IRREDUCIBLE OPERATORS

P. R. Halmos

Dedicated to Marshall Harvey Stone on his 65th birthday.

THEOREM. *On a separable Hilbert space, the set of irreducible operators is a dense* G_δ.

Remarks. The theorem says that the set of irreducible operators is topologically large: most operators are irreducible. (The separability assumption is obviously necessary; on a non-separable space every operator is reducible.) The proof rests on several analytic and algebraic lemmas; they occur in Section 1 below. Section 2 contains a few related remarks on finite-dimensional spaces (everything is easier) and on the set of normal operators (it is topologically small). The topological considerations needed to show that the set of irreducible operators is a G_δ occur in Section 3. Although the principal theorem gives some information about the size of the set of reducible operators, it does not answer all questions about that set. For instance, it is still not known whether the set of reducible operators is dense; Section 4 contains some comments on that subject. Closely related to this whole circle of ideas is the possibility of a topological attack on the problem of invariant subspaces. It is, after all, not inconceivable that the existence of an operator with only trivial invariant subspaces could be proved by showing that the set of all such operators is topologically large. Section 5 contains a result that seems to kill that hope, and it suggests that in fact the set is topologically small. An appendix contains a theorem, a special case of which is used in Section 2; the result (rank is weakly lower semi-continuous) may be of interest in its own right.

Terminology. Hilbert spaces are complex, subspaces are closed linear manifolds, operators are bounded linear transformations, and, when it is not otherwise indicated, all topological concepts (for both vectors and operators) refer to the norm topology. The *commutant* of a set of operators is the set of all those operators that commute with each operator in the given set. An operator is *irreducible* if its commutant contains no projections other than 0 and 1.

Notation. The underlying Hilbert space is H. If E is a subset of H, then $\bigvee E$ is the span of E (the smallest subspace that includes E). If E is an orthonormal basis for H, then $\mathbb{D}(E)$ is the set of all operators that are diagonal with respect to E (that is, the operators for which each element of E is an eigenvector). If A is an operator on H, then

$$\Re A = \frac{1}{2}(A + A^*) \qquad \text{and} \qquad \Im A = \frac{1}{2i}(A - A^*).$$

The set of all those operators A for which both $\Re A$ and $\Im A$ are simple diagonal operators is $\mathbb{D}$. That is: $A \in \mathbb{D}$ if and only if there exist orthonormal bases E and F such that $\Re A \in \mathbb{D}(E)$, $\Im A \in \mathbb{D}(F)$, and all eigenvalues of both $\Re A$ and $\Im A$ have multiplicity 1. If $\mathbb{K}$ is a set of operators, its commutant is $\mathbb{K}'$. The set of irreducible operators is $\mathbb{I}$, the set of reducible operators is $\mathbb{R}$, and the set of scalar multiples of the identity is $\mathbb{C}$.

Received September 22, 1967.
Research supported in part by a grant from the National Science Foundation.

1. DENSITY

LEMMA 1. *A necessary and sufficient condition that an operator* A *be irreducible is that* $\{\mathfrak{R} A\}' \cap \{\mathfrak{I} A\}' = \mathbb{C}$.

Proof. If A is reducible, then there is a non-trivial projection P in $\{A\}'$. Since P is Hermitian, $P \in \{A^*\}'$, and therefore $P \in \{\mathfrak{R} A\}' \cap \{\mathfrak{I} A\}'$; it follows that if $\{\mathfrak{R} A\}' \cap \{\mathfrak{I} A\}' = \mathbb{C}$, then $A \in \mathbb{I}$. If, conversely, $\{\mathfrak{R} A\}' \cap \{\mathfrak{I} A\}' \neq \mathbb{C}$, then that intersection contains a non-trivial projection, and therefore A is reducible; it follows that if $A \in \mathbb{I}$, then $\{\mathfrak{R} A\}' \cap \{\mathfrak{I} A\}' = \mathbb{C}$.

LEMMA 2. *If* E *is an orthonormal basis and* A *is an operator in* $\mathbb{D}(E)$ *with simple spectrum, then* $\{A\}' = \mathbb{D}(E)$.

Proof. A standard and elementary computation.

LEMMA 3. *If* E *and* F *are orthonormal bases, then a necessary and sufficient condition that* $\mathbb{D}(E) \cap \mathbb{D}(F) \neq \mathbb{C}$ *is that there exist non-trivial subsets* E_0 *and* F_0 *of* E *and* F, *respectively, such that* $\bigvee E_0 = \bigvee F_0$.

