Perron - Frobenius eigenfunctions of perturbed stochastic matrices

Rajeeva L. Karandikar
Director
Chennai Mathematical Institute
rlk@cmi.ac.in

Abstract:

Consider a stochastic matrix P for which the Perron-Frobenius Eigenvalue has multiplicity larger than 1 and for $\varepsilon > 0$, let

$$P^{\varepsilon} = (1-\varepsilon)P + \varepsilon Q$$

where Q is a stochastic matrix for which the Perron-Frobenius Eigenvalue has multiplicity 1. Let π^{ε} be the Perron-Frobenius eigenfunction for P^{ε}. We will discuss behavior of π^{ε} as $\varepsilon \to 0$.

This was an important ingredient in showing that if two players repeatedly play Prisoner's Dilemma, without knowing that they are playing a game, and if they play rationally, they end up cooperating. We will discuss this as well in the second half.

The talk will include required background on Markov chains.

Let $P = (p_{ij})$ be a $m \times m$ matrix with $p_{ij} \geq 0$ and $\sum_j p_{ij} = 1$ for all i, j. Clearly, $\lambda = 1$ is an eigenvalue of P with $(1, 1, \ldots, 1)^T$ as an eigenvector. Such matrices are called *Stochastic* matrices.

Perron-Frobenius theorem says that if $p_{ij} > 0$ for all i, j, then multiplicity of eigenvalue 1 is 1, and the left-eigenvector π can be chosen to have all entires positive with $\sum_j \pi_j = 1$. The left-eigenvector so chosen is called Perron-Frobenius eigenvector.

All other eigenvalues λ_k satisfy: $|\lambda_k| < 1$.

In this case, ($p_{ij} > 0$ for all i, j), it follows that

$$(P^n)_{ij} \to \pi_j \text{ for all } i, j.$$

Strong connections with Markov Chains, which we will discuss later. The Perron-Frobenius eigenvector is known as stationary distribution for the Markov chain.

A Stochastic matrix P is said to be *irreducible* if for all i, j, $\exists n \geq 1$ such that

$$(P^n)_{ij} > 0.$$

A Stochastic matrix P is said to be *primitive* if $\exists n \geq 1$ such that

$$(P^n)_{ij} > 0 \text{ for all } i, j.$$

An irreducible P is primitive if and only if

$$\text{g.c.d.}\{n \geq 1 : (P^n)_{ii} > 0\} = 1.$$

Perron-Frobenius theorem is valid verbatim for a primitive stochastic matrix (in Markov chain context, this is called irreducible aperiodic case).

The question we will discuss: *Let P, Q be stochastic matrices and let Q be primitive. For $\varepsilon > 0$, let*

$$P^{\varepsilon} = (1-\varepsilon)P + \varepsilon Q.$$

Then P^{ε} is primitive and let π^{ε} be its Perron-Frobenius eigenvector.

Question : Does π^{ε} converge as $\varepsilon \downarrow 0$ and if so, how do we characterise the limit?

The answer is clear if P is also primitive. What can we say when P is not primitive and geometric multiplicity of eigenvalue 1 is 2 or more for P?

In general, eigenvalues behave well under perturbation but eigenvectors do not, specially in situation like the one here where geometric multiplicity of eigenvalue 1 is 1 for P^{ε}, $\varepsilon > 0$, but for the limit the geometric multiplicity is bigger than 1.

The next result shows that indeed, π^{ε} converges.

Theorem : Let A be an $m \times m$ stochastic matrix. There exist (universal) polynomials $u_1, u_2, \ldots u_m$ of order $m \times m$ with non-negative coefficients such that

$$v_j = u_j(a_{11}, a_{12}, \ldots, a_{1m}, a_{21}, a_{22}, \ldots, a_{2m}, \ldots, a_{m1}, \ldots, a_{mm})$$

$1 \leq j \leq m$, satisfy

$$\sum_j v_j a_{jk} = v_k \text{ for all } k$$

Thus, if $\sum_j v_j = \alpha > 0$, (this can be shown to be true for an irreducible stochastic matrix A) then

$$\pi_j = \frac{1}{\alpha} v_j$$

is the Perron-Frobenius eigenvector.

