



HAL
open science

NNLO positivity bounds on chiral perturbation theory for a general number of flavours

Benjamin Alvarez, Johan Bijnens, Mattias Sjö

► **To cite this version:**

Benjamin Alvarez, Johan Bijnens, Mattias Sjö. NNLO positivity bounds on chiral perturbation theory for a general number of flavours. *Journal of High Energy Physics*, 2022, 2022, pp.159. 10.1007/JHEP03(2022)159 . hal-03500376

HAL Id: hal-03500376

<https://hal.science/hal-03500376>

Submitted on 12 Sep 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

RECEIVED: December 10, 2021

REVISED: March 1, 2022

ACCEPTED: March 1, 2022

PUBLISHED: March 24, 2022

NNLO positivity bounds on chiral perturbation theory for a general number of flavours

Benjamin Alvarez,^a Johan Bijnens^b and Mattias Sjö^b

^a*Aix Marseille Université, Université de Toulon, CNRS, CPT, Marseille, France*

^b*Department of Astronomy and Theoretical Physics, Lund University, Box 43, SE 22100 Lund, Sweden*

E-mail: benjamin-alvarez@univ-tln.fr, bijnens@thep.lu.se, mattias.sjo@thep.lu.se

ABSTRACT: We present positivity bounds, derived from the principles of analyticity, unitarity and crossing symmetry, that constrain the low-energy constants of chiral perturbation theory. Bounds are produced for 2, 3 or more flavours in meson-meson scattering with equal meson masses, up to and including next-to-next-to-leading order (NNLO), using the second and higher derivatives of the amplitude. We enhance the bounds by using the most general isospin combinations possible (or higher-flavour counterparts thereof) and by analytically integrating the low-energy range of the discontinuities. In addition, we present a powerful and general mathematical framework for efficiently managing large numbers of positivity bounds.

KEYWORDS: Chiral Lagrangians, Effective Field Theories

ARXIV EPRINT: [2112.04253](https://arxiv.org/abs/2112.04253)

Contents

1	Introduction	1
2	Chiral perturbation theory	2
2.1	The χ PT Lagrangian	3
3	Scattering amplitudes	4
3.1	Other forms of the amplitude	4
3.2	Structure of the amplitude	5
3.3	Irreducible amplitudes	6
3.4	Eigenstate amplitudes	7
3.5	Crossing symmetry	8
4	Linear constraints	8
4.1	Definition and combination of constraints	9
4.2	Stronger and weaker constraints	9
4.3	Determining the relationship between constraints	10
4.4	Representations and degeneracy	12
5	Positivity bounds	13
5.1	Conditions on a_J	15
5.2	The number of derivatives	16
5.3	The value of t	17
5.4	Independently bounded parameters	17
5.4.1	General number of flavours	18
5.4.2	Two flavours	19
5.4.3	Three flavours	19
5.4.4	The full parameter space	19
5.5	The absence of catastrophic divergences	19
5.6	Integrals above threshold	21
6	Results	22
6.1	Two flavours	22
6.2	Three flavours	25
6.3	Higher number of flavours	28
6.4	Considerations about the integrals	31
6.5	Considerations about a_J	34
7	Conclusions and outlook	38
A	LEC details	40

B	Details and proofs regarding linear constraints	42
B.1	Proof of propositions 4.1 and 4.2	42
B.1.1	The trivial half of the proof	42
B.1.2	Proof that $\mathcal{A}(\omega_0)$ is closed and convex	43
B.1.3	Proof of proposition 4.1 in the $c = 0$ case	44
B.1.4	Proof of proposition 4.2	44
B.1.5	An important corollary	45
B.2	Some mathematical tools	46
B.2.1	The degenerate constraint framework	46
B.2.2	K -faces	49
B.3	Proof and generalisation of proposition 4.3	52
B.3.1	The treatment of degenerate constraints	54
B.3.2	An important corollary	56
B.4	Practical construction of $\mathcal{A}_c(\Omega)$ and $\mathcal{R}(\Omega)$	57
B.4.1	Construction of $\mathcal{A}_c(\Omega)$	57
B.4.2	Proof of lemma B.10	61
B.4.3	Some important corollaries of proposition B.2	64
B.4.4	Construction of $\mathcal{R}(\Omega)$	65
B.4.5	Visualisation of $\mathcal{A}_c(\Omega)$	68
B.5	The duality between $\mathcal{A}_c(\Omega)$ and $\mathcal{B}(\Omega)$	70
B.5.1	Duality with a bounding box	72
B.6	A note on infinite sums of constraints	72
B.7	Mathematical glossary	74
C	The loop integral functions	75
C.1	Expanding around $s = 0$	76
C.2	Expanding around $s = 4$	77
C.3	Integrals above threshold	78
C.3.1	The integral $Z_p^{m,n}$ for $p = 0$	79
C.3.2	Reduction of $Z_p^{m,n}$ to $Z_0^{m,n}$ and $Z_1^{m,n}$	79
C.3.3	Special cases for $n = 0$	80
C.3.4	The treatment of $(z - v)$	81
C.3.5	The completed integral	82

1 Introduction

Chiral perturbation theory (χ PT) is the most widespread theory for low-energy quantum chromodynamics (QCD). It is an effective field theory (EFT) which reformulates the non-perturbative behaviour of low-energy QCD as a perturbative theory of new degrees of freedom, physically interpreted as bound states of quarks. When constructed using n light quark flavours, the degrees of freedom are the $n^2 - 1$ light pseudoscalar mesons: the pions for $n = 2$, with the kaons and eta added for $n = 3$. χ PT was developed by Gasser & Leutwyler [1, 2] based on earlier work by Weinberg [3]; see [4, 5] for modern introductions with further references.

At leading order in the low-energy expansion, the only parameters of χ PT are the meson mass and decay constant, but higher orders introduce a rapidly increasing number of Wilson coefficients or low-energy constants (LECs) which, while in principle derivable from the underlying QCD dynamics, must in practice be seen as unknowns. At next-to-leading order (NLO), the LECs can be measured reasonably well with experimental or lattice methods, although the precision is typically only one or two significant digits. At next-to-next-to-leading order (NNLO), only tentative results are presently available. For a review of LEC measurements, see [6].

All quantum field theories must obey the axioms of unitarity, analyticity and crossing symmetry, and normally do so by construction. However, it turns out that these axioms are not automatically satisfied by EFTs such as χ PT when perturbativity is assumed at a fixed order in the expansion. Therefore, imposing the axioms actually adds new information, typically by placing bounds on the scattering amplitudes. Pioneering work was done by Martin [7] before the development of χ PT as such. Bounds on NLO two-flavour χ PT amplitudes, which in turn translate to bounds on the LECs, were first obtained in [8–10] and extended in [11, 12]. Further improvements were made in [13] and extended to three-flavour χ PT in [14]. There is ongoing research in extending these methods, both specific to χ PT and with broader scope; recent examples include [15–18].

The method of [13, 14], which serves as the basis of our method, is to apply dispersion relations (a consequence of analyticity) to a meson-meson scattering amplitude decomposed into isospin components (for higher flavours, the Clebsch-Gordan decomposition is used). Then, crossing symmetry and the optical theorem (a consequence of unitarity) are applied to give a positivity condition on the decomposed amplitude. With the amplitude calculated in terms of the LECs to some order, this results in bounds on linear combinations of LECs. More recently, stronger bounds have been obtained in [19, 20] by improving this method; put extremely simply, this was done with more sophisticated use of dispersion relations and crossing symmetry, respectively. Put similarly simply, our work instead improves the handling of the isospin decompositions and the LEC bounds themselves, although some improvements similar to [19] are also made. Perhaps more importantly, we perform the first extension to NNLO χ PT with any number of flavours (two flavours was treated in [19]), albeit with the simplification that all mesons have the same mass. The LECs are independent of the chosen masses, although the bounds do depend on the mass. At NLO they depend only on the ratio of the meson mass and the subtraction scale μ , at NNLO also on the ratio of the meson mass and decay constant.

Preliminary results of this work are presented in the Lund University master thesis [21]. Our work is structured as follows: section 2 introduces χ PT and its LECs; section 3 (backed by appendix A) presents the $2 \rightarrow 2$ meson scattering amplitude used to obtain the bounds; section 4 (backed by appendix B) introduces the mathematical framework used to manage them; section 5 (backed by appendix C) presents the method of [13, 14] and the improvements made to it; and section 6 displays the most interesting bounds we obtain, with final remarks given in section 7.

2 Chiral perturbation theory

n -flavour χ PT is based around a non-linear sigma model (NLSM), whose degrees of freedom are the $n^2 - 1$ Nambu-Goldstone bosons that arise when the chiral symmetry $G = \text{SU}(n)_L \times \text{SU}(n)_R$ of n -flavour massless QCD is spontaneously broken into its diagonal subgroup $H = \text{SU}(n)_V$. The Goldstone bosons live in the coset space G/H , which is isomorphic to $\text{SU}(n)$.

The presence of quark masses, electroweak interactions, etc. can be accounted for by including four external $n \times n$ flavour-space matrix fields — s (scalar), p (pseudoscalar), v_μ (vector) and a_μ (axial vector)¹ — into the massless QCD Lagrangian. These additions were introduced in [1, 2], and endow the Nambu-Goldstone bosons with masses and interactions that allow them to accurately model the light pseudoscalar mesons, turning the $\text{SU}(n)$ NLSM into χ PT proper.

The Nambu-Goldstone boson fields can be organised into a $n \times n$ flavour-space matrix field $u(\phi)$ [24, 25]. Under the chiral transformation $(g_L, g_R) \in G$, $u(\phi)$ transforms as

$$u(\phi) \longrightarrow g_R u(\phi) h[g_L, g_R, u(\phi)] = h[g_L, g_R, u(\phi)] u(\phi) g_L^\dagger, \quad (2.1)$$

where $h \in H$ is defined by this transformation. By requiring that G can be made local while leaving the extended QCD Lagrangian invariant, it can be shown that

$$\begin{aligned} \chi &\equiv 2B(s + ip) \longrightarrow g_R \chi g_L^\dagger, \\ \ell_\mu &\equiv v_\mu - a_\mu \longrightarrow g_L \ell_\mu g_L^\dagger - i \partial_\mu g_L g_L^\dagger, \\ r_\mu &\equiv v_\mu + a_\mu \longrightarrow g_R r_\mu g_R^\dagger - i \partial_\mu g_R g_R^\dagger, \end{aligned} \quad (2.2)$$

where B is a constant related to the leading-order (LO) meson decay constant and the $\langle \bar{q}q \rangle$ condensate.

It is possible to rewrite $u(\phi), \chi, \ell_\mu, r_\mu$ in a basis of fields that transform entirely in terms of g_L and g_R , as is done in [2] to derive the NLO χ PT Lagrangian. We instead choose to follow [26–28] and rewrite them in a basis of fields that all transform as $X \rightarrow hXh^\dagger$:

$$\begin{aligned} u_\mu &\equiv i \left[u^\dagger (\partial_\mu - ir_\mu) u - u (\partial_\mu - i\ell_\mu) u^\dagger \right], \\ \chi_\pm &\equiv u^\dagger \chi u^\dagger \pm u \chi^\dagger u, \\ f_\pm^{\mu\nu} &\equiv u F_L^{\mu\nu} u^\dagger \pm u^\dagger F_R^{\mu\nu} u, \end{aligned} \quad (2.3)$$

¹One can add more types of external fields to χ PT. Examples are symmetric or antisymmetric tensors [22, 23]. These extensions are not relevant for this work.

where $F_L^{\mu\nu} \equiv \partial^\mu \ell^\nu - \partial^\nu \ell^\mu - i[\ell^\mu, \ell^\nu]$ and similarly for $F_R^{\mu\nu}$ and r^μ . These transformation properties are conserved under the covariant derivative ∇_μ defined as

$$\nabla_\mu X = \partial_\mu X + [\Gamma_\mu, X], \quad \Gamma_\mu \equiv \frac{1}{2} \left[u^\dagger (\partial_\mu - i r_\mu) u + u (\partial_\mu - i \ell_\mu) u^\dagger \right]. \quad (2.4)$$

2.1 The χ PT Lagrangian

There exists an infinite number of possible Lagrangian terms consistent with the symmetries of χ PT. They can be organised into a power-counting hierarchy in the small energy-momentum scale p , where $u_\mu, \nabla_\mu = \mathcal{O}(p)$ and $\chi_\pm, f_\pm^{\mu\nu} = \mathcal{O}(p^2)$. Thus,

$$\mathcal{L}_{\chi\text{PT}} = \mathcal{L}_2 + \mathcal{L}_4 + \mathcal{L}_6 + \dots, \quad (2.5)$$

where \mathcal{L}_{2n} is $\mathcal{O}(p^{2n})$; odd powers are forbidden by parity. The coefficient of each term in \mathcal{L}_{2n} is a separate LEC.²

The LO Lagrangian is

$$\mathcal{L}_2 = \frac{F^2}{4} \langle u_\mu u^\mu + \chi_+ \rangle, \quad (2.6)$$

where F is a LEC related to the LO meson decay constant, and $\langle \dots \rangle$ indicates a trace over flavour-space indices. The LEC of the χ_+ term is $\frac{BF^2}{4}$ as defined in eq. (2.2). By requiring that the kinetic term is canonically normalised, one can fix $u(\phi) = 1 + \frac{it^a \phi^a}{F\sqrt{2}} + \dots$, where t^a are the generators of $SU(n)$ and Einstein's summation convention is used. The higher-order terms depend on the choice of parametrisation, which influences the computation of amplitudes but not the amplitudes themselves.

The next-to-leading-order (NLO) Lagrangian, which was first determined in [2], is in terms of our basis³

$$\begin{aligned} \mathcal{L}_4 = & \hat{L}_0 \langle u_\mu u_\nu u^\mu u^\nu \rangle + \hat{L}_1 \langle u_\mu u^\mu \rangle^2 + \hat{L}_2 \langle u_\mu u_\nu \rangle \langle u^\mu u^\nu \rangle + \hat{L}_3 \langle (u_\mu u^\mu)^2 \rangle \\ & + \hat{L}_4 \langle u_\mu u^\mu \rangle \langle \chi_+ \rangle + \hat{L}_5 \langle u_\mu u^\mu \chi_+ \rangle + \hat{L}_6 \langle \chi_+ \rangle^2 + \hat{L}_7 \langle \chi_- \rangle^2 + \hat{L}_8 \langle \chi_+^2 + \chi_-^2 \rangle \\ & - i \hat{L}_9 \langle f_+^{\mu\nu} u_\mu u_\nu \rangle + \hat{L}_{10} \langle f_+^{\mu\nu} f_{\mu\nu}^+ - f_-^{\mu\nu} f_{\mu\nu}^- \rangle, \end{aligned} \quad (2.7)$$

where the LECs are \hat{L}_i . The analogous NNLO Lagrangian with 112 LECs K_i was determined in [27]. The 1862-LEC NNNLO Lagrangian, which we do not use here, was determined in [28].

For small n , the Cayley-Hamilton identity reduces the number of independent terms, and consequently the number of LECs. At $n = 3$, it is standard to eliminate \hat{L}_0 ; the remaining LECs are conventionally labelled L_i with i preserved. At $n = 2$, it is customary to also redefine the LECs slightly, resulting in the l_i of the original Gasser-Leutwyler convention [1]. At NNLO, the 112+3 K_i (ordinary+contact terms) are reduced to 90+4 C_i at $n = 3$ and 52+4 c_i at $n = 2$ as detailed in [27]. For more details on the Lagrangians for different n , see [6, 29].

²Some ‘‘terms’’, like the one associated with \hat{L}_{10} in eq. (2.7) below, actually consist of several terms. These transform into each other under the discrete symmetries of the Lagrangian, and must therefore appear with the same LEC.

³There are two additional *contact terms* proportional to $\langle \chi_+^2 - \chi_-^2 \rangle$ and $\langle f_+^{\mu\nu} f_{\mu\nu}^+ + f_-^{\mu\nu} f_{\mu\nu}^- \rangle$. They are needed for renormalisation but make no physical contributions to the amplitudes considered here.

The NLO renormalisation was first carried out in [1, 2], and the extension to NNLO in [29]; for more information on χ PT renormalisation, see [30]. A slightly altered $\overline{\text{MS}}$ scheme is conventionally used, with renormalisation scale $\mu = 0.77 \text{ GeV}$. The renormalised LECs are denoted X_i^r where $X = \ell, L, \hat{L}$, etc. At $n = 2$ flavours it is conventional to use $\bar{\ell}_i$ instead, related to ℓ_i^r through

$$\ell_i^r = \frac{\gamma_i}{32\pi^2} \left[\bar{\ell}_i + \ln \left(\frac{M_{\text{phys}}^2}{\mu^2} \right) \right] \quad (2.8)$$

where M_{phys} is the chosen meson mass and γ_i are coefficients found in [1]. Effectively, eq. (2.8) sets the renormalisation scale to M_{phys} for $\bar{\ell}_i$.

3 Scattering amplitudes

In this section, and in the remainder of the paper, we will restrict ourselves to a simplified version of χ PT. Firstly, we will not include the external (axial) vector fields a_μ, v_μ in the Lagrangian, which essentially amounts to ignoring electroweak corrections to the amplitude. Secondly, we will assume that all mesons have the same mass M_{phys} , as mentioned in the introduction. While this limits the phenomenological applicability of three-flavour χ PT, it is a reasonable approximation that simplifies the procedure for obtaining bounds (see section 5). More importantly, the full NNLO amplitude is currently not available in the general-mass case; available results only cover $\pi\pi$ scattering in two- [30, 31] and three-flavour [32] χ PT, as well as πK scattering [33], and are not expressed in terms of elementary functions. With equal masses, we normalise all Mandelstam variables so that $s + t + u = 4$.

For the general equal-mass n -flavour scattering process $a + b \rightarrow c + d$, there are nine independent flavour structures possible: the six distinct index permutations on $\langle t^a t^b t^c t^d \rangle$ and the three on $\langle t^a t^b \rangle \langle t^c t^d \rangle$. Due to charge conjugation symmetry, a permutation is not independent of its reverse. Thus, the scattering amplitude M may be decomposed as

$$\begin{aligned} M(s, t, u) = & \left[\langle t^a t^b t^c t^d \rangle + \langle t^d t^c t^b t^a \rangle \right] B(s, t, u) \\ & + \left[\langle t^a t^c t^d t^b \rangle + \langle t^b t^d t^c t^a \rangle \right] B(t, u, s) \\ & + \left[\langle t^a t^d t^b t^c \rangle + \langle t^c t^b t^d t^a \rangle \right] B(u, s, t) \\ & + \delta^{ab} \delta^{cd} C(s, t, u) + \delta^{ac} \delta^{bd} C(t, u, s) + \delta^{ad} \delta^{bc} C(u, s, t), \end{aligned} \quad (3.1)$$

where s, t, u are the normalised Mandelstam variables, and crossing symmetry imposes that only two distinct functions B, C are used.⁴ This is the form used in [34], where the functions B, C are given to NNLO for $\text{SU}(n)$ equal-mass χ PT. The NLO results were first obtained in [1, 35].

3.1 Other forms of the amplitude

With two flavours, the traces can be evaluated in terms of Kronecker δ 's, giving⁵

$$M(s, t, u) = \delta^{ab} \delta^{cd} A(s, t, u) + \delta^{ac} \delta^{bd} A(t, u, s) + \delta^{ad} \delta^{bc} A(u, s, t), \quad (3.2)$$

⁴These functions have the symmetries $B(s, t, u) = B(u, t, s)$ and $C(s, t, u) = C(s, u, t)$, which is consistent with the symmetries of the respective flavour structures. Likewise, $A(s, t, u) = A(s, u, t)$ holds in eq. (3.2).

⁵This form can be traced back to the original current-algebra calculation [36] of the $\pi\pi$ amplitude.

which is the form used in [13] (up to reordering the arguments as permitted by the symmetries of A). In terms of the functions above,

$$A(s, t, u) = C(s, t, u) + B(s, t, u) + B(t, u, s) - B(u, s, t), \quad (3.3)$$

the function A was first determined to NLO in [1].

With n flavours, the traces can be evaluated using the anticommutation relation $\{t^a, t^b\} = \frac{2}{n}\delta^{ab} + d^{abc}t^c$ to give⁶

$$M(s, t, u) = d^{abe}d^{cde}B'(s, t, u) + d^{ace}d^{bde}B'(t, u, s) + d^{ade}d^{bce}B'(u, s, t) \\ + \delta^{ab}\delta^{cd}C'(s, t, u) + \delta^{ac}\delta^{bd}C'(t, u, s) + \delta^{ad}\delta^{bc}C'(u, s, t), \quad (3.4)$$

where

$$B'(s, t, u) = \frac{1}{2}[B(s, t, u) + B(t, u, s) - B(u, s, t)], \quad (3.5) \\ C'(s, t, u) = C(s, t, u) + \frac{4}{n}B'(s, t, u).$$

With three flavours, the Cayley-Hamilton theorem⁷ allows for the removal of one term at the expense of symmetry, leaving

$$M(s, t, u) = \delta^{ab}\delta^{cd}A_1(s, t, u) + \delta^{ac}\delta^{bd}A_2(s, t, u) + \delta^{ad}\delta^{bc}A_3(s, t, u) \\ + d^{abe}d^{cde}B_1(s, t, u) + d^{ace}d^{bde}B_2(s, t, u) \quad (3.6)$$

where

$$B_1(s, t, u) = B(t, u, s) - B(u, s, t), \quad B_2(s, t, u) = B(t, u, s) - B(s, t, u), \\ A_1(s, t, u) = C(s, t, u) + B(s, t, u) + \frac{1}{3}B_1(s, t, u), \quad (3.7) \\ A_2(s, t, u) = C(t, u, s) + B(u, s, t) + \frac{1}{3}B_2(s, t, u), \\ A_3(s, t, u) = C(u, s, t) + B(s, t, u) + B(u, s, t) - B(t, u, s).$$

This is the form used in [14].

3.2 Structure of the amplitude

The functions $B(s, t, u)$ and $C(s, t, u)$ consist of one part that is polynomial in the Mandelstam variables and contains the LECs, plus the so-called *unitarity correction* that is non-polynomial in the Mandelstam variables.⁸ The polynomial parts are quadratic at NLO and cubic at NNLO. At NLO, the unitarity correction does not contain any LECs; at NNLO, the unitarity correction depends on the NLO LECs.

⁶The relevant identity is

$$\langle t^a t^b t^c t^d \rangle + \langle t^a t^d t^c t^b \rangle = \frac{1}{2}(d^{abe}d^{cde} + d^{ade}d^{bce} - d^{ace}d^{bde}) + \frac{2}{n}(\delta^{ab}\delta^{cd} + \delta^{ad}\delta^{cb} - \delta^{ac}\delta^{bd}).$$

It is most easily derived by first using $t^a t^b = \frac{1}{n}\delta^{ab} + \frac{1}{2}(d^{abc} + if^{abc})t^c$ repeatedly, and then removing all occurrences of f with the Jacobi-like identity

$$f^{abe}f^{cde} = d^{ace}d^{bde} - d^{bce}d^{ade} + \frac{4}{n}(\delta^{ac}\delta^{bd} - \delta^{ad}\delta^{bc}),$$

which is derived from the observation that $[[t^a, t^b], t^c] = \{\{t^b, t^c\}, t^a\} - \{\{t^c, t^a\}, t^b\}$.

⁷More specifically the $n = 3$ Cayley-Hamilton theorem, recast as the SU(3)-specific identity

$$3(d^{abe}d^{cde} + d^{bce}d^{ade} + d^{cae}d^{bde}) = 2(\delta^{ab}\delta^{cd} + \delta^{bc}\delta^{ad} + \delta^{ca}\delta^{bd}).$$

⁸This split is not uniquely defined, but we adhere to the conventions of [34].

The unitarity correction at NLO depends on the function \bar{J} , which originates in the loop integral as shown in [1]. The NNLO unitarity correction introduces four analogous functions $k_i, i = 1, \dots, 4$ [31, 32, 34, 37]. More details about these functions can be found in appendix C.

The LEC content of the amplitude considered here is more limited than that of the full χ PT Lagrangian. About half of the Lagrangian terms are dropped by not including the external (axial) vector fields, and a significant part of the NNLO Lagrangian cannot appear in a 4-particle process below NNNLO. Also, the number of LECs is reduced by the Cayley-Hamilton theorem in the 2- and 3-flavour case as described in section 2.1. Lastly, $\hat{L}_7, K_{12}, K_{24}, K_{30}, K_{34}, K_{36}, K_{41}$ and K_{42} , i.e. those whose Lagrangian terms contain $\langle \chi_- \rangle$, disappear in the equal-mass limit.⁹ Even with these reductions, there are still 35 (27 at $n = 3$, 18 at $n = 2$) NNLO LECs that are involved in the amplitude at hand, in addition to 8 (7, 4) NLO LECs.

3.3 Irreducible amplitudes

The scattered particles are in the adjoint representation of $SU(n)$. The Clebsch-Gordan decomposition of the initial and final states is therefore¹⁰

$$\text{Adj} \otimes \text{Adj} = R_I + R_S + R_A + R_S^A + R_A^S + R_S^S + R_A^A, \quad (3.8)$$

where R_I is the singlet representation, and the sub(super)scripts on the other representations indicate lower (upper) index pairs that are symmetric (S) or antisymmetric (A). Details on the representations and their dimensions can be found in [34, 38]. From this, it follows that the scattering amplitude can be decomposed in terms of seven corresponding irreducible amplitudes T_J . In terms of eq. (3.1), these are

$$\begin{aligned} R_I : \quad & T_I = 2 \frac{n^2 - 1}{n} [B(s, t, u) + B(t, u, s)] - \frac{2}{n} B(u, s, t) \\ & \quad + (n^2 - 1)C(s, t, u) + C(t, u, s) + C(u, s, t), \\ R_S : \quad & T_S = \frac{n^2 - 4}{n} [B(s, t, u) + B(t, u, s)] - \frac{4}{n} B(u, s, t) \\ & \quad + C(t, u, s) + C(u, s, t), \\ R_A : \quad & T_A = n[B(t, u, s) - B(s, t, u)] + C(t, u, s) - C(u, s, t), \\ R_S^A, R_A^S : \quad & T_{AS} = T_{SA} = C(t, u, s) - C(u, s, t), \\ R_S^S : \quad & T_{SS} = 2B(u, s, t) + C(t, u, s) + C(u, s, t), \\ R_A^A : \quad & T_{AA} = -2B(u, s, t) + C(t, u, s) + C(u, s, t). \end{aligned} \quad (3.9)$$

⁹This can be understood by noting that χ_- has odd parity, so all terms in its expansion contain an odd number of pseudoscalar fields. If the even-parity Lagrangian term contains two traces of odd-parity objects, it can therefore only result in six-point vertices or larger, since the trace of a single field vanishes. Therefore, K_{12}, K_{24} etc. do not appear in the NNLO four-point amplitude, whereas \hat{L}_7 only appears in s, t, u -independent tadpole diagrams. As will be shown in section 5, we only consider s -derivatives of the amplitude, so also \hat{L}_7 disappears for our purposes.

¹⁰[35] contains an intuitive description of how the decomposition is performed.

Only six amplitudes are needed, since T_{SA} and T_{AS} are identical due to crossing symmetry etc., as mentioned in [14].

In $SU(3)$, the R_A^A representation vanishes, so only five amplitudes are needed. In [14], the representations are labelled by their dimensions, which are 1, 8, 8, 10 and 27 in the order they appear in eqs. (3.8) and (3.9).

In $SU(2)$, only R_I , R_A and R_S^S remain and have dimension 1, 3 and 5, respectively. The corresponding amplitudes can be identified with the isospin components T^0, T^1 and T^2 , respectively. In terms of eq. (3.2), they are

$$\begin{aligned} T^0 &= 3A(s, t, u) + A(t, u, s) + A(u, s, t), \\ T^1 &= A(t, u, s) - A(u, s, t), \\ T^2 &= A(t, u, s) + A(u, s, t). \end{aligned} \tag{3.10}$$

This well-known relation can be derived from eqs. (3.3) and (3.9).

3.4 Eigenstate amplitudes

A general amplitude can be expressed as $a_J T^J$, where the index J runs over the representations in the order they appear in eq. (3.8). For a physically applicable scattering process, however, the initial and final states should typically be taken as a product of mass eigenstates such as π, K and η . This corresponds to fixing a_J to a small selection of values so that $T(ab \rightarrow cd) = a_J(ab \rightarrow cd)T^J$. Here, as in [13, 14], we consider only elastic scattering of eigenstates, with $a_J(ab \rightarrow ab) \equiv a_J(ab)$.

With two flavours, where J runs over I, A, SS (alternatively, isospin 0, 1, 2), the eigenstates are¹¹

$$a_J(\pi^0 \pi^0) = \left(\frac{1}{3} \ 0 \ \frac{2}{3} \right), \quad a_J(\pi^0 \pi^\pm) = \left(0 \ \frac{1}{2} \ \frac{1}{2} \right), \quad a_J(\pi^\pm \pi^\pm) = \left(0 \ 0 \ 1 \right), \tag{3.11}$$

and with three flavours, where J runs over I, S, A, AS, SS , they are¹²

$$\begin{aligned} a_J(\pi^0 \pi^0) &= \left(\frac{1}{8} \ \frac{1}{5} \ 0 \ 0 \ \frac{27}{40} \right), & a_J(\pi^\pm \pi^\pm, K^\pm K^\pm, K^0 K^0) &= \left(0 \ 0 \ 0 \ 0 \ 1 \right), \\ a_J(\pi^0 \pi^\pm) &= \left(0 \ 0 \ \frac{1}{3} \ \frac{1}{6} \ \frac{1}{2} \right), & a_J(K^\pm \pi^\pm, K^\pm \pi^\mp, K^0 K^\pm) &= \left(0 \ 0 \ 0 \ \frac{1}{2} \ \frac{1}{2} \right), \\ a_J(K^0 \pi^\pm) &= \left(0 \ \frac{3}{10} \ \frac{1}{6} \ \frac{1}{3} \ \frac{1}{5} \right), & a_J(K \pi^0) &= \left(0 \ \frac{3}{20} \ \frac{1}{12} \ \frac{5}{12} \ \frac{7}{20} \right), \\ a_J(\pi \eta) &= \left(0 \ \frac{1}{5} \ 0 \ \frac{1}{2} \ \frac{3}{10} \right), & a_J(K \eta) &= \left(0 \ \frac{1}{20} \ \frac{1}{4} \ \frac{1}{4} \ \frac{9}{20} \right); \end{aligned} \tag{3.12}$$

see e.g. [13, 14], respectively.¹³ With four or more flavours, χ PT loses its applicability as low-energy QCD since there are only three light quarks in the Standard Model. Therefore, there is little sense in considering eigenstates for n flavours, although we can note that $T(\pi^\pm \pi^\pm \rightarrow \pi^\pm \pi^\pm) = T_{SS}$ regardless of n .

¹¹ a_J is invariant under particle/antiparticle exchange, so $a_J(\pi^+ \pi^+) = a_J(\pi^- \pi^-) \equiv a_J(\pi^\pm \pi^\pm)$. Note, however, that $a_J(\pi^\pm \pi^\mp) \neq a_J(\pi^\pm \pi^\pm)$ — they are instead related by crossing; see eq. (3.14).

¹²Here, π without superscript stands for any of π^\pm or π^0 (and similarly for K) whenever a_J is agnostic about the particular choice. We use $a_J(ab, cd, \dots)$ for $a_J(ab) = a_J(cd) = \dots$

¹³Eq. (3.12) differs from the values given in [14]: there was an error or misprint in $a_J(\pi^0 \pi^\pm)$, and all eigenstates were not included, with $a_J(K^\pm \pi^\pm)$ given as $a_J(K \pi)$.

One of our extensions over previous work is that we use all possible values for the a_J , rather than restricting them to eigenstates (see section 5 for what constitutes “possible”). This can be done without complications, since the mass eigenstates are completely degenerate in the equal-mass limit. However, it is still useful to view those states that remain mass eigenstates in the unequal-mass case as special. Below, by “eigenstate” we will specifically mean scattering between these states. Note that by treating general a_J , we effectively include inelastic scattering such as $a_J(\pi^0\pi^0 \rightarrow \pi^+\pi^-) = \left(\frac{1}{3} \ 0 \ -\frac{1}{3}\right)$. However, it turns out that inelastic scattering is useless for our purposes by invariably failing to satisfy eq. (5.6b). This (in addition to [13, 14]) is why this section has focused mainly on elastic scattering.

3.5 Crossing symmetry

Since all amplitudes can be expressed as $a_J T^J$, crossing symmetry implies that channel crossing must take the form of a linear transformation of a_J . For $s \leftrightarrow u$ crossing, the transformation $T^I(u, t, s) = C_u^{IJ} T^J(s, t, u)$ is given by [21, 38]

$$C_u^{IJ} = \begin{pmatrix} \frac{1}{n^2-1} & 1 & -1 & \frac{4-n^2}{2} & \frac{n^2(n+3)}{4(n+1)} & \frac{n^2(n-3)}{4(n-1)} \\ \frac{1}{n^2-1} & \frac{n^2-12}{2(n^2-4)} & -\frac{1}{2} & 1 & \frac{n^2(n+3)}{4(n+1)(n+2)} & \frac{n^2(3-n)}{4(n-1)(n-2)} \\ \frac{1}{1-n^2} & -\frac{1}{2} & \frac{1}{2} & 0 & \frac{n(n+3)}{4(n+1)} & \frac{n(3-n)}{4(n-1)} \\ \frac{1}{1-n^2} & \frac{2}{n^2-1} & 0 & \frac{1}{2} & \frac{n(n+3)}{4(n+1)(n+2)} & \frac{n(n-3)}{4(n-1)(n-2)} \\ \frac{1}{n^2-1} & \frac{1}{2+n} & \frac{1}{n} & \frac{n-2}{2n} & \frac{n^2+n+2}{4(n+1)(n+2)} & \frac{n-3}{4(n-1)} \\ \frac{1}{n^2-1} & \frac{1}{2-n} & -\frac{1}{n} & \frac{n+2}{2n} & \frac{n+3}{4(n+1)} & \frac{n^2-n+2}{4(n-1)(n-2)} \end{pmatrix} \quad (3.13)$$

which also works at $n = 2, 3$ by removing appropriate rows and columns:

$$C_u^{IJ}|_{\text{SU}(2)} = \frac{1}{6} \begin{pmatrix} 2 & -6 & 10 \\ -2 & 3 & 5 \\ 2 & 3 & 1 \end{pmatrix} \quad C_u^{IJ}|_{\text{SU}(3)} = \begin{pmatrix} \frac{1}{8} & 1 & -1 & -\frac{5}{2} & \frac{27}{8} \\ \frac{1}{8} & -\frac{3}{10} & -\frac{1}{2} & 1 & \frac{9}{8} \\ -\frac{1}{8} & -\frac{1}{2} & \frac{1}{2} & 0 & \frac{27}{40} \\ -\frac{1}{8} & \frac{2}{5} & 0 & \frac{1}{2} & \frac{9}{40} \\ \frac{1}{8} & \frac{1}{5} & \frac{1}{3} & \frac{1}{6} & \frac{7}{40} \end{pmatrix}. \quad (3.14)$$

These versions can be found in [13, 14] respectively.

4 Linear constraints

In this section, we will introduce a mathematical language of *linear constraints*. This formalism is introduced before positivity bounds (see section 5) so that they can be established in full generality. In order to make the handling of the bounds as general and powerful as possible, we dedicate this section to developing some useful mathematical definitions and results.¹⁴

¹⁴In this section employ mathematical notation that, depending on the background of the reader, may not be entirely familiar. We also define new notation for our own purposes. A glossary covering all potentially unfamiliar notation is provided in appendix B.7.

4.1 Definition and combination of constraints

For a set of parameters b_i (e.g. the LECs), a linear constraint takes the general form

$$\alpha_1 b_1 + \alpha_2 b_2 + \dots + \alpha_n b_n - c \geq 0, \tag{4.1}$$

where c, α_i are known coefficients. By treating α_i, b_i as components of vectors, this is equivalent to

$$\boldsymbol{\alpha} \cdot \mathbf{b} \geq c. \tag{4.2}$$

We say that \mathbf{b} lives in the *parameter space*, whereas $\boldsymbol{\alpha}$ lives in the *constraint space*.¹⁵ Since $\boldsymbol{\alpha}$ and c can be rescaled by any positive scalar without changing the inequality, any linear constraint can be described by the pair $\langle \boldsymbol{\alpha}, c \rangle$ with $c \in \{1, 0, -1\}$.

We say that a point \mathbf{b} *satisfies* a constraint $\langle \boldsymbol{\alpha}, c \rangle$ if $\boldsymbol{\alpha} \cdot \mathbf{b} \geq c$. We denote by $\mathcal{B}(\langle \boldsymbol{\alpha}, c \rangle)$ the subset of parameter space that satisfies $\langle \boldsymbol{\alpha}, c \rangle$. For any $\boldsymbol{\alpha}$, it is clear that the origin $\mathbf{b} = \mathbf{0}$ is contained in $\mathcal{B}(\langle \boldsymbol{\alpha}, -1 \rangle)$ but not in $\mathcal{B}(\langle \boldsymbol{\alpha}, 1 \rangle)$, and lies on the boundary of $\mathcal{B}(\langle \boldsymbol{\alpha}, 0 \rangle)$ (except when $\boldsymbol{\alpha} = \mathbf{0}$).

