Jean Goubault-Larrecq

A Cone-Theoretic Krein-Milman Theorem

Research Report LSV-08-18

June 2008

Laboratoire Spécification et Vérification

CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE

Ecole Normale Supérieure de Cachan
61, avenue du Président Wilson
94235 Cachan Cedex France

A Cone-Theoretic Krein-Milman Theorem

Jean Goubault-Larrecq
LSV, ENS Cachan, CNRS, INRIA Futurs projet SECSI
61, avenue du président-Wilson, 94235 Cachan, France
goubault@lsv.ens-cachan.fr

June 5, 2008

Abstract

We prove the following analogue of the Krein-Milman Theorem: in any locally convex T_0 topological cone, every convex compact saturated subset is the compact saturated convex hull of its extreme points.

1 Introduction

The classical Krein-Milman Theorem states that any compact convex subset K of a locally convex topological vector space X is the closed convex hull of its extreme points.

We show that a similar result holds when X is a locally convex topological cone. Remarkably, the only visible modification in the conclusion of the theorem is that K will be the *co*-closed convex hull, where co-closedness denotes closure in the de Groot dual of X. Another important difference is that the right notion of extreme point requires an additional notion of minimality (see Definition 3.1), thus making our result stronger than expected. This being said, our proof follows the lines of the textbook proof of the classical Krein-Milman Theorem (Wikipedia, 2008).

2 Preliminaries

A topological space X is a space equipped with a collection of so-called *open* subsets of X, such that any union of opens is open, and any finite intersection of opens is open. The collection of all the opens is called the *topology* of X. A *closed* subset of X is the complement of an open. Any topological space comes with a *specialization quasi-ordering* $\leq$, defined by $x \leq y$ iff y belongs to every open that contains U. Note that every open subset of X is upward closed in $\leq$. A T_0 *space* is one where $\leq$ is a partial ordering. A subset of X is *saturated* iff it is upward closed.

A subset of X is *compact* iff one can extract a finite subcover from any open cover of it, i.e., if whenever $\bigcup_{i \in I} U_i$ contains Q, where U_i are opens, then there is a finite union of U_i's that contain Q.

In general, every closed subset of a compact subset is compact, but compacts may fail to be closed. We do not require the Hausdorff separation axiom, because non-trivial cones (see below) are never Hausdorff. In a Hausdorff space, every compact subset would be closed. In merely T_0 spaces, the mathematically interesting objects are the compact *saturated* subsets. Note that they are by definition upward closed, while closed subset are *downward* closed instead.

A map $f : X \to Y$ between two topological spaces is *continuous* iff $f^{-1}(V)$ is open in X for every open subset V of Y. The image by a continuous map of a compact set is always compact.

The *upward closure* $\uparrow E = \{x \in X \mid \exists y \in E \cdot y \leq x\}$ is saturated, and compact whenever E is finite. Since every open is upward closed, the upward closure of any compact is compact. The *downward closure* $\downarrow E = \{x \in X \mid \exists y \in E \cdot x \leq y\}$ is closed whenever E is finite. We also write $\uparrow x$ for $\uparrow \{x\}$ and $\downarrow x$ for $\downarrow \{x\}$. The latter is also the closure of $\{x\}$; the *closure* $cl(A)$ of $A \subseteq X$ is by definition the smallest closed set containing A.

A *cpo* is a partially ordered set in which every directed family $(x_i)_{i \in I}$ has a least upper bound $\sup_{i \in I} x_i$. A family $(x_i)_{i \in I}$ is *directed* iff it is non-empty, and any two elements have an upper bound in the family. Any partially ordered set can be equipped with the *Scott topology*, whose opens are the upward closed sets U such that whenever $(x_i)_{i \in I}$ is a directed family that has a least upper bound in U, then some x_i is in U already. The Scott topology is always T_0, and the specialization ordering is the original partial ordering.