Here non-trivial means not empty and not the whole set.

Proof. If $A \in \mathbb{D}(E) \cap \mathbb{D}(F)$ and $A \notin \mathbb{C}$, consider an eigenvalue α_0 of A. Let E_0 and F_0 be the sets of eigenvectors corresponding to α_0 in E and in F, respectively. Clearly, $\bigvee E_0 = \bigvee F_0 = \{f: Af = \alpha_0 f\}$. Conversely, if E_0 and F_0 are non-trivial subsets of E and F with $\bigvee E_0 = \bigvee F_0$, let A be the projection whose range is that common span. Then $A \notin \mathbb{C}$ and $A \in \mathbb{D}(E) \cap \mathbb{D}(F)$.

LEMMA 4. *If* A *is an operator such that both* $\mathfrak{R} A$ *and* $\mathfrak{I} A$ *are diagonal operators with simple spectrum (that is, if* $A \in \mathbb{D}$), *and if* $\mathfrak{I} A$ *has an eigenvector* f *such that* $(e, f) \neq 0$ *for every eigenvector* e *of* $\mathfrak{R} A$, *then* $A \in \mathbb{I}$.

Proof. Let E and F be orthonormal bases such that $\mathfrak{R} A \in \mathbb{D}(E)$ and $\mathfrak{I} A \in \mathbb{D}(F)$. If E_0 and F_0 are subsets of E and F such that $\bigvee E_0 = \bigvee F_0$, then

$$\bigvee(E - E_0) = \left(\bigvee E_0\right)^{\perp} = \left(\bigvee F_0\right)^{\perp} = \bigvee(F - F_0).$$

There is no loss of generality in assuming that $f \in F$ (the simplicity assumption implies that some scalar multiple of f belongs to F), and, in view of the preceding comment, there is no loss of generality in assuming that $f \in F_0$ (since f belongs either to F_0 or to $F - F_0$, and the difference between the two cases is merely notational). Since, by assumption, f has a non-zero projection on every e in E, the only way it can happen that $f \in \bigvee E_0$ is that $E_0 = E$. The desired conclusion follows from Lemmas 3, 2, and 1.

LEMMA 5. $\mathbb{D}$ *is dense.*

Proof. It is sufficient to prove that the set of Hermitian operators in $\mathbb{D}$ is dense in the set of all Hermitian operators. To do this, represent any given Hermitian operator as a multiplication on L^2 over a finite measure space. (This is where the separability assumption comes in.) The multiplier can be uniformly approximated by simple functions. Multiplication by a real-valued simple function is the direct sum of a finite set of real scalars, and consequently it is a diagonal Hermitian operator. A diagonal Hermitian operator can obviously be approximated by one with simple spectrum.

The next auxiliary result is easy, but it has a certain interest in its own right. It is the solution of an approximation problem: given two unit vectors, find a unitary operator that maps one onto the other and is as near as possible to the identity.

LEMMA 6. *If* f *and* g *are unit vectors, then the infimum of* $\|1 - U\|$ *over all unitary operators* U *such that* $Uf = g$ *is equal to* $\|f - g\|$, *and it is always attained.*

The proof shows that U can be chosen so that $Uh = h$ whenever $h \perp \{f, g\}$; this is sometimes useful.

Proof. If $Uf = g$, then $\|1 - U\| \geq \|f - Uf\| = \|f - g\|$, so that the infimum is not smaller than $\|f - g\|$. What follows is the proof that $\|f - g\|$ can always be attained.

If $g = \alpha f$, put $Uf = \alpha f$ and $Uh = h$ whenever $h \perp f$. In all other cases the span of f and g has dimension 2. Define $Uh = h$ whenever $h \perp \{f, g\}$; the problem is thereby reduced to an elementary matrix computation. Choose coordinates so that $f = \langle 1, 0 \rangle$. If $g = \langle \alpha, \beta \rangle$ (with $|\alpha|^2 + |\beta|^2 = 1$, of course), put

$$U = \begin{pmatrix} \alpha & -\beta^* \\ \beta & \alpha^* \end{pmatrix}.$$

Clearly, U is unitary. Since $\|1 - U\|^2 = \|(1 - U)(1 - U^*)\| = \|2 - U - U^*\|$, since $2 - U - U^*$ is $2 - \alpha - \alpha^*$ times the identity, and since $\|f - g\|^2 = 2 - \alpha - \alpha^*$, the proof is complete.

