ESTIMATES ON VOLUMES OF HOMOGENEOUS POLYNOMIAL SPACES

ITAÏ BEN YAACOV

ABSTRACT. In this paper we develop the "local part" of our local/global approach to globally valued fields (GVFs). The "global part", which relies on these results, is developed in a subsequent paper.

We study *virtual divisors* on projective varieties defined over a valued field K, as well as *sub-valuations* on polynomial rings over K (analogous to homogeneous polynomial ideals). We prove a Nullstellensatz-style duality between projective varieties equipped with virtual divisors (analogous to projective varieties over a plain field) and certain sub-valuations on polynomial rings over K (analogous to homogeneous polynomial ideals). Our main result compares the *volume* of a virtual divisor on a variety W, namely its $(\dim W + 1)$-fold self-intersection, with the asymptotic behaviour of the volume of the dual sub-valuation, restricted to the space of polynomial functions of degree m, as $m \to \infty$.

This is work in progress.

CONTENTS

INTRODUCTION

1. DEFINABILITY AND STABILITY IN ALGEBRAICALLY CLOSED (VALUED) FIELDS

Let us recall a few facts from model theory, as it pertains to algebraically closed fields (this is all essentially folklore). By a *formula $\varphi(X)$*, in a tuple of indeterminates X, we mean a Boolean combination of polynomial equalities $f(X) = 0$: if a is a tuple in a field K of the appropriate length, then $\varphi(a)$ is either True or False. We often split the indeterminates in several families, with the notation $\varphi(X, Y)$. We may substitute elements of K for some of the variables, obtaining a *formula with parameters $\varphi(X, b)$*. When $A \subseteq K$, we write $\varphi(A, b)$ for $\{a \in A : \varphi(a, b)\}$ (here and later, $a \in A$ should be understood loosely as "a is a tuple in A of the appropriate length"). When $A = K$ and $b \in K$, the set $\varphi(K, b)$ is called a *definable set* in K, and is just a constructible set defined over K. When no ambiguity may arise, we sometimes omit the parameters b from the notation, saying that $\varphi(X)$ is a formula *over K*.

Algebraically closed fields have *quantifier elimination*: if $\varphi(X, Y)$ is a formula in X, Y, then $\exists Y \, \varphi(X, Y)$ is (equivalent to) a formula $\psi(X)$ (once this holds for formulae without parameters, it also holds for ones with parameters). Indeed, this is just Chevallay's Theorem: a coordinate projection of a constructible set in an algebraically closed field is again a constructible set. Since formulae are closed under negation, the same is true of $\forall Y \, \varphi(X, Y)$, and of any other expression constructed using quantifiers or Boolean operations. (In fact, what we defined as a formula is what is usually called a *quantifier-free* formula, but by quantifier elimination the two notions agree.) Notice that if D is a definable set, defined by $\psi(Y)$, and $\varphi(X, Y)$ is another formula, then $(\exists Y \in D) \, \varphi(X, Y)$ is again a formula (since it is the same as $\exists Y \, (\psi(Y) \wedge \varphi(X, Y))$), and similarly for $(\forall Y \in D) \, \varphi(X, Y)$. One important consequence of quantifier elimination is the following: if $\varphi(X)$ and $\psi(X)$ are two formulae over K, such that $\varphi(K) = \psi(K)$, then $\varphi(L) = \psi(L)$, for any larger field L. Indeed, we may assume that L is algebraically closed, and the formula the formula $\forall X \, \varphi(X) \leftrightarrow \psi(X)$ is true in K, so also in L.

The class of algebraically closed fields is *stable*. This has many equivalent characterisations, one of which is the following: if L/K is an extension of algebraically closed fields, and $\varphi(X, b)$ is a formula with parameters $b \in L$, then the set $\varphi(K, b)$ is definable in K (this is far from being true for any kind of structure: for example, $\mathbf{Q} \subseteq \mathbf{R}$ is an extension of dense linear orders without endpoints, which also have quantifier elimination, but the set $\{q \in \mathbf{Q} : q < \pi\}$ is not definable in $\mathbf{Q}$). An equivalent characterisation of stability is via the existence of a (necessarily unique) *notion of independence* of two structures L and M over a common substructure K, denoted $L \mathop{\smile}\limits_{K} M$, satisfying certain axioms that we do not state here (see for example Pillay [Pil96]). Thus, for example, properties of linear independence over a common subspace imply that the class of vector spaces over some fixed field is stable, and stochastic independence yields stability for probability algebras. In the class of algebraically closed fields, when L and M are subfields of some large ambient field, and K is a common algebraically closed sub-field, we say that $L \mathop{\smile}\limits_{K} M$ if L and M are linearly disjoint over K, i.e., if they generate the algebra $L \otimes_K M$ inside the ambient field (since K is algebraically closed, $L \otimes_K M$ is always an integral domain). It is a general fact that in a stable class of structures, the following are equivalent:

(i) We have $L \mathop{\smile}\limits_{K} M$.

(ii) If $\varphi(X, b)$ is a formula with parameters $b \in L$, then $\varphi(M, b)$ can be defined by a formula $\psi(X, c)$ with a parameter $c \in K$ (so $\psi(X, c)$ also defines, in K, the restriction $\varphi(K, b)$).

Let us see how this is proved for fields. It is enough to consider the special case of a polynomial equality $f(X, b) = 0$, where $f(X, Y) \in \mathbf{Z}[X, Y]$ and $b \in L$. We may assume that f is homogeneous in X, and applying the Veronese map of degree d, we may assume that f is linear in X. But then the set $f(K, b) = 0$ is a linear subspace of K^m, and it is easily definable, say, as the zero set of a family of linear forms with coefficients in K. If M is another extension of K, linearly disjoint of L over K, then the linear form $f(X, b)$ vanishes on $c \in M^m$ if and only if c can be expressed as an M-linear combination of tuples in $f(K, b) = 0$. It follows that the same intersection of zero sets of linear forms over K also defines $f(M, b) = 0$. Going back to our original setting where f is an arbitrary polynomial, if W is an algebraic set defined over K (or indeed, any set definable with parameters in K), and $L \mathop{\smile}\limits_{K} M$, then $f(X, b) = 0$ has a solution in $W(K)$ if and only if it has a solution in $W(M)$ (by quantifier elimination).

We are going to be interested in a similar situation, but in the context of algebraically closed fields equipped with a non-trivial valuation $v \colon K \to \mathbf{R} \cup \{\infty\}$ (equivalently, with a non-trivial, non-Archimedean absolute value). Therefore, from this point onward, a *valuation* will always be in the ordered group $(\mathbf{R}, +, <)$ (not necessarily onto). Fields equipped with such a valuation are studied, from a model-theoretic point of view, in [Ben14]. Following the notation there, we shall write $K \vDash MVF$ to say that K is a complete valued field, and $K \vDash ACMVF$ to say that K is, in addition, algebraically closed with a non-trivial valuation.

The presence of $\mathbf{R}$ as a fixed object requires us to replace classical Boolean logic with real-valued *continuous* logic (see [BU10, BBHU08]). Thus, the set of possible truth values for a formula $\varphi(X)$ is no longer $\{T, F\}$, but some compact interval of $\mathbf{R}$, and formulae are closed under continuous, rather than Boolean, combinations. For this closure property we allow infinite continuous combinations, or, equivalently, finite continuous combinations and uniform limits. We may also allow $[-\infty, \infty]$ as a truth value space, with the notion of uniform convergence corresponding to its compact topology (equivalently, arising from any homeomorphism with a compact interval of $\mathbf{R}$).

Convention 1.1. Breaking with usual terminology, by a *predicate* we mean a map from a set (possibly a Cartesian product of sets) to $[-\infty, \infty]$. A special case of this is the more familiar notion of a *Boolean predicate*, which is a map into $\{0, 1\}$ (or $\{True, False\}$, where *True* is identified with 0 and *False* with 1).

One last complication: the structure cannot be taken to be the field K itself, since it is unbounded, nor the valuation ring, since its ultra-powers are not necessary valuation rings. We take instead the disjoint union of $\mathbf{P}^n(K)$ for all n, which will be denoted $\mathbf{P}(K)$ (each $\mathbf{P}^n$ is a *sort* of $\mathbf{P}(K)$). Let us now define what (quantifier-free) formulae are in this structure.

Notation 1.2. For $a \in K^m$ we let

$$\tilde{v}(a) = \min_{i < m} v(a_i).$$

We shall use this mainly in two contexts, for affine points, and for polynomials, identified with their tuple of coefficients.

Definition 1.3. Let $f \in \mathbf{Z}[X_0, X_1, \ldots]$, where each of X_i is a tuple of $n_i + 1$ indeterminates, such that f is homogeneous of degree d_i in each X_i. Then, for $\xi_i = [x_i] \in \mathbf{P}^{n_i}(K)$ we define

$$vf(\xi_0, \xi_1, \ldots) = v \circ f(x_0, x_1, \ldots) - \sum d_i \tilde{v}(x_i) \in [0, \infty],$$

noting that it depends only on the points $\xi_0, \xi_1, \ldots$ (and not on the representatives).

Each vf is considered an *atomic formula*, and any continuous combinations thereof is a *formula*. Thus, each formula defines a predicate on $\mathbf{P}(K)$ (i.e., on a Cartesian product of sorts of $\mathbf{P}(K)$). Notice that a polynomial $f \in K[X_0, \ldots]$ can be written as $g(b, X_0, \ldots)$, where $g \in \mathbf{Z}[Y, X_0, \ldots]$ is linear in Y and b are the coefficients. In this case

$$vf(\xi_0, \ldots) = vg([b], \xi_0, \ldots) + \tilde{v}(b),$$

so allowing coefficients in K does not change the expressive power.

Fact 1.4 ([Ben14], Theorem 2.4 and Proposition 2.16). *Algebraically closed metric valued fields admit quantifier elimination. In other words, if $\varphi(\xi, \zeta)$ is a formula (where each of ξ and ζ represents a tuple of indeterminate projective points), then $\inf_\zeta \varphi(\xi, \zeta)$ is (equivalent to) a formula in every $K \vDash ACMVF$. Consequently, $\sup_\zeta \varphi(\xi, \zeta)$, as well as any more complex expression constructed using these* quantifiers *and continuous combinations, is a formula.*

Moreover, every algebraic $W \subseteq \mathbf{P}^n$ is a definable set *(in the sense of continuous logic), that is to say that $\inf_{\zeta \in W} \varphi(\xi, \zeta)$ is also a formula over K.*

Notice that [Ben14] uses multiplicative notation, whereas here we use additive notation, but since $e^{-t} \colon [0, \infty] \to [0, 1]$ is a monotone homeomorphism, this changes nothing. Given an incomplete valued field K, it has a unique completion $\widehat{K}$, and the quantifiers inf and sup evaluate the same in K and in $\widehat{K}$. We may therefore assume that all the fields in question are complete (and if an incomplete one arises, replace it tacitly with its completion).

Let us show that the class $ACMVF$ is stable using the characterisation given above (a different argument, by counting types, is given in [Ben14]). In other words, if L/K is an extension of such fields, and $\varphi(X, b)$ is a formula with parameters in L, then its restriction to K is definable, in K, by a formula with parameters there. The fundamental tool for doing this is the notion of a valued vector space.

Definition 1.5. Let $K \vDash MVF$, E a vector space over K. A *valuation* on E is a function $u \colon E \to \mathbf{R} \cup \{\infty\}$ satisfying, for all $x, y \in E$ and $a \in K$:
  (i) $u(ax) = v(a) + u(x)$
  (ii) $u(x + y) \geq u(x) \wedge u(y)$
Its *kernel* is the subspace

$$\ker u = \{x \in E : u(x) = \infty\}.$$

If $\ker u = \{0\}$, then we say that u is *reduced*.

Remark 1.6. One might think that what we call a valuation should be called a *semi-valuation*, reserving the term *valuation* to reduced ones. There are, however, several good reasons for our choice terminology, such as the analogy between ideals and (sub-)valuations on rings, which is explored below.

Definition 1.7. Let E be a valued vector space and $F \subseteq E$ a subspace. We define a *quotient valuation* on either E or E/F (both points of view may be useful) by:

$$(u/F)(x) = u(x + F) = \sup \{u(y) : y \in x + F\}.$$

It is easy to check that the quotient valuation is indeed a valuation. Clearly, $\tilde{v}$ is a reduced valuation on K^m. Conversely:

Lemma 1.8. *Any reduced valuation u on K^m is arbitrarily close, up to a change of coordinates by a triangular matrix, to $\tilde{v}$.*

If u is a reduced valuation on E and $F \subseteq E$ is finite-dimensional, then the quotient valuation on E/F is reduced as well.

Proof. We prove the first item by induction on m, with $m = 0$ being trivial. For $m + 1$, let $(e_i : i \leq m)$ denote a basis of E, and $F = \mathrm{Span}(e_i : i < m)$. Let $\varepsilon > 0$. Applying a change of coordinates to F by a triangular matrix, we may assume that $|u(y) - \tilde{v}(y)| < \varepsilon$ for all $y \in F$. Let $\alpha = u(e_m + F)$. If $\alpha = \infty$, then there exist sequences $(a_{i,n})$ for $i < m$ such that $u(e_m + \sum_{i<m} a_{i,n} e_i) \to \infty$ as $n \to \infty$. By our assumption regarding u on F each of the sequences $(a_{i,n})$ must be Cauchy, and therefore converge to some a_i. But then $u(e_m + \sum_{i<m} a_i e_i) = \infty$, a contradiction.

Therefore $\alpha < \infty$, and we may choose $x \in e_m + F$ such that $\alpha < u(x) + \varepsilon$. Taking $b \in K$ such that $\alpha - \varepsilon < -v(b) < u(x)$, we have $0 < u(bx) \leq \alpha + v(b) < \varepsilon$. Replacing e_m with bx, we complete a triangular change of basis for K^{m+1}, and now $0 < u(e_m) < \alpha < \varepsilon$. Let $y \in E$. If $y \in F$, then we already have $|u(y) - \tilde{v}(y)| < \varepsilon$.

Otherwise, we may assume that $y \in e_m + F$. If $u(y - e_m) \leq 0$, then $u(y) = u(y - e_m)$, and in any case $u(y) \leq \varepsilon$. Therefore

$$u(y) \leq 0 \wedge u(y - e_m) + \varepsilon \leq \tilde{v}(e_m) \wedge \big(\tilde{v}(y - e_m) + \varepsilon\big) + \varepsilon \leq \tilde{v}(y) + 2\varepsilon.$$

Since clearly $u \geq \tilde{v}$, we reach the desired conclusion.

The second assertion follows from our argument that $\alpha < \infty$. ∎

Definition 1.9. Let L/K be an extension.

(i) Say that a tuple $a \in L^m$ is *weakly v-generic* over K, if for every polynomial $h \in K[X]_d$:

$$v\big(h(a)\big) \leq \tilde{v}(h).$$

Notice that this implies that $\tilde{v}(a) \leq 0$.