A cone C is an additive commutative monoid with a scalar multiplication by elements of $\mathbb{R}^+$, satisfying laws similar to those of vector spaces. Precisely, a *cone* C is endowed with an addition $+ : C \times C \to C$, a zero element $0 \in C$, and a scalar multiplication $\cdot : \mathbb{R}^+ \times C \to C$ such that:

$$(x + y) + z = x + (y + z) \quad x + y = y + x \quad x + 0 = x$$
$$(rs) \cdot x = r \cdot (s \cdot x) \quad 1 \cdot x = x \quad 0 \cdot x = 0$$
$$r \cdot (x + y) = r \cdot x + r \cdot y \quad (r + s) \cdot x = r \cdot x + s \cdot x$$

A *topological cone* is a cone equipped with a topology that makes $+$ and $\cdot$ continuous, where $\mathbb{R}^+$ is equipped with its Scott topology (not its usual, metric topology, which we shall almost never use). The opens of $\mathbb{R}^+$ in its Scott topology are the intervals $(r, +\infty)$, $r \in \mathbb{R}^+$, together with $\mathbb{R}^+$ and $\emptyset$. The saturated compacts in $\mathbb{R}^+$ are then the intervals $[r, +\infty)$, $r \in \mathbb{R}^+$, together with $\emptyset$. The specialization ordering of $\mathbb{R}^+$ is the usual ordering $\leq$.

For example, ${\mathbb{R}^+}^n$ is a topological cone. Letting $\overline{\mathbb{R}^+}$ be $\mathbb{R}^+$ plus an additional point at infinity $+\infty$, with $x + (+\infty) = (+\infty) + x = +\infty$, $r \cdot (+\infty) = +\infty$ for every $r > 0$ and $0 \cdot (+\infty) = 0$, then $\overline{\mathbb{R}^+}$ is also a topological cone, and $\overline{\mathbb{R}^+}^n$ as well. Again, $\overline{\mathbb{R}^+}$ is equipped with its Scott topology, whose opens are the intervals $(r, +\infty]$, $r \in \mathbb{R}^+$, together with $\overline{\mathbb{R}^+}$ and $\emptyset$. The saturated compacts in $\mathbb{R}^+$ are then the intervals $[r, +\infty]$, $r \in \mathbb{R}^+$, together with $\emptyset$. Note that $\overline{\mathbb{R}^+}$ is not *regular*, i.e., $x + y = x' + y$ does not entail $x = x'$ (take $y = +\infty$). All vector spaces are regular in this sense.

Other remarkable topological cones include $\langle X \to \mathbb{R}^+\rangle$, the set of bounded continuous map from X to $\mathbb{R}^+$, with the Scott topology. The set of continuous maps from X to $\overline{\mathbb{R}^+}$ is a d-cone, i.e., a cone which is also a cpo. The set of all continuous valuations over a topological space X is a cone, too, for example (Goubault-Larrecq, 2007).

A subset Z of C is *convex* iff $r \cdot x + (1 - r) \cdot y$ is in Z whenever $x, y \in Z$ and $0 \leq r \leq 1$. C is itself *locally convex* iff every point has a basis of convex open neighborhoods, i.e., whenever $x \in U$, U open in C, then there is a convex open V such that $x \in V \subseteq U$.

A map $f : C \to \overline{\mathbb{R}^+}$ is *positively homogeneous* iff $f(r \cdot x) = rf(x)$ for all $x \in C$ and $r \in \mathbb{R}^+$. It is *additive* (resp. *super-additive*, resp. *sub-additive*) iff $f(x + y) = f(x) + f(y)$ (resp. $f(x + y) \geq f(x) + f(y)$, resp. $f(x + y) \leq f(x) + f(y)$) for all $x, y \in C$. It is *linear* (resp. *super-linear*, resp. *sub-linear*) iff its is both positively homogeneous and additive (resp. super-additive, resp. sub-additive).

We shall use the following Separation Theorem (Keimel, 2006, Theorem 9.1): let C be a topological cone, E a non-empty convex subset of C, U an open convex subset of C, disjoint from E; then there is a continuous linear map $f : C \to \overline{\mathbb{R}^+}$ such that $f(x) \leq 1$ for every $x \in E$ and $f(y) > 1$ for every $y \in U$. We shall especially use the following Separation Corollary (Keimel, 2006, Corollary 9.2): let C be a locally convex topological cone, and x, y two elements such that $x \not\leq y$, then there is a continuous linear map $f : C \to \overline{\mathbb{R}^+}$ such that $f(y) < f(x)$.

Finally, we shall use the following Geometric Separation Theorem (Keimel, 2006, Proposition 10.2): let C be a locally convex topological cone, K a compact convex subset of C, F a non-empty closed convex set disjoint from K, then there is a convex open set U containing K and disjoint from F.