Proof of density. Since $\mathbb{D}$ is dense (Lemma 5), it is sufficient to prove that if $\varepsilon > 0$ and $A \in \mathbb{D}$, then there exists an irreducible A_0 in $\mathbb{D}$ with $\|A - A_0\| < \varepsilon$. Write $B = \Re A$ and $C = \Im A$, and let E and F be orthonormal bases such that $B \in \mathbb{D}(E)$ and $C \in \mathbb{D}(F)$.

Consider an arbitrary f in F. The Fourier expansion of f with respect to E may have some zero coefficients; let g be a unit vector obtained from f by varying those coefficients slightly so that $(e, g) \neq 0$ for all e in E. Let the variation be so slight that $\|f - g\| < \varepsilon/2\|C\|$. Lemma 6 yields a unitary operator U such that $Uf = g$ and $\|1 - U\| < \varepsilon/2\|C\|$. If $C_0 = UCU^*$, then C_0 is a diagonal Hermitian operator with simple spectrum and

$$\begin{aligned} \|C - C_0\| &= \|C - UCU^*\| \leq \|C - UC\| + \|UC - UCU^*\| \\ &\leq \|C\| \cdot \|1 - U\| + \|UC\| \cdot \|1 - U^*\| < \varepsilon. \end{aligned}$$

If $A_0 = B + iC_0$, then $A_0 \in \mathbb{D}$. Since C_0 has an eigenvector (namely g) that has a non-zero inner product with every eigenvector of B, it follows from Lemma 4 that A_0 is irreducible. Since, finally,

$$\|A - A_0\| = \|C - C_0\| < \varepsilon,$$

the proof is complete.

R. G. Douglas has observed that in the proof of Lemma 5 it is possible to invoke the von Neumann approximation (each Hermitian operator when suitably perturbed by an operator of arbitrarily small Hilbert-Schmidt norm becomes diagonal) in place of the more obvious norm approximation; the result is that the principal theorem is true in the sense of Hilbert-Schmidt approximation also. The same result was obtained, later but independently, by J. G. Stampfli.

2. FINITE-DIMENSIONALITY AND NORMALITY

If H is finite-dimensional, the theorem has a relatively simple geometric proof, and it can be significantly strengthened.

Here is a possible proof. (a) The operators all whose eigenvalues have algebraic multiplicity 1 are dense. (b) If all the eigenvalues of an operator have algebraic multiplicity 1, then its eigenvectors span H, and consequently there exists a linear basis of H consisting of eigenvectors. (c) By a small perturbation an operator whose eigenvectors span H can be transmuted into another one of the same kind such that some particular element of a basis consisting of eigenvectors is *not* orthogonal to any other. (Consider a basis of eigenvectors, form the span of all but one, note that that one is not in the span, and perturb it slightly, if necessary, so as to push it out of the orthogonal complement of the span.) (d) If a basis consisting of eigenvectors of an operator is such that some element of it is not orthogonal to any other, then that operator is irreducible. (For each reducing subspace, every eigenvector must belong either to it or to its orthogonal complement; whichever one the distinguished "non-orthogonal" vector belongs to must contain all others.)

The strengthening is that $\mathbb{I}$ is not only a G_δ (the proof of this in the general case is in Section 3) but open.

PROPOSITION 1. *On a finite-dimensional Hilbert space the set of reducible operators is closed and nowhere dense.*

Proof. Suppose that A_n is reducible and $A_n \to A$, and, for each n, let P_n be a non-trivial projection that commutes with A_n. Finite-dimensionality implies the compactness of the unit ball in the space of operators. There is, therefore, no loss of generality in assuming that $P_n \to P$, where P is, of course, a projection, and, clearly, $AP = PA$. If $\dim H = k$, then $1 \leq \operatorname{rank} P_n \leq k - 1$; since rank is lower semicontinuous (see the Appendix), it follows that $P \neq 0, 1$. This proves that $\mathbb{R}$ is closed; that it is nowhere dense follows from the already proved density of its complement.

Similar easy techniques give information about the size of the set of normal operators on spaces of arbitrarily large dimension.