The LECs will typically be subject to many linear constraints simultaneously. We will normally use the letter Ω to denote a constraint, either a single one like $\langle \boldsymbol{\alpha}, c \rangle$ or a combination of several such constraints. Given two constraints Ω, Ω' , we write the constraint that imposes both of them simultaneously as $\Omega + \Omega'$. A point \mathbf{b} satisfies $\Omega + \Omega'$ if and only if it satisfies both Ω and Ω' ; thus, the \mathcal{B} notation naturally generalises through $\mathcal{B}(\Omega + \Omega') \equiv \mathcal{B}(\Omega) \cap \mathcal{B}(\Omega')$. For combinations of many constraints, we will generalise $+$ into e.g. $\Omega = \sum_i \langle \boldsymbol{\alpha}_i, c_i \rangle$.

4.2 Stronger and weaker constraints

A hierarchy can be established among the constraints based on how strong (restrictive) they are. For instance, $b_1 \geq 1$ is stronger than $b_1 \geq 0$. We will write the stronger-than relation as $\Omega \geq \Omega'$, which holds if all points that satisfy Ω also satisfy Ω' . Thus, $\Omega \geq \Omega'$ is equivalent to $\mathcal{B}(\Omega) \subseteq \mathcal{B}(\Omega')$. Naturally, we say $\Omega = \Omega'$ if $\mathcal{B}(\Omega) = \mathcal{B}(\Omega')$, and say $\Omega > \Omega'$ if $\Omega \geq \Omega'$ but $\Omega \neq \Omega'$. Just like subset relations, our stronger-than relation is not a total ordering, as there exist many pairs of constraints Ω, Ω' where neither is stronger than the other. From our definitions, it trivially follows that

$$\begin{aligned} (\Omega + \Omega') \geq \Omega, & & (\Omega + \Omega') \geq \Omega', \\ (\Omega + \Omega') = \Omega & \Leftrightarrow & \Omega \geq \Omega', \end{aligned} \tag{4.3}$$

so that if $\Omega \not\geq \Omega'$ and $\Omega' \not\geq \Omega$, their combination $\Omega + \Omega'$ is indeed a new, strictly stronger constraint. Furthermore, we see that, for all $\lambda > 0, \kappa > 1$, and $\boldsymbol{\alpha} \neq \mathbf{0}$,

$$\langle \lambda \boldsymbol{\alpha}, 1 \rangle > \langle \lambda \boldsymbol{\alpha}, 0 \rangle > \langle \lambda \boldsymbol{\alpha}, -1 \rangle, \tag{4.4}$$

$$\langle \kappa \boldsymbol{\alpha}, 1 \rangle < \langle \boldsymbol{\alpha}, 1 \rangle, \quad \langle \lambda \boldsymbol{\alpha}, 0 \rangle = \langle \boldsymbol{\alpha}, 0 \rangle, \quad \langle \kappa \boldsymbol{\alpha}, -1 \rangle > \langle \boldsymbol{\alpha}, -1 \rangle. \tag{4.5}$$

¹⁵We consistently use Roman letters for vectors in parameter space and Greek letters for vectors in constraint space. In general, parameter space may be any finite-dimensional real vector space, with constraint space considered as its dual.

From the $c = 0$ version of eq. (4.5), we see that $\langle \alpha, 0 \rangle$ is not a unique representation of the constraint, since we can freely rescale α without changing it. We may remove this ambiguity by constraining α to be a unit vector.

There exists a constraint Ω_∞ , equivalent to $\langle \mathbf{0}, 1 \rangle$ or e.g. $\langle \alpha, 1 \rangle + \langle -\alpha, 1 \rangle$, that is not satisfied by any point. It follows that $\Omega_\infty \geq \Omega$ and $\Omega_\infty + \Omega = \Omega_\infty$ for any Ω . A constraint that is satisfied by all points, i.e. $\langle \mathbf{0}, -1 \rangle$ or $\langle \mathbf{0}, 0 \rangle$, will be called a *trivial constraint*.

4.3 Determining the relationship between constraints

We will now present a general result, which determines if a given linear constraint $\langle \beta, c \rangle$ is weaker than an arbitrarily complicated constraint Ω . This will serve as the basis for all our uses of linear constraints.¹⁶ For complete proofs, more details, and practical applications, see appendix B.

Consider a set of linear constraints $\langle \alpha_i, c \rangle$ for i in some finite set $I \subset \mathbb{N}$.¹⁷ Note that c is the same for all constraints. Then let $\omega_c \equiv \sum_{i \in I} \langle \alpha_i, c \rangle$;¹⁸ an example of such a constraint is given in figure 1. Then define $\mathcal{A}(\omega_c)$ as the set of all points that can be expressed as

$$\sum_{i \in I} \lambda_i \alpha_i, \quad \lambda_i \geq 0, \quad \sum_{i \in I} \lambda_i \in \Lambda_c, \quad (4.6)$$

where

$$\Lambda_1 = [1, \infty), \quad \Lambda_0 = [0, \infty), \quad \Lambda_{-1} = [0, 1]. \quad (4.7)$$

The shape of $\mathcal{A}(\omega_c)$ is illustrated in figure 2. With these definitions, the following holds:

Proposition 4.1 (determining if constraint is weaker, special case). *Let $\langle \beta, c \rangle$ be a single linear constraint, and let $\omega_c \neq \Omega_\infty$ be defined as above. Then $\langle \beta, c \rangle \leq \omega_c$ if and only if $\beta \in \mathcal{A}(\omega_c)$.*

This is proven in appendix B.1. If ω_c is a single linear constraint, this result reduces down to eq. (4.5). The condition $\omega_c \neq \Omega_\infty$ is necessary, since there exist corner cases where $\omega_c = \Omega_\infty$ but $\mathcal{A}(\omega_c)$ fails to cover the entire constraint space.¹⁹ However, if $\mathcal{A}(\omega_c)$ does cover the entire space, then it is certain that $\omega_c = \Omega_\infty$.

In figure 2, we may note that $\mathcal{A}(\omega_c)$ is closely related to the *convex hull* of the α_i . In essence, $\mathcal{A}(\omega_c)$ is obtained by forming the hull, and then also including all points that give weaker constraints under eq. (4.5). We may also note that the convex hull can be defined as

$$\text{Hull}(\{\alpha_i\}_{i \in I}) = \left\{ \sum_{i \in I} \lambda_i \alpha_i \mid \lambda_i \geq 0, \sum_{i \in I} \lambda_i = 1 \right\}, \quad (4.8)$$

which is very similar to eq. (4.6).

¹⁶Proposition 4.1, along with a version of the notation we use here, was defined in [21], although the proof was completely different. An incorrect version of proposition 4.2 was also presented without proof. To the best of our knowledge, these results are novel, although the relevant literature is vast and lies outside our area of expertise. The closest we have found is [39], although their algorithm requires knowing a point that satisfies Ω , relies on more complicated mathematical machinery, and does not include all the extensions presented further below in section 4.4 and appendix B.

¹⁷It is crucial that only finite combinations of constraints are considered, and it will normally be tacitly assumed that all sets like I are finite. A limited extension to infinite sets is covered in appendix B.6.

¹⁸We use lowercase ω here to emphasise that it is not a general constraint. A similar treatment of general Ω is given below.

¹⁹A trivial example of this is $\omega_1 = \langle \mathbf{0}, 1 \rangle$, where $\mathcal{A}(\omega_1) = \{\mathbf{0}\}$.




Figure 1. A cropped depiction of twelve random two-dimensional constraints $\langle \alpha_i, c \rangle$ for $c = +1$ (left, **yellow**), $c = 0$ (middle, **blue**) and $c = -1$ (right, **red**) illustrated as the parameter-space lines $\alpha_i \cdot \mathbf{b} = c$. The side of the line that is *excluded* by the constraint is hatched. The region $\mathcal{B}(\omega_c)$ for $\omega_c = \sum_i \langle \alpha_i, c \rangle$, i.e. the set of points that satisfy all the constraints, is shaded. The lines corresponding to *relevant* constraints (i.e. those that actually delimit $\mathcal{B}(\omega_c)$; this is more closely defined in section 4.4) are drawn more strongly than the rest.




Figure 2. Examples of the regions $\mathcal{A}(\omega_1)$ (left, **yellow**), $\mathcal{A}(\omega_0)$ (middle, **blue**) and $\mathcal{A}(\omega_{-1})$ (right, **red**) in constraint space, using the same α_i as in figure 1. The α_i are represented as points (\triangle , \square , ∇ , respectively) in the space, and the relevant ones are filled (again, relevancy is defined in section 4.4). The convex hulls (as defined in eq. (4.8)) of the α_i are outlined. For comparison to figure 1, it is helpful to remember that α_i are normal vectors to the lines shown there, and that larger $|\alpha_i|$ correspond to lines passing *closer* to the origin. Note how, given identical α_i , $\mathcal{A}(\omega_0)$ is the union of $\mathcal{A}(\omega_1)$ and $\mathcal{A}(\omega_{-1})$ (this is easy to see from eq. (4.7)) whereas the set of relevant constraints is the intersection of the respective sets.

Now, let us handle the general case. The most general combination of a finite number of linear constraints can be expressed as

$$\Omega = \sum_{i \in I_1} \langle \alpha_i, 1 \rangle + \sum_{i \in I_0} \langle \alpha_i, 0 \rangle + \sum_{i \in I_{-1}} \langle \alpha_i, -1 \rangle, \tag{4.9}$$

where I_c are some disjoint, finite, and possibly empty sets. We may compactly write this as $\Omega = \sum_{i \in I} \langle \alpha_i, c_i \rangle$ where $I \equiv I_{+1} \cup I_0 \cup I_{-1}$ and $c_i = c$ if $i \in I_c$.

Similarly to $\mathcal{A}(\omega_c)$, let $\mathcal{A}_c(\Omega)$ be the set of all points that can be expressed as (recall that $c \in \{1, 0, -1\}$)

$$\sum_{i \in I_{-1}} \lambda_i \alpha_i + \sum_{i \in I_0} \lambda_i \alpha_i + \sum_{i \in I_1} \lambda_i \alpha_i, \quad \lambda_i \geq 0, \tag{4.10}$$

with λ_i constrained by the condition

$$\sum_{i \in I_1} \lambda_i - \sum_{i \in I_{-1}} \lambda_i \geq c. \tag{4.11}$$

An illustration of $\mathcal{A}_c(\Omega)$ can be found in appendix B.4.5. With these definitions, the following holds:

Proposition 4.2 (determining if constraint is weaker, general case). *Let $\langle \beta, c \rangle$ be a linear constraint, and let $\Omega \neq \Omega_\infty$ be defined as above. Then $\langle \beta, c \rangle \leq \Omega$ if and only if $\beta \in \mathcal{A}_c(\Omega)$.*

This is proven in appendix B.1.4. If only one of the I_c is nonempty, this reduces down to proposition 4.1.

While it is not as useful for the purposes of proposition 4.2, one may note that eqs. (4.10) and (4.11) can be more succinctly stated as

$$\mathcal{A}_c(\Omega) = \left\{ \sum_{i \in I} \lambda_i \alpha_i \mid \lambda_i \geq 0, \sum_{i \in I} \lambda_i c_i \geq c \right\}. \tag{4.12}$$

This definition of $\mathcal{A}_c(\Omega)$ works also if c, c_i are not constrained to $\{-1, 0, 1\}$.

4.4 Representations and degeneracy

Checking if \mathbf{b} satisfies Ω becomes computationally expensive if Ω is the combination of many different linear constraints. However, Ω is usually not uniquely determined by how it is expressed as a sum of linear constraints, and it is possible to vastly reduce that redundancy. To that end, we define a *representation* of a constraint Ω as any finite set \mathcal{S} of linear constraints with the property²⁰

$$\Omega = \sum_{\langle \alpha, c \rangle \in \mathcal{S}} \langle \alpha, c \rangle. \tag{4.13}$$

If it is implicit which representation is used for Ω , we may call the $\langle \alpha, c \rangle \in \mathcal{S}$ the *elements of Ω* .

It is clear that there exist *minimal representations*, i.e. representations with the smallest number of elements. As we will see below, there is often a unique minimal representation, which we will label $\mathcal{R}(\Omega)$. However, there is an important exception to this: when $\mathcal{B}(\Omega)$ is contained in a hyperplane. This happens when there are some δ, d such that $\delta \cdot \mathbf{b} = d$ for all $\mathbf{b} \in \mathcal{B}(\Omega)$, or equivalently $\langle \delta, d \rangle + \langle -\delta, -d \rangle \leq \Omega$. We will call Ω *degenerate* if so is the case, and *non-degenerate* otherwise.²¹ With this in mind, we can state the following result:

²⁰Clearly, all Ω also admit representation as a sum of an *infinite* number of constraints. However, we will not consider such representations, and proposition 4.3 below generally only holds if Ω can be expressed as a finite sum. See appendix B.6 for a discussion about infinite sums of constraints.

²¹As defined here, Ω_∞ would be considered a special case of a degenerate constraint. In the closer study of degenerate constraints given in appendix B.2.1, it turns out to be more useful to consider Ω_∞ separately, viewing it as neither degenerate nor non-degenerate.

Proposition 4.3 (finding relevant constraints, non-degenerate case). *If Ω is a non-degenerate constraint, there exists a minimal representation $\mathcal{R}(\Omega)$ that is unique up to the normalisation of its elements. Furthermore, for any representation \mathcal{S} of Ω , the relation $\mathcal{R}(\Omega) \subseteq \mathcal{S}$ is true up to normalisation.*

The elements of $\mathcal{R}(\Omega)$ are exactly those $\langle \alpha, c \rangle \leq \Omega$ for which there is some $\mathbf{b} \in \mathcal{B}(\Omega)$ such that $\alpha \cdot \mathbf{b} = c$ and $\beta \cdot \mathbf{b} > d$ for all $\langle \beta, d \rangle \leq \Omega$ with $\langle \beta, d \rangle \neq \langle \alpha, c \rangle$.

This is proven in appendix B.3. Due to this uniqueness, and the fact that $\mathcal{R}(\Omega)$ is a subset of any representation, we will call the elements of $\mathcal{R}(\Omega)$ the *relevant elements of Ω* , and call all other elements of any representation *irrelevant*, since they can be discarded without altering Ω . A more practical way of finding $\mathcal{R}(\Omega)$, based on proposition 4.2, is given in appendix B.4.4.

When Ω is degenerate, there is typically no unique minimal representation, although there is still a straightforward way to find *some* minimal representation, which we will also label $\mathcal{R}(\Omega)$. This generalisation of proposition 4.3 is discussed in appendix B.2.1, along with a more general method of replacing any degenerate constraint with a non-degenerate analogue in a lower-dimensional space. Note, however, that degenerate constraints are only a corner case with little practical relevance: a small perturbation, e.g. by numerical error, to the elements of a degenerate constraint will either render it non-degenerate, or render it equal to Ω_∞ .

5 Positivity bounds

Equipped with the notion of linear constraints, we are ready to move on to the main topic of this paper: positivity bounds. (For a more detailed version of this derivation, see [13]; various generalisations can be found in e.g. [19, 20].) We start by writing down the fixed- t dispersion relation for the amplitude $a_J T^J$:

$$a_J T^J(s, t) = \frac{1}{2\pi i} \oint_\gamma dz \frac{a_J T^J(z, t)}{z - s}. \tag{5.1}$$

The amplitude has two branch cuts along the real axis: a right-hand cut starting at $z = 4$ corresponding to the s -channel, and a left-hand cut starting at $z = -t$ corresponding to the u -channel. The discontinuity across these cuts is $T(z + i\varepsilon) - T(z - i\varepsilon) = 2i \operatorname{Im} T(z + i\varepsilon)$. For real s in the span $-t < s < 4$, we deform the contour γ as shown in figure 3. We can then reexpress the integral in terms of the discontinuities, which may require derivatives (subtractions) to make the contour at infinity vanish. Using the crossing relation derived in section 3.5 to rewrite the u -channel cut in terms of s , the result is

$$a_J \frac{d^k}{ds^k} T^J(s, t) = \frac{k!}{\pi} \int_4^\infty dz \left[\frac{a_J}{(z - s)^{k+1}} + \frac{(-1)^k a_I C_u^{IJ}}{(z - u)^{k+1}} \right] \operatorname{Im} T^J(z + i\varepsilon, t). \tag{5.2}$$

The Froissart bound [40] shows that the integral converges whenever $k \geq 2$, since $T^J(z + i\varepsilon, t) = \mathcal{O}(s \ln^2 s)$.²² We will discuss specific values for k in section 5.2; here, we keep it general.

²²Note that $T^J(z + i\varepsilon, t)$ on the right-hand side of eq. (5.2) is the exact, non-perturbative amplitude — see e.g. [13]. Indeed, the perturbative χ PT amplitude at any fixed order grows polynomially with s , so it violates the Froissart bound. We can insert the fixed-order perturbative amplitude at the left-hand side thanks to the smallness of s (and t), which guarantees good agreement with the exact one.




Figure 3. The contour integral in the z -plane around $z = s$ used in the dispersion relation.

Above threshold, the partial-wave expansion of the amplitude takes the form²³

$$T^J(s, t) = \sum_{\ell=0}^{\infty} (2\ell + 1) f_{\ell}^J(s) P_{\ell} \left(1 + \frac{2t}{s-4} \right), \quad (5.3)$$

where f_{ℓ}^J are partial wave amplitudes, P_{ℓ} are Legendre polynomials, and the expression in parentheses is the cosine of the scattering angle. The optical theorem then gives

$$\text{Im } f_{\ell}^J(s) = s\beta(s)\sigma_{\ell}^J(s), \quad \beta(s) \equiv \sqrt{1 - \frac{4}{s}}, \quad (5.4)$$

which is positive above threshold since the partial-wave cross-sections σ_{ℓ}^J are always positive. Therefore,

$$\text{Im } T^J(s, t) = \sum_{\ell=0}^{\infty} (2\ell + 1) s\beta(s)\sigma_{\ell}^J P_{\ell} \left(1 + \frac{2t}{s-4} \right) \quad (5.5)$$

is positive above threshold as long as P_{ℓ} is. Since $P_{\ell}(z) \geq 0$ when $z \geq 1$, eq. (5.2) therefore imposes the constraint that, for any $t \in [0, 4]$, $s \in [-t, 4]$ and any representation index J ,

$$a_J \frac{d^k}{ds^k} T^J(s, t) \geq 0 \quad (5.6a)$$

$$\text{if } a_I \left\{ \delta^{IJ} \left[\frac{z-u}{z-s} \right]^{k+1} + (-1)^k C_u^{IJ} \right\} \geq 0 \quad \text{for all } z \geq 4. \quad (5.6b)$$

The region in the s, t plane where this holds is shown in figure 4. Note that $u \in [-4, 4]$, so the expression in square brackets above is always positive.

Up to and including NNLO, the second derivative of T^J is linear in all LECs, so we obtain from eq. (5.6a) an expression of the form

$$\sum_i \alpha_i \hat{L}_i^r + \sum_j \beta_j K_j^r + \gamma \geq 0, \quad (5.7)$$

where the coefficients $\alpha_i, \beta_j, \gamma$ are functions of s, t and a_J , but not of the LECs. This constitutes a linear constraint, and each valid choice of s, t and a_J potentially yields a different constraint. The result of combining these constraints will be that only a limited region in parameter space ($\mathcal{B}(\Omega)$ in the notation of section 4) satisfies the positivity bounds. With some luck, the boundary of this region is close enough to the experimentally measured value to improve on its uncertainty (carefully considering also the uncertainty of the bounds).

²³There is a limited domain of validity for this expansion, but it does not affect the range of s, t used by us. Again, see [13] for details.




Figure 4. The plane of normalised Mandelstam variables. The **red** triangle is the region where the amplitude is real and free from singularities or branch cuts. The positivity conditions eq. (5.6) are valid inside the outlined part. The hatched regions with s, t or u positive are the physical regions for the respective channels.

5.1 Conditions on a_J

If we demand that eq. (5.6b) holds in the entire allowed s, t region, we see that the factor in square brackets can be made arbitrarily large or small by varying s, u, z . Therefore, we obtain the independent conditions $a_J \geq 0$ and $(-1)^k a_I C_u^{IJ} \geq 0$. However, we may apply the dispersion relation independently to each fixed s, t . Then, eq. (5.6b) can be made less restrictive, and a wider range of constraints on the LECs can be generated. This also includes permitting odd k for some s, t .

While we may fix s and t (which in turn fixes u), we must still allow z to cover its entire range. Therefore, finding all valid a_J for given s, t presents some practical issues. We solve this by using the technology of section 4, since eq. (5.6b) is a set of linear constraints on the vector a_I ; we may write it compactly as $\mathbf{a} \cdot \boldsymbol{\beta}^J(z) \geq 0$. Noting that $\left(\frac{z-u}{z-s}\right)^{k+1}$ is monotonic as a function of $z \in [4, \infty)$, we see that it is always possible to write $\boldsymbol{\beta}^J(z) = \mu \boldsymbol{\beta}^J(4) + (1-\mu) \boldsymbol{\beta}^J(\infty)$ for $\mu \in [0, 1]$. By propositions 4.1 and 4.3 (see also proposition B.3), it follows that only $\boldsymbol{\beta}^J(4)$ and $\boldsymbol{\beta}^J(\infty)$ are relevant constraints on \mathbf{a} . Thus, it is sufficient to evaluate eq. (5.6b) at $z = 4$ and $z = \infty$, rather than letting z cover its entire range.

Another practical problem is that the set of allowed a_J is typically unbounded. However, eq. (5.6a) is independent of the magnitude of a_J . The obvious solution is to fix the normalisation of the vector \mathbf{a} , but this is problematic since a linear constraint on \mathbf{a} is not necessarily a linear constraint on $\frac{1}{|\mathbf{a}|} \mathbf{a}$. Instead, we may simply rescale \mathbf{a} so that $\sum_J a_J = 1$. This does not cover all possible \mathbf{a} (for that, we must also look at $\sum_J a_J = 0$ and $\sum_J a_J = -1$), but it turns out that eq. (5.6b) is only satisfied by \mathbf{a} for which this works. Using this, constraints on a_I are shown in figure 5.




Figure 5. Illustration of which a_I, a_A, a_{SS} , normalised so that $a_I + a_A + a_{SS} = 1$, are permitted by eq. (5.6b) for $n = 2, k = 2$ at $t = 4$ and various fixed s . The shaded band is the region permitted by the $z = \infty$ bounds (see the discussion in section 5.1). It is independent of s and extends to infinity. The hatched triangles are the regions permitted by the $z = 4$ bounds for various s as indicated. Thus, the a_J permitted at fixed s, t is the overlap between the triangle and the band. The **orange** rectangle is the region permitted at all s, t . The **blue** points represent the eigenstates, including $a_J(\pi^\pm \pi^\mp) = C_u^{IJ} a_I(\pi^\pm \pi^\pm)$ in addition to those given in eq. (3.11). Not shown is the inelastic scattering $a_J(\pi^0 \pi^0 \rightarrow \pi^+ \pi^-) = (\frac{1}{3} \ 0 \ -\frac{1}{3})$, which never satisfies eq. (5.6b). For $n \geq 2$, the permitted region has an analogous shape, albeit in 4 ($n = 3$) and 5 ($n \geq 4$) dimensions, respectively. Like for $n = 2$, the eigenstate scattering amplitudes are mainly located in the corners of the always-permitted region.

5.2 The number of derivatives

As mentioned before, eq. (5.6) requires $k \geq 2$ to be valid, and $k = 2$ is sufficient; indeed, [41] claims that this value produces the best bounds. However, nothing prevents us from taking more derivatives, and with our generalised methods, we do find new relevant bounds from larger k ; see e.g. figures 6 and 12 below. Also [19] makes use of higher derivatives.

At NLO, the LECs only enter through the second-order polynomial part of the amplitude, so the third and higher derivatives are parameter-independent and do not generate any bounds. This is not the case at NNLO, where the polynomial part is third-order, and where the non-polynomial unitarity correction contains NLO LECs. Therefore, $k = 3$ should yield another set of bounds on the NNLO LECs, and $k \geq 4$ should add bounds on the NLO LECs not obtainable from the NLO-only amplitude.

It also turns out that odd k cannot be used at any order in the 2-flavour case. To see this, look explicitly at eq. (5.6b) at $z = \infty$:

$$\delta^{IJ} - C_u^{IJ} = \frac{1}{6} c_1^I c_2^J \quad \text{where} \quad c_1 = \begin{pmatrix} 2 & 1 & -1 \end{pmatrix}, c_2 = \begin{pmatrix} 2 & 3 & -5 \end{pmatrix} \quad (5.8)$$

Due to what seems to be a coincidence, the matrix factorises into a direct product, and since c_2 has different-sign elements, no nonzero a_I satisfies $(a_I c_1^I) c_2^J \geq 0$ for all J . No such coincidences hinder odd k at 3 or more flavours, and we have explicitly found a_J that satisfy eq. (5.6b) with odd k at 3 and 4 flavours (these turn out to produce very weak bounds, though). Even $k > 2$ remain permitted also at 2 flavours.

5.3 The value of t

There is no immediately obvious reason to favour any specific part of the allowed s, t region when producing bounds. However, one may note that at NLO with $k \geq 2$, the only part of $\frac{d^k}{ds^k} T^J(s, t)$ that depends on s, t is the LEC-independent unitarity correction, which manifests itself as the term labelled γ in eq. (5.7). Therefore, at fixed a_J the most restrictive bound is obtained by minimising γ . It turns out that given s , the magnitude of $\frac{d^k}{ds^k} T^J(s, t)$ tends to increase with t , with minima and maxima always falling along the $t = 4$ line. Therefore, it is expected that all relevant constraints should be found with $t = 4$. While we see no clear *a priori* reason for it to be so, we have verified it by scanning the entire s, t range for bounds; all relevant ones were found at $t = 4$, within numerical uncertainty.

At NNLO, also the α_i in eq. (5.7) may depend on s, t , so the simple argument above does not hold. However, the NNLO corrections are far too small to affect the overall shape of the amplitude, so for NLO LECs, fixing $t = 4$ should remain sufficient. The situation is yet more complicated for the NNLO LECs, since it turns out that certain combinations of them only feature in the amplitude when $t \neq 4$ (see the next section). Therefore, complete NNLO bounds require using the full s, t range.

5.4 Independently bounded parameters

While the parameter space affected by our bounds is technically the full space of (N)NLO LECs, it is of course impractical to work in such a large and redundant space. Many LECs do not receive any bounds at all by not appearing in the Lagrangian relevant for meson-meson scattering (see section 3.2), and others only appear in fixed linear combinations. Specifically, all NLO LECs that appear in the amplitude do so independently, but the NNLO LECs only appear in combinations; consequently, it is not possible to obtain bounds on the individual K_i^r . For instance, we shall see below that the combination $K_4^r - 2K_2^r$ appears in the amplitude and therefore receives bounds,²⁴ but its complement $2K_4^r + K_2^r$, which does not appear, is free to assume any value. Therefore, nothing can be said about the values of K_2^r and K_4^r themselves.

We will therefore reexpress our parameter space in terms of NLO LECs in addition to a new set of independently bounded parameters built from the NNLO LECs. Their form can be deduced from the polynomial parts of the functions B, C described in section 3, which, following [34], are

$$B_P(s, t, u) = \gamma_1 + \gamma_2 t + \gamma_3 t^2 + \gamma_4 (s - u)^2 + \gamma_5 t^3 + \gamma_6 t (s - u)^2 + (\text{NLO}), \quad (5.9a)$$

$$C_P(s, t, u) = \delta_1 + \delta_2 s + \delta_3 s^2 + \delta_4 (t - u)^2 + \delta_5 s^3 + \delta_6 s (t - u)^2 + (\text{NLO}). \quad (5.9b)$$

²⁴This combination is, up to a scale factor, Δ_3 as defined below and explicitly given in appendix A.

Here, γ_i, δ_i are linear combinations of the NNLO LECs, and “(NLO)” contains all NLO LECs, constant terms, etc.^{25,26} Therefore, bounds on the NNLO LECs only come in the form of bounds on γ_i, δ_i . Of course, not all γ_i, δ_i are bounded either — those with $i < 3$ vanish in the second derivative of the amplitude and therefore receive no bounds, and those with $i < 6$ vanish in the third. The remaining combinations are also not necessarily independent, so we will proceed to remould them into a better set of parameters.

5.4.1 General number of flavours

As is discussed in section 5.3, all relevant NLO bounds appear at $t = 4$, so we may expect that this particular t -value is special also at NNLO. Therefore, we express the polynomial parts in terms of s and $\bar{t} \equiv 4 - t$, using $u = \bar{t} - s$:

$$\begin{aligned} B_P(s, t, u) &= s^2\Gamma_1 + \bar{t}s^2(\Gamma_4 - 3\Gamma_3) + \dots, \\ B_P(t, u, s) &= s^2\Gamma_2 - s^3\Gamma_3 + \bar{t}s^2\Gamma_4 + \dots, \\ B_P(u, s, t) &= s^2\Gamma_2 + s^3\Gamma_3 + \bar{t}s^2(\Gamma_4 - 3\Gamma_3) + \dots, \end{aligned} \tag{5.10}$$

where “...” consists of terms that vanish in the second derivative of the amplitude. Here, we have defined the parameters

$$\begin{aligned} \Gamma_1 &= 4\gamma_4 + 16\gamma_6, & \Gamma_2 &= \gamma_3 + \gamma_4 + 8\gamma_6, \\ \Gamma_3 &= \gamma_5 + \gamma_6, & \Gamma_4 &= 3\gamma_5 - \gamma_6, \end{aligned} \tag{5.11}$$

of which Γ_4 only receives bounds when $t \neq 4$ due to the presence of \bar{t} . Similarly,

$$\begin{aligned} C_P(s, t, u) &= s^2\Delta_2 + s^3\Delta_3 + \bar{t}s^2(\Delta_4 - 3\Delta_3) + \dots, \\ C_P(t, u, s) &= s^2\Delta_1 + \bar{t}s^2(\Delta_4 - 3\Delta_3) + \dots, \\ C_P(u, s, t) &= s^2\Delta_2 - s^3\Delta_3 + \bar{t}s^2\Delta_4 + \dots, \end{aligned} \tag{5.12}$$

where Δ_i are defined in terms of δ_i identically to eq. (5.11).

These 8 parameters Δ_i, Γ_i constitute a minimal set of parameters for NNLO bounds with a general number of flavours; explicit expressions are given in appendix A.²⁷ At 2 and 3 flavours, the Cayley-Hamilton identity allows for further reduction of the number of parameters.

²⁵This differs from the convention in [34], where the NLO terms are included in γ_i, δ_i . The “NNLO parts” that we extract here are easily read off from the appendices to that paper.

²⁶[19] use a similar approach in their 2-flavour NNLO bounds, but do not separate NLO and NNLO parts. This results in a smaller parameter space (6 dimensions compared to our 8; see section 5.4.4), but our approach has the benefit of separating the relatively well-determined \bar{t}_i from the much more uncertain NNLO values, allowing for figures such as figure 8. The fact that our parametrisation remains partly redundant is not a major issue, since we always fix some parameters rather than working in the full space.

²⁷While they can technically be considered LECs, we will avoid confusion by referring to the Γ_i, Δ_i as *NNLO parameters*, reserving “LEC” for the coefficients appearing in the standard form of the Lagrangian.

5.4.2 Two flavours

Decomposing the polynomial part of $A(s, t, u)$ using eqs. (3.3), (5.10) and (5.12) reveals

$$\begin{aligned} A_P(s, t, u) &= s^2\Theta_2 - s^3\Theta_3 + \bar{t}s^2\Theta_4 + \dots + \dots, \\ A_P(t, u, s) &= s^2\Theta_1 + \bar{t}s^2(\Theta_4 - 3\Theta_3), \\ A_P(u, s, t) &= s^2\Theta_2 + s^3\Theta_3 + \bar{t}s^2(\Theta_4 - 3\Theta_3) + \dots, \end{aligned} \tag{5.13}$$

where

$$\begin{aligned} \Theta_1 &= 2\gamma_3 - 2\gamma_4 + 4\delta_4 + 16\delta_6, & \Theta_2 &= 4\gamma_4 + 16\gamma_6 + \delta_3 + \delta_4 + 8\delta_6, \\ \Theta_3 &= 2\gamma_5 + 2\gamma_6 - \delta_5 - \delta_6, & \Theta_4 &= 3\gamma_5 - \gamma_6 - 4\delta_6. \end{aligned} \tag{5.14}$$

Θ_4 , like Γ_4 and Δ_4 , is only bounded when $t \neq 4$. Explicit expressions and experiment-based reference values for Θ_i are given in appendix A.

5.4.3 Three flavours

A similar but less elegant simplification is possible in the 3-flavour case, using eq. (3.6):²⁸

$$\begin{aligned} B_{1P}(s, t, u) &= -2s^3\Gamma_3 + 3\bar{t}s^2\Gamma_3, \\ B_{2P}(s, t, u) &= s^2\Xi_1 - s^3\Gamma_3 + 3\bar{t}s^2\Xi_4, \\ A_{1P}(s, t, u) &= s^2\Xi_2 + s^3(\Delta_3 - \frac{2}{3}\Gamma_3) + \bar{t}s^2\Xi_4, \\ A_{2P}(s, t, u) &= s^2\Xi_3 + \frac{2}{3}s^3\Gamma_3 + \bar{t}s^2\Xi_4, \\ A_{3P}(s, t, u) &= s^2\Xi_2 + s^3(2\Gamma_3 - \Delta_3) + \bar{t}s^2(\Xi_4 + 3\Delta_3 - 4\Gamma_3), \end{aligned} \tag{5.15}$$

where

$$\begin{aligned} \Xi_1 &= \gamma_3 - 3\gamma_4 - 8\gamma_6, & \Xi_2 &= 5\gamma_4 + 16\gamma_6 + \delta_3 + \delta_4 + 8\delta_6, \\ \Xi_3 &= \frac{4}{3}\gamma_3 + \frac{16}{3}\gamma_6 + 4\delta_4 + 16\delta_6, & \Xi_4 &= \gamma_5 - 3\gamma_6 - 4\delta_6. \end{aligned} \tag{5.16}$$

Again, expressions and values are given in appendix A.

5.4.4 The full parameter space

Table 1 summarises the parameters affected by our bounds at different orders and number of flavours. Note how the dimension of the space ranges from 2 (NLO 2-flavour) to 16 (NNLO n -flavour, $n \geq 4$). If t is fixed to 4 at NNLO, this is reduced by 1 if $n \leq 3$ and by 2 otherwise; as discussed in section 5.3, $t = 4$ is the only relevant value at NLO. The difficulties associated with large parameter spaces are discussed in section 6.

5.5 The absence of catastrophic divergences

At NLO and above, the coefficient γ in eq. (5.7) diverges in the limit $s \rightarrow 0$ or $u \rightarrow 0$. If the divergence is towards positive infinity, this is not a problem — it simply means that the positivity bound becomes trivial in these limits. However, divergence towards negative

²⁸A more symmetric result would have been obtained by eliminating the $d^{ace}d^{bde}$ term in eq. (3.4) instead, but we choose to follow [14].

Flavours	NLO	NNLO
2	\bar{l}_1, \bar{l}_2	$\underline{\bar{l}_1, \dots, \bar{l}_4}; \quad \Theta_1, \Theta_2, \underline{\Theta_3}, (\Theta_4)$
3	L_1^r, L_2^r, L_3^r	$\underline{L_1^r, \dots, L_6^r, L_8^r}; \quad \Xi_1, \Xi_2, \Xi_3, (\Xi_4), \underline{\Gamma_3}, \underline{\Delta_3}$
≥ 4	$\hat{L}_0^r, \dots, \hat{L}_3^r$	$\underline{\underline{\hat{L}_0^r, \dots, \hat{L}_6^r, \hat{L}_8^r}}; \quad \Gamma_1, \Gamma_2, \underline{\Gamma_3}, (\Gamma_4), \Delta_1, \Delta_2, \underline{\Delta_3}, (\Delta_4)$

Table 1. Summary of the NLO LECs and NNLO parameters appearing in the second derivative of the amplitude. The parameters $\Gamma'_i, \Delta'_i, \Theta'_i, \Xi'_i$ as defined in eqs. (5.11), (5.14) and (5.16). The LECs and parameters that remain in the third derivative are underlined, and those that remain also in the fourth and above are doubly underlined. Parameters that only feature in the amplitude when $t \neq 4$ are placed in parentheses.

infinity would be catastrophic, since no finite LECs could satisfy the positivity condition. If there were some value of a_J for which the divergence is in this direction, the theory would be inconsistent.

The situation becomes more complicated at NNLO, where also α_i diverge. If α_i diverge at the same rate or faster than γ , the positivity conditions remain sensible also in these limits, but if γ were to diverge towards negative infinity faster than α_i , we would again have inconsistencies.

As s approaches 4 from below, the k th derivative of the amplitude diverges as odd powers (up to $2k - 1$) of $1/\delta$, where $s = 4(1 - \delta^2)$; see appendix C.2 for details. Let q^J be the coefficient of the leading divergence $d^k T^J/ds^k$ that contributes to γ . Then consistency requires $a_J q^J \geq 0$ for all valid a_J . Since eq. (5.6a) requires $a_J \geq 0$ in the limit $s \rightarrow 4$, this is satisfied if $q^J \geq 0$. At both NLO and NNLO for any number of flavours n , this turns out to be true for the $s \rightarrow 4$ divergence (this was already noted in [13] for $n = 2$). Also, the divergences of α_i are of equal or lower powers than those of γ .

The same divergence structure appears in the $u \rightarrow 4$ limit, but here the coefficients of the leading divergences are not necessarily positive. However, we may use crossing symmetry to rewrite

$$\frac{d^k}{ds^k} T^I(s, t) = (-1)^k \frac{d^k}{du^k} C_u^{IJ} T^J(u, t). \tag{5.17}$$

Here, we can simply relabel u as s . The coefficient of the leading divergence is here $(-1)^k C_u^{IJ} q^J$, and since eq. (5.6a) requires $(-1)^k a_I C_u^{IJ} \geq 0$, in the limit $u \rightarrow 4$, the fact that $q^J \geq 0$ in the $s \rightarrow 4$ limit guarantees that there are no catastrophic divergences in the $u \rightarrow 4$ limit either.