(ii) Say that a is *v-generic* over K if for every polynomial $h \in K[X]_d$:

$$v\big(h(a)\big) = \tilde{v}(h).$$

Equivalently, if it is weakly v-generic and $\tilde{v}(a) = 0$.

Lemma 1.10. *Let K be a valued field.*

(i) *One can always adjoin to K new v-generic elements.*
(ii) *Assume that $a \in L^m$ is v-generic over K, and $b \in K^m$. Then $a + b \in L^m$ is weakly v-generic over K.*

Proof. It is a standard (and easy) fact that $\tilde{v}$ is multiplicative on $K[X]$, namely, that $\tilde{v}(fg) = \tilde{v}(f) + \tilde{v}(g)$ for all $f, g \in K[X]$. This gives rise to a valuation $w(f/g) = \tilde{v}(f) - \tilde{v}(g)$ on the fraction field $L = K(X)$, for which X are v-generic over K. The rest is immediate. ∎

Proposition 1.11. *Let us fix a degree $d \geq 1$, and let $W \subseteq \mathbf{P}^n$ be an algebraic set defined over K. Then the following are equivalent for a function $\eta \colon W(K) \to \mathbf{R} \cup \{\infty\}$:*

(i) *The function η is a uniform limit on $W(K)$ of functions of the form $\min_{i<N} vg_i\!\restriction_{W(K)}$, where $g_i \in K[X]_d$.*
(ii) *There exists an extension L/K and polynomial $f \in L[X]_d$ such that $\eta = vf\!\restriction_{W(K)}$.*
(iii) *[When $d = 1$] There exists a valuation u on K^{n+1} such that $\eta(\xi) = u(x) - \tilde{v}(x)$ for all $\xi = [x] \in W(K)$.*

Moreover, if η is finite then it is bounded in $\mathbf{R}$, and in the second item one may always take $N = \binom{n+d}{d}$.

Proof. Applying a Veronese map, and replacing n with $\binom{n+d}{d} - 1$, we may assume that $d = 1$. We may further assume that $W(K)$ does not lie in any hyperplane of $\mathbf{P}^n$.

(i) $\Longrightarrow$ (ii). Let us first assume that $\eta = \min_{i<N} vg_i\!\restriction_{W(K)}$ for some $g_i \in K[X]_1$. Let L/K be an extension containing a tuple a of N v-generic elements, and let $f = \sum a_i g_i \in L[X]_1$. Then $\eta = vf\!\restriction_{W(K)}$. Also, a direct calculation shows that $\tilde{v}(f) = \min \tilde{v}(g_i) \leq \inf \eta$.

Let us show that the difference $\inf \eta - \tilde{v}(f)$ is bounded by a constant M which depends only on W, and not on η or the g_i. Indeed, assume not. Then there exists a sequence of $(g_{m,i} : m \in \mathbf{N}, i < N_m) \subseteq K[X]_1$ such that $\min_i \tilde{v}(g_{m,i}) = 0$ for all m, and yet $vg_{m,i}\!\restriction_{W(K)} \geq m$ for all m, i. In particular, there exists a sequence $(g_m) \subseteq K[X]_1$ such that $\tilde{v}(g_m) = 0$ and $\inf_{\xi \in W} vg_i(\xi) \geq m$ holds in K. Let $M = K^{\mathcal{U}} \supseteq K$ be an ultra-power, and let $g \in M[X]_1$ be the limit of the sequence (g_m). Then $\tilde{v}(g) = 0$ and $\inf_{\xi \in W} vg(\xi) \geq m$ for all m, i.e., g vanishes on W. Forgetting the valuation, M/K is an extension of algebraically closed fields. By quantifier elimination for those, there exists $h \in K[X]_1$ which vanishes on W, contradicting our hypothesis.

We can now prove the converse in the general case. Indeed, we assume that η is a uniform limit on $W(K)$ of functions $\eta_k = \min_{i<N_k} vg_{k,i}\!\restriction_{W(K)}$. Say, in particular, that $|\eta - \eta_k| \leq 1$ for all k. Express each η_k as $vf_k\!\restriction_{W(K)}$ where $f_k \in L_k[X]_1$ and $\tilde{v}(f_k) \geq \inf \eta - M - 1$. Let $L = \prod L_k/\mathcal{U}$ be an ultra-product, and $f \in L[X]_d$ the image of the sequence (f_k) (for this to exist we require the common lower bound for $\tilde{v}(f_k)$). It follows that $\eta = vf\!\restriction_{W(K)}$.

(ii) $\Longrightarrow$ (iii). If f is linear, then $u(x) = v \circ f(x)$ is a valuation on K^{n+1}.

(iii) $\Longrightarrow$ (i). Dividing by $\ker u$ we may assume that u is reduced. Then, up to a change of coordinates, it is arbitrarily close to $\tilde{v}$. In other words, $u(x)$ is arbitrarily close to $\min v \circ g_i(x)$ for a family $(g_i : i < n) \subseteq K[X]_1$. It follows that $\min_{i<n} vg_i\!\restriction_{W(K)}$ is as close as desired to η.

Finally, let us prove that if η is finite, then it is bounded. For this we may assume that $\eta = \min_{i<N} vg_i\!\restriction_{W(K)}$ as above. Indeed, if $\sup \eta = \infty$, then there exist $\xi_m \in W(K)$ such that $vg_i(\xi_m) \geq m$ for all i. In an ultra-power M we find $\xi \in W(M)$ on which all the g_i vanish. Forgetting the valuation as above, we find $\xi \in W(K)$ on which all the g_i vanish, so $\eta(\xi) = \infty$, a contradiction. ∎

In particular, for every homogeneous polynomial $f(X) \in L[X]$, the restriction of vf to $\mathbf{P}^n(K)$ is uniformly approximated by formulae with parameters in K, i.e., $vf\restriction_{\mathbf{P}^n(K)}$ is definable in K. The Segre embedding allows us to replace a homogeneous polynomial in several families of indeterminates with one in a single family, and we conclude that the restriction to K of any atomic formula over L is definable in K, and the same follows for every formula. It follows that the theory $ACMVF$ is stable.

Consider now $L, M \vDash MVF$, both embedded in a large valued field, along with a common sub-field $K \vDash ACMVF$. We want to characterise when $L \underset{K}{\downarrow} M$, i.e., when the restriction of a formula over L to M is definable over K. As above, it suffices to consider the case of a formula vf where $f \in L[X]$ is homogeneous in a single family of indeterminates. For this, we require one last tool.

Definition 1.12. If E and F are two valued vector spaces then the *tensor product valuation* on $E \otimes_K F$ as the least one satisfying $u(x \otimes y) \geq u_E(x) + u_F(y)$ for every simple tensor $x \otimes y$. It is sometimes denoted $u_E \otimes u_F$.

Lemma 1.13. *With the hypotheses of Definition 1.12:*

  (i) *The tensor product valuation exists.*
  (ii) *We have $(u_E \otimes u_F)(x \otimes y) = u(x) + u(y)$ for all simple tensors, and more generally,*

$$(u_E \otimes u_F)(z) = \sup\left\{\min\left(u_E(x_i) + u_F(y_i)\right) : z = \sum x_i \otimes y_i\right\},$$

  where the supremum is taken over all presentations of z as a sum of simple tensors.
  (iii) *When $F = K^m$ is equipped with $\tilde{v}$ then $E \otimes_K K^m = E^m$ and $u_E \otimes \tilde{v} = \tilde{u}_E$, i.e.:*

$$(u_E \otimes \tilde{v})(x_0, \ldots, x_{m-1}) = \min_i u_E(x_i).$$

  (This is the direct sum valuation on $E \oplus \cdots \oplus E$.)
  (iv) *If $E_0 \subseteq E$ and $F_0 \subseteq F$ are sub-spaces, then*

$$(u_E\restriction_{E_0}) \otimes (u_F\restriction_{F_0}) = (u_E \otimes u_F)\restriction_{E_0 \otimes_K F_0}.$$

  (v) *We have $\ker(u_E \otimes u_F) = (\ker u_E) \otimes_K (\ker u_F)$.*

Proof. Let us consider first the special case where $E = K^m$ and $F = K^n$ are equipped with $\tilde{v}$, and $E_0 = K^{m_0} \times \{0\}$, $F_0 = K^{n_0} \times \{0\}$. Identifying $E \otimes_K F$ with K^{mn} we have $\tilde{v}(x \otimes y) = \tilde{v}(x) + \tilde{v}(y)$, and $\tilde{v}$ on K^{mn} is least satisfying $\tilde{v}(x_i \otimes y_j) \geq 0$ for x_i and y_j in the respective standard bases, so $\tilde{v} = \tilde{v} \otimes \tilde{v}$ (on the appropriate spaces). In this case, all our assertions are easy to check. The case of any two reduced valuations follows by Lemma 1.8 (since the change of basis is by a triangular matrix, we reduce to the special forms of E_0 and F_0 assumed earlier). In the general case, we divide by the kernels. ∎

If L and M are field extensions of K, then $L \otimes_K M$ is a ring, and it is immediate to check that $u(cd) \geq u(c) + u(d)$ for all $c, d \in L \otimes_K M$. The main theorem of [Ben15] asserts that, when $K \vDash ACMVF$, we have $u(cd) = u(c) + u(d)$, giving rise to a natural valuation $v(c/d) = u(c) - u(d)$ on $\mathrm{Frac}(L \otimes_K M)$. (It is proved using quantifier elimination for $ACVF$, in Boolean logic, but can also be proved using quantifier elimination in $ACMVF$, as well as via other methods, as in Poineau [Poi13].)

Proposition 1.14. *Let $K, L, M \vDash ACMVF$, where both L and M are embedded in some large valued field and K is a common sub-field. Then the following are equivalent:*

  (i) *The compositum LM is $\mathrm{Frac}(L \otimes_K M)$ (as a valued field).*
  (ii) *For every homogeneous polynomial $f \in M[X]$, the restriction of vf to L is definable with parameters in K.*

Proof. In one direction, we may assume, as in the proof of Proposition 1.11, that $f(X)$ is linear. We let $u_1(x) = v \circ f(x)$ for $x = K^N$, and similarly $u_2(x)$ for $x \in L^N$.

Let $x \in L^N$ be given. Let $E \subseteq L$ be the K-vector space generated by the coefficients of x, and let $(y_i : i < \ell)$ be a basis for E. Assume first that (y_i) induces an isomorphism between E and $(K^\ell, \tilde{v})$, and express $x = \sum y_i x_i$, where $x_i \in K^N$. Then $f(x) = \sum y_i f(x_i) \in L \otimes_K M$, and by Lemma 1.13(iii) we have

$$(v_L \otimes u_1)(x) = \min u_1(x_i) = \min v_M(f(x_i)) = (v_L \otimes v_M)(f(x)) = u_2(x).$$

We reduce the general case to this special one (approximately) by Lemma 1.8, so $u_2 = v_L \otimes u_1$.

In particular, we may now reduce to the case where u_1 is a valuation. We then have a family of linear functions g_i over K such that $u_1(x)$ is arbitrarily close to $\min v \circ g_i(x) = \tilde{v}_K(g_i(x) : i < N)$ for $x \in K^N$. Then, for $x \in L^N$, $u_2(x)$ is as close to $\tilde{v}_L(g_i(x) : i < N)$, i.e., to $\min v \circ g_i(x)$ again. Thus vf restricted to L is definable with parameters in K.

For the converse, let us compare the compositum LM (in the ambient valued field) with $\mathrm{Frac}(L \otimes_K M)$. Consider a homogeneous polynomial $f \in M[X]$. The restriction of vf to $\mathbf{P}^n(K)$ is definable, in K (with parameters there), say by $\varphi(\xi)$. In both LM and $\mathrm{Frac}(L \otimes_K M)$, the restriction of vf to $\mathbf{P}^n(L)$ is definable by formulae with parameters in K, say $\psi_1(\xi)$ and $\psi_2(\xi)$. But then $\sup_\xi |\varphi(\xi) - \psi_i(\xi)|$ is a formula with parameters in K, evaluating to zero. Therefore it must also evaluate to zero in L (we are making a non-trivial use of quantifier elimination here), so $\psi_1(\xi) = \psi_2(\xi)$ for all $\xi \in \mathbf{P}^n(L)$. In other words, for $\xi \in \mathbf{P}^n(L)$, $vf(\xi)$ is the same in LM and in $\mathrm{Frac}(L \otimes_K M)$. It follows that the valuation on LM is the one induced from $\mathrm{Frac}(L \otimes_K M)$, as stated. ∎

As in the case of pure fields, we can reduce from arbitrary formulae to ones of the form $vf(\xi)$. We concede that $LM \cong \mathrm{Frac}(L \otimes_K M)$ as valued fields if and only if the restriction of every formula over L to M is definable over K, i.e., if and only if $L \underset{K}{\bigcup} M$ in the sense of model-theoretic stability. Let us conclude with a technical result which will be used later on.

Lemma 1.15. *Assume that $L \underset{K}{\bigcup} M$, where $K, L, M \vDash ACMVF$, and let $a \in L^m$ be (weakly) v-generic over K. Then it is also (weakly) v-generic over M.*

Proof. This can be calculated directly, but let us give an argument using stability. So fix a degree d and let $b \in L^N$ consist of all monomials of degree d in a. Let $f(X) = \sum b_i X_i \in L[X]_1$, and let $\varphi(\xi)$ be the formula, with parameters in K, defining the restriction of vf to $\mathbf{P}(K)$. Since $L \underset{K}{\bigcup} M$, it also defines the restriction of vf to $\mathbf{P}(M)$. By hypothesis we have $\sup_\xi \varphi(\xi) \leq 0$ in K, so also in M. Doing this for all d, we see that a is weakly v-generic over M. It is then v-generic over either field if and only if, in addition, $\tilde{v}(a) = 0$. ∎

2. VIRTUAL DIVISORS AND VIRTUAL CHAINS

Throughout, $K \vDash ACMVF$.

Notation 2.1. If $f \in K[X]_m$, with $m \geq 1$, then for $\xi = [x] \in \mathbf{P}^n(K)$ we write

$$\hat{f}(\xi) = \frac{vf(\xi)}{m} = \frac{v \circ f(x)}{m} - \tilde{v}(x).$$

Definition 2.2. Let $W(K) \subseteq \mathbf{P}^n(K)$ be an algebraic set defined over K. A *virtual divisor* of degree d on $W(K)$ is a function $\eta\colon W(K) \to \mathbf{R}$ which can be expressed as $\hat{f}\restriction_{W(K)}$ for some extension L/K and $f \in L[X]_d$. A *virtual divisor* on $W(K)$ is a uniform limit of virtual divisors (of various degrees), i.e., in the distance:

$$d(\eta, \theta) = \sup_{\xi \in W(K)} |\eta(\xi) - \theta(\xi)|.$$

When η is a virtual divisor of degree d, we define its *d-width* as

$$w_d(\eta) = \sup \eta - \sup\{\tilde{v}(f)/d : f \in L[X]_d, \eta = \hat{f}\restriction_{W(K)}\},$$

as L varies over all possible extensions of K.