Let $S = \{1, 2, \ldots, m\}$. We will be considering directed graphs G on S and for such a graph G, let

$$\theta(G)(A) = \prod_{(j \mapsto k) \in G} a_{kj}$$

Note that $\theta(G)(A)$ is a polynomial in matrix entries with positive coefficients.

A tree rooted at j is a directed connected spanning graph G such that (i) there is no incoming edge into the vertex j, (ii) the incoming degree for all vertices other than j is 1 and (iii) there are no cycles. Let $\Gamma(j)$ denote the set of all trees rooted at j. Let

$$\gamma_j(A) = \sum_{G \in \Gamma(j)} \theta(G)(A)$$

Note that $\gamma_j(A)$ is a polynomial in matrix entries with positive coefficients.

We will prove that

$$\sum_{i \in S, i \neq j} \gamma_i(A) a_{ij} = \gamma_j(A)(1 - a_{jj}) \ \forall j \in S \tag{1}$$

Since A is a stochastic matrix, this will yield

$$\sum_{i \in S} \gamma_i(A) a_{ij} = \gamma_j(A) \ \forall j \in S$$

showing that $\gamma_i(A)$ are the required polynomials.

Let us fix j and we will prove (1). Let Λ be the set of all directed connected spanning graphs on S that have exactly one cycle that contains the vertex j and such that every vertex has incoming degree 1.

We are going to get two ways of computing $\sum_{H \in \Lambda} \theta(H)(A)$ - one method would yield LHS and the other RHS of (1).

Look at the graph of an element of Λ :

(with $S = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$ and $j = 4$) :

directed connected spanning graph on S that has exactly one cycle that contains the vertex 4 and such that every vertex has incoming degree 1.

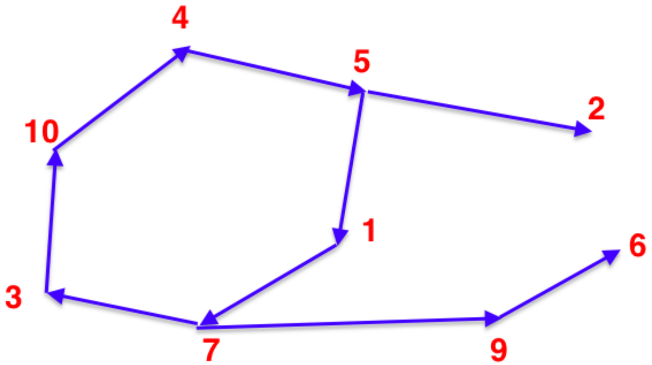

directed spanning graph with exactly one loop containing 4

If we delete the incoming edge to 4, we get a tree rooted at 4.

If we delete the outgoing edge from 4, we get a tree rooted at the other end of the deleted edge, here 5.

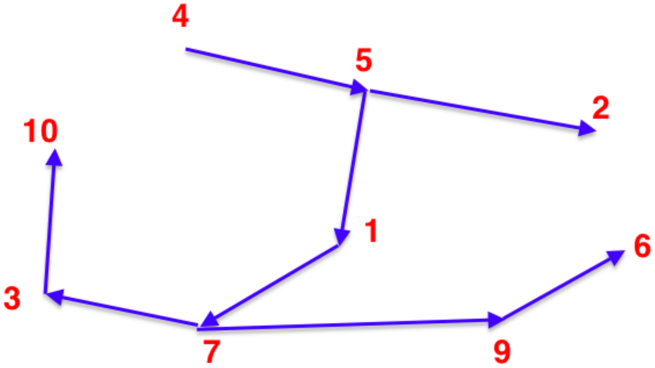

Tree rooted at 4

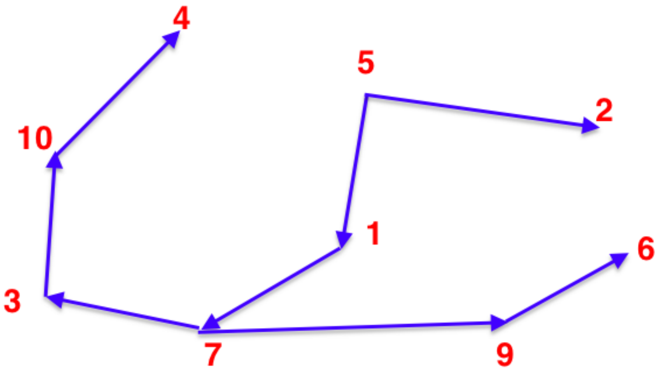

Tree rooted at 5

Also, if we take the tree rooted at 4 and add an edge from any of the other nodes, we will get an element of Λ- *a directed connected spanning graph on* S *that has exactly one cycle that contains the vertex* 4 *and such that every vertex has incoming degree* 1*.*