Since $d\delta/ds = -1/8\delta$, taking another derivative does not change the sign of q^J . Therefore, if no catastrophic divergences appear at the first k where γ diverges, they will not appear at larger k either.²⁹

²⁹The first divergence happens at $k = 1$ for $J \in \{I, S, SS, AA\}$ and at $k = 2$ for $J = A$, regardless of n . There is no divergence for $J = AS$, since $T^{AS}(s, t, u)$ does not contain $\bar{J}(s)$ or $k_i(s)$.

5.6 Integrals above threshold

The right-hand side of eq. (5.2) is, in its standard application, a non-perturbative quantity, about which the only knowledge we have is the fact that it is positive. However, χ PT is a low-energy theory, so its amplitude at any order should be an excellent approximation of the true amplitude for energies sufficiently close to threshold. Taking inspiration from the approach used in [19], we may therefore explicitly evaluate the lowest part of the integral on the right-hand side of eq. (5.2) and subtract it from both sides. Specifically, we define

$$D_k^J(\lambda, v, t) = \frac{k!}{\pi} \int_4^\lambda \frac{dz}{(z-v)^{k+1}} \text{Im} T^J(z + i\varepsilon, t) \quad (5.18)$$

and modify eq. (5.2) to

$$a_J \left[\frac{d^k}{ds^k} T^J(s, t) - D_k^J(\lambda, s, t) - (-1)^k C^{JI} D_k^I(\lambda, u, t) \right] = \frac{k!}{\pi} \int_\lambda^\infty dz [\dots] \text{Im} T^J(z + i\varepsilon, t). \quad (5.19)$$

This has two benefits:

- i) When s, t, u and a_J satisfy the conditions of eq. (5.6), both D_k^J and the right-hand side are positive, so we obtain a stronger positivity bound.
- ii) The right-hand side of eq. (5.19) is positive under a wider range of conditions than that of eq. (5.2), so we obtain more positivity bounds. (This is because the constraint $\mathbf{a} \cdot \boldsymbol{\alpha}^J(4) \geq 0$ is replaced by the weaker $\mathbf{a} \cdot \boldsymbol{\alpha}^J(\lambda) \geq 0$, recalling the notation and discussion in section 5.1). Some of the new bounds are weakened by D_k^J being negative, but they may still contribute.

The size of λ presents a tradeoff: larger values amplify the benefits of using it, but also decrease the accuracy of the relation as the fixed-order χ PT amplitude strays from the exact result. The integral also requires some mathematical machinery; D_k^J is by no means a simple function, but we determine it up to NNLO in appendix C.3 (of course, it could also have been done numerically). By evaluating the NNLO corrections, we obtain a good idea of the accuracy of the NLO result.

An upper bound on λ is provided by [35], which determines the breakdown scale of n -flavour χ PT to be $s \sim (4\pi F)^2/n$. Using the value $F = 92.2(1)$ MeV adopted by [6] along with $M_{\text{phys}} = M_\pi$, this places the breakdown at $\lambda \approx 35$ for $n = 2$ and at $\lambda \approx 25$ for $n = 3$. Thus, we cannot expect sensible results for λ anywhere close to this, and certainly not above it.³⁰

³⁰Note that $(2M_K)^2 \approx 14$ and $(2M_\eta)^2 \approx 16$ are already quite close to the breakdown scale. This offers some motivation as to why the equal-mass approximation is reasonable also at 3 flavours: with unequal, real-world masses, 3-flavour χ PT operates close to the limits of its range of validity even with nonrelativistic particles, which offsets the accuracy gained by increasing the realism of the model.

6 Results

Here, we present the constraints obtained using the methods described in the preceding sections.

Following section 4, we will use the letter Ω to denote each collection of constraints, and $\mathcal{B}(\Omega)$ to denote the sets of parameter-space points that satisfy these. We will compare each $\mathcal{B}(\Omega)$ to a reference point, taken as the central value of the LEC estimates in [6]. These values can also be found in table 2 in appendix A.

In all but the simplest cases, parameter space has too many dimensions to be visualised as a whole. Therefore, we will show lower-dimensional slices, with all omitted parameters set to their reference values. We will primarily show two-dimensional slices, since they are the easiest to understand, although some three-dimensional slices will be needed as well. It is not practical to show an exhaustive set of slices, so we will focus on grouping parameters that are, in some loose sense, related.

As a visual aid and a rough indicator of constraint strength, we define the quantity $\rho(\langle \boldsymbol{\alpha}, c \rangle)$ to be the shortest distance between the reference point and the hyperplane $\boldsymbol{\alpha} \cdot \mathbf{b} = c$.³¹ Alternatively, we may use $\hat{\rho}(\langle \boldsymbol{\alpha}, c \rangle)$, which is the analogous distance if the space where all parameters are rescaled so that their reference values are 1.

6.1 Two flavours

Two-flavour χ PT constraints are in many regards quite simple: there are only 2 parameters at NLO and 7 at NNLO; there is only one reasonable choice for M_{phys} , namely M_π ; and as shown in section 5.2, we do not have to consider odd numbers of derivatives. The NLO constraints have been extensively studied in e.g. [13, 19, 20], whereas the NNLO constraints are novel to this work.

Figure 6 shows constraints obtained using 2 derivatives and various amounts of above-threshold integration. The non-integrated ($\lambda = 4$) constraints are slightly stronger than those in [13], which only considered eigenstate a_J rather than the full space, but the constraints do not come close to the experimental uncertainty of the reference point without using λ that are far too large for the results to be trusted (recall that perturbative breakdown is expected at $\lambda \approx 35$). In [19, 20], comparable bounds are obtained with slightly less extreme λ , but in both cases, λ needs to be rather large to start cutting into the experimental uncertainty. The large discrepancy between the NLO and NNLO versions of the constraints indicate that the bounds are highly unreliable for all but the smallest λ used. Even with $\lambda = 4$, the difference is quite significant.

Figure 7 shows similar NNLO constraints over a more conservative λ range; however, here we display the effects of higher even derivatives (recall again that odd derivatives need not be considered with 2 flavours). Unlike the ones shown in figure 6, these constraints impose *upper* bounds on the LECs as well as lower bounds. Note that in the upper-left part of the plot, the two-derivative bounds are less restrictive than their NLO counterpart. This can partly be seen as an artefact of introducing multiple new parameters and fixing

³¹Note that this refers to distance in the full parameter space, which does not directly correspond to distance in the subspaces shown in the figures.




Figure 6. Two-flavour two-derivative constraints on \bar{l}_1, \bar{l}_2 for various λ , as indicated in the legend. The presentation is similar to figure 1: each version of $\mathcal{B}(\Omega)$ is outlined, with the hatched side indicating the points excluded by the constraints. The bounds from [13] are also drawn. The reference point $\bar{l}_1 = -0.4(6), \bar{l}_2 = 4.3(1)$ is drawn as a dot with an uncertainty region around it. For each set of constraints, the direct NNLO counterpart (i.e. using the same λ, s, t and a_J) is drawn as a dashed outline.




Figure 7. NNLO bounds on \bar{l}_1, \bar{l}_2 with $k = 2, 4, 6$ derivatives, displayed similarly to figure 6 but over a slightly wider part of parameter space. For comparison to figure 6, the corresponding NLO bounds are drawn with weak solid lines. The six-derivative bounds are nearly independent of λ , so all versions are not drawn.




Figure 8. NNLO k -derivative constraints on \bar{l}_3, \bar{l}_4 , displayed like in figure 7, using similarly scaled axes (left) or an extreme scale on the \bar{l}_3 axis (right) to show the full constraint structure. In all cases, the bounds are satisfied by the reference point. The six-derivative bounds (which like in figure 6 are nearly λ -independent) exactly overlap with the $\lambda = 4$ two-derivative bounds and are not drawn.

them to imprecise experimental values (for instance, figure 9 shows that smaller values of \bar{l}_4 strengthen the bounds on \bar{l}_1, \bar{l}_2), but one must keep in mind that switching to a more refined theory can both strengthen and relax the predictions.

Figure 8 shows similar bounds on \bar{l}_3 and \bar{l}_4 , which are the only NLO LECs other than \bar{l}_1, \bar{l}_2 that appear in the NNLO amplitude. The bounds on \bar{l}_3 are extremely weak, since \bar{l}_3 figures in the amplitude with much smaller prefactors than the other LECs. It can be partly understood by noting that the \bar{l}_3 term in the Lagrangian, unlike the other terms, does not contain the field u_μ . Interestingly, the upper bounds on \bar{l}_4 become *weaker* as λ is increased. This does not necessarily contradict the arguments made in section 5.6, due to the complicated NNLO situation where both α and c in a constraint may depend on λ . Nevertheless, it is surprising to see, and does not seem to appear in other bounds, such as those on $\bar{l}_{1,2}$.

Figure 9 combines the \bar{l}_4 bounds with the \bar{l}_1, \bar{l}_2 bounds to form a summary of the effective bounds on the two-flavour NLO LECs obtained in this paper. Unfortunately, only the $k = 2, \lambda \approx 4$ constraints have a shape that is sensible to show in three dimensions.

Lastly, figure 10 shows bounds on the four NNLO parameters Θ_i , using only two derivatives. Here, the bounds are not entirely consistent with the reference point, although not too much meaning should be read into this, as the reference values for the NNLO parameters are little more than educated guesses. It also showcases the phenomenon where some constraints become weaker at larger λ , at least at the particular values at which we have fixed the NLO LECs.

The NNLO parameter bounds are not particularly strong compared to the magnitude of the reference values, but some are still notable. Θ_4 has a strict lower bound, with no values of $\Theta_{1,2,3}$ being permitted if $\Theta_4 \lesssim -2.5$. Θ_3 is bounded from both above and below, and the bounds are fairly independent of the values of the other parameters in a large part




Figure 9. NNLO two-derivative constraints on \bar{l}_1, \bar{l}_2 and \bar{l}_4 with $\lambda = 4$; a cross-section at $\bar{l}_4 = 4.4$ would yield part of figure 6. The constraint surfaces are coloured according to their proximity $1/\rho(\langle \alpha_i, c_i \rangle)$ to the reference point, which is drawn similar to the 2D plots. The “open space” bounded by the constraint surfaces is part of $\mathcal{B}(\Omega)$; the greyed-out region is excluded by the constraints. To clarify its spatial position, the reference point is connected to the boundaries of the plotted region with black dotted lines parallel to the coordinate axes, whenever doing so is possible without intersecting a constraint surface. It is similarly connected to neighbouring surfaces with black arrows. These are orthogonal to the respective surfaces, even though the different scales of the axes makes it not appear so.

of parameter space. Thus, we may write down the tentative single-parameter bounds

$$-1.5 \lesssim 10^3 \Theta_3 \lesssim 1, \quad -2.5 \lesssim 10^3 \Theta_4, \quad (6.1)$$

both of which are satisfied, with a margin of about an order of magnitude, by the reference values in appendix A.

6.2 Three flavours

Three-flavour χ PT bounds cover 3 parameters at NLO and 12 at NNLO, and three choices for M_{phys} (namely M_π, M_K and M_η) present themselves, with no *a priori* indication of which to choose. This would of course be resolved by working with unequal-mass mesons, but this NLO amplitude is far more complicated (see [42] and references therein) and its NNLO counterpart is so far undetermined in a simple analytic fashion; furthermore, unequal masses have implications for the construction of bounds that we do not address here (see [14]).

Figure 11 shows the bounds on L_1^r, L_2^r and L_3^r obtainable at NLO. With $M_{\text{phys}} = M_\pi$, the bounds are consistent with the reference point and qualitatively similar to the two-flavour bounds on \bar{l}_1, \bar{l}_2 . Interestingly, the reference point does not satisfy the bounds at the other choices of M_{phys} , although only barely — the smallest distance between the reference point and $\mathcal{B}(\Omega)$ is $0.5 \cdot 10^{-3}$ for $M_{\text{phys}} = M_K$ and $0.6 \cdot 10^{-3}$ for $M_{\text{phys}} = M_\eta$, which is smaller

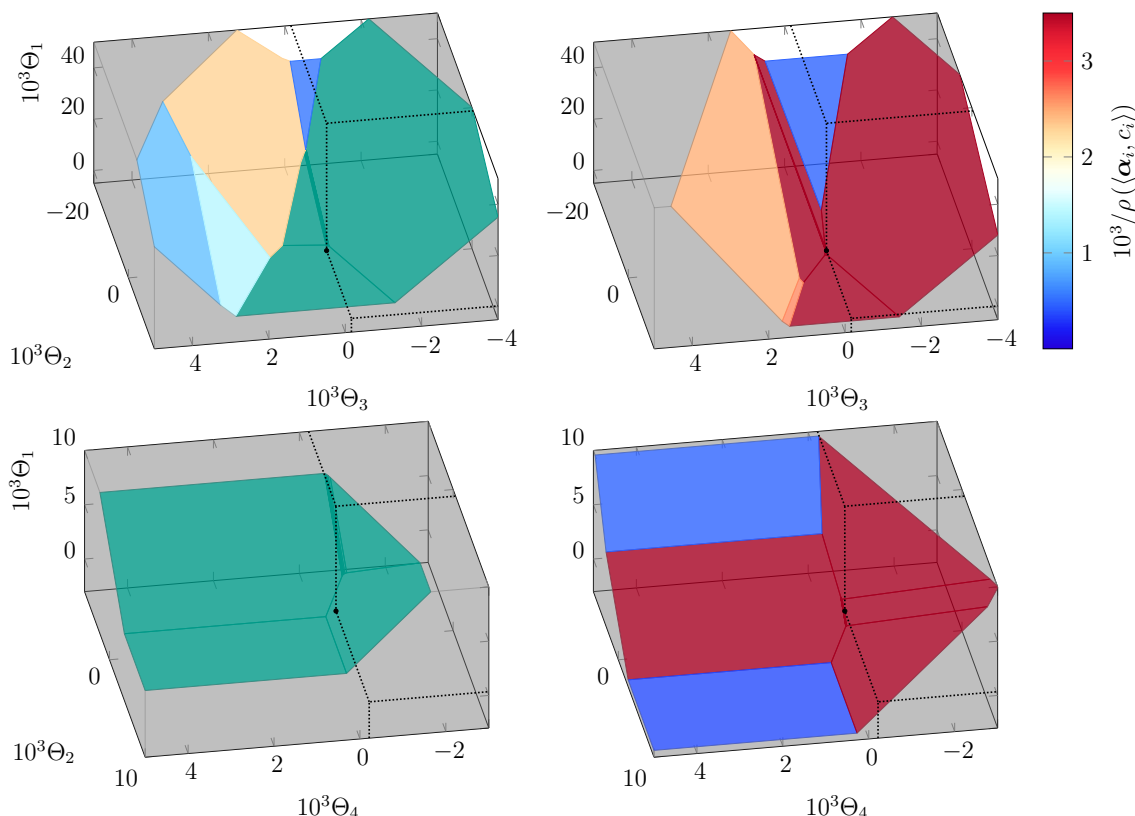


Figure 10. NNLO two-derivative bounds on Θ_i at $\lambda = 4$ (left) and $\lambda = 8$ (right) for $i = 1, 2, 3$ (top) and $1, 2, 4$ (bottom, note different $\Theta_{1,2}$ -axes). These specific i choices are used, for if either Θ_1 or Θ_2 is fixed to its reference value, the remaining constraint is not satisfied by any value of the other Θ_i . Note the rather different scales on the axes; several facets in the top plots (especially for $\lambda = 8$) are in fact nearly perpendicular to the Θ_3 axis, placing upper and lower bounds on Θ_3 that do not depend very strongly on the other parameters. The green facets exclude the reference point. The exclusion is by a rather small amount (no orthogonal arrow can be sensibly drawn), and varying the L_i^r within their uncertainties is sufficient to remedy this. Note how increasing λ to 8 slightly weakens these constraints so that they no longer exclude the reference point.

than the uncertainty in the experimental values (approximately $(1-3)\cdot 10^{-3}$). Therefore, using these M_{phys} does not imply any significant inconsistency.

It is worth noting that the three-flavour constraints are much more sensitive to integration than their two-flavour counterparts. This can be partly understood by noting that integrals scale as roughly the square of the number of flavours (see figure 22).

Figure 12 shows NNLO bounds on the same three LECs. Interestingly, the four-derivative constraints alone confine L_1^r, L_2^r, L_3^r to a bounded region, although it is very large in most directions. However, the bounds are reasonably strict between the two near-parallel faces shown in the figure, so we may write down another double-ended bound, similar to eq. (6.1):

$$-27 \lesssim 4.9L_1^r + 2.8L_2^r + 2.4L_3^r \lesssim 10 \quad \text{for } M_{\text{phys}} = M_\pi, \quad (6.2)$$

where the linear combination of the LECs is chosen to be roughly orthogonal to the bounding faces.

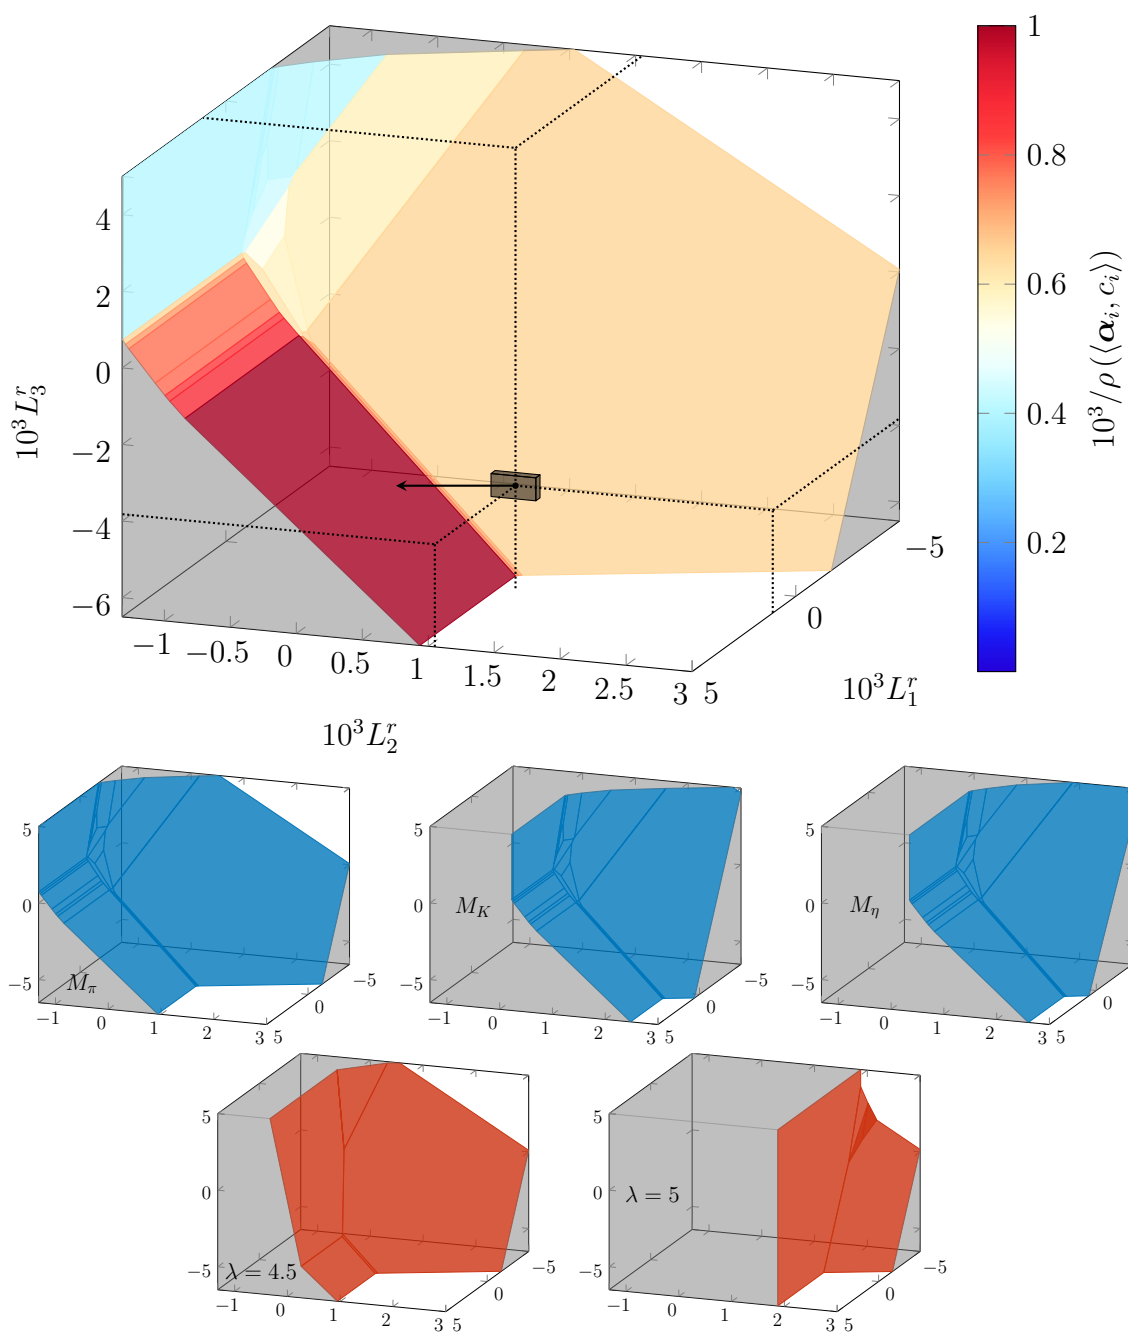


Figure 11. Two-derivative constraints on the three-flavour NLO LECs L_i^r , visualised similarly to previous figures. *Top:* detailed constraints with $M_{\text{phys}} = M_\pi, \lambda = 4$. The visualisation is similar to figure 9; note the different scales on the axes. *Middle:* constraints with (from left to right) $M_{\text{phys}} = M_\pi, M_K$, and M_η , all with $\lambda = 4$. Note that the leftmost figure is just a less detailed version of the top figure, and that the M_K and M_η surfaces only differ from each other by a very small amount. Both these higher-mass constraints exclude the reference point (not drawn), although not by much. *Bottom:* the effect of setting $\lambda = 4.5$ (left) or 5 (right) with $M_{\text{phys}} = M_\pi$, for comparison with the $\lambda = 4$ figure above. The $\lambda = 5$ constraint excludes the reference point (not drawn) rather severely.

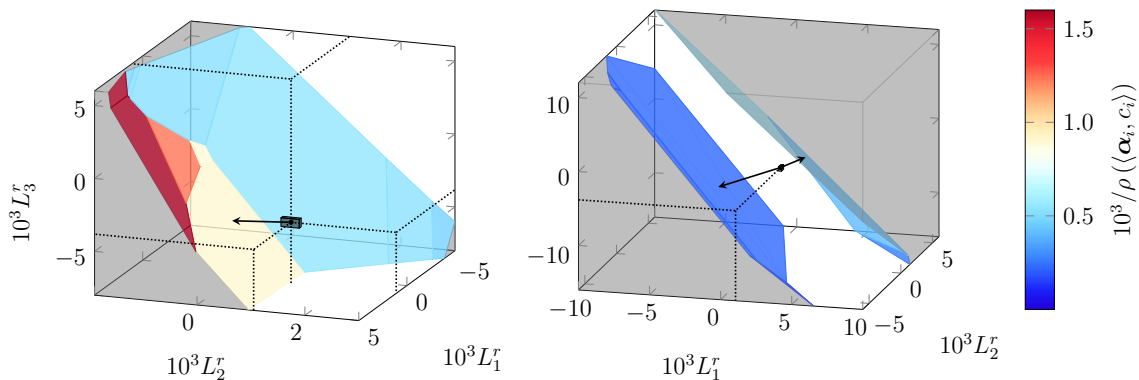


Figure 12. NNLO bounds on L_1^r, L_2^r and L_3^r with $M_{\text{phys}} = M_\pi$ and $\lambda = 4$, using two (left) and four (right) derivatives; three derivatives gives very weak bounds. Thus, the left figure is essentially the NNLO version of figure 11. Its response to larger λ or M_{phys} (not shown) is qualitatively similar to that exhibited at NLO. In the four-derivative case, $\mathcal{B}(\Omega)$ is actually a bounded region. It is a lentil-shaped body whose largest dimension is about two orders of magnitude larger than the region shown in the figure. Note that the axes have been rotated relative to the left figure in order to make the inside of $\mathcal{B}(\Omega)$ reasonably visible.

Figure 13 shows NNLO bounds on the remaining L_i^r . These are rather weak, even though L_6^r obtains a double-ended bound. The weakness is understandable for similar reasons as the weakness of the \bar{l}_3 bounds. With $M_{\text{phys}} = M_K, M_\eta$, the bounds are inconsistent, not only with the reference values for $L_{4,5,6,8}^r$, but with *all* values of these parameters, unless $L_{1,2,3}^r$ are removed from their experimental values by an amount roughly one order of magnitude larger than their stated uncertainty (the NNLO parameters have little effect). This practically renders this version of the theory self-inconsistent, except at experimentally unreasonable points in parameter space. Although the $M_{\text{phys}} = M_\pi$ version remains consistent, only a small amount of integration excludes the reference point, so its validity is quite dubious.

Lastly, figures 14 and 15 show bounds on the NNLO parameters; for the same reason as above, we keep $M_{\text{phys}} = M_\pi$. The former shows some features reminiscent of its SU(2) analogue, figure 10, although there is not a clear double-ended bound on either parameter. The latter is more interesting, since all three parameters are confined to a small bounded region and quite significantly excludes the reference point. The bounds on Γ_3, Δ_3 in particular are fairly independent of each other and Ξ_4 , leading to the single-parameter bounds

$$0.08 \lesssim 10^3 \Gamma_3 \lesssim 0.34, \quad -0.8 \lesssim 10^3 \Delta_3 \lesssim 0.25 \quad \text{for } M_{\text{phys}} = M_\pi. \quad (6.3)$$

6.3 Higher number of flavours

χ PT with more than three flavours are not of direct interest as low-energy QCD, since the large mass of the charm quark makes it entirely invalid as a model of mesons. An arbitrary number of flavours is useful when developing the methods, though, and is interesting in its own right in the context of EFT studies. Furthermore, χ PT has many uses other than QCD (for a review of some of these, see [43]). Besides various numbers of flavours, these




Figure 13. NNLO bounds on L_4^r, L_5^r (left) and L_6^r, L_8^r (right), visualised similarly to figure 6. The uncertainty region around the reference point is hardly visible at this scale. Note how the four-derivative bounds exclude the reference point also for rather small $\lambda > 4$. However, the integrated two-derivative bounds are weaker and are not visible at all in the right figure. The corresponding bounds for $M_{\text{phys}} = M_K, M_\eta$ are entirely inconsistent with the experimental values.




Figure 14. NNLO two-derivative bounds on Ξ_1, Ξ_2, Ξ_3 (those NNLO parameters whose contribution is independent of t) with $\lambda = 4$ and $M_{\text{phys}} = M_\pi$. Higher-derivative bounds are extremely weak and are not shown. The facets are coloured based on their distance to the reference point in the space of NNLO parameters rather than the full parameter space. Like in figure 10, the **green** facet excludes the reference point, although by an extremely small amount, well within the drawn size of the reference point.




Figure 15. NNLO two-derivative bounds on Γ_3 , Δ_3 and Ξ_4 (those NNLO parameters that obtain additional bounds at $t \neq 4$) with $\lambda = 4$ and $M_{\text{phys}} = M_\pi$. Note how they are bounded from all directions and confined to a rather small volume; this makes three-dimensional representation difficult, so we use two-dimensional slices instead. Note how the bounds exclude the reference point. A major source of uncertainty in these bounds is the fixing of the other parameters: as the NLO LECs are varied within their uncertainties, the boundary lines shift in various directions, with the typical amount displacement being roughly 0.2. However, at no values are the bounds consistent with the reference values of Γ_3 , Δ_3 and Ξ_4 .

commonly use different symmetry breaking patterns than $SU(n) \times SU(n) \rightarrow SU(n)$; some of these were treated in [34], so their amplitudes would be a drop-in replacement into our methods. However, that is beyond the scope of this paper, in which we are content to show some basic high-flavour results in “QCD-style” χ PT.

It turns out that the bounds change quite gradually between different $n \geq 4$, to the extent that the difference between e.g. $n = 4$ and $n = 5$ is hardly visible upon first inspection of figures like those used here. Therefore, we have chosen to only display $n = 4$ (the “leading” high-flavour example) and $n = 8$ (a reasonable “very high flavour” example, also of historical technicolour relevance [44]). Understandably, there are no experimental reference values for the high-flavour LECs, so to perform two- or three-dimensional slices of parameter space, we provisionally set $\hat{L}_i^r = L_i^r$ and use the 3-flavour reference values when possible, and use 0 as the reference value for e.g. \hat{L}_0^r . We also retain the distance-to-reference-point colouring of the constraint surfaces to make them more visually distinguishable. Of course, not too much meaning should be read into these distances.

Figures 16 ($n = 4$) and 17 ($n = 8$) show the basic NLO bounds, similarly to figure 11, to which the bounds are qualitatively similar. There is a trend towards weaker constraints as n increases, as can be seen from the amplitude: many of the most important terms go as powers of $1/n$ (cf. eq. (4.5); all constraints have $c = -1$). Note how we wholly abandon the M_{phys} debate and use M_π throughout.

Figures 18, 19 ($n = 4$) and 20, 21 ($n = 8$) show the most interesting NNLO bounds, which are again similar to their lower-flavour analogues. The bounds on $\hat{L}_4^r, \hat{L}_5^r, \hat{L}_6^r$ and \hat{L}_8^r , which are even weaker than the corresponding ones for $n = 3$, are not shown. Likewise, the bounds on Γ_4, Δ_4 are very weak and have been omitted. Note how the prominent lower bound on \hat{L}_1^r in figure 21 breaks the trend of weaker bounds at larger n .




Figure 16. Two-derivative constraints on the 4-flavour NLO LECs L_0^r, \dots, L_3^r with $M_{\text{phys}} = M_\pi$ and $\lambda = 4$. The four possible selections of three LECs are shown, with the top left figure (L_1^r, L_2^r, L_3^r) being almost exactly the 4-flavour equivalent of figure 11 (the LEC and colour ranges differ slightly). No reference point is shown, and the colour information should not be considered very meaningful, as discussed in the text; therefore, the colour bar has been omitted to save space.

Finally, one may ask what happens in the limit $n \rightarrow \infty$. Looking at the amplitudes in [34], we see that they are independent of the LECs at leading order in n : at NLO, the amplitude is $\mathcal{O}(n)$ whereas the LEC parts are $\mathcal{O}(1)$, while at NNLO they are $\mathcal{O}(n^2)$ and $\mathcal{O}(n)$, respectively (this of course indicates convergence problems at high n , in agreement with [35]). Thus, the bounds, expressed in the schematic form of eq. (5.7), will eventually be dominated by γ , and will therefore asymptotically tend towards either $\langle \mathbf{0}, 1 \rangle$ or the trivial $\langle \mathbf{0}, -1 \rangle$; the gradual weakening in figures 16 to 21 suggests the latter.

6.4 Considerations about the integrals

Figure 22 demonstrates the integrals of the components of the amplitude over relevant s, λ ranges. We may note that despite the great complexity seen in appendix C.3, the graphs are typically quite simple and qualitatively similar. There is, however, a very wide range of magnitudes; typically, “lower” J components have larger integrals.




Figure 17. The exact 8-flavour equivalent of figure 16.




Figure 18. NNLO 4-flavour bounds on Γ'_i (left) and Δ'_i (right), analogous to figure 10. Note the different scales on the axes; Γ'_6 and Δ'_6 are bounded by almost parallel faces.




Figure 19. NNLO 4-flavour bounds on L_0^r, \dots, L_3^r presented in the same way as figure 16, but with some axes rotated to give a better view. See figure 12 for the 3-flavour counterpart (no analogous 4-derivative version has been produced due to issues with high-dimensional parameter space).




Figure 20. The 8-flavour equivalent of figure 18. The axes have been zoomed out to adequately view the weaker constraints.




Figure 21. The exact 8-flavour equivalent of figure 19.

It is interesting to note that the ratio between NLO and NNLO integrals is approximately constant in λ , varying only with s and J . This is perhaps unexpected since the NNLO integrand contains terms like $z^3 X(z)$, where X is one of \bar{J} or k_i , whereas the NLO amplitude only contains $z^2 X(z)$. Therefore, we would expect the ratio to grow approximately linearly with λ . However, the $z^3 \bar{J}(z)$ terms are typically suppressed by small numerical coefficients or NLO LECs, and $\bar{J}(z)$ dominates $k_i(z)$, as shown in figure 23. Therefore, this effect does not manifest until λ is much larger than the values relevant to this application, which in practice limits the importance of higher-order corrections to the bounds.

6.5 Considerations about a_J

We have made two innovations in the treatment of a_J :

- (i) not restricting a_J to the physical eigenstates (this was done already in [21]),
- (ii) employing the fixed- s (as opposed to all- s) constraints eq. (5.6b) on a_J .

This section investigates whether these changes actually give any improvements at all — it would be conceivable that the physical eigenstates, which are allowed for all s , were special in a way that guarantees that they generate the strongest bounds.




Figure 22. The integrals $\frac{1}{n^2} [D_k^J(s, t; \lambda) + (-1)^k C^{JI} D_k^I(u, t; \lambda)]$, defined in eq. (5.19), as functions of λ between 4 and 4.5 at $k = 2, t = 4$, for the six possible J . The factor $1/n^2$ roughly cancels the n -dependence, making it possible to show all n at the same vertical scale. Each function is represented for $n = 2$ (green), $n = 3$ (cyan) and $n = 4$ (purple) χ PT, with the latter representing the general n -flavour case reasonably well. Different values of s between 0 and 1.5 (the range of most relevant constraints) in increments of 0.5 are shown with the colour saturation (brighter colour = smaller s). The NLO result is drawn with solid lines, and the NNLO result (with the NLO LECs fixed to their reference values) is drawn with dashed lines. The difference between them is shaded, and the λ ranges are staggered slightly to improve readability. The graphs have been vertically shifted to enhance readability; all integrals are equal to 0 at $\lambda = 4$. Note that the vertical scale differs greatly between different J .




Figure 23. Logarithmic plot of the relative size of the different loop integral functions, covering a rather wide range of λ . Each line represents $\left| \int_4^\lambda X(z) dz \right|$ for the given function X . Note how $\bar{J} \gg k_1 \approx k_3 \gg k_2 \approx k_4$ for small λ , which changes to $\bar{J} \gg k_1 \approx k_2 \approx k_3 \gg k_4$ for large λ . In either case, “ \gg ” is by about two orders of magnitude, and “ \approx ” is within one order of magnitude.

To measure how significantly (i) is used, we consider a_J as a point on the unit sphere and find the angle $\theta(a_J)$ between it and the closest eigenstate point. Thus, larger θ indicates, in a sense, more use of (i). Similarly, we may measure (ii) via the fraction of points on $s \in [-4, 4]$ for which eq. (5.6b) permits a_J .

In terms of the *ad hoc* measures together with $\hat{\rho}$ defined above, figure 24 shows the distribution of NLO bounds for 2, 3 and 4 flavours, and figure 25 shows a geometrically more intuitive version in the 2-flavour case.³² We see that relevant constraints tend to have low $\hat{\rho}$, albeit with many exceptions — the orientation of the constraint is another important factor. The relevant constraints are rather evenly distributed over the θ range permitted by eq. (5.6b), indicating that there is, in this regard, nothing special about the eigenstates, validating (i). Most relevant constraints also occur at a_J that are permitted for very few s (i.e. coloured very pale in the plots), validating (ii).

There are, of course, severe limitations to the analysis in this section. Apart from the roughness of the chosen measures discussed above, it is difficult to assess just how great the benefits of (i) and (ii) are. Stronger constraints are obtained, but not necessarily much stronger: the improvement in figure 6 over [13] is very slight, although it seems that the use of (i) is limited at $n = 2$ but more extensive at e.g. $n = 3$ by comparing the subfigures of figure 24.

³²The points were sampled uniformly over the unit octahedron (i.e. the unit sphere under the 1-norm $\sum_J |a_J| = 1$) and its higher-dimensional analogues in a_J space. This shape was used rather than the unit sphere to preserve the linearity of bounds on a_J ; compare the discussion in section 5.1.