PROPOSITION 2. *On a Hilbert space of dimension greater than* 1, *the set of normal operators is closed and nowhere dense.*

Proof. Closure is obvious; since the mapping $A \to A^*A - AA^*$ is continuous, the set $\{A: A^*A - AA^* = 0\}$ is closed. A closed set is nowhere dense if and only if its complement is dense. Since an irreducible operator can be normal only if the space has dimension 0 or 1, the conclusion for separable spaces follows from the density of the set of irreducible operators.

There is a better, direct proof, independent of the previous density theorem, that works for non-separable spaces just as well as for separable ones. It is sufficient to prove that every normal operator is arbitrarily near non-normal ones. For scalars this is obvious. (Find a non-normal T, and consider $\lambda + \varepsilon T$, where λ is the given scalar and ε is small.) If A is normal but not a scalar, then either $\Re A$ or $\Im A$ is not a scalar; say $\Re A$ is not. Then there exists a Hermitian T that does not commute with $\Re A$ (this is where it is necessary that the dimension be greater than 1). If $\varepsilon > 0$, then $A + i\varepsilon T$ is not normal but can be arbitrarily near to A.

3. TOPOLOGY

This section contains the second half of the proof of the principal theorem, that is, the proof that $\mathbb{I}$ is a G_δ.

Let $\mathbb{P}$ be the set of all those Hermitian operators P on H for which $0 \leq P \leq 1$. Recall that $\mathbb{P}$ is exactly the weak closure of the set of projections. Let $\mathbb{P}_0$ be the subset of those elements of $\mathbb{P}$ that are *not* scalar multiples of the identity. Since $\mathbb{P}$ is a weakly closed subset of the unit ball, it is weakly compact, and hence the weak topology for $\mathbb{P}$ is metrizable. Since the set of scalars is weakly closed, it follows that $\mathbb{P}_0$ is weakly locally compact. Since the weak topology for $\mathbb{P}$ has a countable base, the same is true for $\mathbb{P}_0$, and therefore $\mathbb{P}_0$ is weakly σ-compact. Let $\mathbb{P}_1, \mathbb{P}_2, \cdots$ be weakly compact subsets of $\mathbb{P}_0$ such that $\bigcup_{n=1}^{\infty} \mathbb{P}_n = \mathbb{P}_0$.

It is to be proved that $\mathbb{R}$ is an F_σ (norm topology). Let $\hat{\mathbb{P}}_n$ be the set of all those operators A on H for which there exists a P in $\mathbb{P}_n$ such that $AP = PA$ $(n = 1, 2, 3, \cdots)$; the spectral theorem implies that $\bigcup_{n=1}^{\infty} \hat{\mathbb{P}}_n = \mathbb{R}$.

The proof can be completed by showing that each $\hat{\mathbb{P}}_n$ is (norm) closed. Suppose, therefore, that $A_k \in \hat{\mathbb{P}}_n$ and that $A_k \to A$ (norm). For each k, find a P_k in $\mathbb{P}_n$ such that $A_k P_k = P_k A_k$. Since $\mathbb{P}_n$ is weakly compact and metrizable, there is no loss of generality in assuming that the sequence $\{P_k\}$ is weakly convergent to P, say. (This is the point where it is advantageous to consider all the operators in $\mathbb{P}$, and not just projections; there is no guarantee that P is a projection even if the P_k's are. Note that $P \in \mathbb{P}_n$, so that, in particular, P is not a scalar.)

Assertion: $AP = PA$. This follows from an easy lemma: if $A_k \to A$ (norm) and $P_k \to P$ (weak), then $A_k P_k \to AP$ and $P_k A_k \to PA$ (weak). Indeed:

$$\begin{aligned} |(A_k P_k f, g) - (APf, g)| &\leq |(A_k P_k f, g) - (A P_k f, g)| + |(A P_k f, g) - (APf, g)| \\ &\leq \|A_k - A\| \cdot \|f\| \cdot \|g\| + |((P_k - P)f, A^* g)|. \end{aligned}$$

(It is important that the sequence $\{P_k\}$ is bounded.) These inequalities imply that $A_k P_k \to AP$ (weak); the other order follows from the consideration of adjoints. Once this is done, everything is done: $A \in \hat{\mathbb{P}}_n$, hence $\hat{\mathbb{P}}_n$ is closed (norm), hence $\mathbb{R}$ is an F_σ (norm).