3 Extreme Points, Faces, and the Theorem

Let C be a cone. For any two points x, y of C, let $]x, y[$ be the set of points of the form $r \cdot x + (1 - r) \cdot y$, with $0 < r < 1$. It is tempting to call this the *open line segment* between x and y, however be aware that it is generally not open.

Let Q a subset of a T_0 topological cone C, with specialization ordering $\leq$, which we shall fix for the rest of the section.

Definition 3.1 (Extreme Point) *An* extreme point *of Q is any element $x \in Q$ that is minimal in Q, and such that there are no two distinct points x_1 and x_2 of Q such that $x \in]x_1, x_2[$.*

Compared to the standard definition, we additionally require that x be minimal in Q, that is, that whenever $x' \in Q$ is such that $x' \leq x$, then $x' = x$. In topological

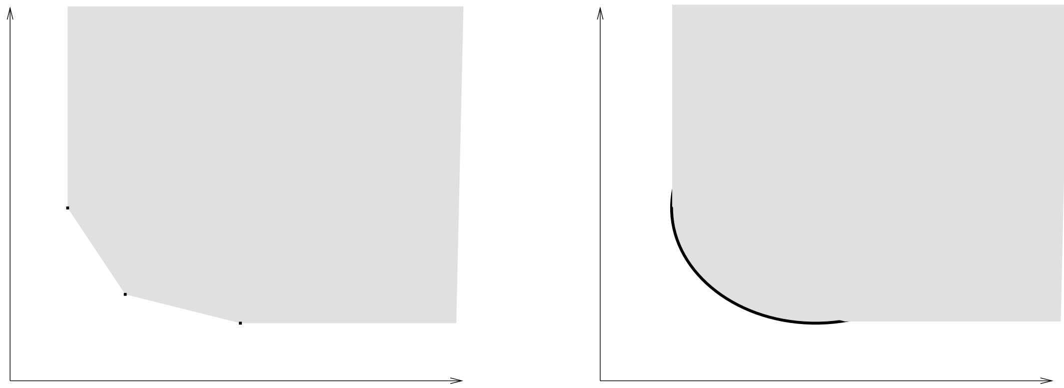

Figure 1: Extreme points

vector spaces, which are Hausdorff, the minimality condition would be vacuously true. However, recall that no non-trivial T_0 cone C can be Hausdorff, or even T_1: since $0 = 0 \cdot x \leq 1 \cdot x \leq x$, 0 is the least element of C, which implies that the specialization ordering $\leq$ does not coincide with equality.

As examples, look at Figure 1, representing convex saturated subsets of the quarter-plane $\mathbb{R}^{+2}$. The polytope on the left has three extreme points, shown as fat dots. The extreme points of the subset on the right are shown as a fat curve.

In the sequel, we shall need the notion of a closed subset of Q. This is by definition the intersection of a closed subset of C with Q.

Definition 3.2 (Face) *Call a* face A *of* Q *any non-empty closed subset of* Q *such that, for any* $x_1, x_2 \in Q$*, if* $]x_1, x_2[$ *intersects* A*, then* $]x_1, x_2[$ *is entirely contained in* A*.*

Lemma 3.3 *If* A *is a closed subset of* Q*, and* $]x_1, x_2[$ *is contained in* A*, then* x_1 *and* x_2 *are also in* A*.*

Proof. Let $[0,1]^m$ denotes $[0,1]$ together with its usual, metric topology. We reserve the notation $[0,1]$ for $[0,1]$ with its topology induced from that of $\mathbb{R}^+$, i.e., whose opens are $(t,1]$ for $t \geq 0$ and $[0,1]$ itself.

We claim that: $(*)$ given $x_1, x_2 \in C$, the map $f_{x_1,x_2} : r \in [0,1]^m \mapsto r \cdot x_1 + (1-r) \cdot x_2$ is continuous from $[0,1]^m$ to C. Indeed, this map is obtained by composing the map $(r,s) \in [0,1]^2 \mapsto r \cdot x_1 + s \cdot x_2$, which is continuous by the definition of a topological cone, and the map $r \in [0,1]^m \mapsto (r, 1-r)$. The latter is continuous, since the inverse image of the basic open set $(t,1] \times (t',1]$ is the open interval $(t, 1-t')$ (taking this to be empty when $t \geq 1-t'$).