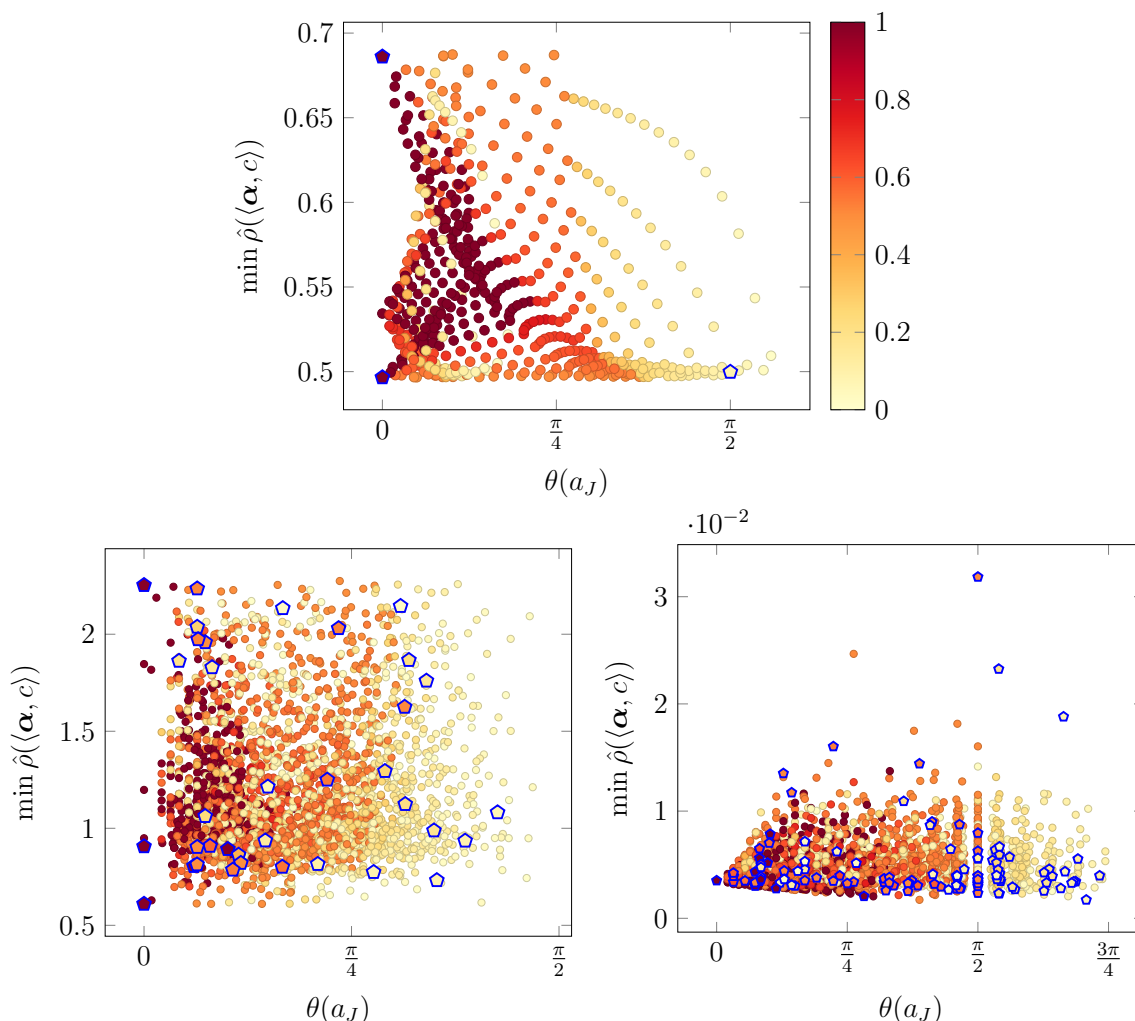


Figure 24. A discrete sample of the full set of a_J permitted by eq. (5.6b), plotted over θ as defined in section 6.5. For each a_J , the closest distance $\hat{\rho}$ between the reference point and any constraint generated at that a_J is indicated with a dot. Relevant constraints are outlined with a blue pentagon. The colour of the dots indicate the fraction of the full range $s \in [-4, 4]$ in which eq. (5.6b) permits that a_J . *Top:* two flavours. The eigenstates, i.e. the points at $\theta = 0$, are, from top to bottom, $\pi^0\pi^+$, $\pi^0\pi^0$ and $\pi^+\pi^+$. *Lower left:* three flavours. The eigenstates are, from top to bottom, $\pi^0\pi^+$, $K\eta$, $\pi^+\pi^+$, $K\pi^0$, $K^\pm\pi^\pm/K^0\pi^\pm/\pi^0\pi^0$ (overlapping) and $\eta\pi$. *Lower right:* four flavours (representing $n > 3$ flavours). Since high-flavour χ PT is not directly applicable to meson physics, we do not use the full set of “physical” eigenstates, but for θ to be defined, we retain a single eigenstate: $\pi^+\pi^+$, whose decomposition (consisting of the SS component only) is uniquely flavour-independent.




Figure 25. The same 2-flavour data as in figure 24, but plotted over a_I, a_A plane with $a_I + a_A + a_{SS} = 1$, like in figure 5, rather than as a function of θ . The colours convey the same meaning as in figure 24, whereas the $\hat{\rho}$ values are indicated by the size of the points: smaller $\hat{\rho}$ (higher proximity to the reference point) corresponds to larger points. Note how the a_J that are permitted at all s can be clearly seen as a uniformly coloured patch, which is also featured in figure 5. It is interesting to note that most relevant constraints come from the $a_{SS} < 0$ region (upper right), which is permitted for significantly fewer s than those points where all a_J are positive.

7 Conclusions and outlook

To recapitulate, our method has been as follows: we scan over the s, t range depicted in figure 4 (or in many cases just the s range, with t fixed to 4, as discussed in section 5.3). At each (s, t) -value, we scan over those a_J that are permitted by eq. (5.6b), and compute the k th derivative of the n -flavour amplitude to either NLO or NNLO at that point, possibly with above-threshold integration up to λ , as described in section 5.6. Through eq. (5.6a), this yields positivity bounds on the LECs (or at NNLO, the parameters derived in section 5.4), which can be handled as linear constraints using the language and methods of section 4 (implemented as discussed in appendix B). In the end, this yields a manageable set of relevant constraints, which can be visualised and interpreted.

Our results at NLO consist of stronger bounds than in [13, 14], whereas the comparison to the more recent works [19, 20] is less clear-cut. As in previous works, most bounds consist

of irregularly shaped and usually infinite regions, although some of the nicer cases allow for more clear-cut bounds such as eqs. (6.1) to (6.3). Although the bounds themselves are highly uncertain, basic compatibility suggests that one may assign uncertainties of roughly 200–1000% to $\Theta_i, \Gamma_i, \Delta_i$ and therefore also to the NNLO LECs, which are not given any error estimates in [6]. Obtaining better error estimates is a possible direction for future work.

The employment of above-threshold integration allows for very strong bounds, but carries the risk of going too far beyond the low-energy limit; the difference between our NLO and NNLO bounds, such as in figure 6, does not inspire much confidence in integrated bounds for λ significantly larger than 4. Our improved handling of a_J , which is evaluated in section 6.5, improves bounds without additional assumptions ([19, 20] also uses assumption-less improvements). Of course, the choice of fixed-order χ PT is itself something that relies on the low-energy limit, although it is easier to motivate than a particular choice of $\lambda > 4$ is.

At NNLO, our methods suffer some practical problems due to the very high dimension of the parameter space, so we have only performed rather coarse scans of the available s, t and a_J ranges. This is remedied by fixing some parameters and focusing on the lower-dimensional space that remains. However, it is important to keep in mind that, even though we may fail to obtain many constraints due to limited scans and technical issues with determining convex hulls (see appendix B.4), it is guaranteed that the constraints we do find are true — the method automatically errs on the side of caution, so to speak. In particular, this means that if the bounds are inconsistent with the experimental values, then the error must lie either with the values, or with the theory itself. In general, the main problem with our method is not its ability to produce bounds, but our ability to rely on the assumptions behind them.

The most prominent example of this is the problems encountered at three flavours, as discussed in section 6.2. The obvious remedy is to replace equal-mass χ PT with the more realistic unequal-mass version, towards which the main hurdle is some so-far-unsolved two-loop integrals. The emergence of distinct mass eigenstates, as well as some other features discussed in [14], could possibly interfere with some of our method innovations, although we are confident that they can be remedied.

Beyond practical improvements and the use of unequal-mass χ PT, a possible step forward is to either go beyond $2 \rightarrow 2$ scattering, or to study NNNLO. The former would allow for bounds on LECs that do not appear in the $2 \rightarrow 2$ amplitude, as well as possibly new bounds on those covered here. While it would be infeasible to manage all the NNNLO LECs in the $2 \rightarrow 2$ amplitude, or even all the NNLO LECs in higher-multiplicity amplitudes, it is not unthinkable that useful results could be obtained by fixing most LECs and studying the rest. In any case, further explorations in these directions are mainly hindered by the lack of available amplitudes; the NLO 2-flavour $2 \rightarrow 4$ amplitude was recently calculated [45], but no higher-order or higher-multiplicity amplitudes are currently known. Furthermore, section 5 would need to be generalised to handle the 9-dimensional kinematic space of 6-particle amplitudes. Lastly, some parts of the NNNLO amplitude would be nonlinear in the NLO LECs, necessitating the development of proposition 4.2-like technology for *nonlinear* constraints, i.e. $\alpha_i b_i + \alpha_{ij} b_i b_j + \dots \geq c$. We believe it possible that at least some of the tools in appendix B can be generalised to handle this, but have not investigated it much.

Lastly, these methods could lend themselves to application on EFTs other than χ PT, e.g. for beyond-the-Standard-Model (BSM) applications. This could be particularly promising if there are no experimentally measured values for the LECs, or if experiments have only yielded bounds. An experimental upper bound coupled with an analytic lower bound could confine the coupling of an unobserved process to a range, or exclude a BSM EFT altogether.

Acknowledgments

Mattias Sjö thanks Torbjörn Lundberg for rewarding discussions resulting in the new treatment of propositions 4.1 and 4.2 as compared to [21]. This work is supported in part by the Swedish Research Council grants contract numbers 2016-05996 and 2019-03779. Colour schemes for the figures, appropriate for colourblind people and monochrome printing, were chosen based on [46, 47].

A LEC details

This section contains further details about the LECs that are introduced in section 2.1, and the NNLO parameters $\Gamma_i, \Delta_i, \Theta_i, \Xi_i$ that are defined in section 5.4. Table 2 contains the values determined in [6] that are used as references in section 6. Naturally, these do not include estimates of parameters that only appear above 3 flavours; similarly, the Γ_i, Δ_i that only appear in the $n \geq 3$ amplitude are just given provisional values based on the 3-flavour data.

Eq. (A.1) shows how the independent NNLO parameters depend on the LECs K_i^r , for a general number of flavours n .³³

$$\frac{1}{8}\Gamma_1 = 6K_1^r - 4K_5^r + K_7^r + K_{11}^r + 2K_{31}^r + nK_8^r + 2nK_{18}^r \quad (\text{A.1a})$$

$$\begin{aligned} \frac{1}{2}\Gamma_2 = & 12K_1^r - 48K_3^r - 12K_5^r + K_7^r + K_{11}^r - 16K_{13}^r + 32K_{17}^r - 8K_{19}^r - 8K_{23}^r - 16K_{28}^r \\ & + 2K_{31}^r - 8K_{33}^r + 16K_{37}^r + nK_8^r - 16nK_{14}^r + 34nK_{18}^r - 4nK_{20}^r \end{aligned} \quad (\text{A.1b})$$

$$\frac{1}{4}\Gamma_3 = K_1^r - 2K_3^r - 2K_5^r \quad (\text{A.1c})$$

$$\frac{1}{8}\Gamma_4 = -3K_3^r - K_5^r \quad (\text{A.1d})$$

$$\frac{1}{16}\Delta_1 = 6K_4^r - 2K_6^r + K_{15}^r + 2K_{29}^r + nK_{16}^r \quad (\text{A.1e})$$

$$\begin{aligned} \frac{1}{4}\Delta_2 = & -48K_2^r + 12K_4^r - 4K_6^r - 16K_9^r + K_{15}^r + 32K_{18}^r + 8K_{20}^r + 16K_{21}^r \\ & + 2K_{29}^r - 8K_{32}^r + 16K_{35}^r - 8K_{38}^r - 16nK_{10}^r + nK_{16}^r + 16nK_{22}^r \end{aligned} \quad (\text{A.1f})$$

$$\frac{1}{8}\Delta_3 = -2K_2^r + K_4^r \quad (\text{A.1g})$$

$$\frac{1}{8}\Delta_4 = -6K_2^r + K_6^r \quad (\text{A.1h})$$

³³These are generated by FORM with minimal post-processing. Common factors have been extracted to make the expressions shorter.

NLO LEC		NNLO parameter ·10 ³	
\bar{l}_1	-0.4(6)	Θ_1	0.34
\bar{l}_2	4.3(1)	Θ_2	0.68
\bar{l}_3	2.9(24)	Θ_3	-0.16
\bar{l}_4	4.4(2)	Θ_4	-0.22
		Ξ_1	0.29
	·10 ³	Ξ_2	0.34
L_1^r	1.11(10)	Ξ_3	0.25
L_2^r	1.05(17)	Ξ_4	-0.008
L_3^r	-3.82(30)	Γ_1	0.008
L_4^r	1.87(53)	Γ_2	-0.71
L_5^r	1.22(06)	Γ_3	-0.10
L_6^r	1.46(46)	Γ_4	-0.22
L_8^r	0.65(07)	Δ_1	-0.032
		Δ_2	-0.83
		Δ_3	-0.048
		Δ_4	-0.14

Table 2. All LECs and parameters covered by the bounds derived in this paper, along with their experimental reference values, taken from the most general fits in [6]. The values in parentheses indicate the uncertainties in the last decimal places. No uncertainties are provided for the NNLO parameters, since none are given for the NNLO LEC values in [6]; these values are little more than educated guesses. The values of Γ_i , Δ_i also depend directly on n for $i < 3$ (see eq. (A.1)); the listed values use $n = 4$.

Eq. (A.2) shows the same for 3 flavours. The application of the $n = 3$ Cayley-Hamilton identity and the numbering of the C_i^r follows [27].

$$\frac{1}{4}\Xi_1 = -6C_1^r + 12C_3^r + 4C_4^r - C_5^r - 3C_6^r - C_8^r + 2C_{10}^r + 6C_{11}^r - 4C_{12}^r - 18C_{13}^r + 4C_{22}^r - 4C_{25}^r \quad (\text{A.2a})$$

$$\frac{3}{8}\Xi_2 = 50C_1^r + 27C_2^r + 22C_3^r - 18C_4^r - 48C_5^r + 54C_6^r + 45C_7^r - 21C_8^r + 24C_9^r + 3C_{10}^r + 9C_{11}^r + 6C_{13}^r + 6C_{16}^r - 4C_{24}^r + 12C_{25}^r - 4C_{26}^r + 4C_{29}^r \quad (\text{A.2b})$$

$$\frac{3}{8}\Xi_3 = 70C_1^r - 18C_2^r + 68C_3^r - 16C_4^r + 43C_5^r + 21C_6^r - 30C_7^r + 25C_8^r - 24C_9^r + 10C_{10}^r + 30C_{11}^r - 8C_{12}^r - 18C_{13}^r + 12C_{16}^r - 4C_{22}^r + 18C_{23}^r + 4C_{24}^r - 2C_{25}^r - 8C_{26}^r + 8C_{29}^r \quad (\text{A.2c})$$

$$\frac{3}{4}\Xi_4 = -38C_1^r + 9C_2^r - 28C_3^r + 12C_4^r - 6C_{16}^r + 6C_{22}^r - 3C_{23}^r - 2C_{24}^r + 4C_{26}^r - 4C_{29}^r \quad (\text{A.2d})$$

$$\frac{1}{4}\Gamma_3 = C_1^r - 2C_3^r - 2C_4^r \quad (\text{A.2e})$$

$$\frac{1}{4}\Delta_3 = 4C_1^r - 5C_2^r + 2C_3^r \quad (\text{A.2f})$$

Lastly, eq. (A.3) shows the same for 2 flavours. Again, we have followed [27].

$$\frac{3}{16}\Theta_1 = 14c_1^r + 40c_2^r - 6c_3^r + 15c_4^r + 30c_5^r - 30c_6^r + 12c_7^r + 9c_{12}^r - 3c_{13}^r - 4c_{14}^r + 4c_{16}^r \tag{A.3a}$$

$$\frac{3}{8}\Theta_2 = 104c_1^r + 22c_2^r - 18c_3^r + 93c_4^r + 15c_5^r + 12c_7^r + 12c_{13}^r - 4c_{14}^r + 4c_{16}^r \tag{A.3b}$$

$$\frac{3}{16}\Theta_3 = 6c_1^r - 9/2c_2^r - 3c_3^r \tag{A.3c}$$

$$\frac{3}{4}\Theta_4 = -14c_1^r - 40c_2^r - 12c_7^r + 3c_{12}^r + 4c_{14}^r - 4c_{16}^r \tag{A.3d}$$

B Details and proofs regarding linear constraints

In this appendix, we prove the propositions stated in section 4 and provide some more details on how they may be applied.³⁴

B.1 Proof of propositions 4.1 and 4.2

As has already been mentioned, proposition 4.1 is a direct consequence of proposition 4.2, obtained by leaving all but one I_c empty. However, directly proving proposition 4.2 is much less straightforward than the following chain of implications,

$$\text{Proposition 4.1 (} c = 0 \text{ only)} \Rightarrow \text{Proposition 4.2} \Rightarrow \text{Proposition 4.1}, \tag{B.1}$$

which we will demonstrate in this section. First, however, we will show an easily accessible partial result (appendix B.1.1), and then prove some properties that are necessary for the main proof (appendix B.1.2).

B.1.1 The trivial half of the proof

One side of proposition 4.2 is easy to prove, namely that $\langle \beta, c \rangle \leq \Omega$ if $\beta \in \mathcal{A}_c(\Omega)$.

Assuming that $\Omega \neq \Omega_\infty$, take any point \mathbf{b} that satisfies Ω . Then by eq. (4.10),

$$\beta \cdot \mathbf{b} = \sum_{i \in I_1} \lambda_i \alpha_i \cdot \mathbf{b} + \sum_{i \in I_0} \lambda_i \alpha_i \cdot \mathbf{b} + \sum_{i \in I_{-1}} \lambda_i \alpha_i \cdot \mathbf{b}, \tag{B.2}$$

and since $\alpha_i \cdot \mathbf{b} \geq c$ for $i \in I_c$, we have

$$\beta \cdot \mathbf{b} \geq \sum_{i \in I_1} \lambda_i - \sum_{i \in I_{-1}} \lambda_i \geq c, \tag{B.3}$$

which uses (and motivates) eq. (4.11). The corresponding result for proposition 4.1 follows immediately.

³⁴Like in section 4, we make use of potentially unfamiliar mathematical notation in this section, so the glossary (appendix B.7) may be helpful.

B.1.2 Proof that $\mathcal{A}(\omega_0)$ is closed and convex

As appendix B.1.3 will show, these properties of $\mathcal{A}(\omega_0)$ are crucial for the main proof. Convexity is easy to show for $\mathcal{A}(\omega_0)$: given any points β_1, β_2 satisfying eq. (4.6), their convex combination

$$\mu\beta_1 + (1 - \mu)\beta_2 = \sum_{i \in I} (\mu\lambda_{1i} + (1 - \mu)\lambda_{2i})\alpha_i, \quad \lambda_{1i} \geq 0, \quad \lambda_{2i} \geq 0 \quad (\text{B.4})$$

satisfies eq. (4.6) as well.

For the proof of closedness, we only need three basic facts: (i) the intersection of closed sets is closed, (ii) the union of a finite number of closed sets is closed, and (iii) for any α and c the set $\mathcal{B}(\langle \alpha, c \rangle)$ is closed.³⁵

Now, we employ *Carathéodory's theorem for convex cones*,³⁶ which for our purposes can be formulated as follows:

Let $I \subset \mathbb{N}$ be finite, and let $\{\alpha_i\}_{i \in I}$ be vectors in \mathbb{R}^D . For any point $\beta \in \mathbb{R}^D$ fulfilling

$$\beta = \sum_{i \in I} \lambda_i \alpha_i, \quad \lambda_i \geq 0, \quad (\text{B.5})$$

there exists a set $J \subseteq I$ with at most D elements such that

$$\beta = \sum_{j \in J} \lambda_j \alpha_j, \quad \lambda_j \geq 0, \quad (\text{B.6})$$

where and all α_j are linearly independent.

Eq. (B.5) is clearly equivalent to eq. (4.6) for $c = 0$.

For any such J , there exists a set of vectors $\{\gamma_k\}$ such that $\{\alpha_j\}_{j \in J} \cup \{\gamma_k\}$ is a basis of \mathbb{R}^D . Let A_J be the invertible matrix whose columns are these basis vectors, and let λ be the vector whose components are λ_j , where $\lambda_j = 0$ if $j \notin J$. Then eq. (B.6) can be rewritten as $\beta = A_J \lambda$, or equivalently $A_J^{-1} \beta = \lambda$. Since $\lambda_j \geq 0$, we therefore obtain the inequalities $\mathbf{a}_j \cdot \beta \geq 0$, where \mathbf{a}_j are the column vectors of A_J^{-1} , and $\mathbf{a}_j \cdot \beta = 0$ if $j \notin J$. In other words,³⁷

$$\beta \in \mathcal{B} \left(\sum_j \langle \mathbf{a}_j, 0 \rangle + \sum_{j \notin J} \langle -\mathbf{a}_j, 0 \rangle \right). \quad (\text{B.7})$$

³⁵This last fact is easy to prove: take a point $\mathbf{b} \notin \mathcal{B}(\langle \alpha, c \rangle)$, i.e. $\alpha \cdot \mathbf{b} < c$. Then for any vector \mathbf{d} such that $|\mathbf{d}| < \varepsilon$,

$$\alpha \cdot (\mathbf{b} + \mathbf{d}) \leq \alpha \cdot \mathbf{b} + |\alpha \cdot \mathbf{d}| < \alpha \cdot \mathbf{b} + \varepsilon |\alpha|,$$

where we used the Cauchy-Schwarz inequality in the last step. For $\varepsilon > 0$ sufficiently small, this is still less than c , so $\mathbf{b} + \mathbf{d} \notin \mathcal{B}(\langle \alpha, c \rangle)$, proving that it is closed. (Note that this also works for $\mathcal{B}(\langle \mathbf{0}, 1 \rangle) = \emptyset$ and $\mathcal{B}(\langle \mathbf{0}, -1 \rangle) = \mathbb{R}^D$.)

³⁶In common mathematical nomenclature, $\mathcal{A}(\omega_0)$ is a convex cone, $\mathcal{A}(\omega_1)$ is an affine cone, and $\mathcal{A}(\omega_{-1})$ is a convex hull. We have chosen not to use these classifications, partly because neither applies to the general $\mathcal{A}_c(\Omega)$.

³⁷Note that β , which normally is part of a linear constraint, is itself constrained here. This is not a problem; in fact, the expression $\alpha \cdot \mathbf{b} \geq c$ can be interpreted both as $\mathbf{b} \in \mathcal{B}(\langle \alpha, c \rangle)$ and as $\alpha \in \mathcal{B}(\langle \mathbf{b}, c \rangle)$. We will return to this symmetric interpretation many times below.

By facts (i) and (iii), this set is always closed. Therefore, the set of all $\beta \in \mathcal{B}(\omega_0)$ associated with the same J is closed. $\mathcal{A}(\omega_0)$ must then be the union of all such sets, but since I is finite, there are finitely many different subsets J , so $\mathcal{A}(\omega_0)$ is the union of a finite number of closed sets. By fact (ii), it is therefore closed. \square

Let us remark that this proof extends to the other cases, so that $\mathcal{A}_c(\Omega)$ is closed and convex for any Ω and c .

B.1.3 Proof of proposition 4.1 in the $c = 0$ case

We will now turn our attention to the statement that $\langle \beta, 0 \rangle \not\leq \omega_0$ if $\beta \notin \mathcal{A}(\omega_0)$, which will complete the proof of proposition 4.1 for $c = 0$.³⁸ We employ the *separating hyperplane theorem*, which can be formulated as follows:

Let \mathcal{X} and \mathcal{Y} be disjoint convex sets, with \mathcal{X} closed and \mathcal{Y} compact. Then there exists a nonzero vector \mathbf{h} and a real number d such that

$$\chi \cdot \mathbf{h} > d \text{ and } \psi \cdot \mathbf{h} < d \tag{B.8}$$

for all $\chi \in \mathcal{X}, \psi \in \mathcal{Y}$. (The set $\{\pi \mid \pi \cdot \mathbf{h} = d\}$ is a hyperplane that separates \mathcal{X} from \mathcal{Y} , hence the name.)

Since $\mathcal{A}(\omega_0)$ is closed and convex, and because the set consisting of the single point $\beta \notin \mathcal{A}(\omega_0)$ is compact and convex, the separating hyperplane theorem implies that there exists $\langle \mathbf{h}, d \rangle$ such that

$$\forall \alpha \in \mathcal{A}(\omega_0), \quad \mathbf{h} \cdot \alpha > d \text{ and } \mathbf{h} \cdot \beta < d. \tag{B.9}$$

Since $\mathbf{0} \in \mathcal{A}(\omega_0)$, we know that $d < 0$. However, we claim that for any $\alpha \in \mathcal{A}(\omega_0)$, we in fact have $\mathbf{h} \cdot \alpha \geq 0$. Indeed, if we assume that there exists $\alpha \in \mathcal{A}(\omega_0)$ such that $\mathbf{h} \cdot \alpha < 0$, then for any $\lambda \geq \frac{|d|}{|\alpha \cdot \mathbf{h}|}$, eq. (4.6) implies that $\lambda \alpha \in \mathcal{A}(\omega_0)$. But $\lambda \alpha \cdot \mathbf{h} \leq -|d|$, which contradicts the fact that $\alpha \cdot \mathbf{h} > d$. Therefore, we have $\alpha \cdot \mathbf{h} \geq 0$ and $\beta \cdot \mathbf{h} < 0$, implying that $\langle \beta, 0 \rangle \not\leq \omega_0$. This proves proposition 4.1. \blacksquare

B.1.4 Proof of proposition 4.2

In order to reduce the general Ω defined in eq. (4.9) to one that can be handled by proposition 4.1 for $c = 0$, we define the “lifted” vector

$$\ell_x(\mathbf{v}) = (v_1, v_2, \dots, v_D, x), \tag{B.10}$$

where D is the dimension of the original vector \mathbf{v} . Then we note that

$$\alpha \cdot \mathbf{b} \geq c \iff \ell_{-c}(\alpha) \cdot \ell_1(\mathbf{b}) = \alpha \cdot \mathbf{b} - c \geq 0. \tag{B.11}$$

³⁸The same method is easy to apply to the $c = -1$ case and, with some slight complications, the $c = 1$ case. With considerable effort, it can also be extended to proposition 4.2. However, we will follow the outline eq. (B.1) and only prove what is necessary.

Thus, any D -dimensional linear constraint $\langle \boldsymbol{\alpha}, c \rangle$ can be lifted into a $(D + 1)$ -dimensional linear constraint $\langle \ell_{-c}(\boldsymbol{\alpha}), 0 \rangle$. We can now show that proposition 4.1 in the lifted space, where we only ever have $c = 0$, is equivalent to proposition 4.2 in the original space. Define

$$\omega_0^\ell = \sum_{i \in I_1} \langle \ell_{-1}(\boldsymbol{\alpha}_i), 0 \rangle + \sum_{i \in I_0} \langle \ell_0(\boldsymbol{\alpha}_i), 0 \rangle + \sum_{i \in I_{-1}} \langle \ell_1(\boldsymbol{\alpha}_i), 0 \rangle + \langle \ell_1(\mathbf{0}), 0 \rangle, \quad (\text{B.12})$$

where the extra constraint $\langle \ell_1(\mathbf{0}), 0 \rangle$ imposes that $\ell_x(\mathbf{b})$ only satisfies ω_0^ℓ if $x \geq 0$. ω_0^ℓ can be thought of as a lifted version of Ω , and as indicated by the notation, it fulfils the definition of ω_0 so that proposition 4.1 for $c = 0$ applies to it.

Now, assume that $\ell_{-c}(\boldsymbol{\beta}) \in \mathcal{A}(\omega_0^\ell)$. Looking at eq. (4.6), we find that

$$\begin{aligned} \boldsymbol{\beta} &= \sum_{i \in I_1} \lambda_i \boldsymbol{\alpha}_i + \sum_{i \in I_0} \lambda_i \boldsymbol{\alpha}_i + \sum_{i \in I_{-1}} \lambda_i \boldsymbol{\alpha}_i + \lambda' \mathbf{0}, \\ -c &= \sum_{i \in I_{-1}} \lambda_i - \sum_{i \in I_1} \lambda_i + \lambda', \quad \lambda_i \geq 0, \lambda' \geq 0. \end{aligned} \quad (\text{B.13})$$

These exactly reproduce eqs. (4.10) and (4.11), so we have shown that $\ell_{-c}(\boldsymbol{\beta}) \in \mathcal{A}(\omega_0^\ell)$ implies $\boldsymbol{\beta} \in \mathcal{A}_c(\Omega)$.

Conversely, let us now assume, $\ell_{-c}(\boldsymbol{\beta}) \notin \mathcal{A}(\omega_0^\ell)$. Proposition 4.1 then implies the existence of some $\ell_x(\mathbf{b}) \in \mathcal{B}(\omega_0^\ell)$ with $x \geq 0$ such that $\ell_{-c}(\boldsymbol{\beta}) \cdot \ell_x(\mathbf{b}) < 0$, and for any $i \in I$, $\ell_{-c_i}(\boldsymbol{\alpha}_i) \cdot \ell_x(\mathbf{b}) \geq 0$. We may moreover assume that $x > 0$, for if $x = 0$, we may choose any $\mathbf{a} \in \mathcal{B}(\Omega)$ and consider $\ell_\varepsilon(\mathbf{b}') = (1 - \varepsilon)\ell_0(\mathbf{b}) + \varepsilon\ell_1(\mathbf{a})$ for $0 < \varepsilon < 1$. By convexity, $\ell_\varepsilon(\mathbf{b}') \in \mathcal{B}(\omega_0^\ell)$, and

$$\ell_{-c}(\boldsymbol{\beta}) \cdot \ell_\varepsilon(\mathbf{b}') = \boldsymbol{\beta} \cdot \mathbf{b} + \varepsilon[\boldsymbol{\beta} \cdot (\mathbf{a} - \mathbf{b}) - c], \quad (\text{B.14})$$

so for ε small enough, $\ell_{-c}(\boldsymbol{\beta}) \cdot \ell_\varepsilon(\mathbf{b}') < 0$. Consequently,

$$\begin{cases} \boldsymbol{\alpha}_i \cdot \mathbf{b} \geq x, & i \in I_1, \\ \boldsymbol{\alpha}_i \cdot \mathbf{b} \geq 0, & i \in I_0, \\ \boldsymbol{\alpha}_i \cdot \mathbf{b} \geq -x, & i \in I_{-1}, \end{cases} \quad \boldsymbol{\beta} \cdot \mathbf{b} < xc. \quad (\text{B.15})$$

Thus, $\frac{1}{x}\mathbf{b}$ is a point that satisfies Ω but not $\langle \boldsymbol{\beta}, c \rangle$, which means that $\ell_{-c}(\boldsymbol{\beta}) \notin \mathcal{A}(\omega_0^\ell)$ implies $\boldsymbol{\beta} \notin \mathcal{A}_c(\Omega)$.

We have now shown that $\ell_{-c}(\boldsymbol{\beta}) \in \mathcal{A}(\omega_0^\ell)$ is equivalent to $\boldsymbol{\beta} \in \mathcal{A}_c(\Omega)$, and since $\langle \ell_{-c}(\boldsymbol{\beta}), 0 \rangle \leq \omega_0^\ell$ is equivalent to $\langle \boldsymbol{\beta}, c \rangle \leq \Omega$, we have therefore proven proposition 4.2 as a consequence of the $c = 0$ version of proposition 4.1. As mentioned before, proposition 4.1 for $c = \pm 1$ follows easily. ■

B.1.5 An important corollary

The following interesting result is a consequence of proposition 4.2:

Corollary B.1 (boundedness of $\mathcal{B}(\Omega)$). *For $\Omega \neq \Omega_\infty$, the region $\mathcal{B}(\Omega)$ is bounded if and only if the origin, $\mathbf{0}$, is in the interior of $\mathcal{A}_{-1}(\Omega)$. This happens if and only if $\{\boldsymbol{\alpha}_i\}_{i \in I}$ spans the full D -dimensional space and there are λ_i such that*

$$\mathbf{0} = \sum_{i \in I} \lambda_i \boldsymbol{\alpha}_i, \quad \lambda_i > 0 \quad (\text{B.16})$$

PROOF. We will show the converse, namely that $\mathcal{B}(\Omega)$ being unbounded is equivalent to $\mathbf{0}$ not being in the interior of $\mathcal{A}_{-1}(\Omega)$.

First assume that $\mathbf{0}$ is not in the interior of $\mathcal{A}_{-1}(\Omega)$. Then for all $\varepsilon > 0$, there exists some β_ε such that $|\beta_\varepsilon| < \varepsilon$ and $\beta_\varepsilon \notin \mathcal{A}_{-1}(\Omega)$. Since $\langle \beta_\varepsilon, -1 \rangle \not\leq \Omega$, there consequently exists a point $\mathbf{b}_\varepsilon \in \mathcal{B}(\Omega)$ such that $\beta_\varepsilon \cdot \mathbf{b}_\varepsilon < -1$. Now, the Cauchy-Schwarz inequality gives

$$1 < |\beta_\varepsilon \cdot \mathbf{b}_\varepsilon| \leq |\beta_\varepsilon| |\mathbf{b}_\varepsilon| < \varepsilon |\mathbf{b}_\varepsilon|. \quad (\text{B.17})$$

Since this holds for arbitrarily small ε , there can be no upper bound on $|\mathbf{b}_\varepsilon|$; therefore, $\mathcal{B}(\Omega)$ is unbounded.

Conversely, assume that $\mathcal{B}(\Omega)$ is unbounded. Then for all $M > 0$, there must exist some $\mathbf{b}_M \in \mathcal{B}(\Omega)$ such that $|\mathbf{b}_M| > M$. Now define

$$\beta_M \equiv -\frac{\mathbf{b}_M}{|\mathbf{b}_M|^{3/2}} \quad \Rightarrow \quad |\beta_M| = \frac{1}{|\mathbf{b}_M|^{1/2}} < \frac{1}{\sqrt{M}}, \quad \beta_M \cdot \mathbf{b}_M = -|\mathbf{b}_M|^{1/2} < -\sqrt{M}. \quad (\text{B.18})$$

For sufficiently large M , the last inequality implies that $\langle \beta_M, -1 \rangle \not\leq \Omega$, so $\beta_M \notin \mathcal{A}_{-1}(\Omega)$. However, the inequality before that tells us that β_M may lie arbitrarily close to the origin. Therefore, $\mathbf{0} \in \mathcal{A}_{-1}(\Omega)$ must lie on the boundary, not the interior, of $\mathcal{A}_{-1}(\Omega)$.

That completes the main proof, but we must also prove condition about the span of $\{\alpha_i\}_{i \in I}$. If the span was lower-dimensional, then there would exist some vector β linearly independent of all α_i , and then clearly $\varepsilon\beta \notin \mathcal{A}_{-1}(\Omega)$ for all $\varepsilon > 0$, implying that $\mathbf{0}$ is not in the interior of $\mathcal{A}_{-1}(\Omega)$. Lastly, we must prove eq. (B.16), which is essentially eq. (4.12) with $\lambda_i > 0$ rather than $\lambda_i \geq 0$. Since $\{\alpha_i\}_{i \in I}$ spans the full space, any vector \mathbf{v} of sufficiently small magnitude satisfies

$$\mathbf{v} \in \text{Hull}(\{\alpha_i\}_{i \in I} \cup \{-\sum_{i \in I} \alpha_i\}). \quad (\text{B.19})$$

Thus, $\mathbf{0}$ is in the interior of $\mathcal{A}_{-1}(\Omega)$ if and only if for sufficiently small $\varepsilon > 0$,³⁹

$$-\varepsilon \sum_{i \in I} \alpha_i \in \mathcal{A}_{-1}(\Omega) \quad \Leftrightarrow \quad -\varepsilon \sum_{i \in I} \alpha_i = \sum_{i \in I} \lambda_i \alpha_i \quad \Leftrightarrow \quad \mathbf{0} = \sum_{i \in I} (\lambda_i + \varepsilon) \alpha_i \quad (\text{B.20})$$

for $\lambda_i \geq 0$, which implies $\lambda_i + \varepsilon > 0$, thereby producing eq. (B.16). ■

B.2 Some mathematical tools

Before moving on with proving proposition 4.3 and deriving further results, we need to establish some tools and terminology that range from useful to crucial in subsequent sections.

B.2.1 The degenerate constraint framework

In this section, we properly define what it means for a constraint to be degenerate, and derive notations and results that are not only useful for the proof of proposition 4.3 and its generalisation, but also for many other things later in this appendix.

³⁹Here, we neglect the condition $\sum_{i \in I} \lambda_i c_i \geq -1$ in the second equality, since it is always possible to multiply both sides by a positive factor to rescale the λ_i appropriately.