By Proposition 1.11, every virtual divisor is a bounded function, definable in the sense of Section 1. The space of virtual divisors on $W(K)$ of a given degree is complete for uniform convergence. Also, by an easy ultra-product argument, the second supremum in the definition of $w_d(\eta)$ is attained as a maximum.

In some sense, all virtual divisors are normalised to have "degree one" (one can imagine a more general definition, bur for our purposes only normalised virtual divisors are needed). The degree of a virtual divisor, as per Definition 2.2, measures its *complexity*. We have $\hat{f} = \widehat{f^m}$ so a virtual divisor of degree d, is also of degree md for every $m \geq 1$.

If η is a virtual divisor on $W(K)$ and L/K is an extension, then η extends naturally to a virtual divisor on $W(L)$, by applying the same definition. Alternatively, if $\eta = \hat{f}\restriction_{W(K)}$ for $f \in M[X]_d$, then we may identify f with a polynomial over $\mathrm{Frac}(L \otimes_K M)$, thus obtaining an extension of η to $\hat{f}\restriction_{W(L)}$. By Proposition 1.14, we obtain the same extension either way. Therefore, we may speak of a *virtual divisor on W* which is *defined over K*, which is the terminology we shall use from here on. By quantifier elimination, if η and θ are virtual divisors on W, say defined over K, then $d(\eta, \theta)$ is the same when calculated in $W(K)$ or in $W(L)$ for any extension L/K.

Lemma 2.3. *Let $W \subseteq \mathbf{P}^n$ be algebraic, defined over K. Then every constant function is a virtual divisor on W, of degree one. The family of virtual divisors on W is closed under translation (by a real), minimum (of finitely many), uniform limit, and finite convex combinations. Moreover, all but convex combinations preserve the degree.*

Proof. Closure under translation, minimum and uniform limit is clear from the definition, and the constant 0 is $\min_{i<n} \widehat{X}_i{\restriction}_{W(K)}$. We can calculate averages as

$$\frac{\hat{f}{\restriction}_{W(K)} + \hat{g}{\restriction}_{W(K)}}{2} = \hat{h}{\restriction}_{W(K)}, \qquad h = f^{\deg g} g^{\deg f}.$$

Closure under convex combinations follows by closure under uniform limits and boundedness. ∎

Lemma 2.4. *Let L/K be an extension, and let $\eta = \hat{f}{\restriction}_{W(K)}$, where $f \in L[X]_d$, be a virtual divisor. Then for every $\xi \in W(L)$ we have*

$$\eta(\xi) \le \hat{f}(\xi).$$

Proof. Let M be a copy of L, let $g \in M[X]_d$ be the corresponding copy of f, and let us embed L and M in $\mathrm{Frac}(L \otimes_K M)$. There is a natural morphism $\varphi \colon L \otimes_K M \to L$, and by definition of the tensor product valuation, it satisfies $v_L(\varphi(z)) \ge (v_L \otimes v_M)(z)$. For $\xi = [x] \in W(L)$ we have $\eta(\xi) = \hat{g}(\xi)$ and $\varphi(g(x)) = f(x)$. Our assertion follows. ∎

Definition 2.5. Let L/K be a valued field extension, and let $\mathfrak{C}$ and $\mathfrak{C}'$ be two Chow forms over L, both in same dimension ℓ. We define the *K-distance* between $\mathfrak{C}$ and $\mathfrak{C}'$ as

$$d_K(\mathfrak{C}, \mathfrak{C}') = \sup_F \left| \frac{v(F \wedge \mathfrak{C})}{\deg F \deg \mathfrak{C}} - \frac{v(F \wedge \mathfrak{C}')}{\deg F \deg \mathfrak{C}'} \right|,$$

as F varies over all families of $(\ell+1)$ non-constant homogeneous polynomials over K, agreeing that $|\infty - \infty| = 0$.

From this point onward we are going to consider chains in projective space as coded by Chow forms. We shall be using the wedge notation for the algebraic intersection of Chow forms with hypersurfaces, as in [Ben].

Lemma 2.6. *Let L/K and M/K be two valued field extensions, and let $W \subseteq \mathbf{P}^n$ be defined over K Let f and g be two homogeneous non-constant polynomials over L, and let $\mathfrak{C}$ and $\mathfrak{C}'$ be two Chow forms over L, both in same dimension ℓ, associated to subsets of W. Then, working in the free amalgam $\mathrm{Frac}(L \otimes_K M)$, we have*

$$d_K(f \wedge \mathfrak{C}, g \wedge \mathfrak{C}') \le d(\hat{f}{\restriction}_{W(K)}, \hat{g}{\restriction}_{W(K)}) + d_K(\mathfrak{C}, \mathfrak{C}').$$

Proof. It will suffice to prove that

$$d_K(f \wedge \mathfrak{C}, g \wedge \mathfrak{C}) \le d(\hat{f}{\restriction}_{W(K)}, \hat{g}{\restriction}_{W(K)}), \qquad d_K(f \wedge \mathfrak{C}, f \wedge \mathfrak{C}') \le d_K(\mathfrak{C}, \mathfrak{C}').$$

We have already observed that, by Proposition 1.14:

$$d(\hat{f}{\restriction}_{W(K)}, \hat{g}{\restriction}_{W(K)}) = d(\hat{f}{\restriction}_{W(M)}, \hat{g}{\restriction}_{W(M)}).$$

The resultant form (in some degrees) $F \mapsto F \wedge \mathfrak{C}$ is just a homogeneous polynomial in the coefficients of $\ell+1$ indeterminate polynomials (in those degrees), so by the same reasoning:

$$d_K(\mathfrak{C}, \mathfrak{C}') = d_L(\mathfrak{C}, \mathfrak{C}').$$

If H is any family of ℓ polynomials over K, then $H \wedge \mathfrak{C}$ splits as $\prod_{i<D} x_i$, where $D = \deg H \deg \mathfrak{C}$ and $[x_i] \in W(M)$. Therefore

$$
\begin{aligned}
d(\hat{f}{\restriction}_{W(K)}, \hat{g}{\restriction}_{W(K)}) &= d(\hat{f}{\restriction}_{W(M)}, \hat{g}{\restriction}_{W(M)}) \\
&\ge \frac{1}{D} \left| \frac{\sum_{i<D} v \circ f(x_i)}{\deg f} - \frac{\sum_{i<D} v \circ g(x_i)}{\deg g} \right| \\
&= \left| \frac{v(H \wedge f \wedge \mathfrak{C})}{\deg H \deg f \deg \mathfrak{C}} - \frac{v(H \wedge g \wedge \mathfrak{C})}{\deg H \deg g \deg \mathfrak{C}} \right|.
\end{aligned}
$$

Similarly,

$$d_K(\mathfrak{C}, \mathfrak{C}') = d_L(\mathfrak{C}, \mathfrak{C}') \ge \left| \frac{v(H \wedge f \wedge \mathfrak{C})}{\deg H \deg f \deg \mathfrak{C}} - \frac{v(H \wedge f \wedge \mathfrak{C}')}{\deg H \deg f \deg \mathfrak{C}'} \right|.$$

The two inequalities, and our assertion, follow. ∎

Definition 2.7. We define a *virtual chain* of dimension ℓ inside W, over a given field K, by taking all Chow forms of dimension ℓ associated to subsets of W, defined over extensions of K, dividing by the kernel of d_K, and completing. The image of a Chow form $\mathfrak{C}$ will be denoted $\mathscr{C} = \widehat{\mathfrak{C}}\!\restriction_K$ (or just $\widehat{\mathfrak{C}}$, if $\mathfrak{C}$ is already over K). Assume that $\mathscr{C}$ is a virtual chain of dimension ℓ inside W, and η is a virtual divisor on W, both given as uniform limits

$$\mathscr{C} = \lim \widehat{\mathfrak{C}}_k\!\restriction_K, \qquad \eta = \lim \hat{f}_k\!\restriction_{W(K)},$$

where $\mathfrak{C}_k$ are over $L \supseteq K$ and f_k are over $M \supseteq K$. Then we define

$$\eta \wedge \mathscr{C} = \lim \widehat{f_k \wedge \mathfrak{C}_k}\!\restriction_K,$$

where $f_k \wedge \mathfrak{C}_k$ is calculated in $\mathrm{Frac}(L \otimes_K M)$. By Lemma 2.6, this is a uniform limit, resulting in a virtual Chow form of dimension $\ell - 1$ inside W, which only depends on η and $\mathscr{C}$.

When iterating this with the same η, we may also write

$$\eta^{\wedge k} \wedge \mathscr{C} = \underbrace{\eta \wedge \ldots \wedge \eta}_{k \text{ times}} \wedge \mathscr{C}.$$

When $\mathscr{C} = \widehat{\mathfrak{C}}_{\mathbf{P}^n}$, for the canonically normalised Chow form of $\mathbf{P}^n$, we omit it, just writing $\eta \wedge \ldots \wedge \theta$ or $\eta^{\wedge k}$. Similarly, the virtual chain $\widehat{g \wedge \mathfrak{C}_{\mathbf{P}^n}}$ will be simply denoted $\hat{g}$, when there is no risk of ambiguity.

By Lemma 2.6,

$$d(\eta \wedge \mathscr{C}, \theta \wedge \mathscr{D})| \leq d(\eta, \theta) + d(\mathscr{C}, \mathscr{D}), \qquad \eta \wedge \theta \wedge \mathscr{C} = \theta \wedge \eta \wedge \mathscr{C}.$$

3. SUB-VALUATIONS

Let us begin with a few general definitions. Throughout, a *ring* is commutative and unital.

Definition 3.1. A *sub-valuation* on a ring A is a function $u\colon A \to \mathbf{R} \cup \{\infty\}$ which satisfies the following properties:

$$u(ab) \geq u(a) + u(b) \qquad \text{(sub-multiplicative)}$$
$$u(a^2) = 2u(a) \qquad \text{(power-multiplicative)}$$
$$u(a + b) \geq u(a) \wedge u(b) \qquad \text{(ultra-metric)}$$
$$u(0) = \infty.$$

(i) We say that u is a *proper* sub-valuation if $u(1) = 0$.
(ii) We say that u is a *valuation* if it is proper and multiplicative: $u(ab) = u(a) + u(b)$.
(iii) We define $\ker u = \{a \in A : u(a) = \infty\}$, and say that u is *reduced* if $\ker u = \{0\}$.
(iv) If $A = \bigoplus A_m$ is a graded ring, and $u(\sum a_m) = \min u(a_m)$ whenever $\sum a_m$ is a decomposition into homogeneous components, then we say that u is a *homogeneous sub-valuation*.

Notice that $u \equiv \infty$ is the unique improper sub-valuation on A. It follows easily from the axioms that $u(a^n) = nu(a)$ for all $n \in \mathbf{N}$ (where $0^0 = 1$ and $0 \cdot \infty = 0$).

We consider general (sub-)valuations as generalisations, in the "valued category", of radical (prime) ideals on A. In particular, a $\{0, \infty\}$-valued (sub-)valuation on A carries exactly the same information as a radical (prime) ideal, namely its kernel.

Fact 3.2 (see Bergman [Ber71]). *For every proper sub-valuation u on a ring A and $a \in A$ there exists a valuation $v \geq u$ on A such that $v(a) = u(a)$. In other words,*

$$u = \inf \{valuation\ v : v \geq u\}.$$

Definition 3.3. Let A be a ring. A pair (a, α), where $a \in A$ and $\alpha \in \mathbf{R} \cup \{\infty\}$, will be called a *condition* (over A), which we may also denote informally as "$u(a) \geq \alpha$". When A is graded and $a \in A$ is homogeneous, we shall say that (a, α) is a *homogeneous condition*.

Let $C \subseteq A \times (\mathbf{R} \cup \{\infty\})$ be a set of conditions such that $\pi(C) = \{a \in A : \exists \alpha\ (a, \alpha) \in C\}$ generates A. We define the sub-valuation *generated* by C, denoted $\langle C \rangle$, to be the least sub-valuation u on A satisfying $u(a) \geq \alpha$ for every condition $(a, \alpha) \in C$.

For the purposes of this definition we identify a subset $C \subseteq A$ (which generates A) with the set of conditions $\{(a, 0) : a \in C\}$, so $\langle C \rangle$ is least such that $u\!\restriction_C \geq 0$.

Lemma 3.4. *With the hypotheses of Definition 3.3, the generated sub-valuation $u = \langle C \rangle$ exists, and can be recovered by (where $\min \varnothing = \infty$):*

$$u(a) = \sup \left\{ \frac{1}{n} \min_i \sum_j \alpha_{ij} : a^n = \sum_i \prod_j b_{ij}, \text{ where } n \geq 1 \text{ and } (b_{ij}, \alpha_{ij}) \in C \right\}.$$

It is proper if and only if, for every $(b_{ij}, \alpha_{ij}) \in C$:

$$\sum_i \prod_j b_{ij} = 1 \qquad \Longrightarrow \qquad \min_i \sum_j \alpha_{ij} \leq 0.$$

Proof. Let us show that u satisfies the ultra-metric inequality. So let $a, b \in A$, and let $\rho < u(a) \wedge u(b)$ be arbitrary. Then there exist n and (c_{ij}, α_{ij}) in C such that $a^n = \sum_i \prod_j c_{ij}$ and $n\rho < \min_i \sum_j \alpha_{ij}$, and similarly (d_{ij}, β_{ij}) for b.

For each $0 \leq k < n$, we may express $a^k b^{n-k} = \sum_i \prod_j e_{kij}$, where $(e_{kij}, \gamma_{kij}) \in C$. Let $\delta = \min_k \min_i \sum_j \gamma_{kij}$, and all we care about is that $\delta > -\infty$. For $\ell \in \mathbf{N}$ we have

$$(a+b)^{\ell n + n - 1} = \sum_{0 \leq m \leq \ell} \left(a^{nm} b^{n(\ell - m)} \sum_{k < n} N_{m,k} a^k b^{n-k} \right),$$

where $N_{m,k}$ is some binomial coefficient. In each each term, express a^n using the c_{ij}, express b^n using the d_{ij}, and the term $a^k b^{n-k}$ using the e_{kij}. We obtain that

$$u(a+b) \geq \frac{\ell n \rho + \delta}{\ell n + n + 1}.$$

Letting $\ell \to \infty$, we obtain $u(a+b) \geq \rho$, as desired.

Everything else is easy. ∎

Every partial map $u' \colon A \dashrightarrow \mathbf{R} \cup \{\infty\}$ can be identified with its graph, which is a set of conditions, so, assuming that $\operatorname{dom} u'$ generates A, we may speak of the generated sub-valuation $\langle u' \rangle$. In particular, a sub-valuation u generates itself. If A is graded and C consists solely of homogeneous conditions, then $\langle C \rangle$ is homogeneous. A sub-valuation u is homogeneous if and only if it is generated by its restriction to homogeneous elements.