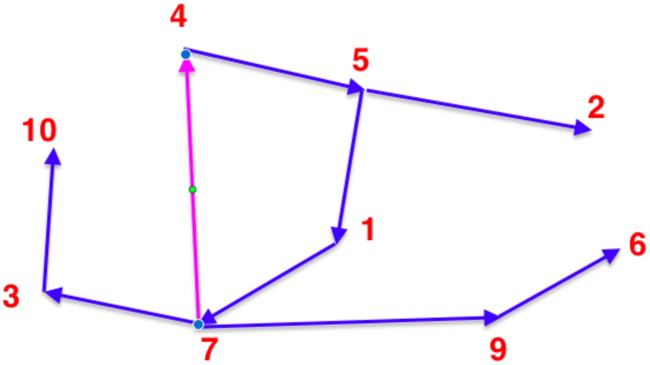
4
5
2
10
1
6
3
7
9

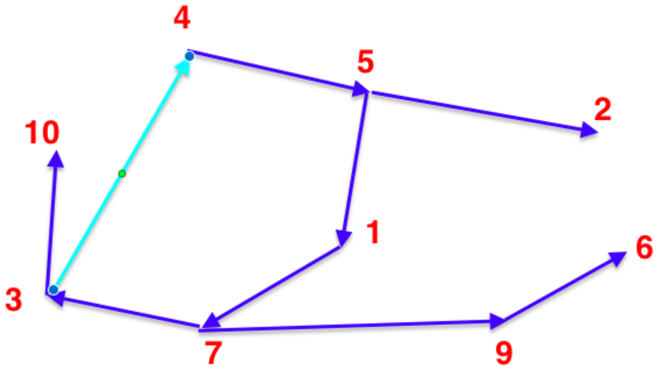
4
5
2
10
1
6
3
7
9

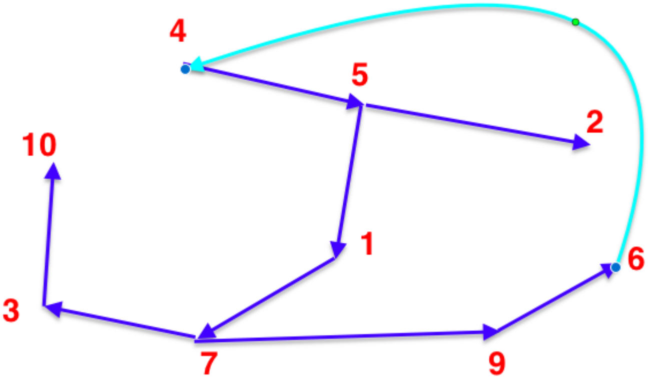
4
5
2
10
1
6
3
7
9

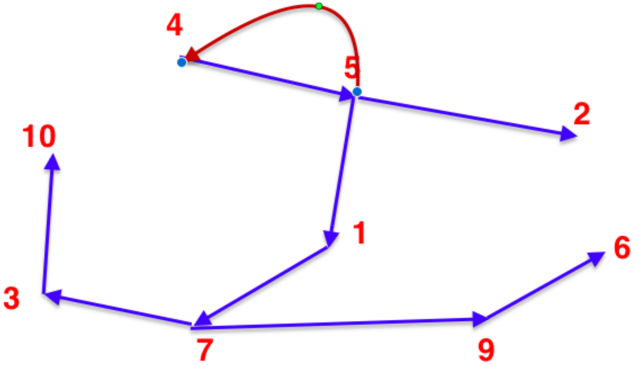
4
5
2
10
1
6
3
7
9

Also, if we take any tree, say rooted at 1, as in next slide, and add an edge from 4 to 1 we will get an element of Λ- *a directed connected spanning graph on* S *that has exactly one cycle that contains the vertex* 4 *and such that every vertex has incoming degree* 1.