In a D -dimensional space, consider an affine subspace E of dimension d . There exists two sets of vectors $\{\mathbf{g}_j\}_{j=1}^d, \{\boldsymbol{\delta}_k\}_{k=1}^{D-d}$ whose union forms an orthonormal basis for \mathbb{R}^D , that, given an arbitrary point $\mathbf{e} \in E$, allow E to be expressed in two complementary ways:⁴⁰

$$E = \left\{ \mathbf{e} + \sum_{j=1}^d x_j \mathbf{g}_j \mid x_j \in \mathbb{R} \right\}, \quad E = \mathcal{B} \left[\sum_{k=1}^{D-d} (\langle \boldsymbol{\delta}_k, \boldsymbol{\delta}_k \cdot \mathbf{e} \rangle + \langle -\boldsymbol{\delta}_k, -\boldsymbol{\delta}_k \cdot \mathbf{e} \rangle) \right]. \quad (\text{B.21})$$

Up to the choice of \mathbf{e} and the basis vectors, any vector \mathbf{b} in parameter space can be uniquely decomposed as

$$\mathbf{b} = \mathbf{e} + \sum_{j=1}^d x_j \mathbf{g}_j + \sum_{k=1}^{D-d} z_k \boldsymbol{\delta}_k. \quad (\text{B.22})$$

We then define $[\mathbf{b}]_E \equiv (x_1, x_2, \dots, x_d)$ and $[\mathbf{b}]^E \equiv (z_1, z_2, \dots, z_{D-d})$. These are d - and $(D-d)$ -dimensional vectors, respectively, and live in spaces separate from the D -dimensional space in which E, \mathbf{e} , etc. live. Note that if $\mathbf{b} \in E$, then $[\mathbf{b}]^E = \mathbf{0}$ and \mathbf{b} is uniquely determined by $[\mathbf{b}]_E$. For constraint-space vectors, we instead make the decomposition

$$\boldsymbol{\alpha} = \sum_{j=1}^d \xi_j \mathbf{g}_j + \sum_{k=1}^{D-d} \zeta_k \boldsymbol{\delta}_k \quad (\text{B.23})$$

and analogously define $[\boldsymbol{\alpha}]_E \equiv (\xi_1, \xi_2, \dots, \xi_d)$ and $[\boldsymbol{\alpha}]^E \equiv (\zeta_1, \zeta_2, \dots, \zeta_{D-d})$.⁴¹ These can form constraints acting on $[\mathbf{b}]_E$ and $[\mathbf{b}]^E$, respectively. As an extension, for any set \mathcal{X} we define

$$[\mathcal{X}]_E \equiv \left\{ [\mathbf{x}]_E \mid \mathbf{x} \in \mathcal{X} \right\}, \quad [\mathcal{X}]^E \equiv \left\{ [\mathbf{x}]^E \mid \mathbf{x} \in \mathcal{X} \right\}. \quad (\text{B.24})$$

We reiterate how important it is to view $\mathcal{X}, [\mathcal{X}]_E$ and $[\mathcal{X}]^E$ as living in three different spaces. There is of course a straightforward mapping between $[\mathcal{X}]_E$ and $\mathcal{X} \cap E$ — indeed, the d -dimensional space can be seen as the vector space underlying the affine subspace E — but the notion of separate spaces makes the proofs below clearer.

Before moving on to constraints, let us make the following definition:

The *dimension* of any nonempty set \mathcal{X} , written $\dim(\mathcal{X})$, is the affine dimension of the smallest (i.e. lowest-dimensional) affine subspace that contains \mathcal{X} .

Equivalently, $\dim(\mathcal{X})$ is the dimension of the affine span of the points in \mathcal{X} .⁴²

⁴⁰Note that we have written \mathbf{g}_j as parameter-space vectors, and $\boldsymbol{\delta}_k$ as constraint-space vectors. This is consistent with their use in eq. (B.21), but eqs. (B.22) and (B.23) are in a sense breaking our conventions by adding vectors of different types. This is of course no problem when both parameter and constraint space are just \mathbb{R}^D , but if we considered constraints in more general spaces, we would have to make appropriate adjustments to our formulae.

⁴¹Our notation does not make the choice of reference point \mathbf{e} explicit, and it is arbitrary for all purposes. Replacing $\mathbf{e} \rightarrow \mathbf{e}'$ simply entails translating all $[\mathbf{b}]_E \rightarrow [\mathbf{b}]_E + [\mathbf{e}' - \mathbf{e}]_E$ while leaving $[\mathbf{b}]^E, [\boldsymbol{\alpha}]_E$ and $[\boldsymbol{\alpha}]^E$ unchanged. Likewise, altering $\{\mathbf{g}_j\}_{j=1}^d, \{\boldsymbol{\delta}_k\}_{k=1}^{D-d}$ just corresponds to a change of basis in the spaces.

⁴²Note that this definition agrees with the usual affine/linear dimension when \mathcal{X} is itself an affine/linear subspace.

This smallest affine subspace is clearly unique, for if it is not, the intersection of all such subspaces is even smaller. For any nonempty convex set $\mathcal{C} \in \mathbb{R}^D$, $D > 0$, the following basic fact holds:

$$\dim(\mathcal{C}) < D \quad \Leftrightarrow \quad \text{int}(\mathcal{C}) = \emptyset, \quad (\text{B.25})$$

where $\text{int}(\mathcal{C})$ is the interior of \mathcal{C} .

Consider now a constraint $\Omega \neq \Omega_\infty$ and define $d_\Omega \equiv \dim[\mathcal{B}(\Omega)]$. We formalise the definition of degeneracy stated in section 4.4 as follows:

A constraint $\Omega \neq \Omega_\infty$ in D -dimensional space is *degenerate* if $d_\Omega < D$ and *non-degenerate* otherwise. Ω_∞ , for which d_{Ω_∞} is undefined, does not fall into either category.

Let E_Ω be the unique d_Ω -dimensional affine subspace that contains $\mathcal{B}(\Omega)$. Given linear constraint $\langle \alpha, c \rangle$, we define

$$\lfloor \langle \alpha, c \rangle \rfloor_\Omega \equiv \langle \lfloor \alpha \rfloor_{E_\Omega}, \quad c - \alpha \cdot \mathbf{e} \rangle, \quad \lceil \langle \alpha, c \rangle \rceil^\Omega \equiv \langle \lceil \alpha \rceil^{E_\Omega}, \quad c - \alpha \cdot \mathbf{e} \rangle. \quad (\text{B.26})$$

For these, the following holds:

Lemma B.1. *For any point \mathbf{b} ,*

(a) *If $\lceil \mathbf{b} \rceil^{E_\Omega} = \mathbf{0}$, then $\lfloor \mathbf{b} \rfloor_{E_\Omega}$ satisfies $\lfloor \langle \alpha, c \rangle \rfloor_\Omega$ if and only if \mathbf{b} satisfies $\langle \alpha, c \rangle$.*

(b) *If $\lfloor \alpha \rfloor_{E_\Omega} = \mathbf{0}$, then $\lceil \mathbf{b} \rceil^{E_\Omega}$ satisfies $\lceil \langle \alpha, c \rangle \rceil^\Omega$ if and only if \mathbf{b} satisfies $\langle \alpha, c \rangle$.*

PROOF. This follows directly from eqs. (B.22), (B.23) and (B.26) and the orthonormality of $\{\mathbf{g}_j\}_{j=1}^d \cup \{\delta_k\}_{k=1}^{D-d}$. \square

Using this, we define for any $\Omega \neq \Omega_\infty$ with representation \mathcal{S}

$$\lfloor \Omega \rfloor \equiv \sum_{\langle \alpha, c \rangle \in \mathcal{S}} \lfloor \langle \alpha, c \rangle \rfloor_\Omega. \quad (\text{B.27})$$

This constraint, which acts on the d_Ω -dimensional space of vectors $\lfloor \mathbf{b} \rfloor_{E_\Omega}$, has three important properties. Firstly, $\lfloor \Omega \rfloor$ is satisfied by $\lfloor \mathbf{b} \rfloor_{E_\Omega}$ if Ω is satisfied by \mathbf{b} , and the converse holds when $\lceil \mathbf{b} \rceil^{E_\Omega} = \mathbf{0}$ (this follows from lemma B.1), so

$$\mathcal{B}(\lfloor \Omega \rfloor) = \lfloor \mathcal{B}(\Omega) \rfloor_{E_\Omega}. \quad (\text{B.28})$$

Secondly, $\lfloor \Omega \rfloor$ is independent of \mathcal{S} as a consequence of this. Lastly, $\lfloor \Omega \rfloor$ is, by construction, non-degenerate.⁴³ Thanks to these properties, $\lfloor \Omega \rfloor$ is key to all further treatment of degenerate constraints.

⁴³If $d_\Omega = 0$, so that $\mathcal{B}(\Omega)$ is a single point, then $\lfloor \Omega \rfloor$ is a zero-dimensional constraint. This is not conceptually a problem for proposition B.1 below: $\lfloor \Omega \rfloor$ is satisfied by $\mathbf{0}$, which is the only point in zero-dimensional space, and $\mathcal{R}(\lfloor \Omega \rfloor) = \emptyset$.

B.2.2 K -faces

In this section, we introduce K -faces, which will be highly useful in subsequent sections; especially corollary B.2 and propositions B.3 and B.4 rely heavily on them. Like several other things introduced here, it is partially based on standard concepts and nomenclature, but has been adapted and extended to fit the context of linear constraints. K -faces are defined as follows:

Let \mathcal{F} be a non-empty convex subset of a closed convex set \mathcal{C} , with $K = \dim(\mathcal{F})$. Then \mathcal{F} is called a K -face of \mathcal{C} if the following holds: for every $\varphi \in \mathcal{F}$, if there exists $\eta_{1,2} \in \mathcal{C}$ such that $\varphi = \mu\eta_1 + (1 - \mu)\eta_2$ with $\mu \in (0, 1)$, then $\eta_{1,2} \in \mathcal{F}$.

The single point in a 0-face of \mathcal{C} is called an *extreme point* or *vertex* of \mathcal{C} ; it is a point that cannot be expressed as a convex combination of any two points in \mathcal{C} distinct from itself.

A 1-face of \mathcal{C} is called an *edge*.

A $(D - 1)$ -face of \mathcal{C} is called a *facet*.

For convex polygons, polyhedra, etc., these definitions agree with the usual concepts of vertices, edges and facets. A number of useful properties of K -faces easily follow from the definitions:

- (i) No K -face of \mathcal{C} is a strict subset of another K -face of \mathcal{C} with the same K .
- (ii) There is a unique $\dim(\mathcal{C})$ -face of \mathcal{C} , namely \mathcal{C} itself.
- (iii) All K -faces of \mathcal{C} (except possibly \mathcal{C} itself) are contained in the boundary of \mathcal{C} .
- (iv) If \mathcal{F} is a K -face of \mathcal{C} , then for all $K' \leq K$ the K' -faces of \mathcal{F} are also K' -faces of \mathcal{C} . Specifically, an edge may have up to two vertices, which are its endpoints.
- (v) If \mathcal{F} is a K -face of \mathcal{C} and $\mathcal{C}' \subseteq \mathcal{C}$, then if $\mathcal{F} \cap \mathcal{C}'$ is non-empty, it is a K' -face ($K' \leq K$) of \mathcal{C}' .

We will now prove some less obvious properties. In the remainder of this section, let \mathcal{C} be any convex set such that there exists a finite set of constraints $\{\langle \mathbf{g}_\ell, \mathbf{c}_\ell \rangle\}_{\ell \in L}$ fulfilling⁴⁴

$$\mathcal{C} = \mathcal{B} \left(\sum_{\ell \in L} \langle \mathbf{g}_\ell, \mathbf{c}_\ell \rangle \right). \tag{B.29}$$

We then begin with the following technical lemma:

Lemma B.2. *Let \mathcal{C} be a convex set defined as in eq. (B.29), and let $\mathcal{F} \subseteq \mathcal{C}$ be non-empty with dimension $K < D$. Then \mathcal{F} is a K -face of \mathcal{C} if and only if there exists $J \subset L$ with $|J| = D - K$ such that $\{\mathbf{g}_j\}_{j \in J}$ are linearly independent and*

$$\forall \varphi \in \mathcal{F}, \forall j \in J, \quad \mathbf{g}_j \cdot \varphi = c_j. \tag{B.30}$$

⁴⁴This class of sets includes all convex hulls of finite sets, all linear and affine subspaces, as well as most other sets we work with, including $\mathcal{A}_c(\Omega)$ as we will prove in appendix B.4. For convenience, we express \mathcal{C} as a subset of constraint space, but all results hold equally well if \mathcal{C} is a subset of parameter space.

Note that there may be $\ell \in L \setminus J$ such that $\mathbf{g}_\ell \cdot \boldsymbol{\varphi} = c_\ell$. All such \mathbf{g}_ℓ are contained in the span of $\{\mathbf{g}_j\}_{j \in J}$, though.

PROOF. If such a J exists, then $\mathcal{F}' \equiv \{\boldsymbol{\varphi} \in \mathbb{R}^D \mid \forall j \in J, \mathbf{g}_j \cdot \boldsymbol{\varphi} = c_j\}$ is clearly an affine subspace of dimension K . Let $\boldsymbol{\varphi} \in \mathcal{F}'$ and $\boldsymbol{\eta}_{1,2} \in \mathcal{C}$, and assume $\boldsymbol{\varphi} = \mu \boldsymbol{\eta}_1 + (1 - \mu) \boldsymbol{\eta}_2$ with $\mu \in (0, 1)$. Then

$$\mathbf{g}_j \cdot \boldsymbol{\varphi} = d = \mathbf{g}_j \cdot [\mu \boldsymbol{\eta}_1 + (1 - \mu) \boldsymbol{\eta}_2] \tag{B.31}$$

and since $\mathbf{g}_j \cdot \boldsymbol{\eta}_{1,2} \geq c_j$, this implies $\mathbf{g}_j \cdot \boldsymbol{\eta}_{1,2} = c_j$. Therefore $\boldsymbol{\eta}_{1,2} \in \mathcal{F}'$, so $\mathcal{F} = \mathcal{F}' \cap \mathcal{C}$ is a K -face.

For the less straightforward converse, let $\omega = \sum_{\ell \in L} \langle \mathbf{g}_\ell, c_\ell \rangle$, with $\mathcal{C} = \mathcal{B}(\omega)$. If ω is degenerate, then \mathcal{C} clearly has no K -faces for $K > d_\omega$, and the unique d_ω -face is \mathcal{C} itself by property (ii). In that case, the proof follows trivially from eq. (B.21). For the remaining cases, we may substitute $\omega \rightarrow \lfloor \omega \rfloor, \mathcal{C} \rightarrow \lfloor \mathcal{C} \rfloor_\omega = \mathcal{B}(\lfloor \omega \rfloor)$ and thus assume without loss of generality that ω is non-degenerate.

Let us then note that J cannot be empty. Indeed, if it were the case, then $\boldsymbol{\varphi} \in \mathcal{F}$ would be in the interior of \mathcal{C} which via property (iii) contradicts the fact that \mathcal{F} is a K -face with $K < D$. Also, $|J| \leq D - K$, since otherwise the set $\{\boldsymbol{\chi} \in \mathcal{C} \mid \forall j \in J, \boldsymbol{\chi} \cdot \mathbf{g}_j = c_j\}$, being the intersection of more than $D - K$ independent hyperplanes, would have dimension less than K . We will then proceed by induction on $D > K$. The result is trivial in $K + 1$ dimensions, since $|J| = 1$ is guaranteed by J being nonempty.

Assume then that the lemma holds in n dimensions, and consider \mathcal{C}, \mathcal{F} in $(n + 1)$ -dimensional space. Since $J \neq \emptyset$, there is some $i \in L$ such that $\forall \boldsymbol{\varphi} \in \mathcal{F}, \mathbf{g}_i \cdot \boldsymbol{\varphi} = c_i$. Then consider

$$E_i \equiv \mathcal{B}(\langle \mathbf{g}_i, c_i \rangle + \langle -\mathbf{g}_i, -c_i \rangle), \quad \lfloor \mathcal{C} \rfloor_{E_i} = \mathcal{B}\left(\sum_{\ell \in L} \lfloor \langle \mathbf{g}_\ell, c_\ell \rangle \rfloor_{E_i}\right). \tag{B.32}$$

By construction, $\lfloor \mathcal{F} \rfloor_{E_i}$ is a K -face of $\lfloor \mathcal{C} \rfloor_{E_i}$. Since these are sets in a n -dimensional space, we know by the induction hypothesis that we have J' with $|J'| = n - K$ such that $\{\lfloor \mathbf{g}_j \rfloor_{E_i}\}_{j \in J'}$ are linearly independent and eq. (B.30) is satisfied. Furthermore, $\lfloor \mathbf{g}_i \rfloor_{E_i} = \mathbf{0}$, which cannot be expressed as a linear combination of $\{\lfloor \mathbf{g}_j \rfloor_{E_i}\}_{j \in J'}$ with nonzero coefficients. Therefore, \mathbf{g}_i is linearly independent of $\{\mathbf{g}_j\}_{j \in J'}$. Thus, $J = J' \cup \{i\}$ has $|J| = n + 1$, $\{\mathbf{g}_j\}_{j \in J}$ linearly independent and satisfies eq. (B.30). This proves that the lemma holds in $n + 1$ dimensions, and completes the induction. \square

Based on this, we can prove two more interesting lemmata:

Lemma B.3. *Let \mathcal{C} be a convex set satisfying eq. (B.29). Let \mathcal{F} be a nonempty convex subset of \mathcal{C} , and let $K = \dim(\mathcal{F})$. Then \mathcal{F} is a K -face of \mathcal{C} if and only if there is some constraint $\langle \mathbf{h}, d \rangle$ such that $\mathcal{C} \subseteq \mathcal{B}(\langle \mathbf{h}, d \rangle)$ and $\{\boldsymbol{\chi} \in \mathcal{C} \mid \mathbf{h} \cdot \boldsymbol{\chi} = d\} = \mathcal{F}$.⁴⁵*

PROOF. First note that this trivially holds when $\{\boldsymbol{\chi} \in \mathcal{C} \mid \mathbf{h} \cdot \boldsymbol{\chi} = d\} = \mathcal{C}$ (compare property (ii)), which may happen even when \mathcal{C} is not contained in any hyperplane if $\langle \mathbf{h}, d \rangle = \langle \mathbf{0}, 0 \rangle$.

⁴⁵Although we do not use it, this holds for $K = D - 1$ for any convex set \mathcal{C} , not just those satisfying eq. (B.29) (this is proven in a later footnote). This is not the case for smaller K : take e.g. the $D = 2$ example $\mathcal{C} = \{(x, y) \in \mathbb{R}^2 \mid y \geq \max(0, x^3)\}$, for which $(0, 0)$ is an extreme point but $\langle \mathbf{h}, d \rangle$ does not exist.

Setting aside the trivial cases, assume that such a $\langle \mathbf{h}, d \rangle$ exists. Then \mathcal{F} is a K -face by essentially the same argument that was made around eq. (B.31).

Conversely, if \mathcal{F} is a K -face then lemma B.2 holds. In the notation of that lemma, let

$$\mathbf{h} = \sum_{\ell \in J} \mathbf{g}_\ell, \quad d = \sum_{\ell \in J} c_\ell, \quad (\text{B.33})$$

so that $\forall \varphi \in \mathcal{F}, \mathbf{h} \cdot \varphi = d$. Let us now consider any $\boldsymbol{\eta} \in \mathcal{C}$ such that $\mathbf{h} \cdot \boldsymbol{\eta} = d$. Let E_J be defined as in eq. (B.32). Then by construction, $[\mathcal{F}]_{E_J}$ has the same dimension as its native space, i.e. $D - K$, and has nonempty interior by eq. (B.25) (recall that $D - K > 0$). Thus, there exists $\varphi \in \mathcal{F}$ such that $[\varphi]_{E_J}$ is in the interior of $[\mathcal{F}]_{E_J}$. For $\varepsilon > 0$ small enough, we therefore have

$$[\boldsymbol{\eta}_\varepsilon]_{E_J} \equiv \varepsilon [\boldsymbol{\eta}]_{E_J} + (1 - \varepsilon) [\varphi]_{E_J} \in [\mathcal{F}]_{E_J}. \quad (\text{B.34})$$

Since by construction $[\boldsymbol{\chi}]^{E_J} = \mathbf{0}$ for all $\boldsymbol{\chi}$ such that $\mathbf{h} \cdot \boldsymbol{\chi} = d$, lemma B.1 implies that

$$\boldsymbol{\eta}_\varepsilon = \varepsilon \boldsymbol{\eta} + (1 - \varepsilon) \varphi. \quad (\text{B.35})$$

By the definition of a K -face, this implies that $\boldsymbol{\eta} \in \mathcal{F}$, which concludes the proof.⁴⁶ \square

Lemma B.4. *Let $\mathcal{C} \subseteq \mathcal{C}'$ be convex sets satisfying eq. (B.29), and let \mathcal{F} be a K -face of \mathcal{C} . Then there exists a K -face \mathcal{F}' of \mathcal{C}' with $\mathcal{F} \subseteq \mathcal{F}'$ if there is some $\varphi \in \mathcal{F}$ and $\varepsilon > 0$ such that $\{\boldsymbol{\chi} \in \mathcal{C}' \mid \varepsilon \geq |\boldsymbol{\chi} - \varphi|\} \subseteq \mathcal{C}$.*

Note how this complements property (v). Note also that $\mathcal{F}' \supset \mathcal{F}$ may exist even if \mathcal{F} fails to satisfy the given conditions.

PROOF. By lemma B.3, there exists $\langle \mathbf{h}, d \rangle$ such that $\{\boldsymbol{\chi} \in \mathcal{C} \mid \mathbf{h} \cdot \boldsymbol{\chi} = d\} = \mathcal{F}$ and $\mathcal{C} \subseteq \mathcal{B}(\langle \mathbf{h}, d \rangle)$. We moreover claim that $\mathcal{C}' \subseteq \mathcal{B}(\langle \mathbf{h}, d \rangle)$. Indeed, let $\varphi \in \mathcal{F}$ and $\varepsilon > 0$ such that $\{\boldsymbol{\chi} \in \mathcal{C}' \mid \varepsilon \geq |\boldsymbol{\chi} - \varphi|\} \subseteq \mathcal{C}$. For any $\boldsymbol{\eta} \in \mathcal{C}'$ it follows from the triangle inequality that

$$\left| \varphi - \left[\left(1 - \frac{\varepsilon}{|\varphi| + |\boldsymbol{\eta}|}\right) \varphi + \frac{\varepsilon}{|\varphi| + |\boldsymbol{\eta}|} \boldsymbol{\eta} \right] \right| = \varepsilon \frac{|\varphi - \boldsymbol{\eta}|}{|\varphi| + |\boldsymbol{\eta}|} \leq \varepsilon, \quad (\text{B.36})$$

so $\left(1 - \frac{\varepsilon}{|\varphi| + |\boldsymbol{\eta}|}\right) \varphi + \frac{\varepsilon}{|\varphi| + |\boldsymbol{\eta}|} \boldsymbol{\eta} \in \mathcal{C}$. In particular,

$$\mathbf{h} \cdot \left[\left(1 - \frac{\varepsilon}{|\varphi| + |\boldsymbol{\eta}|}\right) \varphi + \frac{\varepsilon}{|\varphi| + |\boldsymbol{\eta}|} \boldsymbol{\eta} \right] \geq d \quad \Rightarrow \quad \mathbf{h} \cdot \boldsymbol{\eta} \geq d. \quad (\text{B.37})$$

Therefore, lemma B.3 guarantees that $\mathcal{F}' = \{\boldsymbol{\eta} \in \mathcal{C}' \mid \mathbf{h} \cdot \boldsymbol{\eta} = d\}$ is a K' -face for some K' . Since $\mathcal{F} \subseteq \mathcal{F}'$ we have $K \leq K'$, so it remains to prove that $K \geq K'$. For any $\varphi' \in \mathcal{F}'$, we have $\varphi + \varepsilon \frac{\varphi'}{|\varphi'|} \in \mathcal{F}$ since $\left| \varphi - \left(\varphi + \varepsilon \frac{\varphi'}{|\varphi'|} \right) \right| \leq \varepsilon$. However, since $-\varphi$ belongs to the span of \mathcal{F} , we can deduce that $\varepsilon \frac{\varphi'}{|\varphi'|}$, and therefore also φ' , belongs to the span of \mathcal{F} . Thus, the span of \mathcal{F}' is included in the span of \mathcal{F} , implying that $K \geq K'$, which completes the proof. \square

⁴⁶When $K = D - 1$, the following proof, which does not use lemma B.2 and therefore holds for all convex sets \mathcal{C} , works for the converse: let \mathcal{F}' be the affine span of \mathcal{F} , which is a hyperplane. By a variant of the separating hyperplane theorem (for any disjoint convex sets \mathcal{X} and \mathcal{Y} (no closedness/compactness needed), there exists $\langle \mathbf{h}, d \rangle$ with $\mathbf{h} \cdot \boldsymbol{\chi} \leq d'$ for all $\boldsymbol{\chi} \in \mathcal{X}$ and $\mathbf{h} \cdot \boldsymbol{\psi} \geq d'$ for all $\boldsymbol{\psi} \in \mathcal{Y}$) applied to the disjoint convex sets \mathcal{F}' and $\mathcal{C} \setminus \mathcal{F}$, we immediately find our desired $\langle \mathbf{h}, d \rangle$. This hinges on the properties of hyperplanes: the separating hyperplane must be parallel to \mathcal{F}' , or else they would intersect. Therefore, this does not work if $K < D - 1$.

B.3 Proof and generalisation of proposition 4.3

To simplify this proof, we will introduce the following terminology:

A linear constraint $\langle \alpha, c \rangle$ with $\alpha \neq \mathbf{0}$ supports a point \mathbf{b} if $\alpha \cdot \mathbf{b} = c$.⁴⁷

A linear constraint $\langle \mathbf{0}, c \rangle$ supports no point (this is natural for $c = \pm 1$, but we define it to be so also for $c = 0$; this simplifies most statements expressed in terms of support).

Given Ω , a linear constraint $\langle \alpha, c \rangle$ uniquely supports a point $\mathbf{b} \in \mathcal{B}(\Omega)$ if it supports \mathbf{b} , and there is no other $\langle \beta, d \rangle \leq \Omega$ with $\langle \beta, d \rangle \neq \langle \alpha, c \rangle$ that supports \mathbf{b} .

A constraint representation \mathcal{S} is non-redundant if it contains no trivial constraints, and there are no two elements $\langle \alpha, c \rangle \in \mathcal{S}$ and $\langle \beta, d \rangle \in \mathcal{S}$ such that $\langle \alpha, c \rangle = \langle \beta, d \rangle$.

Reducing a representation to a non-redundant one is of course trivial. Noting that the second paragraph of proposition 4.3 can be reduced to “The elements of $\mathcal{R}(\Omega)$ are exactly those $\langle \alpha, c \rangle \leq \Omega$ that uniquely support a point $\mathbf{b} \in \mathcal{B}(\Omega)$ ”, we will then begin with the following lemma:

Lemma B.5. *Let $\Omega = \sum_{i \in I} \langle \alpha_i, c_i \rangle$ be a non-degenerate constraint, as defined in section 4.4, and let the representation $\{\langle \alpha_i, c_i \rangle\}_{i \in I}$ be non-redundant. Then for any $j \in I$ and $\mathbf{b} \in \mathcal{B}(\Omega)$, $\langle \alpha_j, c_j \rangle$ uniquely supports \mathbf{b} if and only if $\alpha_j \cdot \mathbf{b} = c_j$ and $\alpha_i \cdot \mathbf{b} > c_i$ for all $i \neq j$.*

Note that compared to the definition of unique support, this only concerns the elements of a non-redundant representation rather than all $\langle \beta, c \rangle \leq \Omega$.

PROOF. Assume that $\alpha_j \cdot \mathbf{b} = c_j$ and $\alpha_i \cdot \mathbf{b} > c_i$ for all $i \neq j$, and assume there is some $\langle \beta, c \rangle < \Omega$ that supports \mathbf{b} , i.e. that $\langle \alpha_j, c_j \rangle$ does not support it uniquely. By proposition 4.2, there exist some positive numbers $\{\lambda_i\}_{i \in I}$ such that

$$c = \beta \cdot \mathbf{b} = \sum_{i \in I} \lambda_i \alpha_i \cdot \mathbf{b} \quad \Rightarrow \quad \begin{cases} c = \lambda_j c_j & \text{if } \lambda_i = 0 \text{ for all } i \neq j, \\ c > \sum_{i \in I} \lambda_i c_i \geq c & \text{otherwise.} \end{cases} \quad (\text{B.38})$$

The second case is a contradiction, so $\beta = \lambda_j \alpha_j$ and $c = \lambda_j c_j$ with $\lambda_j \geq 0$. If $c_j = \pm 1$, then either $\beta = -\alpha_j$ and $c = -c_j$ (contradicting non-degeneracy), or $\beta = \mathbf{0}$ (impossible since $\langle \mathbf{0}, \pm 1 \rangle$ supports no point), or $\beta = \alpha_j$ and $c = c_j$ (contradicting non-redundancy). If $c_j = 0$, then $c = 0$ and β is proportional to α_j (again contradicts non-redundancy, via eq. (4.5)). With no non-contradictory cases left, we have proven that $\langle \alpha_j, c_j \rangle$ supports \mathbf{b} uniquely if $\alpha_j \cdot \mathbf{b} = c_j$ and $\alpha_i \cdot \mathbf{b} > c_i$ for all $i \neq j$. The converse is trivial. \square

Let us then prove the following more significant lemma:

Lemma B.6. *Let $\langle \alpha, c \rangle$ be a linear constraint, and let Ω be a constraint such that $\Omega + \langle \alpha, c \rangle$ is non-degenerate. Assume that Ω uses a non-redundant representation \mathcal{S} and that $\mathcal{S} \cup \{\langle \alpha, c \rangle\}$ is a non-redundant representation for $\Omega + \langle \alpha, c \rangle$. Then $\Omega + \langle \alpha, c \rangle > \Omega$ if and only if $\langle \alpha, c \rangle$ uniquely supports some point $\mathbf{b} \in \mathcal{B}[\Omega + \langle \alpha, c \rangle]$.*

⁴⁷This is inspired by the standard concept of a supporting hyperplane.

PROOF. Assume that $\langle \alpha, c \rangle$ uniquely supports some point $\mathbf{b} \in \mathcal{B}(\Omega + \langle \alpha, c \rangle)$, so that $\alpha \cdot \mathbf{b} = c$ and $\beta \cdot \mathbf{b} > d$ for all $\langle \beta, d \rangle \in \mathcal{S}$. Then define

$$\mathbf{b}_\varepsilon \equiv \mathbf{b} - \varepsilon \frac{\alpha}{|\alpha|^2}, \quad \varepsilon > 0 \tag{B.39}$$

so that

$$\alpha \cdot \mathbf{b}_\varepsilon = c - \varepsilon < c, \quad \beta \cdot \mathbf{b}_\varepsilon > d - \varepsilon \frac{\beta \cdot \alpha}{|\alpha|^2}. \tag{B.40}$$

Thus, $\mathbf{b}_\varepsilon \notin \mathcal{B}(\Omega + \langle \alpha, c \rangle)$, but $\mathbf{b}_\varepsilon \in \mathcal{B}(\Omega)$ for sufficiently small ε . Then by definition, $\Omega + \langle \alpha, c \rangle > \Omega$.

Conversely, assume that $\Omega + \langle \alpha, c \rangle > \Omega$. Then there exists some $\mathbf{b} \in \mathcal{B}(\Omega)$ such that $\mathbf{b} \notin \mathcal{B}(\Omega + \langle \alpha, c \rangle)$. Furthermore, since $\Omega + \langle \alpha, c \rangle$ is non-degenerate, the interior of $\mathcal{B}(\Omega + \langle \alpha, c \rangle)$ is non-empty;⁴⁸ therefore, it contains some point \mathbf{n} . These points have the properties

$$\begin{aligned} \alpha \cdot \mathbf{b} < c, & \quad \beta \cdot \mathbf{b} \geq d, \\ \alpha \cdot \mathbf{n} > c, & \quad \beta \cdot \mathbf{n} > d. \end{aligned} \tag{B.41}$$

where again $\langle \beta, d \rangle \in \mathcal{S}$. By the intermediate value theorem, there must therefore exist some $\mu \in (0, 1)$ such that⁴⁹

$$\alpha \cdot [\mu \mathbf{b} + (1 - \mu) \mathbf{n}] = c, \quad \beta \cdot [\mu \mathbf{b} + (1 - \mu) \mathbf{n}] > d, \tag{B.42}$$

which, through lemma B.5, proves that $\langle \alpha, c \rangle$ supports $\mu \mathbf{b} + (1 - \mu) \mathbf{n}$ uniquely. \square

We now move on to proving proposition 4.3. Let a non-degenerate constraint Ω be expressed as

$$\Omega = \sum_{i \in I} \langle \alpha_i, c_i \rangle = \sum_{j \in J} \langle \gamma_j, d_j \rangle \tag{B.43}$$

where $\{\langle \alpha_i, c_i \rangle\}_{i \in I}$ is any representation, whereas $\{\langle \gamma_j, d_j \rangle\}_{j \in J}$ is minimal. A minimal representation must exist, since Ω can be written as a sum of a finite number of constraints. For any $k \in J$, minimality implies that

$$\sum_{j \in J} \langle \gamma_j, d_j \rangle > \sum_{\substack{j \in J \\ j \neq k}} \langle \gamma_j, d_j \rangle, \tag{B.44}$$

but by lemma B.6, $\langle \gamma_k, d_k \rangle$ must then uniquely support some $\mathbf{b} \in \mathcal{B}(\Omega)$. Then, there must also be some $i_k \in I$ such that $\langle \alpha_{i_k}, c_{i_k} \rangle$ also supports \mathbf{b} . To see this, assume that no $\langle \alpha_i, c_i \rangle, i \in I$ supports \mathbf{b} , and consider

$$\mathbf{b}_\varepsilon \equiv \mathbf{b} - \varepsilon \frac{\gamma_k}{|\gamma_k|^2} \quad \Rightarrow \quad \alpha_{i_k} \cdot \mathbf{b}_\varepsilon > c - \varepsilon \frac{\alpha_{i_k} \cdot \gamma_k}{|\gamma_k|^2}, \tag{B.45}$$

⁴⁸Recall that Ω being degenerate is equivalent to $\mathcal{B}(\Omega)$ being contained in a hyperplane, and a convex set has empty interior if and only if it is contained in a hyperplane; see also eq. (B.25). In general, the existence of a point in the interior of $\mathcal{B}(\Omega)$ is the *only* property of non-degenerate constraints used in this proof. Circumventing this requirement is key to proposition B.1 below.

⁴⁹Specifically, $\mu = \frac{\alpha \cdot \mathbf{n} - c}{\alpha \cdot \mathbf{n} - \alpha \cdot \mathbf{b}}$.

so that for sufficiently small ε , $\alpha_i \cdot \mathbf{b}_\varepsilon > c$ for all $i \in I$, and thus $\mathbf{b}_\varepsilon \in \mathcal{B}(\Omega)$. On the other hand,

$$\gamma_k \cdot \mathbf{b}_\varepsilon = d_k - \varepsilon \tag{B.46}$$

implying $\mathbf{b} \notin \mathcal{B}(\Omega)$, a contradiction. Therefore, $i_k \in I$ does exist. By the definition of unique support, we must then have $\langle \gamma_k, d_k \rangle = \langle \alpha_{i_k}, c_{i_k} \rangle$; that is, they are identical up to normalisation. By repeating this argument, we see that a distinct i_k exists for each $k \in J$, so $\{\langle \gamma_j, d_j \rangle\}_{j \in J} \subseteq \{\langle \alpha_i, c_i \rangle\}_{i \in I}$ up to normalisation. Having shown this, we may without loss of generality normalise and re-index $\langle \gamma_j, d_j \rangle$ so that $J \subseteq I$. Then

$$\begin{aligned} \Omega &= \sum_{i \in I} \langle \alpha_i, c_i \rangle \\ &= \sum_{i \in J} \langle \alpha_i, c_i \rangle + \sum_{i \in I \setminus J} \langle \alpha_i, c_i \rangle \\ &= \Omega + \sum_{i \in I \setminus J} \langle \alpha_i, c_i \rangle. \end{aligned} \tag{B.47}$$

By lemma B.6, the sum in the last line can only contain constraints that do not uniquely support any point. Thus, the minimal representation $\mathcal{R}(\Omega) = \{\langle \alpha_i, c_i \rangle\}_{i \in J}$ consists (up to normalisation) of *exactly* those elements of any representation that uniquely support a point. From this it follows that $\mathcal{R}(\Omega)$ consists of exactly all those $\langle \gamma, d \rangle \leq \Omega$ that uniquely support a point, since $\{\langle \alpha_i, c_i \rangle\}_{i \in I}$ could be made to include all $\langle \alpha, c \rangle \leq \Omega$, and from that it follows that $\mathcal{R}(\Omega)$ is unique up to normalisation. ■

B.3.1 The treatment of degenerate constraints

The following result generalises proposition 4.3 to all Ω :

Proposition B.1 (finding relevant constraints, general case). *Let Ω be any constraint in D -dimensional space. Then a minimal representation $\mathcal{R}(\Omega)$ can be determined as follows:*

- (i) *If Ω is non-degenerate, then proposition 4.3 applies. $\mathcal{R}(\Omega)$ is therefore unique up to normalisation, and is a subset of any representation of Ω .*
- (ii) *If $\Omega = \Omega_\infty$, then trivially $\mathcal{R}(\Omega) = \{\langle \mathbf{0}, 1 \rangle\}$. This is unique, but not necessarily a subset of other representations.*
- (iii) *If Ω is degenerate, let E_Ω be the unique d_Ω -dimensional affine subspace that contains $\mathcal{B}(\Omega)$. Let $\{\langle \alpha_i, c_i \rangle\}_{i \in I}$ be any set such that $\{\lfloor \langle \alpha_i, c_i \rangle \rfloor_\Omega\}_{i \in I} = \mathcal{R}(\lfloor \Omega \rfloor)$, with $\lfloor \langle \alpha_i, c_i \rangle \rfloor_\Omega \neq \lfloor \langle \alpha_j, c_j \rangle \rfloor_\Omega$ for all $i \neq j$. Let $\{\sigma_k\}_{k=0}^{d_\Omega}$ be any set of $(D - d_\Omega + 1)$ vectors with the following properties:*

- $\lfloor \sigma_k \rfloor_{E_\Omega} = \mathbf{0}$;
- The dimension of $\text{span}(\{\sigma_k\}_{k=0}^{D-d_\Omega})$ is $D - d_\Omega$;
- There exists a solution to

$$\mathbf{0} = \sum_{k=0}^{D-d_\Omega} \lambda_k \sigma_k, \quad \lambda_k > 0. \tag{B.48}$$

Then for arbitrary $\mathbf{e} \in E_\Omega$,

$$\mathcal{R}(\Omega) = \{ \langle \boldsymbol{\alpha}_i, c_i \rangle \}_{i \in I} \cup \{ \langle \boldsymbol{\sigma}_k, \boldsymbol{\sigma}_k \cdot \mathbf{e} \rangle \}_{k=0}^{D-d_\Omega}. \quad (\text{B.49})$$

This $\mathcal{R}(\Omega)$ is generally not unique, and is not necessarily a subset of any given representation of Ω . However, there is no minimal representation of Ω that is not of this form.