4. REDUCIBILITY

By the principal theorem, the set $\mathbb{R}$ of reducible operators is always an F_σ, and, by Proposition 1, in the finite-dimensional case $\mathbb{R}$ is closed. Could it be that $\mathbb{R}$ is always closed? The answer is no. Reason: on an infinite-dimensional space every operator of finite rank is reducible, so that every compact operator is in the closure of $\mathbb{R}$, but it is easy to construct compact operators (weighted shifts) that are irreducible.

There is another example, which shows that $\mathbb{R}$ is not closed in a more surprising way. For each positive integer n, let $\mathbb{R}_n$ be the set of all operators that have a reducing subspace of dimension n.

PROPOSITION 3. *Every isometry is in the closure of* $\mathbb{R}_1$.

Proof. Observe to begin with that for an operator to be nearly irreducible is the same as to be near to a reducible operator. This auxiliary assertion can be stated as follows, in precise terms: if $\|A - A_0\| < \varepsilon$ and P is a projection such that $A_0 P = PA_0$, then $\|AP - PA\| < 2\varepsilon$; if, conversely, $\|AP - PA\| < \varepsilon$ and $A_0 = PAP + (1 - P)A(1 - P)$, then $A_0 P = PA_0$ and $\|A - A_0\| < 2\varepsilon$. The proof of the first assertion is implied by the inequality

$$\|AP - PA\| \leq \|AP - A_0 P\| + \|A_0 P - PA_0\| + \|PA_0 - PA\| .$$

The proof of the second assertion is implied by

$$\begin{aligned}\|A - A_0\| &= \|(1 - P)AP + PA(1 - P)\| \\ &= \|(AP - PA)P - P(AP - PA)\| \leq 2\|AP - PA\| .\end{aligned}$$

The motivation for the definition of A_0 is the consideration of the matrix of A corresponding to the direct sum decomposition of H into ran P and $(\text{ran } P)^{\perp}$: throw away the off corners.

In view of the preceding paragraph, it is sufficient to prove that if U is an isometry, then there exist projections of rank 1 that nearly commute with U. To prove it, let λ be a number of modulus 1 that is an approximate eigenvalue of U, that is, an element of the approximate point spectrum. It follows, by definition, that corresponding to each positive number ε there is a unit vector e such that $\|Ue - \lambda e\| < \varepsilon$, and hence such that

$$\|U^* e - \lambda^* e\| = \|-\lambda^* U^* (Ue - \lambda e)\| < \varepsilon .$$

If P is the projection onto e, that is, if $Pf = (f, e)e$ for all f, then

$$\begin{aligned}\|(UP - PU)f\| = \|(f, e)Ue - (Uf, e)e\| &\leq \|(f, e)Ue - (f, e)\lambda e\| + \|(f, e)\lambda e - (f, U^* e)e\| \\ &\leq |(f, e)| \cdot \|Ue - \lambda e\| + \|f\| \cdot \|U^* e - \lambda e\| \leq 2\varepsilon \|f\| ,\end{aligned}$$

so that $\|UP - PU\| \leq 2\varepsilon$. The proof is complete.

Since there exist irreducible isometries (for example, the unilateral shift), Proposition 3 implies again that $\mathbb{R}$ is not closed.

Proposition 3 raises the hope that $\mathbb{R}_1$ is dense, but it is not. If H is a Hilbert space, and if A is the operator on $H \oplus H$ with matrix $\begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}$, then A is not in the closure of $\mathbb{R}_1$. (It is easy to see that A is in $\mathbb{R}_2$; in fact, A is the direct sum of operators of rank 2, equal in number to the dimension of H.)

To prove that A is not in the closure of $\mathbb{R}_1$, it is sufficient to prove that the infimum of $\|A - B\|$, as B varies over $\mathbb{R}_1$, is positive. Given B in $\mathbb{R}_1$, let $\langle f, g \rangle$ (with f and g in H) be a reducing eigenvector of B of norm 1; that is,

$$B\langle f, g \rangle = \lambda \langle f, g \rangle, \quad B^*\langle f, g \rangle = \lambda^* \langle f, g \rangle, \quad \|f\|^2 + \|g\|^2 = 1 .$$

Since $A - \lambda = (A - B) + (B - \lambda)$, it follows that

$$\|(A - \lambda)\langle f, g \rangle\| \leq \|A - B\| \quad \text{and} \quad \|(A^* - \lambda^*)\langle f, g \rangle\| \leq \|A - B\| ;$$

since $(A - \lambda)\langle f, g \rangle = \langle -\lambda f, f - \lambda g \rangle$ and $(A^* - \lambda^*)\langle f, g \rangle = \langle g - \lambda^* f, -\lambda^* g \rangle$, it follows that

$|\lambda|^2 \|f\|^2 + \|f - \lambda g\|^2 \leq \|A - B\|^2$ and $\|g - \lambda^* f\|^2 + |\lambda|^2 \|g\|^2 \leq \|A - B\|^2$.