Assume $]x_1, x_2[\subseteq A$, that is, $r \cdot x_1 + (1-r) \cdot x_2$ is in A for all r, $0 < r < 1$. So $f^{-1}_{x_1,x_2}(A)$ contains the open interval $(0,1)$. Remember that A is closed in Q, so we may write A as $F \cap Q$, where F is closed in X. Then $f^{-1}_{x_1,x_2}(F)$ contains $(0,1)$. However, by $(*)$ f_{x_1,x_2} is continuous. Since F is closed, $f^{-1}_{x_1,x_2}(F)$ is closed in $[0,1]^m$. But the only closed set in $[0,1]^m$ containing $(0,1)$ is $[0,1]$ itself. So both 0 and 1 are in $f^{-1}_{x_1,x_2}(F)$, i.e., x_1 and x_2 are in F. Since x_1 and x_2 are in Q, they are in A, too. $\square$

The extreme points of Q are exactly its one-element faces:

Proposition 3.4 *For every point $x \in C$, x is an extreme point of Q if and only if $\{x\}$ is a face of Q.*

Proof. Let x be an extreme point of Q, and let us show that $A = \{x\}$ is a face. First, A is non-empty, and since x is minimal in Q, A is the intersection of the closed set $\downarrow x$ with Q, hence is closed in Q. Assume $x_1, x_2 \in Q$ and $]x_1, x_2[$ intersects A. That is, x is in $]x_1, x_2[$. If x_1 and x_2 were disjoint, this would contradict the fact that x is extreme. So $x_1 = x_2$. Then $]x_1, x_2[= \{x\}$ is indeed contained in A.

Conversely, let $A = \{x\}$ be a one-element face of Q. Since A is closed in Q, we may write $A = F \cap Q$ with F closed. Also, $x \in F$, so F contains the closure $\downarrow x$ of x. So A contains $\downarrow x \cap Q$. Let $x' \in Q$ be such that $x' \leq x$. It follows that $x' \in \downarrow x \cap Q \subseteq A$, so $x' = x$. That is, x is indeed minimal in Q.

Assume x_1, x_2 are two distinct elements of Q such that $x \in]x_1, x_2[$. Since $]x_1, x_2[$ intersects A (at x), $]x_1, x_2[\subseteq A$. By Lemma 3.3, x_1 and x_2 are in A. That is, $x_1 = x$ and $x_2 = x$. This contradicts the fact that x_1 and x_2 are distinct. $\square$

Lemma 3.5 *If Q is convex and non-empty, then Q itself is a face of Q.*

Proof. Q is closed in itself, as the intersection of the closed set X with Q. Whenever $x_1, x_2 \in Q$, $]x_1, x_2[$ is contained in Q by convexity. $\square$

Let $Face(Q)$ denote the collection of all faces of Q, ordered by *reverse* inclusion $\supseteq$. Lemma 3.5 states that $Face(Q)$ has a least element.

Lemma 3.6 *If Q is compact, then $Face(Q)$ is a cpo.*

Proof. Let $(A_i)_{i \in I}$ be a directed family of faces. Since the ordering is $\supseteq$, directedness means that for any $i, i' \in I$, there is a face A_k contained in $A_i \cap A_{i'}$.

One may write A_i as $F_i \cap Q$, where F_i is closed and intersects Q. Let $F = \bigcap_{i \in I} F_i$. This is a closed set. Let U_i be the complement of F_i, U that of F. We claim that F intersects Q. Otherwise Q would be contained in $U = \bigcup_{i \in I} U_i$. Since each U_i is open and Q is compact, Q would be contained in some finite union $\bigcup_{k=0}^{n} U_{i_k}$, i.e., would not intersect $\bigcap_{k=0}^{n} F_{i_k}$. Since the family $(A_i)_{i \in I}$ is directed (for $\supseteq$), there is a face A_j that is contained in every A_{i_k}, $0 \leq k \leq n$. In particular, A_j is contained in $\bigcap_{k=0}^{n} F_{i_k}$, hence would not intersect Q. This contradicts the fact that, as a face, A_j is non-empty and contained in Q.