Definition 3.5. Let (B, u_B) be a sub-valued ring, and $\varphi \colon B \to A$ a ring morphism, making A a B-algebra.

(i) A *B-sub-valuation* on A is a sub-valuation u_A which satisfies, in addition, $u_A \circ \varphi \geq u_B$ (i.e., $u_A(\varphi(b)) \geq u_B(b)$ for all $b \in B$).

(ii) Let C be a set of conditions over A (in the sense of Definition 3.3) such that $\pi(C)$ generates A as a B-algebra. Then the *B-sub-valuation generated by C*, denoted $\langle C \rangle_B$, is the sub-valuation on A generated by $C \cup \{(\varphi(b), u_B(b)) : b \in B\}$.

Notice that if $B = K$ is a valued field, then any K-sub-valuation on A is, in particular, a K-vector space valuation, and must agree with v_K on the image of K in A.

Definition 3.6. Assume that A is a finitely generated B-algebra. Fix a finite generating tuple $c \in A^k$. For $a \in A$, define $\deg_c a$ to be the least degree of a polynomial $f \in B[X]$ such that $a = f(c)$. We say that a sequence of B-sub-valuations u_m *converges uniformly* to u, if for every $\varepsilon > 0$ there exists N such that for all $m \geq N$ and $a \in A$:

$$|u(a) - u_m(a)| \leq \varepsilon \deg_c a, \qquad \text{where } |\infty - \infty| = 0.$$

It is easy to check that this does not depend on the choice of generating tuple.

Definition 3.7. Continuing Definition 3.5, assume that A is a finitely generated B-algebra. Let u be a B-sub-valuation on A.

(i) We say that u is *finitely generated as a B-sub-valuation* (or *finitely generated over B*) if is generated, as a B-sub-valuation, by finitely many conditions. Notice that this implies, in particular, that A is finitely generated as a B-algebra.

(ii) We say that u is *almost finitely generated as a B-sub-valuation* (or *almost finitely generated over B*) if it is a uniform limit of a sequence of finitely generated B-sub-valuations.

(iii) More precisely, let $C = \{(a_i, \alpha_i) : i \in \mathbf{N}\}$ be a countable family of conditions over A, such that $\pi(C) = \{a_i : i \in \mathbf{N}\}$ generates A over B. Then $\{a_i : i < N\}$ generates A over B for some N, and we say that C *almost finitely generates u* as a B-sub-valuation if

$$\langle (a_i, \alpha_i) : i < N + n \rangle \xrightarrow[n \to \infty]{} u \qquad \text{uniformly.}$$

Notice that:

(i) Definition 3.7(iii) does not depend on the choice of enumeration of the conditions.
(ii) A B-sub-valuation on A is (almost) finitely generated if and only if it (almost finitely) generated by its restriction to a finite (countable) subset of A.
(iii) It is homogeneous if and only if the finite (countable) generating set can be taken to consist of homogeneous elements.

Let us now restrict the scope a little. As usual, we let $W \subseteq \mathbf{P}^n$ be algebraic, defined over $K \vDash ACMVF$. We let $I(W)$ denote the radical homogeneous ideal associated to W, and let $K[W] = K[X]/I(W)$, which inherits the structure of a graded K-algebra (it is the graded K-algebra of regular functions on the associated line bundle). Notice that any reduced finitely generated graded K-algebra can be presented in this manner, so even from a purely algebraic point of view this is essentially the general case.

Remark 3.8. A K-sub-valuation on $K[X]$ need not necessarily be almost finitely generated. In other words, we do not expect to have an analogue of the Hilbert Basis Theorem regarding Noetherianity. Indeed, assume that $(a_i : i \in \mathbf{N}) \subseteq K$ are such that $v(a_i - a_j) = 0$ for all $i \neq j$. Then the sub-valuation u generated by the conditions $u\big(\prod_{i<m}(X - Ya_i)\big) \geq m^2$ is not almost finitely generated.

The algebra $K[X]$ has a canonical set of generators over K, giving rise to a canonical notion of distance on proper sub-valuations:

$$d(u, w) = \sup_f \frac{|u(f) - w(f)|}{\deg f},$$

as f varies over all non-constant homogeneous polynomials (again, $|\infty - \infty| = 0$). Then $u_n \to u$ uniformly, in the sense of Definition 3.6, if and only if $d(u_n, u) \to 0$. If u and w have distinct kernels then clearly $d(u, w) = \infty$. As we show below, if u and w are homogeneous and almost finitely generated over K then the converse holds as well.

Example 3.9. The simplest valuation on $K[X]$ is $\tilde{v}$. It is homogeneous, and (finitely) generated by the conditions $\tilde{v}(X_i) \geq 0$.

Conversely (but this is far from being the general case):

Proposition 3.10. *Assume u is a reduced K-sub-valuation on $K[X]$ which is generated by a family C of* linear *conditions, i.e., of the form (λ, α) (or $u(\lambda) \geq \alpha$), where $\lambda \in K[X]_1$. Then up to a change of coordinates by a triangular matrix, u is arbitrarily close to $\tilde{v}$. Consequently, u is an almost finitely generated valuation.*

Proof. By Lemma 1.8. ∎

Hilbert's Nullstellensatz can be stated as a bijection between (homogeneous) radical ideals and (projective) zero-sets. The analogous statement in our setting asserts the existence of an isometric bijection between almost finitely generated homogeneous sub-valuations u on $K[X]$ (analogous to homogeneous radical ideals) and projective zero-sets $W \subseteq \mathbf{P}^n$ augmented with a virtual divisor η. Moreover, W is always the zero-set of $\ker u$ (this is just Hilbert's Nullstellensatz).

Definition 3.11. Let $W \subseteq \mathbf{P}^n$ be Zariski closed, and $\eta \colon W \to \mathbf{R}$ any function. For homogeneous $f \in K[X]_m$ we define

$$\eta_m^*(f) = \inf_{\xi \in W} \big(vf(\xi) - m\eta(\xi)\big) = m \inf_{\xi \in W} \big(\hat{f}(\xi) - \eta(\xi)\big).$$

We extend this to non-homogeneous polynomials by

$$\eta^*(f) = \min_m \eta_m^*(f_m),$$

where $f = \sum_m f_m$ is the decomposition into homogeneous components.

Conversely, let $u \colon K[W] \to \mathbf{R} \cup \{\infty\}$ be any function. For $\xi \in W$ we define

$$u^*(\xi) = \inf_f \left(\hat{f}(\xi) - \frac{u(f)}{\deg f} \right),$$

as f varies over non-constant homogeneous polynomials.

Theorem 3.12 (Virtual divisor Nullstellensatz). *Let $W \subseteq \mathbf{P}^n$ be Zariski closed, and let u be a reduced homogeneous sub-valuation on $K[W]$, almost finitely generated, say by the homogeneous conditions $C = \big\{(f_k, \gamma_k) : k \in \mathbf{N}\big\}$.*

(i) *We have*

$$u^* = \inf_k \left(\hat{f}_k - \frac{\gamma_k}{\deg f_k} \right)$$

uniformly, i.e.,

$$\min_{k<m} \left(\hat{f}_k - \frac{\gamma_k}{\deg f_k} \right) \xrightarrow[m\to\infty]{} u^* \qquad uniformly.$$

In particular, u^* is a virtual divisor on W. Conversely, every virtual divisor on W arises in this manner, and moreover, with all $\gamma_k = 0$.

(ii) *The double dual u^{**} is equal to u. Consequently, for any virtual divisor η on W we have $\eta^{**} = \eta$.*

(iii) *This bijection between virtual divisors on W and reduced almost finitely generated homogeneous sub-valuations on $K[W]$ is isometric:*

$$d(u,w) = d(u^*,w^*).$$

Proof. Let

$$u_m = \langle (f_k, \gamma_k) : k < m \rangle_K, \qquad \eta_m = \min_{k<m} \left(\hat{f}_k - \frac{\gamma_k}{\deg f_k} \right).$$

Clearly, $u_m^* \le \eta_m$, and for m large enough, the sequence $(f_k : k < m)$ generates $K[W]$ as a K-algebra, and its linear members generate $K[W]_1$ as a K-vector space. If $\xi \in W$ and $u_m^*(\xi) < \rho$, then there exists a homogeneous polynomial f which does not vanish at ξ, such that $\frac{u(f)}{\deg f} > \hat{f}(\xi) - \rho$. The same remains true if we replace f with f^n (for $n \ge 1$), so we may assume that

$$f = \sum_i \prod_j g_{ij}, \qquad \frac{\min_i \sum_j \delta_{ij}}{\deg f} + \rho > \hat{f}(\xi),$$

where $(g_{ij}, \delta_{ij}) \in \{(f_k, \gamma_k) : k < m\}$. But then, unwinding the definition of $\hat{f}$, there must exist a pair $(f_k, \gamma_k) = (g_{ij}, \delta_{ij})$ such that $\frac{\gamma_k}{\deg f_k} + \rho > \hat{f}_k(\xi)$, so $\rho > \eta_m(\xi)$, and $u_m^* = \eta_m$. It follows immediately from the definitions that $d(u^*, w^*) \le d(u, w)$, so $\eta_m \to u^*$ uniformly. Since the $v(K^\times)$ is dense in $\mathbf{R}$, we may take all the γ_k to be zero. By Proposition 1.11, each η_m is a virtual divisor, and so is u^*, and conversely, every virtual divisor arises in this fashion, proving the first item.

For the rest, let us indeed assume that $\gamma_k = 0$ for all k. It is immediate from Definition 3.11 that u^{**} is a homogeneous sub-valuation on $K[W]$, and that $u^{**} \ge u$. Assume that for some $g \in K[W]_d$ we have $u(g) < \rho$. By Fact 3.2, there exists a valuation w on $K[W]$ such that $w \ge u$ and $w(g) < \rho$. Then $\ker w$ is a prime ideal of $K[W]$ containing $I(W)$, and w induces a valuation v_L on the field $L = \operatorname{Frac}(K[X]/\ker w)$, extending v_K.

Let $\xi = [x] \in \mathbf{P}^n(L)$, where $x \in L^{n+1}$ is the image of X. Then $v \circ f_k(x) = w(f) \ge u(f) \ge 0$, $\xi \in W(L)$, and u^* extends to $W(L)$ by definability, so:

$$u^*(\xi) = \inf_k \hat{f}_k(\xi) = \inf_k \frac{w(f_k)}{\deg f_k} - \tilde{v}(x) \ge -\tilde{v}(x).$$

Therefore

$$vg(\xi) - du^*(\xi) \le vg(\xi) + d\tilde{v}(x) = w(g) < \rho.$$

By quantifier elimination in ACMVF, there exists $\zeta \in W(K)$ such that

$$vg(\zeta) - du^*(\zeta) < \rho,$$

so $u^{**}(g) < \rho$ as well, proving that $u^{**} = u$. Since every virtual divisor η on W can be expressed as u^*, we also have $\eta^{**} = u^{***} = u^* = \eta$, proving the second item.

It is again immediate from the definitions that $d(\eta^*, \theta^*) < d(\eta, \theta)$ for any two virtual divisors. Therefore $d(u,w) = d(u^{**}, w^{**}) \le d(u^*, w^*) \le d(u, w)$, completing the proof. ∎

Let us recall an easy corollary of the classical Nullstellensatz (we prefer to work in projective space, but one can also give an affine version).

Fact 3.13. *For every $n, N \in \mathbf{N}$ there exists $m \in \mathbf{N}$ such that the following holds. Let K be an algebraically closed field, let $I \subseteq K[X]$ be homogeneous, generated by polynomials of degree at most N, and let g be homogeneous of degree at most N vanishing on $V(I) \subseteq \mathbf{P}^n(K)$. Then $g^m \in I$.*

Proof. Let $D = \dim K[X]_{\leq N}$, which only depends on n and N. Then one can generate I with fewer than D polynomials, all of degree at most N. In particular, the fact that g vanishes on $V(I)$ can be expressed in first-order logic.

Assume that no such m exists. Taking an ultra-product of counter-examples for all m, we obtain an algebraically closed field K, a finite family of homogeneous polynomials of degree at most N generating an ideal I, and g homogeneous of degree at most N, which vanishes on $V(I)$. By Hilbert's Nullstellensatz, there exists M such that $g^M \in I$. Viewed as a property of the coefficients, the same must hold in infinitely many of the aforementioned counter-examples. In particular, it holds in the mth counter-example for some $m \geq M$, a contradiction. ∎

The corresponding corollary of Theorem 3.12 is the following.

Corollary 3.14. *For every* $n, N \in \mathbf{N}$ *there exists* $m \in \mathbf{N}$ *such that the following holds. Let K be an algebraically closed valued field, and let $W \subseteq \mathbf{P}^n(K)$ be defined by polynomials of degree at most N. Let u be a reduced homogeneous subvaluation on $K[W]$, almost finitely generated by a family C of homogeneous conditions (f, γ), where $\deg f \leq N$. Let*

$$c = \sup\left\{ \left|\gamma - \tilde{v}(f)\right| : (f, \gamma) \in C \right\}.$$ *Then for every g homogeneous of degree at most N:*

$$u(g) - \frac{c}{N} \leq \frac{1}{m}\sup\left\{ \min_i v(a_i) + \sum_j \gamma_{ij} : g^m = \sum_i a_i \prod_j f_{ij},\ \text{where } a_i \in K \text{ and } (f_{ij}, \gamma_{ij}) \in C \right\} \leq u(g). \quad (1)$$

Proof. The second inequality always holds, so we only prove the first. Possibly re-scaling the valuation, we may assume that $c = 1$. As in the proof of Fact 3.13, let $D = \dim K[X]_{\leq N}$. Then a sub-valuation arbitrarily close to u can be generated by a subset of C of size at most D, so we may add the hypothesis that $|C| \leq D$. Similarly, we require that $\tilde{v}(f) = 0$, and so $\gamma \in [-1, 1]$, for every $(f, \gamma) \in C$, and that $\tilde{v}(g) = 0$.

By Theorem 3.12, value $u(g)$ is uniformly definable in g and C (i.e., in the coefficients of g, viewed as a point in $\mathbf{P}$, and, for each $(f, \gamma) \in C$, in f, viewed in the same manner, and in the real constant γ). Letting $u_m(g)$ denote the middle expression in (1), it similarly definable in g and C.