The latter inequalities imply that

$$1 + 2|\lambda|^2 - 4\,\Re\,\lambda^*(f, g) \leq 2\|A - B\|^2,$$

and hence that $1 + 2|\lambda|^2 - 2|\lambda| \leq 2\|A - B\|^2$. Since $1 - 2|\lambda|^2 - 2|\lambda| \geq 1/2$ for all λ, the proof is complete.

There is at least one question along these lines that seems to be of interest and that is unanswered: is $\mathbb{R}$ dense?

5. INVARIANCE

An operator is *transitive* if $\{0\}$ and H are the only subspaces it leaves invariant. The problem of invariant subspaces is to decide whether there exist transitive operators on Hilbert spaces of dimension greater than 1. (The word "transitive" was suggested by W. B. Arveson. Its present meaning is not identical with its meanings in group theory and ergodic theory, but it is in close harmony with them.) Let $\mathbb{T}$ be the set of transitive operators. One possible approach to the problem is to try to prove that $\mathbb{T}$ is not empty by proving that it is topologically large, *i. e.*, that it is (or includes) a dense G_δ. As it stands, this is doomed to failure: $\mathbb{T}$ is not dense.

PROPOSITION 4. *If* U *is a non-invertible isometry, and if* $\|U^* - A\| < 1$, *then* $\ker A \neq \{0\}$.

Proof. Since $\|U^* - A\| < 1$, it follows that $\|U^* U - AU\| < 1$, *i. e.*, that $\|1 - AU\| < 1$. From this, in turn, it follows that AU is invertible, and hence that ran A = H. If it were true that $\ker A = \{0\}$, then it would follow from the closed graph theorem that A in invertible. Since AU is already known to be invertible, it would then follow that U is invertible, and this contradicts the assumption.

Since there exist non-invertible isometries (for example, the unilateral shift), Proposition 4 implies that $\mathbb{T}$ is not dense.

There is an open question along these lines that has at least some curiosity value: is the complement of $\mathbb{T}$ topologically large? The following result is pertinent.

PROPOSITION 5. *The set of all operators with an eigenvalue is dense.*

Proof. The result is an easy consequence of the existence of approximate eigenvalues (compare the proof of Proposition 3). Given an operator A, let λ be an approximate eigenvalue of A; it follows that corresponding to each positive number ε there is a unit vector e such that $\|Ae - \lambda e\| < \varepsilon$. If P is the projection onto e, that is, if $Pf = (f, e)e$ for all f, and if $A_0 = A - (1 - P)AP$, then a direct verification proves that e is an eigenvector of A_0 with eigenvalue (Ae, e). Since $|(Ae, e) - \lambda| < \varepsilon$, it follows that

$$\|(A - A_0)f\| = \|(1 - P)APf\| \leq |(f, e)| \cdot \|(1 - P)Ae\|$$
$$\leq \|f\| \cdot \|Ae - (Ae, e)e\| \leq 2\varepsilon \|f\|,$$

so that $\|A - A_0\| \leq 2\varepsilon$. The proof is complete.

(In matrix terms the proof could have been phrased this way: choose an approximate eigenvector e for A with approximate eigenvalue λ; use e as the first term of an orthonormal basis; replace the first column of the resulting matrix for A by $\langle \lambda, 0, 0, \cdots \rangle$; the new matrix A_0 has the eigenvalue λ and is near to A.)

APPENDIX

THEOREM. *Rank is weakly lower semicontinuous.*

The rank of an operator is the dimension of the closure of its range. The statement means that for each operator A_0 there exists a weak neighborhood N of A_0 such that rank $A \geq$ rank A_0 for all A in N. Equivalently, in terms of convergence: if $\{A_n\}$ is a net that converges weakly to A_0, then $\liminf_n$ rank $A_n \geq$ rank A_0.