The non-uniqueness in case (iii) comes about in two ways. Firstly, $[\boldsymbol{\alpha}]^{E_\Omega}$ is arbitrary for $[\langle \boldsymbol{\alpha}, c \rangle]_\Omega \in \mathcal{R}([\Omega])$. Secondly, there is clearly freedom in the choice of $\{\boldsymbol{\sigma}_k\}_{k=0}^{D-d_\Omega}$. Given $\{\boldsymbol{\delta}_k\}_{k=1}^{D-d_\Omega}$ as defined above eq. (B.21), a straightforward choice is

$$\{\boldsymbol{\sigma}_k\}_{k=1}^{D-d_\Omega} = \{\boldsymbol{\delta}_k\}_{k=1}^{D-d_\Omega}, \quad -\boldsymbol{\sigma}_0 = \sum_{k=1}^{D-d_\Omega} \boldsymbol{\sigma}_k. \quad (\text{B.50})$$

PROOF. We only need to prove case (iii). For brevity, we will omit some sub/superscripts: $[\langle \boldsymbol{\alpha}, c \rangle]$ should be read as $[\langle \boldsymbol{\alpha}, c \rangle]_\Omega$, $[\mathbf{b}]$ as $[\mathbf{b}]^{E_\Omega}$, and so on.

Let $\{ \langle \boldsymbol{\alpha}_i, c_i \rangle \}_{i \in J}$ be a minimal representation of Ω , and subdivide it as

$$J_1 \equiv \{ i \in J \mid [\langle \boldsymbol{\alpha}_i, c_i \rangle] \neq \langle \mathbf{0}, 0 \rangle \}, \quad J_2 \equiv \{ i \in J \mid [\langle \boldsymbol{\alpha}_i, c_i \rangle] = \langle \mathbf{0}, 0 \rangle \}, \quad (\text{B.51})$$

for which the following holds:

Lemma B.7. $\{ [\langle \boldsymbol{\alpha}_i, c_i \rangle] \}_{i \in J_1} = \mathcal{R}([\Omega])$ (up to normalisation).

PROOF. By proposition 4.3, it is clear that $\{ [\langle \boldsymbol{\alpha}_i, c_i \rangle] \}_{i \in J} \supseteq \mathcal{R}([\Omega])$, since it is a representation of $[\Omega]$. Then, for some $i \in J$, consider

$$\Omega' \equiv \sum_{\substack{j \in J \\ j \neq i}} \langle \boldsymbol{\alpha}_j, c_j \rangle. \quad (\text{B.52})$$

Let us assume that $[\langle \boldsymbol{\alpha}_i, c_i \rangle] \notin \mathcal{R}([\Omega])$, so that $[\Omega'] = [\Omega]$. However, $\Omega' \neq \Omega$, since $\{ \langle \boldsymbol{\alpha}_i, c_i \rangle \}_{i \in J}$ is minimal. Thus, by lemma B.1, Ω' must be satisfied by some point $\mathbf{a} \notin E_\Omega$. For each $\mathbf{b} \in \mathcal{B}(\Omega)$, consider then $x\mathbf{a} + (1-x)\mathbf{b}$ for $x \in (0, 1]$. This point satisfies Ω' but not Ω , since it lies outside E_Ω . From this, we conclude that it does not satisfy $\langle \boldsymbol{\alpha}_i, c_i \rangle$. Thus, the continuous function

$$f_{\mathbf{b}}(x) = \boldsymbol{\alpha}_i \cdot [x\mathbf{a} + (1-x)\mathbf{b}] \quad (\text{B.53})$$

has $f_{\mathbf{b}}(x) < c_i$ for $x > 0$. However, $f_{\mathbf{b}}(0) \geq c_i$ since \mathbf{b} satisfies $\langle \boldsymbol{\alpha}_i, c_i \rangle$, and this is only consistent with continuity if $f_{\mathbf{b}}(0) = c_i$, i.e. that $\boldsymbol{\alpha}_i \cdot \mathbf{b} = c_i$, for all $\mathbf{b} \in \mathcal{B}(\Omega)$. Then $\mathcal{B}(\Omega) \subseteq \mathcal{B}(\langle \boldsymbol{\alpha}_i, c_i \rangle + \langle -\boldsymbol{\alpha}_i, -c_i \rangle)$, so it follows from eq. (B.21) that $\boldsymbol{\alpha}_i$ is a linear combination of $\{\boldsymbol{\delta}_k\}_{k=1}^{D-d_\Omega}$, i.e. that $[\boldsymbol{\alpha}_i] = \mathbf{0}$. Eq. (B.26) then shows that $i \in J_2$, completing the proof. \square

Now, let

$$\Omega = \sum_{i \in J_1} \langle \boldsymbol{\alpha}_i, c_i \rangle + \sum_{i \in J_2} \langle \boldsymbol{\alpha}_i, c_i \rangle. \quad (\text{B.54})$$

Lemma B.7 connects the J_1 part with $\mathcal{R}([\Omega])$, so it remains to study the J_2 part. We claim that eq. (B.54) holds true if and only if

$$\mathcal{B} \left(\sum_{i \in J_2} [\langle \alpha_i, c_i \rangle] \right) = \{\mathbf{0}\}. \tag{B.55}$$

It follows immediately from lemma B.7 that

$$\mathbf{0} \in \mathcal{B} \left(\sum_{i \in J_2} [\langle \alpha_i, c_i \rangle] \right). \tag{B.56}$$

Assume then that this set also contains some $[\mathbf{v}] \neq \mathbf{0}$. Since $[\Omega]$ is non-degenerate, there exists $[\mathbf{b}] \in \text{int} [\mathcal{B}([\Omega])]$; choosing $[\mathbf{b}] = \mathbf{0}$, it follows that $\forall i \in J_1, \alpha_i \cdot \mathbf{b} > c_i$. Therefore, for $\varepsilon > 0$ small enough we have $\forall i \in J_1, \alpha_i \cdot (\mathbf{b} + \varepsilon \mathbf{v}) \geq c_i$. Since $[\mathbf{b} + \varepsilon \mathbf{v}] = \varepsilon [\mathbf{v}]$, we also have $\forall i \in J_2, \alpha_i \cdot (\mathbf{b} + \varepsilon \mathbf{v}) \geq c_i$. Thus, $\mathbf{b} + \varepsilon \mathbf{v} \in \mathcal{B}(\Omega)$, but since $[\mathbf{v}] \neq \mathbf{0}$, we have $\mathbf{b} + \varepsilon \mathbf{v} \notin E_\Omega$, a contradiction. Along with its trivial converse, this proves the equivalence between eqs. (B.54) and (B.55).

What is then the minimal set $\{\langle \alpha_i, c_i \rangle\}_{i \in J_2}$ that produces eq. (B.55)? Since $\{\mathbf{0}\}$ is a bounded set, corollary B.1 states that $\{[\alpha_i]\}_{i \in J_2}$ spans the full $(D - d_\Omega)$ -dimensional space, and that there are $\lambda_i > 0, i \in J_2$ such that

$$\mathbf{0} = \sum_{i \in J_2} \lambda_i \alpha_i. \tag{B.57}$$

The span condition requires $|J_2| \geq D - d_\Omega$, but in order for there to be a nontrivial linear combination equal to zero, we must in fact have $|J_2| \geq D - d_\Omega + 1$. It is easy to see that this bound is sufficient (for details, see the proof of corollary B.1), so the minimal set must have $|J_2| = D - d_\Omega + 1$. Identifying $\{\alpha_i\}_{i \in J_2}$ with $\{\sigma_k\}_{k=0}^{D-d_\Omega}$, we see that we have just derived all conditions stated in the proposition, so the proof is complete. ■

B.3.2 An important corollary

The following interesting result, which is also our first use of the K -faces defined in appendix B.2.2, is a consequence of proposition 4.3:

Corollary B.2 (facet supported by relevant element). *Let $\Omega \neq \Omega_\infty$ be a non-degenerate constraint, and let $\langle \alpha, c \rangle \leq \Omega$. Then $\langle \alpha, c \rangle \in \mathcal{R}(\Omega)$ (up to normalisation) if and only if the set $\mathcal{F} \equiv \{\mathbf{b} \in \mathcal{B}(\Omega) \mid \alpha \cdot \mathbf{b} = c\}$ is a facet of $\mathcal{B}(\Omega)$.⁵⁰*

PROOF. It is guaranteed via lemma B.3 that \mathcal{F} is a K -face; we only need to show that $K = D - 1$ so that it is a facet.

⁵⁰It may seem obvious that this condition is equivalent to uniquely supporting a point, but is in fact rather subtle, and crucially depends on Ω having a finite representation. For instance, if $\mathcal{B}(\Omega)$ were a closed unit ball, then every point \mathbf{b} on its surface would be uniquely supported by the constraint $\langle -\mathbf{b}, -1 \rangle$. For such a constraint, $\mathcal{F} = \{\mathbf{b}\}$, which is not a facet in $D > 1$ dimensions.

Assume $\langle \alpha, c \rangle \in \mathcal{R}(\Omega)$. Let $\langle \beta, d \rangle$ be any constraint that supports all of \mathcal{F} , and construct the constraints

$$\Omega' \equiv \sum_{\langle \alpha', c' \rangle \in \mathcal{R}(\Omega)} \langle \alpha', c' \rangle + \langle \beta, d \rangle, \quad \Omega'' \equiv \sum_{\langle \alpha', c' \rangle \in \mathcal{R}(\Omega)} \langle \alpha', c' \rangle + \langle -\beta, d \rangle. \quad (\text{B.58})$$

Then it is clear that $\mathcal{B}(\Omega) = \mathcal{B}(\Omega') \cup \mathcal{B}(\Omega'')$, so at least one of Ω', Ω'' must be non-degenerate. Without loss of generality, assume Ω' is non-degenerate. Since $\langle \alpha, c \rangle$ does not uniquely support any point in $\mathcal{B}(\Omega')$ (all points supported by it are also supported by $\langle \beta, d \rangle$), proposition 4.3 (or rather lemma B.6) gives

$$\Omega' = \sum_{\substack{\langle \alpha', c' \rangle \in \mathcal{R}(\Omega) \\ \langle \alpha', c' \rangle \neq \langle \alpha, c \rangle}} \langle \alpha', c' \rangle + \langle \beta, d \rangle. \quad (\text{B.59})$$

Recall that $\langle \alpha, c \rangle$, viewed as an element of Ω , uniquely supports some point $\mathbf{b} \in \mathcal{F}$. Therefore, $\alpha' \cdot \mathbf{b} > c'$ for all $\langle \alpha', c' \rangle \in \mathcal{R}(\Omega) \setminus \{\langle \alpha, c \rangle\}$. By eq. (B.59), $\langle \beta, d \rangle$ then uniquely supports $\mathbf{b} \in \mathcal{B}(\Omega')$, but since $\langle \alpha, c \rangle \leq \Omega'$ and $\alpha \cdot \mathbf{b} = c$, the definition of unique support gives $\langle \alpha, c \rangle = \langle \beta, d \rangle$. Thus, the only constraints that support all of \mathcal{F} are, up to normalisation, $\langle \alpha, c \rangle$ and $\langle -\alpha, -c \rangle$, confirming via lemma B.2 that $\dim(\mathcal{F}) = D - 1$ (compare also eq. (B.21)). The converse is trivial. \square

B.4 Practical construction of $\mathcal{A}_c(\Omega)$ and $\mathcal{R}(\Omega)$

This section describes how to leverage propositions 4.2 and 4.3 for the practical management of linear constraints. The results stated here, namely propositions B.2 and B.3 along with their corollaries, double as the algorithms which we used in practice to obtain the results presented in section 6.⁵¹

B.4.1 Construction of $\mathcal{A}_c(\Omega)$

When using proposition 4.2 to determine if $\langle \beta, c \rangle \leq \Omega$ for some β, c and Ω , the rather indirect definition in eq. (4.10) is of little practical use. Instead, we will take the approach of finding a constraint that is satisfied by β if and only if $\beta \in \mathcal{A}_c(\Omega)$.

To understand why such a constraint exists and has a finite representation, note that eq. (4.11) along with $\lambda_i \geq 0$ are nothing more than an obfuscated set of linear constraints on the set $\mathcal{A}(\Omega)$. Provided a finite set of α_i , a considerable amount of linear algebra will determine a finite representation this way. Here, however, we present a simpler method in which the only complicated operation is the determination of the convex hull of a set of points. Highly efficient algorithms for determining convex hulls exist; we use the QuickHull algorithm [48] and the associated `qhull` implementation.⁵²

The key to the construction is that, given a finite set of points $\{\beta_j\}_{j \in J}$, a side-effect of the QuickHull algorithm is the creation of a set of constraints $\{\langle \mathbf{n}_\ell, r_\ell \rangle\}_{\ell \in L}$ such that

$$\mathcal{B}(\sum_{\ell \in L} \langle \mathbf{n}_\ell, r_\ell \rangle) = \text{Hull}(\{\beta_j\}_{j \in J}), \quad (\text{B.60})$$

⁵¹The implementation code is available from Mattias Sjö upon request.

⁵²For up-to-date information about `qhull`, see <http://www.qhull.org/>.

since \mathbf{n}_ℓ, r_ℓ are the normals and offsets of the facets of the hull.⁵³ This is, by construction, a minimal representation. If we choose β_j such that $\text{Hull}(\{\beta_j\}_{j \in J})$ is a suitable subset of $\mathcal{A}_c(\Omega)$, we will see that it is possible to write a simple rule that selects a subset $M \subseteq L$ such that

$$\mathcal{B}(\sum_{\ell \in M} \langle \mathbf{n}_\ell, r_\ell \rangle) = \mathcal{A}_c(\Omega). \quad (\text{B.61})$$

In order to do this, let Ω be given as in eq. (4.9). Importantly, assume without loss of generality that $\langle \alpha_i, -1 \rangle = \langle \mathbf{0}, -1 \rangle$ for some $i \in I_{-1}$, but that $\alpha_i \neq \mathbf{0}$ for all $i \in I_0$.⁵⁴ Then, define the following sets of points:

$$\begin{aligned} \mathcal{P}_\pm(\Omega) &= \{\alpha_i \mid i \in I_{\pm 1}\}, \\ \mathcal{Z}_\pm^{(p)}(\Omega) &= \{\alpha_i + (p-1)\alpha_j \mid i \in I_{\pm 1}, j \in I_0\}, \\ \mathcal{N}_\pm^{(p)}(\Omega) &= \{p\alpha_i + (p-1)\alpha_j \mid i \in I_{\pm 1}, j \in I_{\mp 1}\} \end{aligned} \quad (\text{B.62})$$

for integer $p > 1$. Also define

$$\mathcal{H}_\pm^{(p)}(\Omega) = \text{Hull}[\mathcal{P}_\pm(\Omega) \cup \mathcal{Z}_\pm^{(p)}(\Omega) \cup \mathcal{N}_\pm^{(p)}(\Omega)], \quad (\text{B.63})$$

where Hull denotes the convex hull; see eq. (4.8). From now on, we will often drop the “(Ω)” for brevity. In terms of these, we have the following result:

Proposition B.2 (constructing $\mathcal{A}_c(\Omega)$). *Let Ω be a constraint, and arbitrarily select an integer $p > 1$. Construct a minimal representation $\{\langle \mathbf{n}_\ell, r_\ell \rangle\}_{\ell \in L_\pm^{(p)}}$ such that⁵⁵*

$$\mathcal{B}\left[\sum_{\ell \in L_\pm^{(p)}} \langle \mathbf{n}_\ell, r_\ell \rangle\right] = \mathcal{H}_\pm^{(p)}(\Omega), \quad (\text{B.64})$$

with $\mathcal{H}_\pm^{(p)}(\Omega)$ defined as in eq. (B.63). Let $M_\pm^{(p)} \subseteq L_\pm^{(p)}$ be the set of all ℓ for which $\langle \mathbf{n}_\ell, r_\ell \rangle$ supports at least one point $\pi \in \mathcal{P}_\pm$. Then

$$\mathcal{B}\left[\sum_{\ell \in M_\pm^{(p)}} \langle \mathbf{n}_\ell, r_\ell \rangle\right] = \mathcal{A}_{\pm 1}(\Omega), \quad \mathcal{B}\left[\sum_{\ell \in M_\pm^{(p)}} \langle \mathbf{n}_\ell, 0 \rangle\right] = \mathcal{A}_0(\Omega). \quad (\text{B.65})$$

Note that $\mathcal{A}_0(\Omega)$ can be constructed from either $M_+^{(p)}$ or $M_-^{(p)}$. The exception is when $I_1 = \emptyset$, in which case the construction of $\mathcal{A}_1(\Omega)$ fails; $\mathcal{A}_1(\Omega) = \emptyset$ trivially, and $\mathcal{A}_0(\Omega)$ can only be constructed from $M_-^{(p)}$.

⁵³qhull uses a different sign convention, but the conversion to the format given above is trivial.

⁵⁴For maximum efficiency, the representation used for Ω should otherwise have as few elements as possible. The best easily accessible one is

$$\mathcal{R}\left(\sum_{i \in I_1} \langle \alpha_i, c_i \rangle\right) \cup \mathcal{R}\left(\sum_{i \in I_0} \langle \alpha_i, c_i \rangle\right) \cup \mathcal{R}\left(\sum_{i \in I_{-1}} \langle \alpha_i, c_i \rangle\right) \cup \langle \mathbf{0}, -1 \rangle,$$

where the minimal representations are determined with proposition B.3 applied to corollary B.4.

⁵⁵To avoid clutter, we do not indicate any p -dependence on $\langle \mathbf{n}_\ell, r_\ell \rangle$, but one should bear in mind that they may be entirely different constraints for different p (and different \pm). To remember this, it can be useful to think of $L_\pm^{(p)}$ as disjoint sets for different p, \pm .

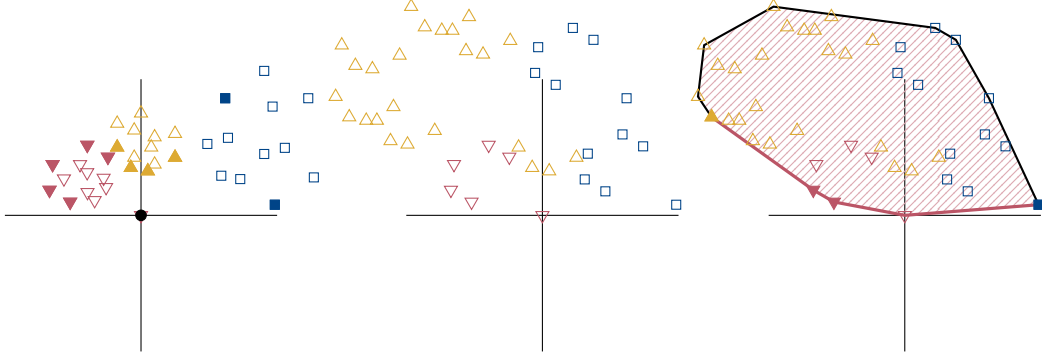


Figure 26. Example of the construction of $\mathcal{A}_{-1}(\Omega)$ as described in proposition B.2. *Left:* three sets of points α_i for $i \in I_{-1}$ (\blacktriangledown), $i \in I_0$ (\blacksquare) and $i \in I_1$ (\blacktriangle). The sets are the same up to rescaling and rotation about the origin, and the addition of $\langle \mathbf{0}, -1 \rangle$. The relevant elements of the respective ω_c are marked as filled points, and the rest are left empty; compare to the similar sets in figure 2. *Middle:* the sets \mathcal{P}_- (\blacktriangledown), $\mathcal{Z}_-^{(2)}$ (\square) and $\mathcal{N}_-^{(2)}$ (\blacktriangle) constructed using only the relevant elements, as remarked above eq. (B.62). *Right:* the convex hull thereof (red), i.e. $\mathcal{H}_-^{(2)}$. The segments of its boundary correspond to $\langle \mathbf{n}_\ell, r_\ell \rangle$, $\ell \in L_-^{(2)}$. The ones for which $\ell \in M_-^{(2)}$ are highlighted in red. Removing the other segments leaves behind the unbounded region $\mathcal{A}_{-1}(\Omega)$, which can be seen in figure 27 below. The relevant elements of Ω have been marked as filled points. They follow from proposition B.3, although it can be intuitively seen that they alone influence the shape of $\mathcal{A}_{-1}(\Omega)$.

This result (along with corollaries B.3 and B.4 and proposition B.3 below) outlines the procedure we use in practice to obtain minimal representations. An example of this construction for $c = -1$ is given in figure 26.

PROOF. For brevity, we will write

$$\mathcal{U}_\pm^{(p)} \equiv \sum_{\ell \in M_\pm^{(p)}} \langle \mathbf{n}_\ell, r_\ell \rangle. \quad (\text{B.66})$$

The goal is then to show that $\mathcal{B}(\mathcal{U}_\pm^{(p)}) = \mathcal{A}_{\pm 1}(\Omega)$, which we will do by showing that $\mathcal{B}(\mathcal{U}_\pm^{(p)}) \supseteq \mathcal{A}_{\pm 1}(\Omega)$ followed by $\mathcal{B}(\mathcal{U}_\pm^{(p)}) \subseteq \mathcal{A}_{\pm 1}(\Omega)$. We will then show the $c = 0$ case as a consequence of the others. First, however, we will establish some lemmata.

Lemma B.8. $r_\ell \geq 0$ for all $\ell \in M_+^{(p)}$, and $r_\ell \leq 0$ for all $\ell \in M_-^{(p)}$.

PROOF. The latter inequality is trivial, since $\mathbf{0} \in \mathcal{H}_-^{(p)}$. In the $M_+^{(p)}$ case, recall that by definition, $\mathbf{n}_\ell \cdot \alpha_{i_\ell} = r_\ell$ for some $i_\ell \in I_1$. We must also have $r_\ell \leq \mathbf{n}_\ell \cdot (p\alpha_{i_\ell} + \mathbf{0}) = pr_\ell$ since $(p\alpha_{i_\ell} + \mathbf{0}) \in \mathcal{N}_+^{(p)}$, but $r_\ell \leq pr_\ell$ for $p > 1$ implies that $r_\ell \geq 0$. \square

Lemma B.9. $\mathbf{n}_\ell \cdot \alpha_i \geq c_i |r_\ell|$ for all $i \in I, \ell \in M_\pm^{(p)}$.

PROOF. For $\ell \in M_\pm^{(p)}$, let $i_\ell \in I_{\pm 1}$ be such that $\mathbf{n}_\ell \cdot \alpha_{i_\ell} = r_\ell$. Then consider two specific points in $\mathcal{Z}_\pm^{(p)}$ and $\mathcal{N}_\pm^{(p)}$:

$$\forall j \in I_0, \quad \mathbf{n}_\ell \cdot (\alpha_{i_\ell} + (p-1)\alpha_j) \geq r_\ell \quad \Rightarrow \quad \mathbf{n}_\ell \cdot \alpha_j \geq 0 \quad (\text{B.67})$$

$$\forall j' \in I_{\mp 1}, \quad \mathbf{n}_\ell \cdot (p\alpha_{i_\ell} + (p-1)\alpha_{j'}) \geq r_\ell \quad \Rightarrow \quad \mathbf{n}_\ell \cdot \alpha_{j'} \geq \frac{1-p}{p-1} r_\ell = -r_\ell. \quad (\text{B.68})$$

This, together with lemma B.8, implies lemma B.9. \square

Now for the main proof. Let $\beta \in \mathcal{A}_{\pm 1}(\Omega)$ (note that this excludes the exceptional case $\mathcal{A}_1(\Omega) = \emptyset$). Using eq. (4.12), it can therefore be written

$$\beta = \sum_{i \in I} \lambda_i \alpha_i, \quad \lambda_i \geq 0, \quad \sum_{i \in I} \lambda_i c_i \geq c, \quad (\text{B.69})$$

so for all $\ell \in M_{\pm}^{(p)}$, lemma B.9 gives

$$\mathbf{n}_\ell \cdot \beta \geq |r_\ell| \sum_{i \in I} \lambda_i c_i \geq |r_\ell| c. \quad (\text{B.70})$$

Since $|r_\ell|c = r_\ell$ for $c = \pm 1$ by lemma B.8, this means that all $\beta \in \mathcal{A}_{\pm 1}(\Omega)$ satisfy $\langle \mathbf{n}_\ell, r_\ell \rangle$ for all $\ell \in M_{\pm}^{(p)}$, thereby proving that $\mathcal{A}_{\pm 1}(\Omega) \subseteq \mathcal{B}(\mathcal{U}_{\pm}^{(p)})$.

For the converse, we first note a direct consequence of eqs. (4.10) and (B.62),

$$\mathcal{H}_{\pm}^{(p)}(\Omega) \subseteq \mathcal{A}_{\pm 1}(\Omega), \quad \mathcal{H}_{\pm}^{(p)} \subseteq \mathcal{H}_{\pm}^{(q)} \quad \text{if } q \geq p. \quad (\text{B.71})$$

Then, the proof hinges on the following deceptively simple result:

Lemma B.10. *For any $\beta \in \mathcal{B}(\mathcal{U}_{\pm}^{(p)})$, there is some q such that $\beta \in \mathcal{H}_{\pm}^{(q)}$.*

We will postpone its lengthy proof until after the main proof is complete. Lemma B.10, along with eq. (B.71), shows that

$$\mathcal{B}(\mathcal{U}_{\pm}^{(p)}) \subseteq \bigcup_{q=p}^{\infty} \mathcal{H}_{\pm}^{(q)} \subseteq \mathcal{A}_{\pm 1}(\Omega). \quad (\text{B.72})$$

This proves that $\mathcal{B}(\mathcal{U}_{\pm}^{(p)}) = \mathcal{A}_{\pm 1}(\Omega)$.

For \mathcal{A}_0 , let us turn to eq. (4.12), which lets us straightforwardly generalise $\mathcal{A}_c(\Omega)$ to non-integer c . One easily finds the following generalisation of eq. (4.4):

$$\langle \alpha, 0 \rangle \leq \Omega \quad \Leftrightarrow \quad \exists \varepsilon > 0, \langle \alpha, +\varepsilon \rangle \leq \Omega \quad \Leftrightarrow \quad \forall \varepsilon > 0, \langle \alpha, -\varepsilon \rangle \leq \Omega, \quad (\text{B.73})$$

from which it follows that

$$\bigcup_{n=1}^{\infty} \mathcal{A}_{+\frac{1}{n}}(\Omega) = \mathcal{A}_0(\Omega) = \bigcap_{n=1}^{\infty} \mathcal{A}_{-\frac{1}{n}}(\Omega). \quad (\text{B.74})$$

Using that we have proven proposition B.2 for $c = \pm 1$, which generalises to all $c \neq 0$ by rescaling, we have for all positive integers n

$$\mathcal{A}_{\pm \frac{1}{n}}(\Omega) = \mathcal{B} \left(\sum_{\ell \in M_{\pm}^{(p)}} \langle \mathbf{n}_\ell, \frac{1}{n} r_\ell \rangle \right), \quad (\text{B.75})$$

so by extension,

$$\begin{aligned} \bigcup_{n=1}^{\infty} \mathcal{A}_{\pm \frac{1}{n}}(\Omega) &= \bigcup_{n=1}^{\infty} \mathcal{B} \left(\sum_{\ell \in M_{\pm}^{(p)}} \langle \mathbf{n}_\ell, \frac{1}{n} r_\ell \rangle \right), \\ \bigcap_{n=1}^{\infty} \mathcal{A}_{\pm \frac{1}{n}}(\Omega) &= \bigcap_{n=1}^{\infty} \mathcal{B} \left(\sum_{\ell \in M_{\pm}^{(p)}} \langle \mathbf{n}_\ell, \frac{1}{n} r_\ell \rangle \right). \end{aligned} \quad (\text{B.76})$$

With the sign of r_ℓ given by lemma B.8, we have

$$\begin{aligned} \bigcup_{n=1}^{\infty} \mathcal{B}\left(\sum_{\ell \in L'_+} \langle \mathbf{n}_\ell, \frac{1}{n} r_\ell \rangle\right) &= \mathcal{B}\left(\sum_{\ell \in L'_+} \langle \mathbf{n}_\ell, 0 \rangle\right), \\ \bigcap_{n=1}^{\infty} \mathcal{B}\left(\sum_{\ell \in L'_-} \langle \mathbf{n}_\ell, \frac{1}{n} r_\ell \rangle\right) &= \mathcal{B}\left(\sum_{\ell \in L'_-} \langle \mathbf{n}_\ell, 0 \rangle\right), \end{aligned} \quad (\text{B.77})$$

so by eqs. (B.74) and (B.76), both of these sets are equal to $\mathcal{A}_0(\Omega)$, which completes the proof. \blacksquare

B.4.2 Proof of lemma B.10

This lemma is the key to proving proposition B.2, and relies on several other lemmata that we will now establish. They also serve to elucidate some aspects of the proposition and its proof; for instance, lemma B.14 explains why the choice of p is arbitrary.

Lemma B.11. *Let $\langle \mathbf{m}, s \rangle$ be a constraint such that $\mathcal{H}_\pm^{(p)} \in \mathcal{B}(\langle \mathbf{m}, s \rangle)$, and which supports a point $\alpha_i + (p-1)\alpha_j \in \mathcal{Z}_\pm^{(p)}$ or a point $p\alpha_i + (p-1)\alpha_j \in \mathcal{N}_\pm^{(p)}$. Then $\langle \mathbf{m}, s \rangle$ either supports $\alpha_i \in \mathcal{P}_\pm$ for that same $i \in I_{\pm 1}$, or supports no point in \mathcal{P}_\pm at all.*

PROOF. Assume $\langle \mathbf{m}, s \rangle$ supports some $\pi \in \mathcal{P}_\pm$, but that it does not support α_i . Then

$$\mathbf{m} \cdot (p\alpha_i + (p-1)\alpha_j) = s \quad \Rightarrow \quad \mathbf{m} \cdot (p\pi + (p-1)\alpha_j) < s, \quad (\text{B.78})$$

since $\mathbf{m} \cdot \alpha_i > \mathbf{m} \cdot \pi = s$. This is a contradiction, since

$$p\pi + (p-1)\alpha_j \in \mathcal{N}_\pm^{(p)} \subseteq \mathcal{H}_\pm^{(p)} \subseteq \mathcal{B}(\langle \mathbf{m}, s \rangle). \quad (\text{B.79})$$

The argument for $\alpha_i + (p-1)\alpha_j \in \mathcal{Z}_\pm^{(p)}$ is the same. \square

Lemma B.12. *Let $\langle \mathbf{m}, s \rangle$ be a constraint with $\mathcal{H}_\pm^{(p)} \subseteq \mathcal{B}(\langle \mathbf{m}, s \rangle)$, and let it support at least one point in \mathcal{P}_\pm . Then if it supports $\alpha_i + (p-1)\alpha_j \in \mathcal{Z}_\pm^{(p)}$, it also supports $\alpha_i + (q-1)\alpha_j \in \mathcal{Z}_\pm^{(q)}$ for all $q > 1$. Likewise, if it supports $p\alpha_i + (p-1)\alpha_j \in \mathcal{N}_\pm^{(p)}$, it also supports $q\alpha_i + (q-1)\alpha_j \in \mathcal{N}_\pm^{(q)}$ for all $q > 1$.⁵⁶*

PROOF. Assume $\mathbf{m} \cdot [p\alpha_i + (p-1)\alpha_j] = s$. Then by lemma B.11, $\langle \mathbf{m}, s \rangle$ also supports α_i , and thus $\mathbf{m} \cdot (\alpha_i + \alpha_j) = 0$. Adding $(q-p)(\alpha_i + \alpha_j)$ therefore gives $\mathbf{m} \cdot [q\alpha_i + (q-1)\alpha_j] = s$. The argument for $\alpha_i + (p-1)\alpha_j \in \mathcal{Z}_\pm^{(p)}$ is the same. \square

Lemma B.13. *$\dim(\mathcal{H}_\pm^{(q)})$, as defined in appendix B.2.1, is independent of q for $q > 1$. Furthermore, for any $J_c \subseteq I_c$, $\dim[\mathcal{J}_q(J_{\pm 1}, J_0, J_{\mp 1})]$ is independent of q for $q > 1$, where*

$$\mathcal{J}_q(J_{\pm 1}, J_0, J_{\mp 1}) \equiv \{\alpha_i\}_{i \in J_{\pm 1}} \cup \{\alpha_i + (q-1)\alpha_j\}_{\substack{i \in J_{\pm 1} \\ j \in J_0}} \cup \{q\alpha_i + (q-1)\alpha_j\}_{\substack{i \in J_{\pm 1} \\ j \in J_{\mp 1}}}. \quad (\text{B.80})$$

⁵⁶Lemmas B.12 and B.13 actually work for all $q \neq 1$. The only places where proposition B.2 actually requires $p > 1$ rather than $p < 1$ are in lemmas B.8 and B.9 and in eq. (B.71).

PROOF. The main statement is actually a special case of the “furthermore” statement, since $\mathcal{P}_\pm \cup \mathcal{Z}_\pm^{(q)} \cup \mathcal{N}_\pm^{(q)} = \mathcal{J}_q(I_{\pm 1}, I_0, I_{\mp 1})$, and since for any set \mathcal{X} , $\dim[\text{Hull}(\mathcal{X})] = \dim(\mathcal{X})$ because convex combinations are a special case of affine combinations.

By definition, the affine span of $\mathcal{J}_q(J_{\pm 1}, J_0, J_{\mp 1})$ has the same dimension as the linear span of $\{\alpha - \beta \mid \alpha, \beta \in \mathcal{J}_q(J_{\pm 1}, J_0, J_{\mp 1})\}$; according to eq. (B.80), this set consists of

$$\begin{cases} \alpha_i - \alpha_j & \text{for } i, j \in J_{\pm 1}, \\ (q-1)\alpha_j & \text{for } j \in J_0, \\ (q-1)(\alpha_i - \alpha_j) & \text{for } i \in J_{\pm 1} \cup J_0 \cup J_{\mp 1} \text{ and } j \in J_0 \cup J_{\mp 1}, \\ \text{various linear combinations of the above.} \end{cases} \quad (\text{B.81})$$

The linear span is unaffected by the inclusion of extra linear combinations or nonzero scale factors, so as long as $(q-1) \neq 0$ we can drop these and be left with the span of

$$\{\alpha_i - \alpha_j \mid i, j \in J_{\pm 1} \cup J_0 \cup J_{\mp 1}\} \cup \{\alpha_j \mid j \in J_0\}, \quad (\text{B.82})$$

which is clearly q -independent. □

Lemma B.14. *If $\dim(\mathcal{H}_\pm^{(q)}) = D$ for some $q > 1$, then $\mathcal{U}_\pm^{(q)}$ is non-degenerate for all $q > 1$ and $\mathcal{U}_\pm^{(p)} = \mathcal{U}_\pm^{(q)}$ for all $p, q > 1$.*

This makes it quite clear why p is arbitrary in proposition B.2.

PROOF. By lemma B.13, $\dim(\mathcal{H}_\pm^{(q)})$ equals D for all $q > 1$ if it does for some $q > 1$. Then $\mathcal{U}_\pm^{(q)}$ is non-degenerate, since $\mathcal{H}_\pm^{(p)} \subseteq \mathcal{B}(\mathcal{U}_\pm^{(p)})$ has dimension D .

Now for the converse. Given $\ell \in M_\pm^{(p)}, p > 1$, define

$$\mathcal{G}_\ell^{(p)} = \left\{ \pi \in \mathcal{P}_\pm \cup \mathcal{Z}_\pm^{(p)} \cup \mathcal{N}_\pm^{(p)} \mid \mathbf{n}_\ell \cdot \pi = r_\ell \right\}. \quad (\text{B.83})$$

By the definition of $M_\pm^{(p)}$ along with lemma B.11, there are some $J_c^\ell \subseteq I_c$ such that $\mathcal{G}_\ell^{(p)} = \mathcal{J}_p(J_{\pm 1}^\ell, J_0^\ell, J_{\mp 1}^\ell)$ as defined in eq. (B.80). Then by lemma B.12, that same $\langle \mathbf{n}_\ell, r_\ell \rangle$ also supports $\mathcal{G}_\ell^{(q)} \equiv \mathcal{J}_q(J_{\pm 1}^\ell, J_0^\ell, J_{\mp 1}^\ell)$ for all $q > 1$. From this and eq. (B.71), it follows that $\langle \mathbf{n}_\ell, r_\ell \rangle \leq \mathcal{U}_\pm^{(q)}$.