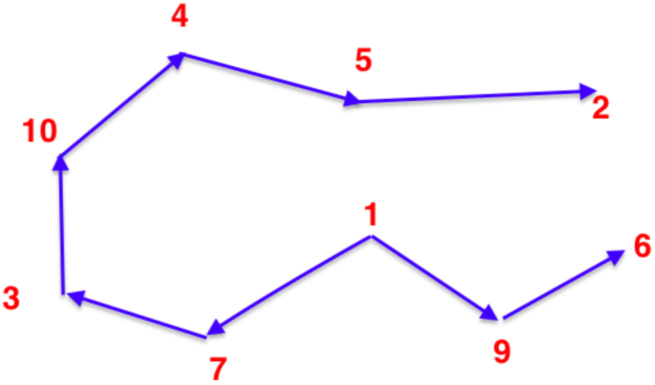
4
5
2
10
1
6
3
9
7

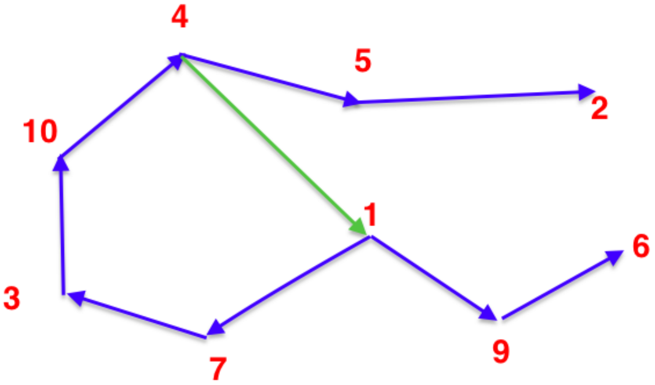

4
5
2
10
1
6
3
9
7

This discussion leads to :

The mapping f from $\cup_{k \in S, k \neq j} \Gamma(k)$ into Λ defined by:

$$f(G) = G \cup \{j \mapsto k\} \;\; \text{for } G \in \Gamma(k)$$

(*i.e.* take $k \neq j$, $G \in \Gamma(k)$ and add the directed edge $j \mapsto k$ to G)
is one-one onto Λ.

The mapping g from $\Gamma(j) \times (S - j)$ into Λ defined by: for

$$g(G, k) = G \cup \{k \mapsto j\} \quad G \in \Gamma(j) \text{ and } k \in (S - j)$$

(*i.e.* take $G \in \Gamma(j)$, $k \in (S - \{j\})$ and add the directed edge $k \mapsto j$ to G)

is one-one onto Λ.

Thus

$$\sum_{H\in\Lambda}\theta(H)(A) = \sum_{k\in S,k\neq j}\left(\sum_{G\in\Gamma(k)}\theta(f(G))(A)\right)$$
$$= \sum_{k\in S,k\neq j}\left(\sum_{G\in\Gamma(k)}\theta(G)(A)a_{kj}\right)$$
$$= \sum_{k\in S,k\neq j}\gamma_k(A)a_{kj}$$

The RHS above is the LHS of (1)

and

$$\sum_{H\in\Lambda}\theta(H)(A) = \sum_{k\in S,\,k\neq j}\left(\sum_{G\in\Gamma(j)}\theta(g(G,k))(A)\right)$$
$$= \sum_{k\in S,\,k\neq j}\left(\sum_{G\in\Gamma(j)}\theta(G)(A)a_{jk}\right)$$
$$= \sum_{k\in S,\,k\neq j}\gamma_j(A)a_{jk}$$
$$= \gamma_j(A)(1-a_{jj})$$

The RHS above is the RHS of (1). This completes the proof.

Returning to $P^{\varepsilon} = (1-\varepsilon)P + \varepsilon Q$ we see that each $\gamma_j(P^{\varepsilon})$ is a polynomial in ε with positive coefficients. Hence in

$$\pi^{\varepsilon} = \frac{\gamma_j(P^{\varepsilon})}{\sum_k \gamma_k(P^{\varepsilon})}$$

the smallest power of ε with non-zero coefficient in the numerator is larger than the the smallest power of ε with non-zero coefficient in the denominator. Thus

$$\lim_{\varepsilon \to 0} \pi^{\varepsilon} \text{ exists.}$$

How does one characterize the limit π^* of π^ε? It is clearly one of the eigenvectors of P corresponding to eigenvalue 1 since

$$\pi^* P = \pi^*.$$

Introduction to Markov chains:

Given a $m \times m$ stochastic matrix $P = ((p_{ij})$ and a $1 \times m$ vector $a = (a_1, a_2, \ldots a_m)$, we can construct a stochastic process X_k with values in $S = \{1, 2, \ldots, m\}$ such that:

$$\text{Prob}(X_0 = i_0, X_1 = i_1, X_2 = i_2, \ldots X_n = i_n) = a_{i_0} p_{i_0 i_1} p_{i_1 i_2} \ldots p_{i_{n-1} i_n}.$$

Here, a is the initial distribution, or distribution of X_0 , and P is called the *transition probability* matrix since

$$\text{Prob}(X_{k+1} = j \mid X_k = i) = p_{ij} \ \textit{for } i, j \in S$$

Moreover,

$$\text{Prob}(X_{k+1} = j \mid (X_0, X_1, \ldots, X_{k-1}) \in B, X_k = i) = p_{ij}$$

for any $B \subseteq S^k$. This is called the Markov property of the process $\{X_n\}$, which is called a Markov Chain. It can be seen that

$$\text{Prob}(X_{k+n} = j \mid X_k = i) = (P^n)_{ij}$$

If $\pi P = \pi$ and we construct the chain with π as the initial distribution, then

$$\text{Prob}(X_n = j) = \pi_j$$

for all $n \geq 1$, for all $j \in S$. Thus the distribution of X_n is stationary and thus π is also called stationary initial distribution.

P is irreducible here means that for all $i, j \in S$, $\exists n \geq 1$ such that

$$\text{Prob}(X_n = j \mid X_0 = i) > 0.$$

The period d_i of $i \in S$ is defined as

$$d_i = \text{g.c.d.}\ \{n \geq 1 : (P^n)_{ii} > 0.$$

For an irreducible chain, $d_i = d_j$ for all $i, j \in S$. The chain is called aperiodic if $d_i = 1$ for all i.

For an irreducible aperiodic chain (when P matrix is primitive),

$$\lim_{n\to\infty} \text{Prob}(X_n = j \mid X_0 = i) = \lim_{n\to\infty} (P^n)_{ij} = \pi_j.$$

and as a consequence

$$\lim_{n\to\infty} \text{Prob}(X_n = j) = \lim_{n\to\infty} \sum_{i\in S} a_i (P^n)_{ij} = \pi_j.$$

Thus irrespective of the initial distribution, the distribution of X_n converges to π.

Heuristics: Returning to $P^{\varepsilon} = (1-\varepsilon)P + \varepsilon Q$

Consider the case where P^n converges to a matrix$R$ and Q^n converges to S. Both R and S are stochastic matrices. Let

$B^{\varepsilon} = (1-\varepsilon)R + \varepsilon Q$

$C^{\varepsilon} = (1-\varepsilon)P + \varepsilon S$

$D^{\varepsilon} = (1-\varepsilon)R + \varepsilon S$

and let θ^{ε}, ξ^{ε} and η^{ε} be the Perron-Frobenious eigenvectors for B^{ε}, C^{ε}, D^{ε} respectively.

Through extensive numerical computation, I discovered that limits π^* of π^{ε} and θ^* of θ^{ε} are the same and limits xi^* of ξ^{ε} and η^* of η^{ε} are the same, while in general $\pi^* \neq \xi^*$.

Using the large number of examples, I kept coming up with conjectures to characterize π^*.

Led me to a conjecture that if RQR is irreducible and aperiodic (or primitive) and thus has a unique stationary distribution (Perron-Frobenius eigenvector) $\tilde{\pi}$, then

$$\pi^* = \tilde{\pi}.$$

I also had a probabilistic proof, justifying the intuitive reasoning.

But then my Economist friend told me that this was the toy case and their real interest is when P is the probability transition kernel for $[0, 1]$-valued Markov Chain (in discrete time).

The Theory of Markov Chains in discrete time with an uncountable set as its state space is not that well studied.

Restatement of the problem:

Suppose the state space S is a compact metric space.

A *probability transition kernel* Γ on S is a mapping from $S \times \mathscr{B}(S)$ into $[0,1]$ such that

For each $A \in \mathscr{B}(S)$, $x \mapsto \Gamma(x,A)$ is a Borel measurable function on S

For each $x \in S$, $A \mapsto \Gamma(x,A)$ is a probability measure on $(S, \mathscr{B}(S))$.