Assume that our assertion is false, and again take and ultra-product of counter-examples. We may assume that for every m, the set of all multiples of m is large in the ultra-filter. Since u_m increases with the divisibility relation, we obtain g and C such that $u(g) \geq u_m(g) + 1/N$ for all m, which is impossible by definition of $u(g)$. ∎

Corollary 3.15. *Let η_K be a virtual divisor on $W(K)$, and let $u_K = \eta_K^*$ be the corresponding sub-valuation on $K[W]$. Let L/K be an extension. Let η_L be the extension of η_K to $W(L)$ by definability, and let $u_L = \langle u_K \rangle_L$ be the sub-valuation on $L[W]$ generated by u_K. Then $\eta_L^* = u_L = u_K \otimes v_L$. In other words, the duality commutes with scalar extensions.*

Proof. Let u_K be almost finitely generated by $C = \{ f_k : k \in \mathbf{N} \}$ over K, with f_k homogeneous, so $\eta_K = \inf_k \hat{f}_k\!\restriction_{W(K)}$ uniformly. Then u_L is almost finitely generated by C over K, and $\eta_L = \inf_k \hat{f}_k\!\restriction_{W(L)}$ uniformly, so indeed $u_L = \eta_L^*$.

It is also clear that $u_K \otimes v_L \leq u_L$, since being a K-algebra sub-valuation is a stronger property than begin a K-vector space valuation. If they are not equal, then there exists some $g \in L[W]_d$ and $\varepsilon > 0$ such that $(u_K \otimes v_L)(g) + 2\varepsilon < u_L(g)$. We can find a basis $(f_i : i < N)$ for $K[W]_d$, with respect to which u_K is ε-close to $\tilde{v}$. This is also a basis for $K[W]_d \otimes_K L$, with respect to which $u_K \otimes v_L$ is ε-close to $\tilde{v}$. We can express $g = \sum b_i f_i$, with $b_i \in L$ so $u_L(g) > (u_K \otimes v_L)(g) + 2\varepsilon \geq \tilde{v}(b) + \varepsilon$. In other words, in L we have:

$$\sup_{[z] \in \mathbf{P}^{N-1}} \inf_{[x] \in W} \left(v\left(\sum z_i f_i(x)\right) - d\tilde{v}(x) - \tilde{v}(z) \right) > \varepsilon.$$

By quantifier elimination, the same holds in K. In other words, there are $a \in K^N$ such that, if $h = \sum a_i f_i \in K[W]_d$, then $u_K(h) > \tilde{v}(a) + \varepsilon$, contradicting the choice of (f_i). (Notice that this implies, in particular, that the natural surjective map $K[W] \otimes_K L \to L[W]$ is bijective.) ∎

Corollary 3.16. *Let $W \subseteq V \subseteq \mathbf{P}^n$ be algebraic sets and η a virtual divisor on V. Then for homogeneous $f \in K[X]$ we have*

$$(\eta\!\restriction_W)^*(f) = \lim_{m \to \infty} \frac{\eta^*\left(f^m + I(W)\right)}{m}. \quad (2)$$

Proof. Say that $\eta = \inf_k \hat{f}_k\!\restriction_V$ uniformly, so η^* is the least sub-valuation on $K[V]$ satisfying $\eta^*(f_k) \geq 0$. Let u denote the right hand side, extended to non-homogeneous polynomials in the usual way. It is easy to check that the limit is increasing with respect to the divisibility relation $m \mid m'$, and to deduce that u is indeed a homogeneous sub-valuation on $K[W]$. Pulled back to a sub-valuation on $K[V]$, it is clearly the least one

satisfying $u \geq \eta^*$ and $\ker u \supseteq I(W)$. In other words, u is the least sub-valuation on $K[W]$ satisfying $u(f_k) \geq 0$ for all k. On the other hand, $\eta\!\restriction_W = \inf_k \hat{f}_k\!\restriction_W$, so $u = (\eta\!\restriction_W)^*$. ∎

4. STATEMENT OF THE MAIN RESULT (AND A FEW EASY CASES)

In this section we state our main result, Theorem 4.3. It relates the self-intersection of a virtual divisor η on a projective variety, on the one hand, with the asymptotic behaviour of the volume of the dual η^*_m, on the other hand, as the degree m goes to infinity.

Let E be a vector space. For any k, the symmetric group $\mathfrak{S}_k$ acts naturally on $E^{\otimes k}$, and we may define the exterior power of E by

$$\bigwedge^k E = \left\{ x \in E^{\otimes k} : \sigma(x) = (\operatorname{sgn} \sigma)x \text{ for all } \sigma \in \mathfrak{S}_k \right\}.$$

In particular, for a finite family of vectors $\mathbf{x} \in E^k$ we shall write

$$\mathbf{x}^\wedge = x_0 \wedge \cdots \wedge x_{k-1} = \sum_{\sigma \in \mathfrak{S}_k} (\operatorname{sgn} \sigma) x_{\sigma(0)} \otimes \cdots \otimes x_{\sigma(k-1)} \in \bigwedge^k E.$$

Given a valuation u on E, this gives rise to a value $u^{\otimes k}(\mathbf{x}^\wedge)$. When there is no risk of ambiguity, we may simply denote this by $u(\mathbf{x}^\wedge)$.

Definition 4.1. When $\mathbf{x}$ is a basis of E, we call $u(\mathbf{x}^\wedge)$ the *volume* of (E, u) (or of u), relative to $\mathbf{x}$, denoted $\operatorname{vol}_{\mathbf{x}} E$ or $\operatorname{vol}_{\mathbf{x}} u$, depending on context.

The volume can be viewed as a (particularly simple) special case of Definition 2.7. Indeed, we may assume that $E = K^m$, equipped with the standard basis $\mathbf{e}$. Then $\eta([x]) = u(x) - \tilde{v}(x)$ is a virtual divisor on $\mathbf{P}^{m-1}$, and

$$\eta^{\wedge m} = u(\mathbf{e}^\wedge).$$

In order to state our main result we need one last technical definition.

Definition 4.2. Let K_0 be a field, T an infinite family of indeterminates, and $L \supseteq K_0(T)$. We shall call a valuation on $K_0(T)$, or any larger field, *good* (with respect to K_0, T), if
   (i) either the tuple T is v-generic over K_0,
   (ii) or the restriction to $K_0(T)$ is the P-adic valuation for some irreducible polynomial $P \in K_0[T]$.

Theorem 4.3. *Let K_0 be an algebraically closed field, T an infinite family of indeterminates, and $K = K_0(T)^a$. Let $d \in \mathbf{N}$, and let $W \subseteq \mathbf{P}^n(K)$ be a projective variety defined over K_0, of dimension ℓ, with Chow form $\mathfrak{C}_W$. Then there exist a tuples a in K, a sequence $\gamma_m \to 0$, and for infinitely many $m \in \mathbf{N}$ there exists a basis B_m for $K[W]_m$, such that for any good valuation on K, and any virtual divisor η on $W(K)$ of degree d, defined over K_0, with d-width $w_d(\eta)$ as per Definition 2.2, we have*

$$\frac{\operatorname{vol}_{B_m} \eta^*}{m \dim K[W]_m} \geq -\frac{\eta^{\wedge \ell+1} \wedge \widehat{\mathfrak{C}}_W}{\ell+1} + \gamma_m \big(\tilde{v}(a) - w_d(\eta) \big). \tag{3}$$

Let us start with a few easy observations.

Lemma 4.4. *Let K be a valued field, $W \subseteq \mathbf{P}^n$ a variety of dimension ℓ, and let η and θ be two virtual divisors on W. Let $\mathfrak{C}_W$ be a Chow form for W, and let B be a basis for $K[W]_m$. Then*

$$\left| \frac{\operatorname{vol}_B \eta^*}{m \dim K[W]_m} - \frac{\operatorname{vol}_B \theta^*}{m \dim K[W]_m} \right| \leq d(\eta, \theta), \qquad \left| \frac{\eta^{\wedge \ell+1} \wedge \widehat{\mathfrak{C}}_W}{\ell+1} - \frac{\theta^{\wedge \ell+1} \wedge \widehat{\mathfrak{C}}_W}{\ell+1} \right| \leq d(\eta, \theta). \tag{4}$$

If $\theta = \eta + \alpha$ for some $\alpha \in \mathbf{R}$, then

$$\frac{\operatorname{vol}_B \theta^*}{m \dim K[W]_m} = \frac{\operatorname{vol}_B \eta^*}{m \dim K[W]_m} - \alpha, \qquad \frac{\theta^{\wedge \ell+1} \wedge \widehat{\mathfrak{C}}_W}{\ell+1} = \frac{\eta^{\wedge \ell+1} \wedge \widehat{\mathfrak{C}}_W}{\ell+1} + \alpha. \tag{5}$$

Consequently, it will suffice to prove Theorem 4.3 in the case where $\sup \eta = 0$.

Proof. If $g \in K[W]_m$, then $\left| \eta^*(g) - \theta^*(g) \right| \leq md(\eta, \theta)$. Therefore, if $z \in K[X]_m^{\otimes \dim K[X]_m}$ (and in particular if $z = B^\wedge$), then $\left| \eta^*(z) - \theta^*(z) \right| \leq m \dim K[W]_m d(\eta, \theta)$. The second inequality in (4) is just Lemma 2.6. If $\theta = \eta + \alpha$, then $\theta^*_m = \eta^*_m - m\alpha$, whence the identity in (5). We may express $\eta = \hat{f}\!\restriction_{W(K)}$ and $\theta = \widehat{af}\!\restriction_{W(K)}$, where $v(a) = \alpha \deg f$, and the second identity follows.

Let $\eta = \hat{f}\!\restriction_{W(K)}$, where $f \in L[X]_d$ for some extension L/K, such that $\tilde{v}(f)$ is maximal possible. By (5) we may reduce to the case where $\sup \eta = 0$. ∎

We are going to prove Theorem 4.3 by calculating the volume of $K[X]_m$ with respect to particularly convenient bases. For this we require a few basic properties of the volume function.

Lemma 4.5. *Let E be a valued vector space over K, and let $\mathbf{x} \in E^m$ be a basis.*

(i) *If $\mathbf{y} = \mathbf{x}A$ is any other basis (we consider bases as rows of vectors) then $\mathrm{vol}_{\mathbf{y}}\, E = \mathrm{vol}_{\mathbf{x}}\, E + v(\det A)$.*

(ii) *Let L/K be an extension, and $E_L = E \otimes_K L$ be the extension of scalars. Then, identifying $\mathbf{x}$ with $\mathbf{x} \otimes 1 \in E_L^m$, the volume remains unchanged: $\mathrm{vol}_{\mathbf{x}}\, E = \mathrm{vol}_{\mathbf{x}}\, E_L$.*

(iii) *We have*

$$\mathrm{vol}_{\mathbf{x}}\, E = \sum_{i<m} u(x_i + F_i), \qquad F_i = \mathrm{Span}(x_j : j < i).$$

(iv) *Let $\mathbf{y} \subseteq \mathbf{x}$ generate $F \subseteq E$. Then*

$$\mathrm{vol}_{\mathbf{x}}\, E = \mathrm{vol}_{\mathbf{y}}\, F + \mathrm{vol}_{\mathbf{x} \smallsetminus \mathbf{y}}(E/F),$$

where by $\mathbf{x} \smallsetminus \mathbf{y}$ we actually mean the image in E/F.

Proof.   (i) Since $\mathbf{y}^{\wedge} = (\det A)\mathbf{x}^{\wedge}$.

(ii) We may identify $(E_L)^{\otimes m}$ with $(E^{\otimes m})_L$ as valued vector spaces over L (by reduction to $\tilde{v}$), and then $(\mathbf{x} \otimes 1)^{\wedge} = \mathbf{x}^{\wedge} \otimes 1$.

(iii) By (the proof of) Lemma 1.8, together with the fact that the volume of $\tilde{v}$ with respect to the standard basis is zero.

(iv) Follows from the previous item.   ■

Now to the construction of a basis, given a virtual divisor.

Lemma 4.6. *Let η be a virtual divisor of degree d on $\mathbf{P}^n$, defined over K, with $\sup \eta = 0$ and $\inf \eta = -\alpha$. Let $m \geq n$ be given.*

(i) *In an extension L/K, one can find $f \in L[X]_d$ weakly v-generic over K such that $\eta = \hat{f}\!\restriction_{\mathbf{P}^n(K)}$, in which case $\tilde{v}(f) = -d\alpha$, so $\alpha = w_d(f)$. Moreover, one can construct a family $F = (f_i : i < m) \subseteq L[X]_d$ of polynomials are isomorphic to f over K, and are independent over K in the sense of Section 1, i.e., such that if $L_i = K(f_i)$ is the subfield generated by K and the coefficients of f_i, then their compositum is $\mathrm{Frac}(L_0 \otimes_K \cdots \otimes_K L_{m-1})$, as a valued field.*

(ii) *Let us fix such a family. Then each f_i is weakly v-generic over $K(f_j : i \neq j)$.*

(iii) *Let us also fix $dX \in L^{n+1}$ v-generic over $K(F)$, which we use as direction for Hasse derivation, as in [Ben, Section 4]. Then dX is algebraically generic over $K(F)$, and F is good in the sense of [Ben, Section 4].*

Proof. For (i), choose f such that $\eta = \hat{f}\!\restriction_{\mathbf{P}^n(K)}$. Adding v-generic elements (over K) to the coefficients of f does not change η and makes f weakly v-generic. Since the variety here is the entire space, it is easy to see that $\tilde{v}(f) = d \inf \eta$. The construction of a field containing an independent family is essentially given in the statement. Items (ii) and (iii) are immediate.   ■

Lemma 4.7. *Continuing Lemma 4.6, assume that $h \in K[X]_D$ is an additional polynomial, weakly v-generic over the prime field, so Fh is again a good family (possibly $D = 0$ and $h = 1$). Let $\psi \subseteq \left[(Fh)^{\partial n}\right]$ be a good set. Then, identifying $K[X]_{md+D-n}$ with $K^{\binom{md+D}{n}}$ via the basis $\Phi_{Fh,\psi}$, we have $\eta^* = \tilde{v} + O\big(\alpha - \tilde{v}(h)\big)$. In other words, there exists a constant c which depends only on n, d and D (but not on m, F, or h) such that for all choices of coefficients $a_{\xi} \in L$:*

$$\left| \eta^* \left(\sum_{\xi \in \psi} a_{\xi} \varphi_{Fh,\psi,\xi} \right) - \min_{\xi \in \psi} v(a_{\xi}) \right| \leq c\big(\alpha - \tilde{v}(h)\big).$$

Proof. Let us estimate $\eta^*(\varphi_{Fh,\psi,\xi})$ for $\xi \in \psi$. There exists some set $G \in \binom{F}{n}$ such that $\xi \in \psi_{Gh} = \psi \cap \left[(Gh)^{\partial n}\right]$, and we have $\varphi_{Fh,\psi,\xi} = \varphi_{Gh,\psi,\xi} \prod(F \smallsetminus G)$ and $\deg \varphi_{Gh,\psi,\xi} = D + nd - n$. The coefficients of $\varphi_{Gh,\psi,\xi}$ are algebraic over dX and the coefficients of Gh, with finitely many possibilities for the irreducible polynomials (since all possibilities for G are isomorphic). The valuations of the coefficients of these irreducible polynomials are $O\big(\alpha - \tilde{v}(h)\big)$ in absolute value, by weak v-genericity. By Newton's polygon we deduce that $\tilde{v}(\varphi_{Gh,\psi,\xi}) = O\big(\alpha - \tilde{v}(h)\big)$ as well. For any $\zeta \in \mathbf{P}^n(L)$, we have, by Lemma 2.4:

$$v\varphi_{Fh,\psi,\xi}(\zeta) \geq (m-n)d\eta(\zeta) + v\varphi_{Gh,\psi,\xi}(\zeta). \tag{6}$$