Recall now that $\langle \mathbf{n}_\ell, r_\ell \rangle \in \mathcal{R}(\mathcal{U}_\pm^{(p)})$, so corollary B.2 implies that $\dim[\mathcal{G}_\ell^{(p)}] = D - 1$.⁵⁷ Then by lemma B.13, $\dim[\mathcal{G}_\ell^{(q)}] = D - 1$ for all $q > 1$, so again by corollary B.2, $\langle \mathbf{n}_\ell, r_\ell \rangle \in \mathcal{R}(\mathcal{U}_\pm^{(q)})$. By doing this for all $\ell \in M_\pm^{(p)}$ and repeating with p and q exchanged, we see that $\mathcal{R}(\mathcal{U}_\pm^{(q)}) = \mathcal{R}(\mathcal{U}_\pm^{(p)})$ up to normalisation. This implies $\mathcal{U}_\pm^{(q)} = \mathcal{U}_\pm^{(p)}$. □

Thanks to lemma B.14, we will drop the “ (p) ” superscript on \mathcal{U}_\pm from now on. However, we face the problem that \mathcal{U}_\pm may be degenerate, which would make lemma B.14 inapplicable. It can be circumvented by using the notion of induced constraints developed in appendix B.2.1: just substitute

$$\mathcal{U}_\pm \rightarrow [\mathcal{U}_\pm], \quad \mathcal{A}_{\pm 1}(\Omega) \rightarrow [\mathcal{A}_{\pm 1}(\Omega)]_{E_{\mathcal{U}_\pm}}, \quad (\text{B.84})$$

⁵⁷What corollary B.2 calls “ \mathcal{F} ” is here $\text{Hull}(\mathcal{G}_\ell^{(p)})$, but $\dim[\mathcal{X}] = \dim[\text{Hull}(\mathcal{X})]$ for any set \mathcal{X} since convex combinations are a special case of affine combinations.

since it follows from lemma B.1 (along with $\mathcal{A}_{\pm 1}(\Omega) \subseteq E_{\mathcal{U}_{\pm}}$ which we proved earlier) that⁵⁸

$$[\mathcal{A}_{\pm 1}(\Omega)]_{E_{\mathcal{U}_{\pm}}} \supseteq \mathcal{B}([\mathcal{U}_{\pm}]) \iff \mathcal{A}_{\pm 1}(\Omega) \supseteq \mathcal{B}(\mathcal{U}_{\pm}). \quad (\text{B.85})$$

Thus, we may for the remainder assume that \mathcal{U}_{\pm} is non-degenerate.

We are now, at long last, ready to prove lemma B.10 itself. Consider $\mathcal{H}_{\pm}^{(q)}$ for arbitrary q , and presume $\beta \notin \mathcal{H}_{\pm}^{(q)}$. Select some $\alpha \in \mathcal{H}_{\pm}^{(q)}$, and draw the line segment joining α and β . It must intersect the boundary of $\mathcal{H}_{\pm}^{(q)}$ in some point γ , which is supported by one or more $\langle \mathbf{n}_{\ell}, r_{\ell} \rangle, \ell \in L_{\pm}^{(q)}$, at least one of which is not satisfied by β . We therefore have some $\ell \in L_{\pm}^{(q)}$ such that

$$\mathbf{n}_{\ell} \cdot \alpha \geq r_{\ell}, \quad \mathbf{n}_{\ell} \cdot \beta < r_{\ell}, \quad \mathbf{n}_{\ell} \cdot \gamma = r_{\ell}. \quad (\text{B.86})$$

If $\ell \in M_{\pm}^{(q)}$, then $\beta \notin \mathcal{B}(\mathcal{U}_{\pm})$; this is where the q -independence of \mathcal{U}_{\pm} proven in lemma B.14 is crucial. Thus, we can assume that $\ell \notin M_{\pm}^{(q)}$.

The magnitude of γ is bounded from above by $|\gamma| \leq \max(|\alpha|, |\beta|)$. We will now attempt to prove that $|\gamma| > \max(|\alpha|, |\beta|)$, which leads to a contradiction, proving that γ does not exist and consequently that $\beta \in \mathcal{H}_{\pm}^{(q)}$.

For any $\ell \in L_{\pm}^{(q)}$, let $\mathcal{G}_{\ell}^{(q)} \subseteq (\mathcal{P}_{\pm} \cup \mathcal{Z}_{\pm}^{(q)} \cup \mathcal{N}_{\pm}^{(q)})$ be defined as in eq. (B.83), and let us think about the structure of $\mathcal{G}_{\ell}^{(q)}$. Each of its elements is of the form $\alpha_i + (q-1)\alpha'$, where $i \in I_{\pm 1}$ and α' may take the form $\mathbf{0}, \alpha_j$ or $\alpha_i + \alpha_j$ depending on whether the element is part of $\mathcal{P}_{\pm}, \mathcal{Z}_{\pm}^{(q)}$ or $\mathcal{N}_{\pm}^{(q)}$. This observation allows us to write the Minkowski sum

$$\text{Hull}(\mathcal{G}_{\ell}^{(q)}) = \mathcal{U}_{\ell}^{(q)} + (q-1)\mathcal{V}_{\ell}^{(q)}, \quad (\text{B.87})$$

where $\mathcal{U}_{\ell}^{(q)} \subseteq \text{Hull}(\mathcal{P}_{\pm})$. $\mathcal{V}_{\ell}^{(q)}$ is the convex hull of a subset of $\{\alpha_i\}_{i \in I}$; the details are messy and unimportant, so we will not write it explicitly, but the important thing is that it only depends on q and ℓ through the specific choice of subset. This means that when counted over the infinitely many choices of q and ℓ , there is only a finite number of distinct $\mathcal{V}_{\ell}^{(q)}$ that appear: at most as many as there are subsets of I .

Now, focus on the case $\ell \in L_{\pm}^{(q)} \setminus M_{\pm}^{(q)}$, where we find that $\mathbf{0} \notin \mathcal{V}_{\ell}^{(q)}$: otherwise, $\text{Hull}(\mathcal{G}_{\ell}^{(q)}) \cap \text{Hull}(\mathcal{P}_{\pm}) \neq \emptyset$, which straightforwardly leads to a contradiction of the definition of $M_{\pm}^{(q)}$. We may also observe that $\mathcal{V}_{\ell}^{(q)}$ (and $\mathcal{U}_{\ell}^{(q)}$) are closed sets, being the convex hulls of finite sets of points.

Consider then the finite set

$$\left\{ \mathcal{V} \mid \mathcal{V}_{\ell}^{(q)} = \mathcal{V} \text{ for some } q \geq 2, \ell \in L_{\pm}^{(q)} \setminus M_{\pm}^{(q)} \right\}, \quad (\text{B.88})$$

and let \mathcal{W}_{\pm} be the union of all elements of this set. By the observations we have made about $\mathcal{V}_{\ell}^{(q)}$, this is a closed set (being the finite union of closed sets) that does not contain $\mathbf{0}$. Consequently, there is some $m > 0$ such that \mathcal{W}_{\pm} contains no vector of magnitude less than m .

⁵⁸For practical applications, there is the additional problem that `qhull` does not function properly when its output would be degenerate. This has not been a problem for us, and is of course no issue for the present proof, but if needed, one could identify the affine subspace E containing all points defined in eq. (B.62), apply $[\dots]_E$ to them, and work entirely in the lower-dimensional space where there are no degeneracies.

Recall now that $\gamma \in \mathcal{G}_\ell^{(q)}$ for some $\ell \in L_\pm^{(q)} \setminus M_\pm^{(q)}$. Thus, eq. (B.87) gives

$$\gamma = \pi + (q-1)\eta, \quad \pi \in \text{Hull}(\mathcal{P}_\pm), \eta \in \mathcal{W}_\pm, \quad (\text{B.89})$$

and this is true no matter the value of q and no matter which α, β are used to obtain γ . As derived above, $|\eta| \geq m$, and since $\text{Hull}(\mathcal{P}_\pm)$ is a bounded set, $|\pi| \leq M$ for sufficiently large M . For q sufficiently large that $(q-1)m > M$, the triangle inequality then gives

$$|\gamma| \geq (q-1)m - M, \quad (\text{B.90})$$

which can be made arbitrarily large by further increasing q , thereby providing the desired contradiction and completing the proof. \square

B.4.3 Some important corollaries of proposition B.2

We can refine the treatment of $\mathcal{A}_0(\Omega)$ with the following:⁵⁹

Corollary B.3 (construction of $\mathcal{A}_0(\Omega)$). *Let $\Omega, M_\pm^{(p)}$, etc. be as in proposition B.2. Let $N_\pm^{(p)} \subseteq M_\pm^{(p)}$ be the set of those ℓ for which $\langle \mathbf{n}_\ell, r_\ell \rangle$ also supports at least one point in $(\mathcal{Z}_\pm^{(p)} \cup \mathcal{N}_\pm^{(p)})$. If $\sum_{\ell \in M_\pm^{(p)}} \langle \mathbf{n}_\ell, 0 \rangle$ is a non-degenerate constraint,⁶⁰ then*

$$\mathcal{A}_0(\Omega) = \mathcal{B}\left(\sum_{\ell \in N_\pm^{(p)}} \langle \mathbf{n}_\ell, 0 \rangle\right). \quad (\text{B.91})$$

As in proposition B.2, the $M_+^{(p)}$ construction does not work when $I_1 = \emptyset$.

PROOF. For ℓ contained in $M_\pm^{(p)}$ but not in $N_\pm^{(p)}$, lemma B.9 more specifically gives

$$\forall i \in I_{\pm 1}, \mathbf{n}_\ell \cdot \alpha_i \geq r_\ell, \quad \forall i \in I_0, \mathbf{n}_\ell \cdot \alpha_i > 0, \quad \forall i \in I_{\mp 1}, \mathbf{n}_\ell \cdot \alpha_i > -r_\ell \quad (\text{B.92})$$

with lemma B.8 dictating the sign of r_ℓ . Then if $\mathbf{n}_\ell \cdot \beta = 0$ for $\beta \in \mathcal{A}_0(\Omega)$, a look at eq. (4.10) tells us that all $\lambda_i, i \in I_0 \cup I_1$ must be zero due to the above inequalities, and eq. (4.11) then implies that also $\lambda_i, i \in I_{-1}$ must be zero. Thus, the only point $\beta \in \mathcal{A}_0(\Omega)$ supported by $\langle \mathbf{n}_\ell, 0 \rangle$ is the trivial $\beta = \mathbf{0}$, which is supported by $\langle \alpha, 0 \rangle$ for all $\alpha \neq \mathbf{0}$. Therefore, $\langle \mathbf{n}_\ell, 0 \rangle$ does not uniquely support any point, so by proposition 4.3 (which requires non-degeneracy), it can be omitted. \blacksquare

⁵⁹An intuitive understanding of the construction of $\mathcal{A}_0(\Omega)$ can be gained by noting that in a sense, one can make $c \rightarrow 0$ by “zooming out” on constraint space. This is the principle that is formalised in the end of the proof of proposition B.2. Each facet of the body $\mathcal{A}_c(\Omega)$, consisting of the points supported by one $\langle \mathbf{n}_\ell, r_\ell \rangle$, is thereby shifted so that it passes through the origin (hence $r_\ell \rightarrow 0$), and if it was bounded, it shrinks down to a point. Corollary B.3 identifies those bounded facets and removes them. It is easy to see why it works: all facets of $\mathcal{H}_\pm^{(p)}$ are bounded, and only by being adjacent to a facet that is removed in the restriction $L_\pm^{(p)} \rightarrow M_\pm^{(p)}$ can a facet become unbounded. The removed facets are those that only support points in $\mathcal{Z}_\pm^{(p)} \cup \mathcal{N}_\pm^{(p)}$, so their neighbours are the ones that support at least one point in these sets.

⁶⁰This is not equivalent to $\sum_{\ell \in M_\pm^{(p)}} \langle \mathbf{n}_\ell, r_\ell \rangle$ being non-degenerate. Consider as a counterexample $\Omega = \langle \binom{1}{1}, -1 \rangle + \langle \binom{1}{-1}, -1 \rangle + \langle \binom{1}{0}, 0 \rangle$, for which $\mathcal{A}_0(\Omega)$ is contained in a hyperplane (i.e. a line) whereas $\mathcal{A}_{-1}(\Omega)$ is not.

One also easily finds the following simplification:

Corollary B.4 (construction of $\mathcal{A}_c(\omega_c)$). *For those ω_c covered by proposition 4.1, proposition B.2 reduces down to the following:*

$c = +1$: $\mathcal{H}_+^{(p)} = \text{Hull}(\{\alpha_i, p\alpha_i\}_{i \in I})$. $M_+^{(p)}$ consists of those ℓ for which $r_\ell \geq 0$.

$c = 0$: $\mathcal{H}_-^{(p)} = \text{Hull}(\{\mathbf{0}\} \cup \{p\alpha_i\}_{i \in I})$. $N_-^{(p)} = M_-^{(p)}$ consists of those ℓ where $r_\ell = 0$.

$c = -1$: $\mathcal{H}_-^{(p)} = \text{Hull}(\{\mathbf{0}\} \cup \{\alpha_i\}_{i \in I})$. All $\langle \mathbf{n}_\ell, r_\ell \rangle$ are kept, since $M_-^{(p)} = L_-^{(p)}$. ■

An alternative to using proposition B.2 is to apply corollary B.4 to ω_0^ℓ as defined in appendix B.1.4, and then “unlifting” the result. This requires the treatment of a much smaller number of points, which makes the QuickHull algorithm run faster; on the other hand, lifting increases the dimension, which makes the QuickHull algorithm run slower and be less numerically stable. The time complexity of the `qhull` implementation suggests that asymptotically, lifting should be the faster method, but since `qhull` is vastly more efficient in 2 and 3 dimensions, not lifting should be preferable when the number of dimensions is small. In practice, we only used proposition B.2 directly without lifting.

B.4.4 Construction of $\mathcal{R}(\Omega)$

The marriage of propositions 4.3 and B.2 makes for a practical way of determining the minimal representation $\mathcal{R}(\Omega)$ of any non-degenerate constraint Ω . As remarked before, degenerate constraints are of little practical relevance, although if needed, the degenerate case can be covered by adapting proposition B.1.

When forming a convex hull, `qhull` produces a list of its vertices and readily checks if two vertices form the endpoints of an edge. Based on that, we devise the following:

Proposition B.3 (constructing $\mathcal{R}(\Omega)$). *Let Ω be a non-degenerate constraint, and let $\mathcal{H}_\pm^{(p)}$, $M_\pm^{(p)}$, etc. be defined as in proposition B.2. Then the unique minimal representation $\mathcal{R}(\Omega)$ is determined as follows:*

- For $i \in I_{\pm 1}$, $\langle \alpha_i, \pm 1 \rangle \in \mathcal{R}(\Omega)$ if and only if α_i is an extreme point of $\mathcal{A}_{\pm 1}(\Omega)$. (The exception is $\langle \mathbf{0}, -1 \rangle$, which is of course not in $\mathcal{R}(\Omega)$.)

Equivalently, $\langle \alpha_i, \pm 1 \rangle \in \mathcal{R}(\Omega)$ if and only if $\alpha_i \in \mathcal{P}_\pm$ is an extreme point of $\mathcal{H}_\pm^{(p)}$ for some $p > 1$, with the same exception.

- For $j \in I_0$, $\langle \alpha_j, 0 \rangle \in \mathcal{R}(\Omega)$ if and only if there is some $i \in I_{\pm 1}$ such that the ray $\{\alpha_i + \lambda \alpha_j \mid \lambda \geq 0\}$ is an edge of $\mathcal{A}_{\pm 1}(\Omega)$ that contains no point in $\mathcal{N}_\pm^{(p)}(\Omega)$. (This is up to normalisation; several equivalent $\langle \alpha_i, 0 \rangle$ may satisfy this condition.)

Equivalently, $\langle \alpha_j, 0 \rangle \in \mathcal{R}(\Omega)$ if and only if there is some $i \in I_{\pm 1}$ such that the line segment between α_i and $\alpha_i + (p-1)\alpha_j$ is an edge of $\mathcal{H}_\pm^{(p)}(\Omega)$ that contains no point in $\mathcal{N}_\pm^{(p)}(\Omega)$. (This breaks the normalisation ambiguity: if several $\langle \alpha_j, 0 \rangle$ are equivalent under eq. (4.5), then only the one with the largest $|\alpha_j|$ will form the endpoint of their edge and be included in $\mathcal{R}(\Omega)$.)

Note how this can be applied to corollary B.4 without modification. Illustrations can be found in figures 26 to 29.

PROOF. We will prove that the stated conditions are equivalent to $\langle \alpha_i, c_i \rangle$ uniquely supporting a point; the rest follows from proposition 4.3. We will focus on proving the conditions based on $\mathcal{H}_\pm^{(p)}$ rather than $\mathcal{A}_{\pm 1}(\Omega)$; that they are equivalent follows easily from property (v) of K -faces along with lemma B.4 ($\alpha_i \in \mathcal{P}_\pm$ serves as the point φ).⁶¹

Consider first $i \in I_{\pm 1}$. If $\langle \alpha_i, \pm 1 \rangle \in \mathcal{R}(\Omega)$, then there is some \mathbf{h} uniquely supported by $\langle \alpha_i, \pm 1 \rangle$; in particular, for any $\beta \in \mathcal{A}_{\pm 1}(\Omega)$, $\beta \neq \alpha$ we have $\beta \cdot \mathbf{h} > \alpha_i \cdot \mathbf{h} = \pm 1$. Since by eq. (B.72)

$$\mathcal{A}_{\pm 1}(\Omega) = \bigcup_{p=2}^{\infty} \mathcal{H}_\pm^{(p)} \quad (\text{B.93})$$

we have that for any $p > 1$

$$\left\{ \boldsymbol{\eta} \in \mathcal{H}_\pm^{(p)} \mid \boldsymbol{\eta} \cdot \mathbf{h} = \pm 1 \right\} = \{ \alpha_i \}. \quad (\text{B.94})$$

From lemma B.3, we conclude that α_i is an extreme point of $\mathcal{H}_\pm^{(p)}$.

Conversely, if α_i is an extreme point of $\mathcal{H}_\pm^{(p)}$, then eq. (B.94) holds. In particular, if $j \in I_{\mp 1}$ then $\mathbf{h} \cdot (p\alpha_i + (p-1)\alpha_j) > \pm 1$ so $\alpha_j \cdot \mathbf{h} > \mp 1$. Similarly, for any $j \in I_0$ we have $\alpha_j \cdot \mathbf{h} > 0$. It then follows from lemma B.5 and proposition 4.3 that $\langle \alpha_i, \pm 1 \rangle \in \mathcal{R}(\Omega)$.

Consider then $j \in I_0$. If the line segment between α_i and $\alpha_i + (p-1)\alpha_j$ is an edge of $\mathcal{H}_\pm^{(p)}$, then by lemma B.3 there exists $\langle \mathbf{h}, d \rangle$ such that

$$\left\{ \boldsymbol{\eta} \in \mathcal{H}_\pm^{(p)} \mid \boldsymbol{\eta} \cdot \mathbf{h} = d \right\} = \{ \alpha_i + \lambda \alpha_j \mid \lambda \in [0, (p-1)] \} \quad (\text{B.95})$$

and $\mathcal{H}_\pm^{(p)} \subseteq \mathcal{B}(\langle \mathbf{h}, d \rangle)$. We then immediately find that $\mathbf{h} \cdot \alpha_i = d$, $\mathbf{h} \cdot \alpha_j = 0$.

What about the other $\alpha_k, k \in I$? If $k \in I_{\mp}$ then $p\alpha_i + (p-1)\alpha_k \in \mathcal{N}_\pm^{(p)}$ which is by assumption not supported by $\langle \mathbf{h}, d \rangle$, so it follows that $\mathbf{h} \cdot \alpha_k > -d$. In the same way, for any $k \in I_0$ we have $\alpha_i + (p-1)\alpha_k \in \mathcal{Z}_\pm^{(p)}$, implying that $\mathbf{h} \cdot \alpha_k \geq 0$. If $\mathbf{h} \cdot \alpha_k = 0$ then it follows that α_k is collinear with α_j , so either $\langle \alpha_i, 0 \rangle = \langle \alpha_k, 0 \rangle$ or $\langle \alpha_i, 0 \rangle = \langle -\alpha_k, 0 \rangle$ under eq. (4.5); the latter is excluded by non-degeneracy. Therefore, we can without loss of generality assume $\mathbf{h} \cdot \alpha_k > 0$ for all $k \in I_0 \setminus \{j\}$ by omitting equivalent constraints. Lastly, if $k \in I_\pm$ then $\mathbf{h} \cdot \alpha_k \geq d$; here, the non-strict inequality is unavoidable but does not pose a problem.

If $i \in I_+$, then the fact that $p\alpha_i \in \mathcal{N}_+^{(p)}$ implies that $d > 0$. If $i \in I_-$, we can use the fact that $\mathbf{0} \in \mathcal{P}_-$ to prove that $d \leq 0$, with equality only if $\alpha_i = \mathbf{0}$. Combining this with the previous paragraph, we conclude that

$$\forall k \in I \setminus \{i, j\}, \mathbf{h} \cdot \alpha_k \geq c_k |d| \quad (\text{B.96})$$

⁶¹There is one subtlety for $\langle \alpha_j, 0 \rangle \in \mathcal{R}(\Omega)$: the line segment between α_i and $\alpha_i + (p-1)\alpha_j$ may fail to be an edge of $\mathcal{H}_\pm^{(p)}$ even though $\{ \alpha_i + \lambda \alpha_j \mid \lambda \geq 0 \}$ is an edge of $\mathcal{A}_{\pm 1}(\Omega)$, if it is contained in the line segment between α_i and some $\alpha_{i'}, i' \in I_{\pm 1}$. This can be remedied by using sufficiently large p . For $\langle \alpha_i, \pm 1 \rangle \in \mathcal{R}(\Omega)$, all p work equally well, as should be apparent from the proof.

(equality only possible if $k \in I_{\pm}$). For sufficiently small $\varepsilon > 0$, we may replace it by $|d| \rightarrow |d| \mp \varepsilon$ without invalidating the above inequality; this also guards against problems when $d = 0$. Then

$$\frac{1}{|d| \mp \varepsilon} \mathbf{h} \cdot \boldsymbol{\alpha}_i > \pm 1, \quad \frac{1}{|d| \mp \varepsilon} \mathbf{h} \cdot \boldsymbol{\alpha}_j = 0, \quad \forall k \in I \setminus \{i, j\}, \quad \frac{1}{|d| \mp \varepsilon} \mathbf{h} \cdot \boldsymbol{\alpha}_k > c_k. \quad (\text{B.97})$$

Thus, $\langle \boldsymbol{\alpha}_j, 0 \rangle$ uniquely supports $\frac{1}{|d| \mp \varepsilon} \mathbf{h}$ by lemma B.5.

Conversely, assume $\boldsymbol{\alpha}_j$ uniquely supports a point \mathbf{b} . Then

$$\mathbf{b} \cdot \boldsymbol{\alpha}_j = 0, \quad \forall i \in I \setminus j, \quad \mathbf{b} \cdot \boldsymbol{\alpha}_i > c_i. \quad (\text{B.98})$$

In particular,

$$\mathcal{F} = \{\boldsymbol{\beta} \in \mathcal{A}_0(\Omega) \mid \mathbf{b} \cdot \boldsymbol{\beta} = 0\} = \{\lambda \boldsymbol{\alpha}_j \mid \lambda \geq 0\}. \quad (\text{B.99})$$

(Non-degeneracy ensures $\lambda \geq 0$). By lemma B.3 and proposition B.2, \mathcal{F} is an edge of $\mathcal{A}_0(\Omega) = \mathcal{B}(\sum_{\ell \in M_p^{(\pm)}} \langle \mathbf{n}_\ell, 0 \rangle)$. Therefore, lemma B.2 guarantees that there exists $Q_{\pm} \subseteq M_{\pm}^{(p)}$ such that

$$\forall \ell \in Q_{\pm}, \quad \mathbf{n}_\ell \cdot \boldsymbol{\alpha}_j = 0, \quad \forall \ell \in M_{\pm}^{(p)} \setminus Q_{\pm}, \quad \mathbf{n}_\ell \cdot \boldsymbol{\alpha}_j > 0 \quad (\text{B.100})$$

where $\{\mathbf{n}_\ell\}_{\ell \in Q_{\pm}}$ contains a subset of $D-1$ linearly independent vectors. Now, consider the set

$$\mathcal{C} \equiv \mathcal{B}\left(\sum_{\ell \in Q_{\pm}} \langle \mathbf{n}_\ell, r_\ell \rangle\right) \supseteq \mathcal{A}_{\pm 1}(\Omega). \quad (\text{B.101})$$

This set must have at least one edge, since the intersection of $D-1$ independent hyperplanes is a line. By lemma B.2, there is then a subset $J_{\pm} \subseteq Q_{\pm}$ with $|J_{\pm}| = D-1$ and $\{\mathbf{n}_\ell\}_{\ell \in J_{\pm}}$ linearly independent, that describes that edge according to eq. (B.30) (if there are several, we choose one arbitrarily). By construction, \mathcal{C} has no extreme points: $\{\mathbf{n}_\ell\}_{\ell \in Q_{\pm}}$ has no subset of D linearly independent vectors. Therefore, the chosen edge can be written like

$$\mathcal{F}_{\pm} \equiv \left\{ \boldsymbol{\beta} \in \mathbb{R}^D \mid \forall \ell \in J_{\pm}, \quad \mathbf{n}_\ell \cdot \boldsymbol{\beta} = r_\ell \right\}. \quad (\text{B.102})$$

We now claim that

$$\mathcal{F}_{\pm} \cap \mathcal{A}_{\pm 1}(\Omega) = \mathcal{F}_{\pm} \cap \mathcal{B}\left(\sum_{\ell \in M_{\pm}^{(p)}} \langle \mathbf{n}_\ell, r_\ell \rangle\right) \neq \emptyset. \quad (\text{B.103})$$

All points in \mathcal{F}_{\pm} are satisfied by $\langle \mathbf{n}_\ell, r_\ell \rangle$ for $\ell \in Q_{\pm}$ by the arguments made above about \mathcal{C} . For $\ell \notin Q_{\pm}$, we note that for any $\boldsymbol{\beta} \in \mathcal{F}$, we also have $\boldsymbol{\beta} + \lambda \boldsymbol{\alpha}_j \in \mathcal{F}$ for all $\lambda \in \mathbb{R}$, so since $\mathbf{n}_\ell \cdot \boldsymbol{\alpha}_j > 0$, $\mathbf{n}_\ell \cdot (\boldsymbol{\beta} + \lambda \boldsymbol{\alpha}_j) \geq r_\ell$ for sufficiently large λ no matter what $\mathbf{n}_\ell \cdot \boldsymbol{\beta}$ is. Therefore, at least some subset of \mathcal{F}_{\pm} is satisfied by all $\langle \mathbf{n}_\ell, r_\ell \rangle$, $\ell \in M_{\pm}^{(p)} \setminus Q_{\pm}$, which proves eq. (B.103). By then considering the constraint

$$\langle \mathbf{h}, d \rangle \equiv \left\langle \sum_{\ell \in J_{\pm}} \mathbf{n}_\ell, \sum_{\ell \in J_{\pm}} r_\ell \right\rangle, \quad (\text{B.104})$$

it follows from lemma B.3 that $\mathcal{F}_\pm \cap \mathcal{A}_{\pm 1}(\Omega)$ is a K -facet of $\mathcal{A}_{\pm 1}(\Omega)$.

Let us now note that $\mathcal{F}_\pm \cap \mathcal{A}_{\pm 1}(\Omega)$ is an edge, not an extreme point, since for any $\beta \in \mathcal{F}_\pm$ and $\lambda \geq 0$,⁶²

$$\beta + \lambda \alpha_j \in \mathcal{F}_\pm \cap \mathcal{A}_{\pm 1}(\Omega). \tag{B.105}$$

From eqs. (B.71), (B.93) and (B.103), there is some $q > 1$ such that $\mathcal{F}_\pm \cap \mathcal{H}_\pm^{(p)} \neq \emptyset$ for all $p \geq q$, and by property (v) of K -facets, $\mathcal{F}_\pm \cap \mathcal{H}_\pm^{(p)}$ is an edge of $\mathcal{H}_\pm^{(p)}$. Unlike $\mathcal{A}_{\pm 1}(\Omega)$, $\mathcal{H}_\pm^{(p)}$ is always a compact set, so all its edges must be line segments with exactly two endpoints, which are extreme points of $\mathcal{H}_\pm^{(p)}$ by property (iv). Thus, $\mathcal{F}_\pm \cap \mathcal{H}_\pm^{(p)}$ is the line segment between two extreme points of $\mathcal{H}_\pm^{(p)}$, and these must clearly be elements of $\mathcal{P}_\pm \cup \mathcal{Z}_\pm^{(p)} \cup \mathcal{N}_\pm^{(p)}$.⁶³

Then, let us prove that at least one of these endpoints of $\mathcal{F}_\pm \cap \mathcal{H}_\pm^{(p)}$ is in \mathcal{P}_\pm by considering the alternatives. For $\mathcal{N}_\pm^{(p)}$, i.e. if $p\alpha_i + (p-1)\alpha_k$ is one endpoint for some $i \in I_{\pm 1}, k \in I_{\mp 1}$, then it follows from lemma B.11 that $\alpha_i \in \mathcal{F}_\pm \cap \mathcal{H}_\pm^{(p)}$, but also (see the proof of lemma B.12) that $(p-1)(\alpha_i + \alpha_k) = \lambda \alpha_j$ for some λ . This contradicts the fact that λ uniquely supports a point, since any point in $\mathcal{B}(\Omega)$ it supports is also supported by α_i and α_k . For $\mathcal{Z}_\pm^{(p)}$, i.e. if $\alpha_i + (p-1)\alpha_k$ is one endpoint for some $i \in I_{\pm 1}, k \in I_0$, then it similarly follows that $\alpha_i \in \mathcal{F}_\pm \cap \mathcal{H}_\pm^{(p)}$, and that $\alpha_k = \lambda \alpha_j$, where $\lambda > 0$ by non-degeneracy.

Thus, $\mathcal{N}_\pm^{(p)}$ is excluded and $\mathcal{Z}_\pm^{(p)}$ can account for at most one endpoint of $\mathcal{F}_\pm \cap \mathcal{H}_\pm^{(p)}$, since the line segment between $\alpha_i + \lambda(p-1)\alpha_j$ and $\alpha_{i'} + \lambda'(p-1)\alpha_j$ cannot contain both α_i and $\alpha_{i'}$, which are clearly contained in $\mathcal{F}_\pm \cap \mathcal{H}_\pm^{(p)}$. Hence, at least one of the endpoints is α_i for some $i \in I_{\pm 1}$, and the other may either be $\alpha_{i'}, i' \in I_{\pm 1}$ or $\alpha_i + (p-1)\alpha_j$ (possibly after exchanging $\langle \alpha_j, 0 \rangle$ for an equivalent constraint). The former case was covered in the first paragraph of this proof, and is removed by considering sufficiently large p ; the latter completes our proof. ■

B.4.5 Visualisation of $\mathcal{A}_c(\Omega)$

We provided figure 2 for illustration along with the statement of proposition 4.1, since the shapes of $\mathcal{A}(\omega_c)$ are quite simple to interpret, and provide some insight into the result. The same cannot be said for the general $\mathcal{A}_c(\Omega)$, however, so we have put off a similar display until now. Figure 26 showed a single example in great detail, and now figures 27 to 29 illustrate $\mathcal{A}_c(\Omega)$ in a similar manner to how figure 2 illustrated $\mathcal{A}(\omega_c)$. For legibility, we have omitted the supporting sets $\mathcal{P}_\pm, \mathcal{Z}_\pm^{(p)}$ and $\mathcal{N}_\pm^{(p)}$, but if one wishes, it is not difficult to imagine them in the figures like in figure 26 to make sense of the shapes.

⁶²It also follows from this that $\mathcal{F}_\pm \cap \mathcal{A}_{\pm 1}(\Omega)$ is not a line segment. That it is a ray rather than a full line is not difficult to prove from non-degeneracy, and it is then possible to show that $\mathcal{F}_\pm \cap \mathcal{A}_{\pm 1}(\Omega) = \{\alpha_i + \lambda \alpha_j \mid \lambda \geq 0\}$ which completes the proof, but we choose an easier path.

⁶³This is made rigorous by comparing eq. (B.63) with the *Krein-Milman theorem*:

Let \mathcal{C} be a compact convex set and \mathcal{E} be the set of its extreme points. Then $\mathcal{C} = \text{Hull}(\mathcal{E})$, and if $\mathcal{C} = \text{Hull}(\mathcal{E}')$ then $\mathcal{E}' \supseteq \mathcal{E}$.




Figure 27. Example of the region $\mathcal{A}_{-1}(\Omega)$ (red). The α_i , identical (up to $c_i = 0$ rescaling) to those in figure 26, are marked with ∇ for $i \in I_{-1}$, \square for I_0 and \triangle for I_1 , and have their convex hulls outlined. The α_i that are identified as relevant from the construction of $\mathcal{A}_{\pm 1}(\Omega)$, as per proposition B.3, are filled; the rest are left empty. Note how the $c_i = +1$ ones can strictly speaking only be deduced from $\mathcal{A}_1(\Omega)$ (figure 28 below), although their relevance can be easily seen in “nice” cases such as this. In the large figure, the full $\mathcal{A}_{-1}(\Omega)$ is drawn. The smaller figures demonstrate the effect of omitting α_i from the definition of Ω , with i in various combinations of I_1, I_0 and I_{-1} . The only- I_{-1} figure (bottom left) is thus analogous to $\mathcal{A}(\omega_{-1})$ in figure 2. Note how in the only- I_1 figure (bottom right), proposition B.2 still works since $\mathcal{P}_{-1} = \{\mathbf{0}\}$ rather than being empty.

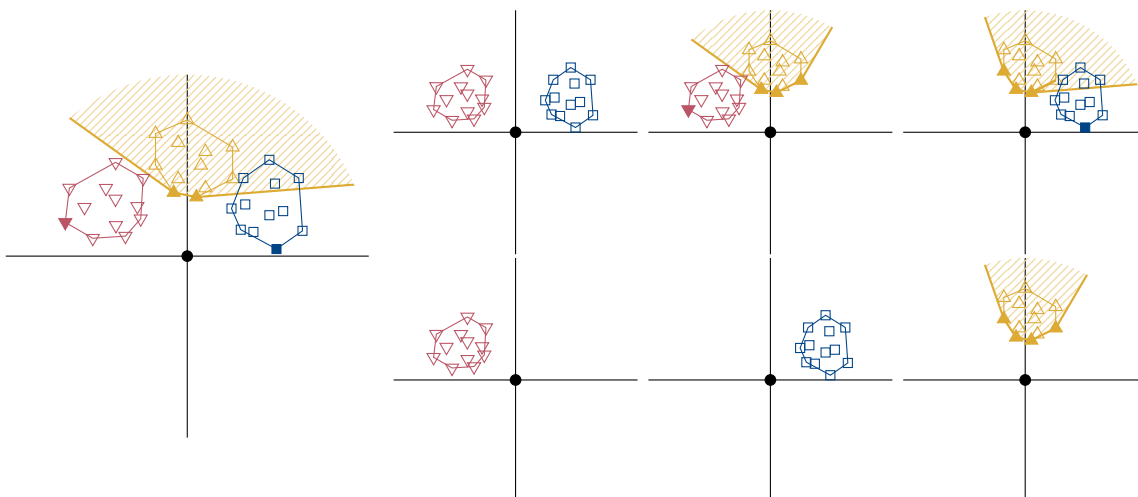


Figure 28. Example of the region $\mathcal{A}_1(\Omega)$ (yellow), with the same α_i and analogous presentation as in figure 27. Note how some relevant $\alpha_i, i \notin I_1$ lie outside $\mathcal{A}_1(\Omega)$; this is also possible for $\mathcal{A}_0(\Omega)$, but not for $\mathcal{A}_{-1}(\Omega)$, since it contains all α_i . As in figure 27, the smaller figures illustrate the effect of omitting some α_i . The only- I_1 figure (bottom right) is thus analogous to $\mathcal{A}(\omega_1)$ in figure 2. Note how $\mathcal{A}_1(\Omega) = \emptyset$ when $i \in I_1$ are omitted, as can be seen from eq. (4.11) with $\sum_{i \in I_1} \lambda_i = 0$ (this is the exception to proposition B.2).




Figure 29. Example of the region $\mathcal{A}_0(\Omega)$ (blue), with the same α_i as in figures 27 and 28. $\mathcal{A}_{\pm 1}(\Omega)$ are outlined to demonstrate how $\mathcal{A}_0(\Omega)$ can be constructed by taking their unbounded facets and shifting them so they pass through the origin, as discussed in appendix B.4.3. Note that unlike in figures 27 and 28, not all relevant α_i influence the shape of $\mathcal{A}_0(\Omega)$. As in figures 27 and 28, the smaller figures illustrate the effect of omitting some α_i . The only- I_0 figure (bottom centre) is thus analogous to $\mathcal{A}(\omega_0)$ in figure 2. Note how in the only- I_{-1} figure (bottom left), $\mathcal{A}_0(\Omega) = \{\mathbf{0}\}$ since $\mathcal{A}_{-1}(\Omega)$ has no unbounded faces (this is one of the exceptions to corollary B.3).

B.5 The duality between $\mathcal{A}_c(\Omega)$ and $\mathcal{B}(\Omega)$

As has been used many times above, parameter space and constraint space are dual in the sense that points in one correspond to hyperplanes in the other. This extends to $\mathcal{B}(\Omega)$ and $\mathcal{A}_c(\Omega)$, which are similarly related in ways we will explore in this section.

A taste of this duality can be found in the following result, which will also be useful further on:

Corollary B.5 (duality). *With $\Omega \neq \Omega_\infty$, let \mathcal{U}_\pm be defined as in eq. (B.66), so that $\mathcal{B}(\mathcal{U}_\pm) = \mathcal{A}_{\pm 1}(\Omega)$. Then $\mathcal{B}(\Omega) = \mathcal{A}_1(\mathcal{U}_+) \cap \mathcal{A}_{-1}(\mathcal{U}_-)$ if $I_1 \neq \emptyset$, and $\mathcal{B}(\Omega) = \mathcal{A}_{-1}(\mathcal{U}_-)$ otherwise.*

Loosely, one can think of this as “proposition B.2 is its own inverse”: applying it to Ω gives \mathcal{U}_\pm , and applying it to \mathcal{U}_\pm gives Ω .