Let $A = F \cap Q$. We have just shown that A was non-empty. For any $x_1, x_2 \in Q$, if $]x_1, x_2[$ intersects A, then $]x_1, x_2[$ also intersects each A_i, $i \in I$. Since each A_i is a face, $]x_1, x_2[$ is contained in each A_i, hence also in A. So A is a face. Clearly, $A = \bigcap_{i \in I} A_i$, and is therefore the largest face contained in all A_i's. So A is the least upper bound of $(A_i)_{i \in I}$ in $\supseteq$. $\square$

Lemma 3.7 *If Q is compact, then every face of Q contains a minimal face (for set inclusion).*

Proof. In any cpo, any element is below some maximal element. This is an easy consequence of Zorn's Lemma. Now apply Lemma 3.6. □

Lemma 3.8 *Let K be a non-empty compact subset of C, and f any continuous map from C to $\overline{\mathbb{R}^+}$. Then the greatest lower bound $a = \inf_{z \in K} f(z)$ is attained, i.e., there is a element $z \in K$ such that $f(z) = a$.*

Proof. The image $f(K)$ of K under f is compact, so its upward closure $\uparrow f(K)$ is, too. Also, since K is non-empty, $\uparrow f(K)$ is non-empty. However, the non-empty saturated compacts in $\overline{\mathbb{R}^+}$ are the intervals of the form $[r, +\infty]$, $r \in \mathbb{R}^+$. Let r be such that $\uparrow f(K) = [r, +\infty]$. Then there is a $z \in K$ such that $f(z) = r$, and clearly $r = a$. So the greatest lower bound a is attained. □

Lemma 3.9 *If Q is compact, then for any face A of Q, and any linear continuous map $f : C \to \overline{\mathbb{R}^+}$, the greatest lower bound $\inf_{z \in A} f(z)$ is attained, write it $\min_A f$. Then the set $\text{argmin}_A f = \{x \in A \mid f(x) = \min_A f\}$ is a face of Q.*

Proof. Let $a = \inf_{z \in A} f(z)$. Since A is a face, in particular the intersection of a closed set and the compact set Q, A is compact. We conclude that a is attained by Lemma 3.8.

Since $a = \min_A f$ is attained, it also follows that $\text{argmin}_A f$ is non-empty.

Since A is a face, we may write A as $F \cap Q$, where F is closed. Let F' be the set $\{x \in X \mid f(x) \leq a\}$. This is the inverse image of the closed set $[0, a]$ by the continuous map f, and is therefore closed. Clearly, $\text{argmin}_A f = A \cap F' = (F \cap F') \cap Q$. Since $F \cap F'$ is closed in C, $\text{argmin}_A f$ is closed in Q.

For any $x_1, x_2 \in Q$ such that $]x_1, x_2[$ intersects $\text{argmin}_A f$, let $x = r \cdot x_1 + (1-r) \cdot x_2$ be in $\text{argmin}_A f$, $0 < r < 1$. Since $]x_1, x_2[$ intersects A at x, and A is a face, $]x_1, x_2[$ is contained in A. By Lemma 3.3, x_1 and x_2 are in A, so in particular $f(x_1) \geq a$, $f(x_2) \geq a$. Also, since f is linear, $f(x) = rf(x_1) + (1-r)f(x_2)$. However, $f(x) = a$ since x is in $\text{argmin}_A f$, so necessarily $f(x_1) = a$ and $f(x_2) = a$. This entails that for any element $y = s \cdot x_1 + (1-s) \cdot x_2$, $f(s) = sf(x_1) + (1-s)f(x_2) = sa + (1-s)a = a$, i.e., that $]x_1, x_2[$ is contained in $\text{argmin}_A f$. So $\text{argmin}_A f$ is a face of Q. □

Lemma 3.10 *Let C be a locally convex T_0 topological cone, Q a compact subset of C. Any minimal face of Q (for set inclusion) is of the form $\{x\}$, with x extreme in Q.*

Proof. Let A be a minimal face of Q, and assume it contains two distinct elements x and y. Since C is T_0, we may assume by symmetry that $x \not\leq y$. By the Separation Corollary (Keimel, 2006, Corollary 9.2) cited in Section 2, there is a continuous linear map $f : C \to \overline{\mathbb{R}^+}$ such that $f(y) < f(x)$. By Lemma 3.9, $\text{argmin}_A f$ is a face of Q. It is clearly contained in A, even strictly: since $f(y) < f(x)$, x cannot be in $\text{argmin}_A f$. So A is not minimal, contradiction.