Therefore,

$$\eta^*(\varphi_{Fh,\psi,\xi}) = \inf_{\zeta \in \mathbf{P}^n(L)} \big(v\varphi_{Fh,\psi,\xi}(\zeta) - (D + md - n)\eta(\zeta) \big)$$

$$\geq \inf_{\zeta \in \mathbf{P}^n(L)} \big(v\varphi_{Gh,\psi,\xi}(\zeta) - (D + nd - n)\eta(\zeta) \big)$$

$$\geq \tilde{v}(\varphi_{Gh,\psi,\xi}) = O\big(\alpha - \tilde{v}(h) \big).$$

For the converse inequality, observe that equality holds in (6) for $\zeta = \xi$ (since $\xi \perp_K F \smallsetminus G$). If $g = \sum_{\xi \in \psi} a_\xi \varphi_{Fg,\psi,\xi}$, then:

$$\eta^*(g) \leq \min\{ vg(\xi) - (D + md - n)\eta(\xi) : \xi \in \psi \}$$

$$= \min\{ v(a_\xi) + v\varphi_{Fh,\psi,\xi}(\xi) - (D + md - n)\eta(\xi) : \xi \in \psi \}$$

$$= \min \left\{ v(a_\xi) + v\varphi_{Gh,\psi,\xi}(\xi) - (D + nd - n)\eta(\xi) : G \in \binom{F}{n}, \ \xi \in \psi_{Gh} \right\}$$

$$= \min_{\xi \in \psi} v(a_\xi) + O\big(\alpha - \tilde{v}(h) \big).$$

Together, these two inequalities conclude the proof. ∎

Lemma 4.8. *Continuing* Lemma 4.6, *let* $h \in K[X]_D$. *Then for any* $G \in \binom{F}{n}$, *and for some constant* $C = C(n,d)$, *we have*

$$\eta^*(h) = \min\{ vh(\xi) - D\eta(\xi) : \xi \in [G^{\partial n}] \}, \tag{7}$$

$$vh(\xi) \leq \tilde{v}(h) + CDw_d(\eta) \qquad \forall \xi \in [F^{\partial n}]. \tag{8}$$

Proof. Let $h_k = h^{kd} X_0^{n(d-1)}$, of degree $m_k d - n$, where $m_k = kD + n$. In what follows, for $O(\cdot)$ notation, we consider all the data fixed, with the exception of k, which may vary. On the one hand, since η^* is a sub-valuation, we have

$$\eta^*(h_k) \geq n(d-1)\eta^*(X_0) + dk\eta^*(h) = \frac{m_k d}{D}\eta^*(h) + O(1).$$

On the other hand, in $\eta^*(h) = \inf_\xi vh(\xi) - D\eta(\xi)$, the infimum is obtained along some sequence of ξ_ℓ. For each ℓ there is some $i \leq n$ such that $\widehat{X}_i(\xi_\ell) = 0$ (i.e., $\tilde{v}(x_\ell) = v(x_{\ell,i})$). Possibly passing to a sub-sequence and permuting coordinates, we may assume that $\widehat{X}_0(\xi_\ell) = 0$ throughout. But then,

$$\eta^*(h_k) \leq \min_\ell \big[vh_k(\xi_\ell) - (m_k d - n)\eta(\xi_\ell) \big] = \frac{m_k d}{D}\eta^*(h) + O(1).$$

It follows that

$$\lim_{k \to \infty} \frac{D\eta^*(h_k)}{m_k d} = \eta^*(h).$$

Fixing k for a while, let $F = F_k = (f_i : i < m_k)$, let $\psi \subseteq [F^{\partial n}]$ be good, and let $h_k = \sum_{\xi \in \psi} a_\xi \varphi_{F,\psi,\xi}$. For each $\xi = [x] \in [F^{\partial n}]$ there is some $G_\xi \in \binom{F}{n}$ such that $\xi \in [G_\xi^{\partial n}]$. Let us identify F and G_ξ with their products, so $F = (F/G_\xi)G_\xi$. Since $\partial_j G_\xi(x)$ vanish for $j < n$, we have $\partial_n F(x) = (F/G_\xi)(x)\partial_n G_\xi(x)$, and therefore $v\partial_n F(\xi) = kDd\eta(\xi) + v\partial_n G_\xi(\xi)$. Therefore, letting k vary and keeping everything else fixed, we have

$$\eta^*(h_k) = \min_{\xi \in \psi} v(a_\xi) + O(1)$$

$$= \min_{\xi \in \psi} \big(vh_k(\xi) - v\partial_n F_k(\xi) \big) + O(1)$$

$$= \min_{\xi \in \psi} \big(vh_k(\xi) - kDd\eta(\xi) - v\partial_n G_\xi(\xi) \big) + O(1).$$

The constant in $O(1)$ does not depend on the choice of good set ψ, so we may take the minimum over $\xi \in [F^{\partial n}]$. But then, by symmetry over K, we have for $G = \{f_0 : i < n\}$ (of any other n-family of the f_i):

$$\eta^*(h_k) = \min_{\xi \in [G^{\partial n}]} \big(vh_k(\xi) - kDd\eta(\xi) - v\partial_n G(\xi) \big) + O(1)$$

$$= \min_{\xi \in [G^{\partial n}]} \big(m_k dvh(\xi)/D - m_k d\eta(\xi) \big) + O(1)$$

$$= \frac{m_k d}{D} \min_{\xi \in [G^{\partial n}]} \big(vh(\xi) - D\eta(\xi) \big) + O(1).$$

Therefore,

$$\lim_{k\to\infty}\frac{D\eta^*(h_k)}{m_k d} = \min_{\zeta\in[G^{\partial n}]}\big(vh(\zeta) - D\eta(\zeta)\big).$$

This proves (7).

For (8), let $\xi \in [F^{\partial n}]$, so $\xi \in [G^{\partial n}]$ for some $G \in \binom{F}{n}$, and let $g = \prod G$. Let $N = n!\binom{nd}{n}$, which is the cardinal of the multi-set $[G^{\partial n}]$, and let $C = \sum_{i<n}\frac{dnN}{nd-i}$. We have $\tilde{v}(g) = n\tilde{v}(f_0) = -dnw_d(\eta)$. For $i < n$, the degree of $\partial_i g$ in $G^{\partial n}$ is $N/(nd-i)$, so the total degree of g in $G^{\partial n}$ is C/dn. It follows that $\tilde{v}(G^{\partial n}) \geq -Cw_d(\eta)$.

Let $A = K(dX, G)$ (i.e., the K-algebra generated by dX and the coefficients of G), and let $B = K[dX, \lambda_{i,j} : i < n, j < d]$, where the $\lambda_{i,j}$ are indeterminate linear forms (again, adjoining $\lambda_{i,j}$ means adjoining its coefficients). We equip B with the valuation $v_B = \tilde{v}$, i.e., we make the (coefficients of) $\lambda_{i,j}$ v-generic over K. Finally, we define a map $s\colon A \to B$ fixing $K[dX]$ and sending $f_i \mapsto \prod_j \lambda_{i,j}$ (we assume, as we may, that $G = (f_i : i < n)$). Since each f_i is weakly v-generic over everything else, and $\tilde{v}(\lambda_{i,j}) \geq 0$, we have $v_A \leq v_B \circ s$. In addition, $s(G^{\partial n})$ is then $n!$-power of the product of formal intersections of n among the $\lambda_{i,j}$, i.e., a product of N many v-generic algebraic points (times some scalar which is a power of dX and can be ignored). Therefore,

$$v(h \wedge G^{\partial n}) \leq v_B\left(h \wedge \left(\prod_{i,j}\lambda_{i,j}\right)^{\partial n}\right) = N\tilde{v}(h).$$

Since $vh(\zeta) \geq \tilde{v}(h)$,

$$\begin{aligned}
vg(\xi) &\leq \sum_{\zeta\in[G]^{\partial n}} vg(\zeta) - (N-1)\tilde{v}(h) \\
&= v(h \wedge G^{\partial n}) - D\tilde{v}(G^{\partial n}) - (N-1)\tilde{v}(h) \\
&\leq N\tilde{v}(h) + CDw_d(\eta) - (N-1)\tilde{v}(h) \\
&= \tilde{v}(h) + CDw_d(\eta),
\end{aligned}$$

as claimed.   ■

We can now start proving our main result, in several steps. The first step is presented mostly for expository purposes, in order to present the main ingredients in a simpler setting.

Lemma 4.9. *Theorem 4.3 holds when $W = \mathbf{P}^n$.*

Moreover, there is no need to adjoin the indeterminates T or assume that the valuation is good.

Proof. We may assume that $\mathfrak{C}_{\mathbf{P}^n}$ is the one with the canonical normalisation (otherwise just multiply some vector in each basis by an appropriate scalar). We will show that (3) holds for degrees of the form $md - n$, for $m \geq 2n$, where $B_{md-n} = \mathfrak{M}_{md-n}$ is the set of monomials. Throughout, let $m \geq n$ and $N = \dim K[X]_{md-n} = \binom{md}{n}$.

By Lemma 4.4, we may assume that $\sup \eta = 0$, and re-scaling, we may assume that $w_d(\eta) = 1$. Let L, dX, and $F = (f_i : i \leq m)$ be as per Lemma 4.6. In particular, F is good, so let $\psi \subseteq [F^{\partial n}]$ be a good set. Identifying $L[X]_{md-n}$ with L^N via the basis $\Phi_{F,\psi}$, we have $\eta^* = \tilde{v} + O(1)$, and so $\mathrm{vol}_{\Phi_{F,\psi}}\eta^* = O(N) = O(m^n)$.

By [Ben, Theorem 4.15]:

$$\det\Phi_{F,\psi} = \mathfrak{d}_{F,\psi} = \prod_{H\in\binom{F}{n+1}} H^\wedge \prod_{G\subseteq F, |G|\leq n} \widehat{\mathfrak{d}}_{G,\psi}.$$

Since $m \geq 2n$, we have $\binom{m}{k} \leq \binom{m}{n}$ for $k \leq n$, so we have fewer than $(n+1)\binom{m}{n} = O(m^n)$ factors of the form $\widehat{\mathfrak{d}}_{G,\psi}$. Each $\widehat{\mathfrak{d}}_{G,\psi}$ is algebraic over dX and the coefficients of G, with only finitely many possibilities for its irreducible polynomial over them. Since all the f_i are $O(1)$-generic over K, we have $|v(\widehat{\mathfrak{d}}_{G,\psi})| = O(1)$ (by Newton's polygon). Therefore

$$v(\det\Phi_{F,\psi}) = \sum_{H\in\binom{F}{n+1}} v(H^\wedge) + O(m^n) = d^{n+1}\binom{m}{n+1}\eta^{\wedge n+1} + O(m^n).$$

We have $\frac{(md-n)N}{n+1} = \binom{md}{n+1} = d^{n+1}\binom{m}{n+1} + O(m^n)$, and $\eta^{\wedge n+1} = O(1)$ (since $-1 \leq \eta \leq 0$), so:

$$d^{n+1}\binom{m}{n+1}\eta^{\wedge n+1} = \frac{(md-n)N}{n}\eta^{\wedge n+1} + O(m^n).$$

We conclude that

$$
\begin{aligned}
\operatorname{vol}_{\mathfrak{M}_{md-n}} \eta^* &= \operatorname{vol}_{\Phi_{F,\psi}} \eta^* - v(\det \Phi_{F,\psi}) \\
&= -d^{n+1} \binom{m}{n+1} \eta^{\wedge n+1} + O(m^n) \\
&= -\frac{(md-n)N}{n+1} \eta^{\wedge n+1} + O(m^n).
\end{aligned}
$$

Dividing by $(md-n)N$ we obtain the desired estimate. ∎

In the second step, we consider a hypersurface. The argument follows a similar path, with several added technical complications.

Lemma 4.10. *Theorem 4.3 holds when $W \subseteq \mathbf{P}^n$ is a hypersurface.*
Moreover, there is no need to adjoin the indeterminates T or assume that the valuation is good.

Proof. We have $\mathfrak{C}_W = g \wedge \mathfrak{C}_{\mathbf{P}^n}$, with $g \in K_0[X]_D$ being the irreducible polynomial defining W. Notice that then $\tilde{v}(g) = \tilde{v}(\mathfrak{C}_W)$, and we shall identify $\widehat{g}$ with $\widehat{\mathfrak{C}_W}$.

For $m \geq D$ we have $K[W]_m = K[X]_m / gK[X]_{m-D}$, for which we construct a basis B_m as follows. The set of monomials $\mathfrak{M}_m$ is a basis for $K[X]_m$, while $g\mathfrak{M}_{m-D}$ is a basis for $gK[X]_{m-D}$. We define $B_m^0 \subseteq K[X]_m$ to be any set which completes $g\mathfrak{M}_{m-D}$ to a basis for $K[X]_m$, such that the transition matrix to $\mathfrak{M}_m$ has determinant ± 1. Its image B_m in $K[W]_m$ is a basis there. We also let the tuple a be $(1, g)$, so $\tilde{v}(a) = 0 \wedge \tilde{v}(g)$.

By Lemma 4.4, we may assume that $\sup \eta = 0$, and re-scaling, we may assume that $w_d(\eta) = 1$. It is then easy to construct θ of degree d on $\mathbf{P}^n$ such that $\eta = \theta|_W$, $\sup \theta = 0$ and $w_d(\theta) = -\inf \theta = w_d(\eta)$. Then θ^* is a sub-valuation on $K[X]$, and on $K[W]_m$ we have $\theta_m^*/gK[X]_{m-D} \leq \eta_m^*$ (since $\theta_m^* \leq \eta_m^*$ by definition, and $gK[X]_{m-D} \subseteq \ker \eta_m^*$). By Lemma 4.5(iv):

$$
\operatorname{vol}_{B_m} \eta_m^* \geq \operatorname{vol}_{B_m} \left(\theta_m^* / gK[X]_{m-D} \right) = \operatorname{vol}_{\mathfrak{M}_m} \theta_m^* - \operatorname{vol}_{g\mathfrak{M}_{m-D}} \theta_m^*.
$$

On the other hand, $\eta^{\wedge n} \wedge \widehat{g} = \theta^{\wedge n} \wedge \widehat{g}$.