The possible values of rank in this context are the non-negative integers, together with ∞; no distinction is made among different infinite cardinals. Were such a distinction to be made, the result would become false. Here is an example. Let H be a non-separable Hilbert space with an orthonormal basis $\{e_j\}$. Let D be the set of all countable subsets of the index set, ordered by inclusion: for each n in D, write A_n for the projection onto $\bigvee \{e_j\colon j \in n\}$. Since for each f_0 in H there exists an n_0 in D such that $f_0 \perp e_j$ whenever $j \notin n_0$, it follows that $A_n \to 1$ (not only weakly, but, in fact, strongly). Since rank $A_n = \aleph_0$ and rank $1 > \aleph_0$, the cardinal version of semicontinuity is false.

LEMMA 1. *If* $\{e_1, \cdots, e_n\}$ *is an orthonormal set and* $\|f_i - e_i\| < 1/\sqrt{n}$ $(i = 1, \cdots, n)$, *then the set* $\{f_1, \cdots, f_n\}$ *is linearly independent.*

Proof. If $\xi_i \neq 0$ for at least one i, then

$$\left\|\sum_{i=1}^{n} \xi_i (f_i - e_i)\right\| \leq \sum_{i=1}^{n} |\xi_i| \cdot \|f_i - e_i\| < \left(\sum_{i=1}^{n} |\xi_i|\right) \cdot (1/\sqrt{n}) \leq \sqrt{n}\sqrt{\sum_{i=1}^{n} |\xi_i|^2}\Big/\sqrt{n},$$

and therefore

$$\left\|\sum_{i=1}^{n} \xi_i f_i\right\| \geq \left\|\sum_{i=1}^{n} \xi_i e_i\right\| - \left\|\sum_{i=1}^{n} \xi_i (f_i - e_i)\right\| > \sqrt{\sum_{i=1}^{n} |\xi_i|^2} - \sqrt{\sum_{i=1}^{n} |\xi_i|^2}.$$

LEMMA 2. *Rank is strongly lower semicontinuous.*

Proof. To prove: if rank $A_0 \geq n$ $(= 1, 2, 3, \cdots)$, then there exists a strong neighborhood N of A_0 such that rank $A \geq n$ for all A in N. Let $\{e_1, \cdots, e_n\}$ be an orthonormal set in ran A_0; find $f_1, \cdots, f_n$ such that $A_0 f_i = e_i$ $(i = 1, \cdots, n)$. Write

$$N = \{A\colon \|Af_i - A_0 f_i\| < 1/\sqrt{n},\ i = 1, \cdots, n\};$$

it follows from Lemma 1 that, for each A in N, the set $\{Af_1, \cdots, Af_n\}$ is linearly independent.

Proof of the theorem. To prove: if rank $A_0 \geq n$ $(= 1, 2, 3, \cdots)$, then there exists a weak neighborhood N of A_0 such that rank $A \geq n$ for all A in N. Let $\{e_1, \cdots, e_n\}$ be an orthonormal set in ran A_0; find $f_1, \cdots, f_n$ such that $A_0 f_i = e_i$ $(i = 1, \cdots, n)$. Write

$$N = \{A: |((A - A_0)f_j, e_i)| < \varepsilon;\ i, j = 1, \cdots, n\},$$

where ε is an as yet unspecified positive number. Given A in N, write $\alpha_{ij} = (Af_j, e_i)$. Note that $(A_0 f_j, e_i) = (e_j, e_i) = \alpha_{ij}$. Since (by Lemma 2) rank is strongly lower semicontinuous, and since for finite-dimensional spaces all the usual operator topologies coincide, it follows that if the matrix α is sufficiently near the identity matrix (that is, if ε is sufficiently small), then rank $\alpha \geq n$, and therefore rank $\alpha = n$. In other words, if ε is sufficiently small, then α is invertible. This implies that if $\sum_{j=1}^{n} \xi_j A f_j = 0$, so that

$$\sum_{j=1}^{n} \alpha_{ij} \xi_j = \sum_{j=1}^{n} (Af_j, e_i) \xi_j = 0,$$

then $\xi_1 = \cdots = \xi_n = 0$, so that the set $\{Af_1, \cdots, Af_n\}$ is linearly independent.

The University of Michigan
Ann Arbor, Mich. 48104