A probability measure μ on S is an invariant measure for P if

$$\int \Gamma(x,A)\, d\mu(x) = \mu(A) \;\; \forall A \in \mathscr{B}(S).$$

Suppose P, Q are probability transition kernels on S and $P^{\varepsilon} = (1-\varepsilon)P + \varepsilon Q$.
Suppose that for $0 < \varepsilon < 1$, P^{ε} admits a unique invariant measure, π^{ε}. Then does π^{ε} converge, and if it does converge to π^*, how does one characterize π^*?

Let me give a proof of the discrete state space case which can be scaled to the compact state space case under appropriate conditions:

It can be shown that for any stochastic matrix P,

$$\frac{1}{n}\sum_{t=1}^{n} P^t$$

converges (as $n \to \infty$) to a stochastic matrix, say R.
Moreover, for $0 < \lambda < 1$, writing $K_\lambda = (I - \lambda P)^{-1}$, one has

$$\lim_{\lambda\to 1}(1-\lambda)K_\lambda \to R$$

or

$$\lim_{\varepsilon\to 0}\varepsilon(I-(1-\varepsilon)P)^{-1} \to R \qquad (2)$$

Theorem : Suppose that either (i) the matrix RQR is primitive, namely, the eigenvalue 1 has geometric multiplicity 1 for RQ or (ii) the matrix QR is primitive. Then π^* is this unique eigenvector (of RQR and/or QR).

Proof : Since

$$\pi^{\varepsilon}((1-\varepsilon)P+\varepsilon Q)=\pi^{\varepsilon}$$

it follows that

$$\pi^{\varepsilon}\varepsilon Q=\pi^{\varepsilon}(I-(1-\varepsilon)P)$$

From $\pi^{\varepsilon}\varepsilon Q = \pi^{\varepsilon}(I-(1-\varepsilon)P)$ we conclude

$$\pi^{\varepsilon} Q[\varepsilon(I-(1-\varepsilon)P)^{-1}] = \pi^{\varepsilon}$$

Taking limit as $\varepsilon \downarrow 0$ and using (2) we conclude

$$\pi^* QR = \pi^*$$

Since π^* obviously satisfies $\pi^* P = \pi^*$ and hence $\pi^* R = \pi^*$, it follows that

$$\pi^* RQR = \pi^*.$$

Coming to the case of compact metric space S as the state space, need some definitions:

A probability transition kernel Γ is said to be *strongly Feller* if for all bounded continuous functions f on S, $x \mapsto \int f(u)\Gamma(x, du)$ is continuous.

A probability transition kernel Γ is said to be *open set irreducible* if for all open sets U in S and all $x \in S$, $\sum_{n=1}^{\infty} \Gamma^n(x, U) > 0$

For probability transition kernels Γ, Λ on S, $\Gamma * \Lambda$ is defined by

$$[\Gamma * \Lambda](x, A) = \int \Lambda(u, A)\Gamma(x, du).$$

For a probability transition kernel Γ, $\Gamma^{(n)}$ are defined inductively by $\Gamma^{(1)} = \Gamma$ and for $k \geq 1$, $\Gamma^{(k+1)} = \Gamma^{(k)} * \Gamma$, *i.e.*

$$\Gamma^{(k+1)}(x, A) = \int \Gamma^{(k)}(u, A)\Gamma(x, du).$$

Suppose P, Q are probability transition kernels on S such that

(i) Q is strongly Feller and open set irreducible.

(ii) $\exists R$ - a probability transition kernel on S such that for all bounded continuous functions f on S,

$$\lim_n \frac{1}{n}\sum_{k=1}^{n}\int f(u)\, dP^{(k)}(x, du) \to \int f(u) R(x, du).$$

(iii) The kernel $Q * R$ admits a unique invariant probability measure π^*.

Let π^{ε} be the unique invariant probability measure for $P^{\varepsilon} = (1-\varepsilon)P + \varepsilon Q$.
Then π^{ε} converges to π^{*} in the sense of weak convergence of measures.

Prisoner's Dilemma Consider the following 2×2 game:

	C	D
C	(σ, σ)	$(0, \theta)$
D	$(\theta, 0)$	(δ, δ)

where $\theta > \sigma > \delta > 0$.