We will show that (3) holds for degrees of the form $md + D - n$, for $m \geq n$. Throughout, let $M = \dim K[W]_{md+D-n} = \binom{md+D}{n} - \binom{md}{n}$. Then we want to show that:

$$
\frac{\operatorname{vol}_{\mathfrak{M}_{md+D-n}} \theta^* - \operatorname{vol}_{g\mathfrak{M}_{md-n}} \theta^*}{(md+D-n)M} + \frac{\theta^{\wedge n} \wedge \widehat{g}}{n} \geq \frac{\gamma}{m} \left(v(a) - w_d(\eta) \right). \tag{9}
$$

Let L, dX, and $F = (f_i : i < m)$ correspond to θ as per Lemma 4.6.

Let us see how (9) changes when we multiply g by a constant $b \in K$. We have

$$
\frac{\theta^{\wedge n} \wedge \widehat{bg}}{n} = \frac{\theta^{\wedge n} \wedge \widehat{g} + v(b)/D}{n} = \frac{\theta^{\wedge n} \wedge \widehat{g}}{n} + \frac{v(b)}{nD},
$$

$$
\operatorname{vol}_{bg\mathfrak{M}_{md-n}} \theta^* = \operatorname{vol}_{g\mathfrak{M}_{md-n}} \theta^* + \binom{md}{n} v(b).
$$

Observe that

$$
\begin{aligned}
(md+D-n)M &= (n+1)\binom{md+D}{n+1} - (n+1)\binom{md}{n+1} - D\binom{md}{n} \\
&= nD\binom{md}{n} + (n+1)\left[\binom{md+D}{n+1} - \binom{md}{n+1} - D\binom{md}{n} \right] \\
&= nD\binom{md}{n} + (n+1) \sum_{i<D} \left[\binom{md+i}{n} - \binom{md}{n} \right],
\end{aligned}
$$

so for some constants c, c':

$$
(md+D-n)M - nD\binom{md}{n} \in [0, cm^{n-1}], \qquad \frac{(md+D-n)M - nD\binom{md}{n}}{nD(md+D-n)M} \in [0, c'/m]. \tag{10}
$$

Now choose b such that $v(b) = -\tilde{v}(g) - CD$, where $C = C(n,d)$ is as per Lemma 4.8. If $v(b) \leq 0$, then the left hand side of (9) decreases, and $\tilde{v}(g) \geq -CD$, so $\tilde{v}(a)$ in the right hand side changes by at most CD. And if $v(b) \geq 0$, then $\tilde{v}(g) \leq -CD$, and now the left hand side increases by at most $v(b)O(m^{-1})$, and $\tilde{v}(a)$ in the right hand side increases by $v(b)$. In either case, choosing γ large enough, we may reduce to the case where $\tilde{v}(g) = -CD$.

For the purpose of calculating the left hand side, we may adjoin to K new v-generic elements, so let $g_0 \in K[X]_D$ be v-generic over K_0, and let $h = g + g_0$. Then, by Lemma 4.8, we have $vg(\xi) = vh(\xi)$ for each

$\xi \in [F^{\partial n}]$, so $v(g\chi)(\xi) = v(h\chi)(\xi)$ for every polynomial $\chi \in K[X]_{md-n}$ and every $\xi \in [F^{\partial n}]$, and therefore $\theta^*(g\chi) = \theta^*(h\chi)$. We conclude that $\mathrm{vol}_{g\mathfrak{M}_{md-n}} \theta^* = \mathrm{vol}_{h\mathfrak{M}_{md-n}} \theta^*$, and since $[f_0 \wedge \ldots \wedge f_{n-1}] \subseteq [F^{\partial n}]$, also $\theta^{\wedge n} \wedge \widehat{g} = \theta^{\wedge n} \wedge \widehat{h}$. We may therefore replace g with h, and assume that g is also weakly v-generic (over the prime field).

We have thus reduced to the hypotheses Lemma 4.7. In particular, Fg is a good family, and let us fix a good set $\psi \subseteq [(Fg)^{\partial n}]$. Then

$$\frac{\det \Phi_{Fg,\psi}}{\det \Phi_{F,\psi}} = \prod_{H \in \binom{F}{n}} g \wedge H^{\wedge} \prod_{G \subseteq F, \, |G| \leq n-1} \widehat{\mathfrak{d}}_{Gg,\psi}.$$

As in the proof of Lemma 4.9 we have $v(\widehat{\mathfrak{d}}_{Gg,\psi}) = O(1)$, so

$$\begin{aligned}
v(\det \Phi_{Fg,\psi} / \det \Phi_{F,\psi_F}) &= \sum_{H \in \binom{F}{n}} v(g \wedge H^{\wedge}) + O(m^{n-1}) \\
&= Dd^n \binom{m}{n} \theta^{\wedge n} \wedge \widehat{g} + O(m^{n-1}) \\
&= D \binom{md}{n} \theta^{\wedge n} \wedge \widehat{g} + O(m^{n-1}) \\
&= (md + D - n)M \frac{\theta^{\wedge n} \wedge \widehat{g}}{n} + O(m^{n-1}).
\end{aligned}$$

Notice that $g\Phi_{F,\psi} \subseteq \Phi_{Fg,\psi}$, so let $\Phi' = \Phi_{Fg,\psi} \smallsetminus g\Phi_{F,\psi} = (\varphi_{Fg,\psi,\xi} : \xi \in \psi \smallsetminus \psi_F)$, generating $E \subseteq L[X]_{md+D-n}$. By Lemma 4.7, the three valuations θ^*_{md+D-n}, $\theta^*_{md+D-n}/gL[X]_{md-n}$ and $\tilde{v}$ (with respect to the basis Φ') agree on E up to $O(1)$. It follows that

$$\mathrm{vol}_{\Phi_{Fg,\psi}} \theta^* - \mathrm{vol}_{g\Phi_{F,\psi}} \theta^* = \mathrm{vol}_{\Phi'}(\theta^*/gL[X]_{md-n}) = O(M) = O(m^{n-1}).$$

We conclude that

$$\begin{aligned}
\mathrm{vol}_{\mathfrak{M}_{md+D-n}} \theta^* - \mathrm{vol}_{g\mathfrak{M}_{md-n}} \theta^* &= \mathrm{vol}_{\Phi_{Fg,\psi}} \theta^* - \mathrm{vol}_{g\Phi_{F,\psi}} \theta^* - v(\det \Phi_{Fg,\psi} / \det \Phi_{F,\psi}) \\
&= -(md + D - n)M \frac{\theta^{\wedge n} \wedge \widehat{g}}{n} + O(m^{n-1}).
\end{aligned}$$

Dividing by $(md + D - n)M$, we obtain (9). $\blacksquare$

5. GENERIC PROJECTIONS OF VARIETIES

We shall prove Theorem 4.3 by reducing the general case to that of a hypersurface, via a generic projection. A projective variety of dimension $\nu - 1$ can always be projected onto a hypersurface in $\mathbf{P}^{\nu}$, so let us study the behaviour of virtual divisors under such a projection.

Throughout, let $\nu \leq n$, $X = (X_0, \ldots, X_n)$ and $Y = (Y_0, \ldots, Y_\nu)$. As in Theorem 4.3, we let K_0 be a small algebraically closed field, and $K = K_0(T)^a$ for some infinite family of indeterminates T. Let $W \subseteq \mathbf{P}^n(K)$ be a projective variety of dimension $\nu - 1$ defined over K_0, with Chow form $\mathfrak{C}_W$.

For $i \leq \nu$, let $\mu_i = \sum_{j \leq m} T_{i,j} X_j$, where $T_{i,j}$ are distinct members of the family T. We are going to work in a large ambient algebraically closed field L containing K. For any polynomial $f \in L[Y]$ define $f_\mu(X) = f(\mu X) \in L[X]$. Define

$$P(Y) = \mathfrak{C}_W \left(\frac{\mu_i}{Y_i} - \frac{\mu_\nu}{Y_\nu} : i < \nu \right) \prod_{i \leq \nu} Y_i^{\deg W}. \tag{11}$$

Recall that by [Ben, Lemma 2.8], this is a polynomial in $K[Y]_{\deg W}$, invariant, up to sign, under permutations of Y. In addition, by [Ben, Fact 2.10(v)], for a point $\xi = [x] \in \mathbf{P}^n$ algebraically independent over K_0 from K, we have $P_\mu(x) = 0$ if and only if $\xi \in W$.

Lemma 5.1.   (i) *If $\xi = [x] \in W$, then $\mu x \neq 0$, so $\pi\xi = [\mu x]$ defines a map $\pi \colon W \to \mathbf{P}^\nu$.*
  (ii) *This map is generically injective, namely if $\xi \in W$ is generic over K, then the fibre of ξ is a singleton.*
  (iii) *Let $U = \pi W \subseteq \mathbf{P}^\nu$ denote the image. The polynomial P defined above is irreducible, and $V(P) = U \subseteq \mathbf{P}^\nu$. In particular, $\deg U = \deg W$ as well.*

(iv) *There exists a homogeneous polynomial $R \in K[Y]$ such that, for every homogeneous $f \in L[X]$ there exists $g \in L[Y]$ such that $g_\mu - f R_\mu \in I(W)$. In particular, π is injective outside $V(R_\mu)$. Moreover, there exists a tuple a in K such that, if L is equipped with a good valuation, then*

$$\tilde{v}(g) \geq \tilde{v}(R) + \tilde{v}(f) + \tilde{v}(a) \deg f.$$

(v) *For any homogeneous polynomials $f = (f_i : i < v)$ in $L[Z]$ we have*

$$f \wedge P = f_\mu \wedge \mathfrak{C}_W.$$

Proof. The intersection of $v + 1$ (or even just v) generic hyperplanes with W is empty, so $\xi = [x] \in W$ implies $\mu x \neq 0$, whence (i). Given $v = [y] \in \mathbf{P}^n$ such that $y_v \neq 0$, its fibre is the common zero-set of the forms

$$\lambda_i^v = \mu_i - \frac{y_i}{y_v} \mu_v, \qquad i < v.$$

Therefore $P(y) = 0$ if and only if $v \in U$. Since P is invariant (up to sign) under permutations, we may always assume that $y_v \neq 0$, so $U = V(P)$.

Consider now a point $\xi \in W$ generic over K, or equivalently, generic over K_0 and algebraically independent from K. The linear forms $\lambda^{\pi\xi}$ are then generic over $K_0(\xi)$, or even $K_0(\xi, \mu_v)$, modulo the constraint that they vanish at ξ, so $W \cap V(\lambda^{\pi\xi}) = \{\xi\}$, and the fibre is a singleton, proving (ii).

Since W is irreducible, so is U, so $P = P_0^r$ for some irreducible $P_0 \in K[Y]$. Intersecting U with $v - 1$ generic hyperplanes we obtain $\deg P_0$ distinct generic points. Pulling back to $\mathbf{P}^n$, we obtain the intersection of W with $v - 1$ generic hyperplanes, consisting of $\deg W$ distinct generic points. Since the fibre of a generic point of W is a singleton, we must have $\deg U = \deg P_0 = \deg W$, i.e., $r = 1$ and P is irreducible, proving (iii).

For (iv), let us consider the variety $\widehat{W} \subseteq \mathbf{P}^{v+n+1}$ consisting of all $[x, \mu x]$ such that $[x] \in W$. The map π factors naturally via $W \to \widehat{W} \to U$. Let $\xi = [x] \in W$ be again generic over K. The linear forms $(\lambda_i^{\pi\xi} : i < v - 1)$ are generic over K_0. Considering them as linear forms in X, Y, the intersection $\lambda_0^{\pi\xi} \wedge \ldots \wedge \lambda_{v-2}^{\pi\xi} \wedge \mathfrak{C}_{\widehat{W}}$ splits into distinct points, on of which is $\hat{\xi} = [x, \mu x]$. By [Ben, Fact 2.10(iv)], $K(\hat{\xi})$ is separable over $K_0(\lambda_i^{\pi\xi} : i < v - 1)$, and *a fortiori* over $K(\pi\xi)$. On the other hand, since π is generically injective, $K(\hat{\xi})$ is purely inseparable over $K(\pi\xi)$. We conclude that $\hat{\xi}$ is rational over $K(\pi\xi)$. We can therefore express each $x_j/\mu_v x$ as a zero-degree rational function over K in μx.

We also claim that for each $j \leq n$ there exists a homogeneous polynomial $X_j^d + \ldots \in K[X_j, Y]$ satisfied by $(x_j, \mu x)$. Indeed, if not, then $[1, 0]$ specialises $[x_j, \mu x]$ over K, so there exists $x' \neq 0$ such that $[x', 0]$ specialises $[x, \mu x]$ over K. In particular, $[x'] \in W$ and $\mu x' = 0$, contradicting (i). This proves our claim, which implies that for every $f(X) \in K[X]_d$ there exists $g \in K[X, Y]_d$ such that $f(X) - g(X, \mu X) \in I(W)$ and $\deg_X g \leq d_0$, where d_0 only depends on W. Together with the previous paragraph, we obtain (iv). The moreover part follows from the argument.

For (v), if $f^\wedge = 0$, i.e., $\dim V(f) > 0$, then $f \wedge P = 0 = f_\mu \wedge \mathfrak{C}_W$. We may therefore assume that $f^\wedge \neq 0$, in which case it splits into finitely many points. Then $f \wedge P = 0$ if and only if one of these points is in the image of W, i.e., if and only if the fibre over one of these points intersects W. But $V(f_\mu)$ is the union of these fibres, so $f \wedge P = 0$ if and only if $f_\mu \wedge \mathfrak{C}_W = 0$. Thinking of f as indeterminate polynomials, both $f \wedge P$ and $f_\mu \wedge \mathfrak{C}_W$ are polynomials in these indeterminates. They are irreducible by [Ben, Proposition 2.15], and have the same zeros. Therefore, they only differ by a scalar coefficient. To find this coefficient, consider first the case where $f_i = Y_i$, so $f_{i,\mu} = \mu_i$. Then

$$f \wedge P = (-1)^{vD} P \wedge Z_0 \wedge \ldots \wedge Z_{v-1} = P(0, \ldots, 0, 1) = \mathfrak{C}_W(\mu_i : i < v) = f_\mu \wedge \mathfrak{C}_W.$$

We obtain the general case, for arbitrary degrees, by specialising f_i to $Y_i^{\deg f_i}$. $\blacksquare$

Definition 5.2. Let L be equipped with a valuation, and let θ be a virtual divisor on $U(K)$, and define

$$\mu_0 = \min_i \hat{\mu}_i \restriction_{W(K)}, \qquad \pi^*\theta = \theta \circ \pi + \mu_0.$$

We call $\pi^*\theta$ the *pull-back* of θ to W.

Clearly, μ_0 is a virtual divisor (or degree one) on W, also defined by $\mu_0(\xi) = \tilde{v}(\mu x) - \tilde{v}(x)$. If $\theta = \hat{f}\restriction_{U(K)}$ where $f \in L[X]_d$, and $\xi = [x] \in W(K)$, then

$$\hat{f}_\mu(\xi) = \frac{v(f(\mu x))}{d} - \tilde{v}(x) = \hat{f}(\pi\xi) + \mu_0(\xi) = \theta(\pi\xi) + \mu_0(\xi).$$

Therefore $\pi^*\theta$ is also a virtual divisor on W, of the same degree as θ.