It follows that we can write A as $\{x\}$; then x is extreme by Proposition 3.4. □

Finally, we get the following cone-theoretic version of the Krein-Milman Theorem:

Theorem 3.11 *Let C be a locally convex T_0 topological cone, Q a convex compact saturated subset of C. Then Q is the smallest convex compact saturated subset of C containing the extreme points of Q.*

Proof. If Q is empty, this is clear, so assume $Q \neq \emptyset$. Clearly, Q is convex, compact, saturated, and contain all the extreme points of Q.

The hard direction is the converse. Let Q' be any convex compact saturated subset of C containing the extreme points of Q, and assume by contradiction that there is a point x in Q that is not in Q'.

Note that the closed subset $\downarrow x$ does not intersect Q', otherwise there would be an $x' \in Q'$ with $x' \leq x$, entailing $x \in Q'$, since Q' is upward closed. Also, $\downarrow x$ is convex: this follows from the fact that $+$ and $\cdot$ are continuous, hence monotonic with respect to the underlying specialization orderings. We may therefore apply the Geometric Separation Theorem (Keimel, 2006, Proposition 10.2) stated in Section 2, with $K = Q'$, $F = \downarrow x$. So there is a convex open set U containing Q' and disjoint from $\downarrow x$, i.e., not containing x. We now use Keimel's Separation Theorem (Keimel, 2006, Theorem 9.1), as stated in Section 2, with $E = \{x\}$ and U just given: there is a continuous linear map $f : C \to \overline{\mathbb{R}^+}$ such that $f(x) \leq 1$, and $f(y) > 1$ for every $y \in U$. Let $a = \inf_{z \in Q'} f(z)$. By Lemma 3.8, a is attained, say $a = f(z_0)$; since $Q' \subseteq U$, $f(z_0) > 1$, so $a > 1$.

By Lemma 3.5, $A = Q$ is a face of Q. By Lemma 3.9, $\mathrm{argmin}_A f$ is then a smaller face of Q. Any element z of $\mathrm{argmin}_A f$ is such that $f(z) = \min_A f = \min_{y \in Q} f(y) \leq f(x) \leq 1$. Since $f(z) \geq a > 1$ for every $z \in Q'$, we conclude that $\mathrm{argmin}_A f$ does not intersect Q'. Let A be a minimal face of Q contained in $\mathrm{argmin}_A f$. This exists by Lemma 3.7. By Lemma 3.10, A is of the form $\{y\}$, with y extreme in Q. Since $y \in A \subseteq \mathrm{argmin}_A f$ and $\mathrm{argmin}_A f$ does not intersect Q', y is an extreme point not contained in Q': contradiction. □

We now compare this to the classical Krein-Milman Theorem. In the introduction, we stated it as: any compact convex subset K of a locally convex topological vector space X is the closed convex hull of its extreme points. The closed convex hull of K is, by definition, the smallest closed convex set containing K.

The parallel with Theorem 3.11 arises from de Groot duality. Given any topological space X, the *cocompact topology* on X is the smallest topology containing the complements of the compact saturated subsets Q of X. The resulting topological space X^{d} is the *de Groot dual* of X. Clearly, every compact saturated subset of X is closed in X^{d}. Then, calling co-closed any subset of X that is closed in X^{d}, Theorem 3.11 implies that Q is the *co-closed* convex hull of the set of extreme points of Q, i.e., the smallest co-closed convex set containing the extreme points of Q.

Acknowledgments

I had many discussions with Roberto Segala about the possible use of the classical Krein-Milman Theorem in establishing certain theorems of probability theory. Adapting it to the cone case was prompted by these discussions, and the need to get a version of it that would work in domain theory.

References

Goubault-Larrecq, J. (2007). Une introduction aux capacités, aux jeux et aux prévisions. `http://www.lsv.ens-cachan.fr/~goubault/ProNobis/pp_1_8.pdf`. 641 pages.

Keimel, K. (2006). Topological cones: Foundations for a domain-theoretical semantics combining probability and nondeterminism. *Electronic Notes in Theoretical Computer Science*, 155:423–443.

Wikipedia (2008). Krein-Milman theorem. `http://en.wikipedia.org/wiki/Krein-Milman_theorem`.