substituting $\lambda = -2Py$ gives

$$Qu + B'Py = 0.$$

Substituting the linear constraint $y = Ax + Bu$ into above equation gives

$$Qu + B'P(Ax + Bu) = 0$$

$$(Q + B'PB)u + B'PAx = 0$$

which is the first-order condition for maximizing L w.r.t. u .

Thus, the optimal choice of u must satisfy

$$u = -(Q + B'PB)^{-1}B'PAx,$$

which follows from the definition of the first-order conditions for Lagrangian equation.

9.1.3 3.

Rewriting our problem by substituting the constraint into the objective function, we get

$$v(x) = \max_u \{ -(Ax + Bu)'P(Ax + Bu) - u'Qu \}.$$

Since we know the optimal choice of u satisfies $u = -(Q + B'PB)^{-1}B'PAx$, then

$$v(x) = -(Ax + Bu)'P(Ax + Bu) - u'Qu \quad \text{with } u = -(Q + B'PB)^{-1}B'PAx$$

To evaluate the function

$$\begin{aligned} v(x) &= -(Ax + Bu)'P(Ax + Bu) - u'Qu \\ &= -(x'A' + u'B')P(Ax + Bu) - u'Qu \\ &= -x'A'PAx - u'B'PAx - x'A'PBu - u'B'PBu - u'Qu \\ &= -x'A'PAx - 2u'B'PAx - u'(Q + B'PB)u \end{aligned}$$

For simplicity, denote by $S := (Q + B'PB)^{-1}B'PA$, then $u = -Sx$.

Regarding the second term $-2u'B'PAx$,

$$\begin{aligned} -2u'B'PAx &= -2x'S'B'PAx \\ &= 2x'A'PB(Q + B'PB)^{-1}B'PAx \end{aligned}$$

Notice that the term $(Q + B'PB)^{-1}$ is symmetric as both P and Q are symmetric.

Regarding the third term $-u'(Q + B'PB)u$,

$$\begin{aligned} -u'(Q + B'PB)u &= -x'S'(Q + B'PB)Sx \\ &= -x'A'PB(Q + B'PB)^{-1}B'PAx \end{aligned}$$

Hence, the summation of second and third terms is $x'A'PB(Q + B'PB)^{-1}B'PAx$.

This implies that

$$\begin{aligned} v(x) &= -x'A'PAx - 2u'B'PAx - u'(Q + B'PB)u \\ &= -x'A'PAx + x'A'PB(Q + B'PB)^{-1}B'PAx \\ &= -x'[A'PA - A'PB(Q + B'PB)^{-1}B'PA]x \end{aligned}$$

Therefore, the solution to the optimization problem $v(x) = -x'\tilde{P}x$ follows the above result by denoting $\tilde{P} := A'PA - A'PB(Q + B'PB)^{-1}B'PA$.

Footnotes

[1] Suppose that $\|S\| < 1$. Take any nonzero vector x , and let $r := \|x\|$. We have $\|Sx\| = r\|S(x/r)\| \leq r\|S\| < r = \|x\|$. Hence every point is pulled towards the origin.

References

- [1] K Jänich. *Linear Algebra*. Springer Undergraduate Texts in Mathematics and Technology. Springer, 1994.