Lemma 5.3. *Assume that L is equipped with a valuation, and let θ be a virtual divisor on U.*

(i) *We have*

$$\theta^{\wedge\nu} \wedge \widehat{P} = (\pi^*\theta)^{\wedge\nu} \wedge \widehat{\mathfrak{C}}_W.$$

(ii) *If $u = \theta^*$ and $w = (\pi^*\theta)^*$, then $u(g) = w(g_\mu)$ for homogeneous $g \in K[U]$.*

Proof. Item (i) is just Lemma 5.1(v). For (ii), let $g \in L[U]_d$. Then

$$u(g) = d \inf_U (\hat{g} - \theta) = d \inf_W (\hat{g}_\mu - \mu_0 - \theta \circ \pi) = d \inf_W (\hat{g}_\mu - \pi^*\theta) = w(g_\mu).$$

∎

Lemma 5.4. *If the valuation on L is good, then $\mu_0 = 0$. In particular, $\pi^*\theta = \theta \circ \pi$ for any virtual divisor θ on U.*

Proof. It will suffice to prove that if $\xi = [x] \in W$ and $\tilde{v}(x) = 0$, then $\tilde{v}(\mu x) = 0$, in each of the two cases described in Definition 4.2. We let $\overline{\xi} = [\overline{x}] \in \mathbf{P}^n(\overline{K})$, so in the residue field we have $\operatorname{trdeg}_{\overline{K_0}} \overline{K_0}(\overline{\xi}) \le \nu - 1$. The linear forms μ_i are over the valuation ring, and descend to the residue field. If there is any $I \subseteq \nu + 1$ of size ν such that $(\mu_i : i \in I)$ are v-generic over K_0, then the residues $(\overline{\mu}_i : i \in I)$ are generic over $\overline{K_0}$, and so $\overline{\mu x} \ne 0$ as in the argument for Lemma 5.1(i).

In case (i), the entire family μ is v-generic. In case (ii), we assume that $K_0(T)$ is equipped with a P-adic valuation for some prime $P \in K_0[T]$. If some of the coefficients of some μ_i, say μ_ν appear in the polynomial P, then $\mu_{<\nu}$ are v-generic over K_0 (in fact, the valuation on $K_0(\mu_{<\nu})$ is trivial). Otherwise, again the entire family μ is v-generic over K_0, concluding the proof. ∎

Push virtual divisors from W to U is more delicate. Indeed, if η is a virtual divisor on W, one could imagine letting $u = \eta^*$ and defining a sub-valuation $\pi_* u(f) = u(f_\mu)$ on $K[U]$. However, it is not at all clear why $\pi_* u$ should be almost finitely generated. Similarly, there may be a (non-generic) fibre which is not a singleton, and on which η need not be constant, so η need not be a limit of $\pi^*\theta$. We therefore follow a somewhat longer route.

Lemma 5.5. *There exists a non-zero tuples a in K, which depends only on W, with the following property. For every d there exists D, which depends only on d and W, such that for any good valuation on K and any virtual divisor η of degree d on W, defined over K_0, there exists a virtual divisor θ of degree D on U, defined over K, satisfying:*

$$\sup \theta = \sup \eta, \qquad \pi^*\theta \ge \eta, \qquad w_D(\theta) \le w_d(\eta) - \tilde{v}(a),$$

$$\theta^{\wedge\nu} \wedge \widehat{P} \le \eta^{\wedge\nu} \wedge \widehat{\mathfrak{C}}_W + \frac{w_d(\eta) - \tilde{v}(a)}{d}.$$

Proof. Let a and $R \in K[Y]_e$ be as per Lemma 5.1(iv), and we may assume that $\tilde{v}(R) = 0$. Let $f \in L[X]_d$ such that $\eta = \hat{f}\restriction_{W(K_0)}$, chosen such that $K_0(f)$ is independent from K over K_0, in the sense of Section 1. Then the extension of η to $W(K)$ is $\hat{f}\restriction_{W(K)}$. We may assume that $\sup \eta = 0$ and $\tilde{v}(f) = -d w_d(\eta)$. Let $m \in \mathbf{N}$ be fixed later, and let $D = md + e$. Let $g \in L[Y]_D$ be as per Lemma 5.1(iv), so $g_\mu = f^m R_\mu$ on W and $\tilde{v}(g) \ge m\tilde{v}(f) + md\tilde{v}(a) = md\tilde{v}(a) - md w_d(\eta)$. Let $h_0 \in L[Y]_D$ be v-generic over $K(g)$, let $h = g + h_0$, and let $\theta = \hat{g}\restriction_{U(K)} \wedge 0 = \hat{h}\restriction_{U(K)}$. Then θ is a virtual divisor on U of degree D, and

$$w_D(\theta) \le -\tilde{v}(h)/D \le w_d(\eta) - \tilde{v}(a).$$

For $\xi \in W$ we have

$$\pi^*\theta(\xi) = 0 \wedge \frac{md\eta(\xi) + e\hat{R}(\pi\xi)}{D} \ge \frac{md}{D}\eta(\xi) \ge \eta(\xi).$$

For the last inequality, let $k < \nu$. Then

$$\widehat{R}_\mu \wedge \eta^{\wedge k} \wedge (\pi^*\theta)^{\wedge\nu-k-1} \wedge \widehat{\mathfrak{C}}_W \le \frac{\tilde{v}(R_\mu \wedge \mathfrak{C}_W)}{e \deg W},$$

$$\eta^{\wedge k+1} \wedge (\pi^*\theta)^{\wedge\nu-k-1} \wedge \widehat{\mathfrak{C}}_W \ge -(k+1)w_d(\eta) - (\nu - k - 1)w_D(\theta) \ge \nu\tilde{v}(a) - \nu w_d(\eta).$$

Since $\pi^*\theta = \hat{h}_\mu\restriction_{W(K)} \le \hat{g}_\mu\restriction_{W(K)}$:

$$\eta^{\wedge k} \wedge (\pi^*\theta)^{\wedge\nu-k} \wedge \widehat{\mathfrak{C}}_W \le \eta^{\wedge k} \wedge (\pi^*\theta)^{\wedge\nu-k-1} \wedge \hat{g}_\mu \wedge \widehat{\mathfrak{C}}_W$$

$$= \frac{md}{D}\eta^{\wedge k+1} \wedge (\pi^*\theta)^{\wedge\nu-k-1} \wedge \widehat{\mathfrak{C}}_W + \frac{e}{D}\hat{R}_\mu \wedge \eta^{\wedge k} \wedge (\pi^*\theta)^{\wedge\nu-k-1} \wedge \widehat{\mathfrak{C}}_W$$

$$\le \eta^{\wedge k+1} \wedge (\pi^*\theta)^{\wedge\nu-k-1} \wedge \widehat{\mathfrak{C}}_W + \frac{e}{D}\left(\frac{\tilde{v}(R_\mu \wedge \mathfrak{C}_W)}{e \deg W} + \nu w_d(\eta) - \nu\tilde{v}(a)\right).$$

Therefore

$$\theta^{\wedge\nu} \wedge \widehat{P} = (\pi^*\theta)^{\wedge\nu} \wedge \widehat{\mathfrak{C}}_W \leq \eta^{\wedge\nu} \wedge \widehat{\mathfrak{C}}_W + \frac{\nu}{D \deg W}\tilde{v}(R_\mu \wedge \mathfrak{C}_W) + \frac{\nu^2 e}{D}\big(w_d(\eta) - \tilde{v}(a)\big).$$

Extending a with the inverse of some non-zero coefficient appearing in $R_\mu \wedge \mathfrak{C}_W$, and choosing m large enough, the last inequality follows. ∎

6. PROOF OF THE MAIN RESULT

We can now prove our main result, Theorem 4.3. We are given $d \in \mathbf{N}$ and $W \subseteq \mathbf{P}^n$ of dimension $\nu - 1$ defined over K_0, and we work in $K = K_0(T)^a$. We seek a tuple a in K, and for each $\gamma > 0$ some large enough m and a basis B_m for $K[W]_m$, such that for any good valuation on K (i.e., such that T is v-generic over K_0) and any η of degree d defined over K_0, we have

$$\frac{\mathrm{vol}_{B_m}\eta^*}{m \dim K[W]_m} \geq -\frac{\eta^{\wedge\nu} \wedge \widehat{\mathfrak{C}}_W}{\nu} + \gamma\big(\tilde{v}(a) - w_d(\eta)\big).$$

If $W = \mathbf{P}^n$, this is Lemma 4.9 (we may also consider $\mathbf{P}^n$ as a hyperplane in $\mathbf{P}^{n+1}$, so Lemma 4.9 is not strictly necessary). Otherwise, we let $\nu = \dim W + 1 \leq n$. Let $U \subseteq \mathbf{P}^\nu(K)$ be the generic projection of W, defined over K by a polynomial P, as in Section 5. Let D be as per Lemma 5.5, and let a' denote the corresponding tuple in K (called a there).

By Lemma 4.10, there is a tuple a'' in K, and for arbitrarily large m we have bases C_m for $K[U]_m$ such that for any valuation on K and virtual divisor θ on U of degree D defined over K_1, we have

$$\frac{\mathrm{vol}_{C_m}\theta^*}{m \dim K[U]_m} \geq -\frac{\theta^{\wedge\nu} \wedge \widehat{P}}{\nu} + \big(\tilde{v}(a'') - w_D(\theta)\big)O(m^{-1}).$$

We define B_m to consist of $(C_m)_\mu = \{h_\mu : h \in C_m\}$, completed with some monomials to a basis of $K[W]_m$, and let a consist of all products of members of a' and a'' (we may assume that 1 appears in each one).

Consider now a good valuation on K and η of degree d on W, and as usual we may assume that $\sup \eta = 0$. We choose θ on U of degree D as per Lemma 5.5. Notice that $\eta \leq 0$ implies that $\tilde{v} \leq \eta^*$, and $\pi^*\theta \geq \eta$ implies $(\pi^*\theta)^* \leq \eta^*$. Therefore

$$\mathrm{vol}_{C_m}\theta^* = \mathrm{vol}_{(C_m)_\mu}(\pi^*\theta)^* \leq \mathrm{vol}_{(C_m)_\mu}\eta^* \leq \mathrm{vol}_{B_m}\eta^*.$$

We therefore have, for some constant δ:

$$\begin{aligned}
\frac{\mathrm{vol}_{B_m}\eta^*}{m \dim K[U]_m} &\geq \frac{\mathrm{vol}_{C_m}\theta^*}{m \dim K[U]_m} \\
&\geq -\frac{\theta^{\wedge\nu} \wedge \widehat{P}}{\nu} + \frac{\delta}{m}\big(\tilde{v}(a'') - w_D(\theta)\big) \\
&\geq -\frac{\eta^{\wedge\nu} \wedge \widehat{\mathfrak{C}}_W}{\nu} + \frac{\tilde{v}(a') - w_d(\eta)}{d\nu} + \frac{\delta}{m}\big(\tilde{v}(a'') + \tilde{v}(a') - w_d(\eta)\big) \\
&\geq -\frac{\eta^{\wedge\nu} \wedge \widehat{\mathfrak{C}}_W}{\nu} + \left(\frac{1}{d\nu} + \frac{\delta}{m}\right)\big(\tilde{v}(a) - w_d(\eta)\big).
\end{aligned}$$

We may always replace d with a multiple, and therefore assume it is large enough, and then choosing m large enough we have:

$$\frac{\mathrm{vol}_{B_m}\eta^*}{m \dim K[U]_m} \geq -\frac{\eta^{\wedge\nu} \wedge \widehat{\mathfrak{C}}_W}{\nu} + \frac{\gamma}{2}\big(\tilde{v}(a) - w_d(\eta)\big).$$

Finally, observe that $\dim K[U]_m \leq \dim K[W]_m$ are both $\deg W\binom{m}{\nu-1} + O(m^{\nu-2})$, and since $\eta \geq -w_d(\eta)$, we have $-\frac{\eta^{\wedge\nu} \wedge \widehat{\mathfrak{C}}_W}{\nu} \leq w_d(\eta)$. Therefore, for large enough m:

$$\frac{\mathrm{vol}_{B_m}\eta^*}{m \dim K[W]_m} \geq -\frac{\eta^{\wedge\nu} \wedge \widehat{\mathfrak{C}}_W}{\nu} + \gamma\big(\tilde{v}(a) - w_d(\eta)\big),$$

concluding the proof.

REFERENCES

[BBHU08] Itaï BEN YAACOV, Alexander BERENSTEIN, C. Ward HENSON, and Alexander USVYATSOV, *Model theory for metric structures*, Model theory with applications to algebra and analysis. Vol. 2, London Math. Soc. Lecture Note Ser., vol. 350, Cambridge Univ. Press, Cambridge, 2008, pp. 315–427, `doi:10.1017/CBO9780511735219.011`.

[Ben]   Itaï BEN YAACOV, *The Vandermonde determinant identity in higher dimension*, research notes, `arXiv:1405.0993`.

[Ben14]   _____, *Model theoretic properties of metric valued fields*, Journal of Symbolic Logic **79** (2014), no. 3, 655–675, `doi:10.1017/jsl.2014.16`, `arXiv:0907.4560`.

[Ben15]   _____, *Tensor products of valued fields*, Bulletin of the London Mathematical Society **47** (2015), no. 1, 42–46, `doi:10.1112/blms/bdu092`, `arXiv:1302.1381`.

[Ber71]   George M. BERGMAN, *A weak Nullstellensatz for valuations*, Proceedings of the American Mathematical Society **28** (1971), 32–38, `doi:10.2307/2037749`.

[BU10]   Itaï BEN YAACOV and Alexander USVYATSOV, *Continuous first order logic and local stability*, Transactions of the American Mathematical Society **362** (2010), no. 10, 5213–5259, `doi:10.1090/S0002-9947-10-04837-3`, `arXiv:0801.4303`.

[Pil96]   Anand PILLAY, *Geometric stability theory*, Oxford Logic Guides, vol. 32, The Clarendon Press Oxford University Press, New York, 1996, Oxford Science Publications.

[Poi13]   Jérôme POINEAU, *Les espaces de Berkovich sont angéliques*, Bulletin de la Société Mathématique de France **141** (2013), no. 2, 267–297, `arXiv:1105.0250`.

ITAÏ BEN YAACOV, UNIVERSITÉ DE LYON, UNIVERSITÉ CLAUDE BERNARD – LYON 1, INSTITUT CAMILLE JORDAN, CNRS UMR 5208, 43 BOULEVARD DU 11 NOVEMBRE 1918, 69622 VILLEURBANNE CEDEX, FRANCE

URL: `http://math.univ-lyon1.fr/~begnac/`