PROOF. By definition and proposition B.2,

$$\begin{aligned} \mathbf{b} \in \mathcal{B}(\Omega) &\Leftrightarrow \forall \langle \alpha, c \rangle \leq \Omega, \quad \alpha \cdot \mathbf{b} \geq c, \\ \mathbf{b} \in \mathcal{A}_{\pm 1}(\mathcal{U}_{\pm 1}) &\Leftrightarrow \mathcal{A}_{\pm 1}(\Omega) \subseteq \mathcal{B}(\langle \mathbf{b}, \pm 1 \rangle). \end{aligned} \tag{B.106}$$

Looking at eq. (4.12), we see that the latter is equivalent to

$$\forall \{\lambda_i\}_{i \in I} \text{ with } \lambda_i \geq 0 \text{ and } \sum_{i \in I} \lambda_i c_i \geq \pm 1, \quad \mathbf{b} \cdot \sum_{i \in I} \lambda_i \alpha_i \geq \pm 1. \tag{B.107}$$

This is true if $\forall i \in I, \mathbf{b} \cdot \alpha_i \geq c_i$, but it only implies $\forall i \in I \setminus I_{\mp 1}, \mathbf{b} \cdot \alpha_i \geq c_i$, as some straightforward algebra shows. By using $\mathcal{A}_1(\mathcal{U}_+) \cap \mathcal{A}_{-1}(\mathcal{U}_-)$, we avoid this shortcoming and complete the proof. When $I_1 = \emptyset$, we cannot use \mathcal{U}_+ but also do not need it, since $I \setminus I_1 = I$. \square

This relation can be extended into a geometric duality between $\mathcal{B}(\Omega)$ and $\mathcal{R}_c(\Omega)$: the vertices of one correspond directly to the facets of the other.⁶⁴

Proposition B.4 (precise duality relations). *Let $M_{\pm}^{(p)}$, etc. be defined as in proposition B.2, and let \mathcal{U}_{\pm} be defined as in eq. (B.66). Let $\pm 1 = 1$ if $I_1 \neq \emptyset$ and $\pm 1 = -1$ otherwise, as in corollary B.5. If Ω and \mathcal{U}_{\pm} are non-degenerate, then the following correspondences hold:*

- (i) α_i for $i \in I_{\pm 1}$ is a vertex of $\mathcal{A}_{\pm 1}(\Omega)$ if and only if $\langle \alpha_i, \pm 1 \rangle \in \mathcal{R}(\Omega)$.
- (ii) $\frac{1}{|r_{\ell}|} \mathbf{n}_{\ell}$ is a vertex of $\mathcal{B}(\Omega)$ if and only if $\ell \in M_{\pm}^{(p)}$ and $r_{\ell} \neq 0$.
- (iii) $\mathcal{F} \equiv \{\beta \in \mathcal{A}_{\pm 1}(\Omega) \mid \mathbf{n}_{\ell} \cdot \beta = r_{\ell}\}$ is a facet of $\mathcal{A}_{\pm 1}(\Omega)$ if and only if $\ell \in M_{\pm}^{(p)}$.
- (iv) $\mathcal{F} \equiv \{\mathbf{b} \in \mathcal{B}(\Omega) \mid \alpha \cdot \mathbf{b} = c\}$ is a facet of $\mathcal{B}(\Omega)$ if and only if $\langle \alpha, c \rangle \in \mathcal{R}(\Omega)$.

All relations are exhaustive: there is no vertex of $\mathcal{A}_{\pm 1}(\Omega)$ that is not covered by correspondence (i), etc. There is the single exception that $\mathbf{0}$ may be a vertex without corresponding to a facet.⁶⁵ Also, the exceptions to proposition B.2 apply.

This was used to obtain the visualisations of $\mathcal{B}(\Omega)$ in section 6. An illustrative example is given in figure 30.

Note how proposition B.4 is rather negligent of $c = 0$ constraints, partly because they are complicated to handle, and partly because an exact zero is an unlikely thing when constraints are generated with numerical inaccuracies. For similar reasons, we do not consider the degenerate case.

PROOF. Correspondence (i) is just proposition B.3.

Correspondence (ii), which is the most useful correspondence for visualisation, follows by applying proposition B.3 to \mathcal{U}_{\pm} instead. To do this, we first apply the normalisation

$$\langle \mathbf{n}_{\ell}, r_{\ell} \rangle \rightarrow \begin{cases} \left\langle \frac{1}{|r_{\ell}|} \mathbf{n}_{\ell}, \pm 1 \right\rangle & \text{if } r_{\ell} \neq 0, \\ \langle \mathbf{n}_{\ell}, 0 \rangle & \text{otherwise,} \end{cases} \quad (\text{B.108})$$

recalling lemma B.8. Granted that \mathcal{U}_{\pm} is non-redundant, proposition B.3 therefore states that the extreme points (vertices) of $\mathcal{A}_{\pm 1}(\mathcal{U}_{\pm})$, except $\mathbf{0}$, are exactly $\frac{1}{|r_{\ell}|} \mathbf{n}_{\ell}$, since $\left\langle \frac{1}{|r_{\ell}|} \mathbf{n}_{\ell}, \pm 1 \right\rangle$ are exactly the relevant elements of \mathcal{U}_{\pm} with $r_{\ell} \neq 0$. The correspondence then follows from corollary B.5, as long as we can ensure that the intersection does not introduce any new vertices that are not vertices of $\mathcal{A}_{\pm 1}(\mathcal{U}_{\pm})$. Such a vertex would, by lemma B.3, be uniquely supported by some constraint that is unaccounted for by proposition B.3, so it does not exist.

Correspondences (iii) and (iv) both follow from proposition 4.3, since each facet contains all points that are uniquely supported by some relevant constraint. It follows from

⁶⁴This is closely related to the concept of *dual polytopes*, where a polytope is a D -dimensional generalisation of a polygon or polyhedron; $\mathcal{B}(\Omega)$ and $\mathcal{A}_c(\Omega)$ are polytopes, if the definition is relaxed to permit unbounded polytopes.

⁶⁵The origin is a vertex of $\mathcal{B}(\Omega)$ if $I_1 = \emptyset$ and $\{\alpha_i\}_{i \in I_0}$ is a basis of \mathbb{R}^D . It is never a vertex of (or even contained in) $\mathcal{A}_1(\Omega)$, but is always a vertex of $\mathcal{A}_{-1}(\Omega)$ if Ω is non-degenerate.




Figure 30. Example of how $\mathcal{A}_{\pm 1}\Omega$ are dual to $\mathcal{B}(\Omega)$. A different, simpler Ω than in previous figures is used. *Left:* constraint space, with $\mathcal{A}_{-1}(\Omega)$ constructed in a similar way to figure 26. The points α_i are marked with filled dots and labelled; points in $\mathcal{P}_-(\Omega)$, $\mathcal{Z}_-^{(2)}(\Omega)$ and $\mathcal{N}_-^{(2)}(\Omega)$ are marked with empty dots. The facets corresponding to constraints $\langle n_\ell, r_\ell \rangle$ are also labelled. *Middle:* constraint space, with $\mathcal{A}_1(\Omega)$ constructed similarly. *Right:* parameter space, drawn similarly to figure 1. $\mathcal{B}(\Omega)$ is shaded; the six constraints $\langle \alpha_i, c_i \rangle$ and vertices $v_\ell = n_\ell \frac{\pm 1}{r_\ell}$ are labelled. Vertices from $\mathcal{A}_{-1}(\Omega)$ are marked with \blacktriangledown , and those from $\mathcal{A}_1(\Omega)$ with \blacktriangle . Note that two vertices sharing an edge of $\mathcal{B}(\Omega)$ corresponds to two constraints supporting a common point in $\mathcal{A}_{\pm 1}\Omega$. In the case of $c = 0$ edges, it corresponds to supporting $(\alpha_i + p\alpha_j) \in \mathcal{Z}_+^{(p)}(\Omega)$ and $(\alpha_{i'} + p\alpha_j) \in \mathcal{Z}_-^{(p)}(\Omega)$ for a common $j \in I_0$. Note also how corollary B.5 applies (compare figure 2).

corollary B.2 and lemma B.3 that a point \mathbf{b} is uniquely supported by some constraint if and only if \mathbf{b} is contained in some $(D - 1)$ -face, but not in any K -face for $K < D - 1$. It also follows from the definition of a K -face that all K -faces contain at least one point that is not contained in any K' -face, $K' < K$. Thus, every facet gives rise to a relevant constraint, and no constraint can be relevant if it does not give rise to a facet. ■

B.5.1 Duality with a bounding box

It would be possible to extend proposition B.4 to relate the n_ℓ with $r_\ell = 0$ to unbounded edges of $\mathcal{B}(\Omega)$ by equating them to the ray $\left\{ \frac{1}{|r_k|} n_k + \lambda n_\ell \mid \lambda \geq 0 \right\}$ for some k (see proposition B.3). However, it is not obvious how to find k , and it is not guaranteed that we exhaustively cover the unbounded edges this way. In either case, unbounded edges are far less pleasant to deal with than vertices, even though they do provide shape information that the vertices alone cannot provide.

Unbounded edges can be wholly avoided by artificially introducing $2D$ extra constraints that constrain $\mathcal{B}(\Omega)$ to a D -dimensional bounding box. All unbounded facets of $\mathcal{B}(\Omega)$ are cropped, and proposition B.4, correspondence (ii) seamlessly provides all the vertices that define the intersections between $\mathcal{B}(\Omega)$ and the walls of the box. This method, demonstrated in figure 31, was used extensively in section 6.

Everything mentioned here can of course be applied equally well to $\mathcal{R}_{\pm 1}(\Omega)$.

B.6 A note on infinite sums of constraints

It is extremely important that we only ever consider sums of a finite number of constraints, not only for the validity of our proofs, but also for the validity of the propositions themselves. Consider as a counterexample the following countably infinite sum of one-dimensional




Figure 31. Example of how $\mathcal{B}(\Omega) = \langle \alpha_1, 1 \rangle + \langle \alpha_2, 1 \rangle$, which would normally have a single vertex v_1 and two unbounded edges corresponding to $r_k = 0$ constraints, is made more manageable by imposing an artificial set of bounding-box constraints (\blacktriangledown), giving rise to some points in $\mathcal{N}_+^{(2)}(\Omega)$ (\blacktriangledown). Like in figure 30, $\mathcal{A}_1(\Omega)$ is shown on the left, and $\mathcal{B}(\Omega)$ is shown on the right.

constraints:

$$\omega_{-1} = \sum_{k=1}^{\infty} \left\langle 1 - \frac{1}{k}, -1 \right\rangle \Rightarrow \mathcal{A}(\omega_{-1}) = [0, 1), \quad \mathcal{B}(\omega_{-1}) = [-1, \infty). \quad (\text{B.109})$$

Here, $\mathcal{A}(\omega_{-1})$ is not a closed set. On the other hand, $\mathcal{B}(\omega_{-1})$ is closed, since for any $\varepsilon > 0$, the point $-(1 + \varepsilon)$ fails to satisfy the element $\left\langle 1 - \frac{1}{k}, -1 \right\rangle$ for sufficiently large k (specifically, larger than $\frac{1+\varepsilon}{\varepsilon}$). Thus, $\mathcal{B}(\omega_{-1}) = \mathcal{B}(\langle 1, -1 \rangle)$, so $\langle 1, -1 \rangle \leq \omega_{-1}$ even though $1 \notin \mathcal{A}(\omega_{-1})$. Thus, proposition 4.1, and therefore proposition 4.2, fails.

There is, however, a straightforward generalisation. We let eqs. (4.9) and (4.10) define Ω and $\mathcal{A}_c(\Omega)$ as before, but now with infinite sums permitted, so that $\{\alpha_i\}_{i \in I_c}$ may be any subset of \mathbb{R}^D . Then we have the following:

Proposition B.5 (determining if constraint is weaker, generalised). *Let $\langle \beta, c \rangle$ be a linear constraint, and let $\Omega \neq \Omega_\infty$ be a possibly infinite sum of linear constraints. Then $\langle \beta, c \rangle \leq \Omega$ if and only if $\beta \in \text{cl}[\mathcal{A}_c(\Omega)]$.*

The simple introduction of the closure solves all issues with infinite combinations, such as the counterexample above. Note how the only part of appendices B.1.2 to B.1.4 that relies on the finiteness of I is the proof that $\mathcal{A}_c(\Omega)$ is closed. The closure is by definition closed, so the remaining arguments in these sections remain valid when applied to $\text{cl}[\mathcal{A}_c(\Omega)]$ instead.

The only part of the proof of proposition 4.2 that does not immediately carry over is appendix B.1.1. However, by the limit definition of closure, for any $\beta \in \text{cl}[\mathcal{A}_c(\Omega)]$ there exists a sequence β_n in $\mathcal{A}_c(\Omega)$ such that $\lim_{n \rightarrow \infty} \beta_n = \beta$. Since the function $\beta \rightarrow \beta \cdot \mathbf{b}$ is continuous, it follows that for any \mathbf{b} , $\lim_{n \rightarrow \infty} \beta_n \cdot \mathbf{b} = \beta \cdot \mathbf{b}$. Thus, since the arguments of appendix B.1.1 hold for all β_n , they also hold for $\beta \in \text{cl}[\mathcal{A}_c(\Omega)]$. ■

Corollary B.1 does not require any adjustment, since the interior of any convex set in \mathbb{R}^D is equal to the interior of its closure.

The generalisation of proposition 4.3, for which counterexamples abound, is less straightforward, partly because there does not necessarily exist a minimal representation if all

representations are infinite. However, the notion of a minimal representation is mainly motivated as being the most practical format of a constraint, so it is not very useful in the infinite case, which is anyway only of theoretical interest. For the same reason, it is not relevant to our study to attempt to generalise any of the propositions presented in this appendix.

B.7 Mathematical glossary

The table below contains a list of notations and terms that may be unfamiliar to some readers, depending on their background. (We have chosen to employ such notation, since it makes some things much more brief and expressive, even though it necessitates this table.)

CONCEPT	DESCRIPTION
\square, \blacksquare	End of proof. We use \blacksquare for main proofs and \square for lemmata.
\forall	The universal quantifier. Informally, short for “for all”.
\exists	The existence quantifier. Informally, short for “there exists”.
$x \in \mathcal{X}$	The object x is contained in (is an element of) the set \mathcal{X} . We typically denote sets using a calligraphic font, but use ordinary italics for sets of indices.
\emptyset	The empty set. The set that contains no elements.
$ \mathcal{X} , \mathbf{v} , s $	Cardinality or magnitude. For a set, this indicates its cardinality (number of elements). For a vector or scalar, the same notation indicates its magnitude $\sqrt{\mathbf{v} \cdot \mathbf{v}}$ or $\sqrt{s^2}$.
$\mathcal{X} \subseteq \mathcal{Y}$	\mathcal{X} is a subset of \mathcal{Y} . All elements of \mathcal{X} are also contained in \mathcal{Y} . Just $\mathcal{X} \subset \mathcal{Y}$ means the same, except that $\mathcal{X} \neq \mathcal{Y}$; that is, \mathcal{Y} has at least one element not contained in \mathcal{X} .
$\mathcal{X} \cup \mathcal{Y}$	The union of two sets. The set that contains all elements contained in either \mathcal{X} or \mathcal{Y} , or both. Clearly, $\mathcal{X} \subseteq \mathcal{X} \cup \mathcal{Y}$ and $\mathcal{Y} \subseteq \mathcal{X} \cup \mathcal{Y}$.
$\mathcal{X} \cap \mathcal{Y}$	The intersection of two sets. The set that contains all elements contained in <i>both</i> \mathcal{X} or \mathcal{Y} . Clearly, $\mathcal{X} \supseteq \mathcal{X} \cap \mathcal{Y}$ and $\mathcal{Y} \supseteq \mathcal{X} \cap \mathcal{Y}$.
Disjoint sets	Said of two sets \mathcal{X}, \mathcal{Y} if $\mathcal{X} \cap \mathcal{Y} = \emptyset$.
$\mathcal{X} \setminus \mathcal{Y}$	Relative complement. Consists of all elements of \mathcal{X} that are not also elements of \mathcal{Y} . Has the properties $(\mathcal{X} \setminus \mathcal{Y}) \cap \mathcal{Y} = \emptyset$, $(\mathcal{X} \setminus \mathcal{Y}) \cup \mathcal{Y} = \mathcal{X}$.
Complement	The complement of \mathcal{X} consists of all elements not contained in \mathcal{X} ; equal to $\mathcal{U} \setminus \mathcal{X}$ where \mathcal{U} is the implicit “universal” set, e.g. \mathbb{R}^D for sets of D -dimensional vectors.
\bigcup, \bigcap	The union/intersection of many sets. Used similarly to \sum, \prod .
$\{x \mid A, B, \dots\}$	Set-builder notation. Denotes the set of all objects x for which all conditions A, B, \dots are true. $\{x \mid x \in \mathcal{X}, A, \dots\}$ is more compactly written $\{x \in \mathcal{X} \mid A, \dots\}$.
$\{a, b, c\}, \{a_i\}_{i \in I}, \{a_n\}_{n=0}^N$	Various shorthands used to define sets. The latter two are equivalent to $\{a_i \mid i \in I\}$ and $\{a_n \mid 0 \leq n \leq N\}$, respectively.
Open set	A subset $\mathcal{X} \subseteq \mathbb{R}^D$ is open if for each $\chi \in \mathcal{X}$, there is some $\varepsilon > 0$ such that $\xi \in \mathcal{X}$ for all $ \xi - \chi < \varepsilon$. This is for subsets of \mathbb{R}^D ; definitions of openness exist for more general sets, but we do not use them.
Closed set	A set whose complement is open.
Finite set	A set that contains a finite (or zero) number of elements. Many intuitive properties, such as the existence of a smallest subset with a given property, are only guaranteed for finite sets.

CONCEPT	DESCRIPTION
Bounded set	A set $\mathcal{X} \subseteq \mathbb{R}^D$ is bounded if there is some M such that $ \chi < M$ for all $\chi \in \mathcal{X}$.
Compact set	A set that is closed and bounded. This is for subsets of \mathbb{R}^D ; we do not use the more general versions. Many properties of finite sets carry over to compact sets, such as having a (not necessarily unique) element that is the minimum or maximum of some property.
$\text{cl}(\mathcal{X})$	Closure. The smallest closed set that has \mathcal{X} as a subset. Equal to \mathcal{X} itself if it is closed.
$\text{int}(\mathcal{X})$	Interior. The largest open set that is a subset of \mathcal{X} . Equal to \mathcal{X} itself if it is open.
Boundary	The boundary of \mathcal{X} is $\text{cl}(\mathcal{X}) \setminus \text{int}(\mathcal{X})$. Note that a non-closed set does not necessarily contain its boundary, and that <i>boundary</i> and <i>boundedness</i> are unrelated concepts.
(a, b) , $[a, b)$, $[a, b]$	Open, half-open and closed intervals. Denotes the range between a and b . A square bracket indicates that the endpoint is included in the interval, a parenthesis that it is not. Note the imperfect agreement with the concept of open and closed sets: (a, b) is open and $[a, b]$ is closed for finite a, b , but $[a, b)$ is neither open nor closed. Also, intervals such as $[a, \infty)$ and $(-\infty, \infty)$ are closed.
Linear/affine/convex combination	A sum of vectors of the form $\sum_i a_i \mathbf{v}_i$. It is a linear combination for all a_i . It is an affine combination if $\sum_i a_i = 1$. It is a convex combination if $\sum_i a_i = 1$ and $a_i \geq 0$ for all i .
Convex set	A set that contains all convex combinations of its elements.
$\text{Hull}(\mathcal{X})$	Convex hull. The set of all convex combinations of the elements of \mathcal{X} . Equivalently, the smallest convex set that contains \mathcal{X} .
Linear/affine span	Of a set of vectors: the set of all linear/affine combinations thereof. The linear span is often just called “span”.
Linear/affine dimension	The linear/affine dimension of a set $\mathcal{X} \subseteq \mathbb{R}^D$ is the smallest number of vectors whose linear/affine span contains \mathcal{X} . The dimension is $\leq D$. When it is clear from context, just “dimension” is often used.
Linear/affine subspace	A subset of \mathbb{R}^D that contains all linear/affine combinations of its elements. Equivalently, the linear/affine span of some set of vectors.
Hyperplane	An affine subspace of affine dimension $D - 1$. Generalises the notion of a plane in 3D space and a line in 2D space. Given a nonzero vector \mathbf{v} , the set $\{\mathbf{x} \mid \mathbf{v} \cdot \mathbf{x} = u\}$ is a distinct hyperplane for each u . The intersection of K hyperplanes in \mathbb{R}^D with linearly independent \mathbf{v} vectors is an affine subspace of dimension $D - K$.
$\mathcal{A} + \mathcal{B}$	Minkowski sum. The set $\{\mathbf{a} + \mathbf{b} \mid \mathbf{a} \in \mathcal{A}, \mathbf{b} \in \mathcal{B}\}$ (assuming the elements of \mathcal{A}, \mathcal{B} support addition).
$x\mathcal{A}$	The set $\{x\mathbf{a} \mid \mathbf{a} \in \mathcal{A}\}$ (assuming the elements of \mathcal{A} support scalar multiplication).

C The loop integral functions

This appendix contains details on the functions appearing in the NLO and NNLO amplitudes, which originate in loop integrals. Appendices C.1 and C.2 contain expansions important to their numerical evaluation, and appendix C.3 contains the derivation of their analytic integrals over the Mandelstam variables.

Using the conventions of [34],⁶⁶ the function \bar{J} can be defined as

$$\bar{J} = \pi_{16}(\beta^2 h + 2), \quad h = \frac{1}{\beta} \ln \frac{\beta - 1}{\beta + 1}, \quad \beta = \sqrt{1 - \frac{4}{s}}, \quad (\text{C.1})$$

where $\pi_{16} \equiv 1/16\pi^2$. To reduce clutter, we define $\hat{J} \equiv \bar{J}/\pi_{16}$. Similarly, with $\hat{k}_i \equiv k_i/\pi_{16}^2$ for brevity, the additional functions at NNLO are defined as

$$\begin{aligned} \hat{k}_1 &= \beta^2 h^2, & \hat{k}_3 &= \frac{\beta^2 h^3}{s} + \frac{\pi^2 h}{s} - \frac{\pi^2}{2}, \\ \hat{k}_2 &= \beta^4 h^2 - 4, & \hat{k}_4 &= \frac{1}{s\beta^2} \left[\frac{\hat{k}_1}{2} + \frac{\hat{k}_3}{3} + \hat{J} + \frac{(\pi^2 - 6)s}{12} \right]. \end{aligned} \quad (\text{C.2})$$

\bar{J} and k_i are real below threshold, and are finite as $s \rightarrow 4$ from below. However, β is not real when $0 < s < 4$, which poses a problem for numerical evaluation. This can be remedied by defining $\tilde{\beta} = -i\beta$ and rewriting h as

$$h = \frac{2 \tan^{-1}(1/\tilde{\beta})}{\tilde{\beta}}, \quad (\text{C.3})$$

which can be evaluated for $0 < s < 4$ using only real numbers.

The functions have further numerical problems. β diverges at $s = 0$, which leaves \bar{J} and k_i with removable singularities there. These are rendered harmless with a series expansion, as shown below. As $s \rightarrow 4$ from below, h diverges while \bar{J} and k_i stay finite. Reliable evaluation of this limit also requires series expansion. The derivatives of \bar{J} and k_i diverge in this limit (starting at the second derivative for k_2 and the first derivative for the others), which also necessitates series expansion for reliable handling.

C.1 Expanding around $s = 0$

Near $s = 0$, we make an expansion in $s = 4\varepsilon^2$, where ε may be complex. To make the NNLO unitarity corrections numerically well-behaved in all cases used by us, an $\mathcal{O}(\varepsilon^8) = \mathcal{O}(s^4)$ expansion is needed. For β and $1/\beta$, it is

$$\beta = \frac{\sqrt{\varepsilon^2 - 1}}{\varepsilon} = \frac{i}{\varepsilon} \left(1 - \frac{\varepsilon^2}{2} - \frac{\varepsilon^4}{8} - \frac{\varepsilon^6}{16} - \frac{5\varepsilon^8}{128} + \mathcal{O}(\varepsilon^{10}) \right), \quad (\text{C.4})$$

$$\frac{1}{\beta} = \frac{\varepsilon}{\sqrt{\varepsilon^2 - 1}} = -i\varepsilon \left(1 + \frac{\varepsilon^2}{2} + \frac{3\varepsilon^4}{8} + \frac{5\varepsilon^6}{16} + \frac{35\varepsilon^8}{128} + \mathcal{O}(\varepsilon^{10}) \right). \quad (\text{C.5})$$

By using the expansion

$$h = -2 \sum_{n=1}^{\infty} \frac{1}{(2n-1)\beta^{2n}}, \quad (\text{C.6})$$

⁶⁶We use β rather than σ for consistency with eq. (5.4).

it follows that

$$h = 2\varepsilon^2 + \frac{4}{3}\varepsilon^4 + \frac{16}{15}\varepsilon^6 + \frac{32}{35}\varepsilon^8, \quad (\text{C.7a})$$

$$\hat{J} = \frac{2}{3}\varepsilon^2 + \frac{4}{15}\varepsilon^4 + \frac{16}{105}\varepsilon^6 + \frac{32}{35}\varepsilon^8 \quad (\text{C.7b})$$

$$\hat{k}_1 = -4\varepsilon^2 - \frac{4}{3}\varepsilon^4 - \frac{32}{45}\varepsilon^6 - \frac{16}{35}\varepsilon^8 \quad (\text{C.7c})$$

$$\hat{k}_2 = -\frac{8}{3}\varepsilon^2 - \frac{28}{45}\varepsilon^4 - \frac{64}{63}\varepsilon^6 - \frac{592}{175}\varepsilon^8 \quad (\text{C.7d})$$

$$\hat{k}_3 = \left(\frac{\pi^2}{3} - 2\right)\varepsilon^2 + \left(\frac{4\pi^2}{15} - 2\right)\varepsilon^4 + \left(\frac{8\pi^2}{45} - \frac{28}{15}\right)\varepsilon^6 + \frac{328}{189}\varepsilon^8 \quad (\text{C.7e})$$

$$\hat{k}_4 = \left(1 - \frac{\pi^2}{9}\right)\varepsilon^2 + \left(\frac{19}{15} - \frac{\pi^2}{15}\right)\varepsilon^4 + \left(\frac{464}{315} - \frac{16\pi^2}{105}\right)\varepsilon^6 + \left(\frac{820}{567} - \frac{16\pi^2}{105}\right)\varepsilon^8, \quad (\text{C.7f})$$

where we have omitted “ $+\mathcal{O}(\varepsilon^{10})$ ” for brevity. These functions are real, so only even powers of ε appear. Note also how all numerical coefficients stay roughly order 1.

C.2 Expanding around $s = 4$

As $s \rightarrow 4$ from below, we expand $s = 4(1 - \delta^2)$ with $\delta > 0$, and again need an eighth-order expansion at NNLO. The expansion of β and $1/\beta$ is

$$\beta = \frac{\delta}{\sqrt{\delta^2 - 1}} = -i\delta \left(1 + \frac{1}{2}\delta^2 + \frac{3}{8}\delta^4 + \frac{5}{16}\delta^6 + \frac{35}{128}\delta^8 + \mathcal{O}(\delta^{10})\right), \quad (\text{C.8})$$

$$\frac{1}{\beta} = \frac{\sqrt{\delta^2 - 1}}{\delta} = \frac{i}{\delta} \left(1 - \frac{1}{2}\delta^2 - \frac{1}{8}\delta^4 - \frac{1}{16}\delta^6 - \frac{5}{128}\delta^8 + \mathcal{O}(\delta^{10})\right). \quad (\text{C.9})$$

By using the expansion

$$h = -\frac{i\pi}{\beta} + 2i \sum_{n=0}^{\infty} \frac{\beta^{2n}}{2n+1}, \quad (\text{C.10})$$

it follows that (suppressing “ $+\mathcal{O}(\delta^{10})$ ”)

$$h = \frac{\pi}{\delta} - 2 \left[1 + \frac{1}{3}\delta^2 + \frac{2}{15}\delta^4 + \frac{8}{105}\delta^6 + \frac{16}{315}\delta^8 \right] - \frac{\pi\delta}{2} \left[1 + \frac{1}{4}\delta^2 + \frac{1}{8}\delta^4 + \frac{5}{64}\delta^6 \right], \quad (\text{C.11a})$$

$$\hat{J} = 2 \left[1 + \delta^2 + \frac{2}{3}\delta^4 + \frac{8}{15}\delta^6 + \frac{16}{35}\delta^8 \right] - \pi\delta \left[1 + \frac{1}{2}\delta^2 + \frac{3}{8}\delta^4 + \frac{5}{16}\delta^6 \right], \quad (\text{C.11b})$$

$$\hat{k}_1 = -4 \left[\frac{\pi^2}{4} + \delta^2 + \frac{1}{3}\delta^4 + \frac{8}{45}\delta^6 + \frac{4}{35}\delta^8 \right] + \pi\delta \left[4 + \frac{2}{3}\delta^2 + \frac{3}{10}\delta^4 + \frac{5}{28}\delta^6 \right], \quad (\text{C.11c})$$

$$\hat{k}_2 = -4 \left[1 - \delta^4 - \frac{4}{3}\delta^6 - \frac{68}{45}\delta^8 \right] - \pi\delta \left[4\delta^2 + \frac{14}{3}\delta^4 + \frac{149}{30}\delta^6 \right] + \pi^2 \left[\delta^2 + \delta^4 + \delta^6 + \delta^8 \right], \quad (\text{C.11d})$$

$$\hat{k}_3 = 2 \left[\delta^2 + \delta^4 + \frac{14}{15}\delta^6 + \frac{164}{189}\delta^8 \right] - \pi\delta \left[3 + \frac{5}{2}\delta^2 + \frac{259}{120}\delta^4 + \frac{3229}{1680}\delta^6 \right] + \pi^2 \left[\frac{1}{2} + \frac{2}{3}\delta^2 + \frac{8}{15}\delta^4 + \frac{16}{35}\delta^6 + \frac{128}{315}\delta^8 \right], \quad (\text{C.11e})$$

$$\hat{k}_4 = -\frac{1}{3} \left[2 + \delta^2 + \delta^4 + \frac{128}{105}\delta^6 + \frac{844}{945}\delta^8 \right] + \frac{\pi\delta}{4} \left[1 + \frac{17}{18}\delta^2 + \frac{311}{360}\delta^4 + \frac{2227}{2800}\delta^6 \right] + \frac{\pi^2}{3} \left[\frac{1}{12} - \frac{2}{15}\delta^2 - \frac{4}{35}\delta^4 - \frac{32}{315}\delta^6 + \frac{187}{1260}\delta^8 \right] - \frac{\pi^3\delta}{4} \left[\frac{121}{1536}\delta^6 \right]. \quad (\text{C.11f})$$

Since δ is real, odd powers of δ are permitted. Note again how the coefficients stay roughly order 1. Due to $d\delta/ds = -1/8\delta$, the derivative of the linear terms diverges as $\delta \rightarrow 0$. Since

\hat{k}_2 lacks a linear term, its derivative remains finite. Since the derivative changes the powers in steps of two, only odd negative powers appear in derivatives of any order.

C.3 Integrals above threshold

Here, we seek to analytically determine the function $D_k^J(\lambda, v, t)$ defined in eq. (5.19). In the relevant z range, $\beta(z)$ remains real while $h(z)$ obtains an imaginary part.⁶⁷ The most convenient form of h is

$$h(z) = \frac{H(\beta) + i\pi}{\beta}, \quad H(\beta) \equiv \ln \frac{1-\beta}{1+\beta} = -\frac{1}{2} \tanh^{-1}(\beta), \quad (\text{C.12})$$

where we take the branch with positive imaginary part. We make the following easily verifiable observations about the real function $H(\beta)$:

$$H' = \frac{-2}{1-\beta^2} = -\frac{1}{2}z, \quad (\text{C.13a}) \quad \beta^2 = 1 + \frac{2}{H'}, \quad (\text{C.13c})$$

$$H'' = -\beta H'^2, \quad (\text{C.13b}) \quad dz = -2H'' d\beta = 2\beta H'^2 d\beta. \quad (\text{C.13d})$$

$D_k^J(\lambda, v, t)$ will be a linear combination of integrals of the form $\int_4^\lambda \frac{z^n dz}{(z-v)^{k+1}} \text{Im } X(z)$, where X is one of \bar{J} and k_i . We will ignore the denominator for now, and show later how to reduce all integrals to the form $\int z^n \text{Im } X(z) dz$. This will involve a wide range of values for n , so it is easiest to treat general n and then read off the special cases. Reading from eqs. (C.1) and (C.2), we find that above threshold,

$$\begin{aligned} \text{Im } \hat{J} &= \pi\beta, & \text{Im } \hat{k}_1 &= 2\pi H, & \text{Im } \hat{k}_2 &= 2\pi\beta^2 H, \\ \text{Im } \hat{k}_3 &= \frac{3\pi H^2}{z\beta}, & \text{Im } \hat{k}_4 &= \frac{1}{z\beta^2} \left[H + \frac{H^2}{z\beta} + \beta \right]. \end{aligned} \quad (\text{C.14})$$

Using eq. (C.13), we see that everything can be expressed in terms of the functions

$$Z_p^{m,n}(\beta) \equiv \int \zeta_p^{m,n}(\beta) d\beta, \quad \zeta_p^{m,n}(\beta) \equiv \frac{H^m H^n}{\beta^p}. \quad (\text{C.15})$$

Thanks to eq. (C.13b), the family of functions $\zeta_p^{m,n}$ is closed under derivatives:

$$\frac{d}{d\beta} \zeta_p^{m,n} = m \zeta_p^{m-1,n+1} - n \zeta_{p-1}^{m,n+1} - p \zeta_{p+1}^{m,n}. \quad (\text{C.16})$$

It is therefore our hope that the highly nontrivial integral $Z_p^{m,n}$ can mostly be expressed as a sum of $\zeta_{p'}^{m',n'}$, plus some special cases (for instance, $\zeta_p^{m,0}$ is not the derivative of another $\zeta_{p'}^{m',n'}$). We will therefore attempt to find recurrence relations on m, n, p that allow Z to be reduced to ζ 's and a few special cases.

⁶⁷We will drop the dependence on z from now on; everything in this section implicitly depends on it unless otherwise specified.

C.3.1 The integral $Z_p^{m,n}$ for $p = 0$

First, we note that n can be assumed non-negative, since

$$Z_p^{m,n} = \frac{1}{2}[Z_{p-2}^{m,n+1} - Z_p^{m,n+1}], \quad (\text{C.17})$$

according to eq. (C.13c).⁶⁸ Then, before treating general p , we consider $Z^{m,n} \equiv Z_0^{m,n}$ for $n, m \geq 0$. The integral of eq. (C.16) gives

$$\begin{aligned} \zeta^{m-1,n-1} &= (m-1)Z^{m-2,n} - (n-1)Z_{-1}^{m-1,n}, \\ \zeta_{-1}^{m,n-1} &= mZ_{-1}^{m-1,n} - (n-1)Z_{-2}^{m,n} + Z^{m,n-1}, \end{aligned} \quad (\text{C.18})$$

where the second line allows for the removal of $Z_{-1}^{m-1,n}$ in the first. Eq. (C.17) can be invoked to turn $Z_{-2}^{m,n}$ into $2Z^{m,n-1} + Z^{m,n}$, and after extracting $Z^{m,n}$, we get a recurrence relation where both m and n decrease:

$$Z^{m,n} = \frac{3-2n}{n-1}Z^{m,n-1} - \frac{\zeta_{-1}^{m,n-1}}{n-1} + \frac{m}{(n-1)^2}[(m-1)Z^{m-2,n} - \zeta^{m-1,n-1}]. \quad (\text{C.19})$$

This is valid for all $m \geq 0, n > 1$. The $n = 1$ case is covered by the trivial identity⁶⁹

$$Z^{m,1} = \frac{\zeta^{m+1,0}}{m+1}. \quad (\text{C.20})$$

$Z^{m,0}$ requires some more thought and will be treated later (see eq. (C.26)).

C.3.2 Reduction of $Z_p^{m,n}$ to $Z_0^{m,n}$ and $Z_1^{m,n}$

We can integrate and restructure eq. (C.16) into the recurrence relation

$$Z_p^{m,n} = \frac{m}{p-1}Z_{p-1}^{m-1,n+1} - \frac{n}{p-1}Z_{p-2}^{m,n+1} - \frac{\zeta_{p-1}^{m,n}}{p-1}, \quad p \neq 1. \quad (\text{C.21})$$

This allows any $p > 0$ to be reduced to the cases $p = 0$ and $p = 1$. When p is negative, (C.17) furnishes the simpler relation

$$Z_p^{m,n} = Z_{p+2}^{m,n} + 2Z_{p+2}^{m,n-1}. \quad (\text{C.22})$$

This reduces any $p < 0$ to $p = 0$ and $p = 1$, as long as n stays positive. In the $n = 0$ case, we again turn to eq. (C.16):

$$Z_p^{m,0} = (p+1)Z_{p+2}^{m,-1} - mZ_{p+1}^{m-1,0} - \zeta_{p+1}^{m,-1}. \quad (\text{C.23})$$

Using eq. (C.17) to treat the $n = -1$ terms, we arrive at

$$Z_p^{m,0} = \frac{p+1}{p-1}Z_{p+2}^{m,0} + \frac{2m}{p-1}Z_{p+1}^{m-1,0} + \frac{\zeta_{p+1}^{m,0} - \zeta_{p-1}^{m,0}}{p-1}, \quad (\text{C.24})$$

which works for $n = 0, p \neq 1$.⁷⁰

⁶⁸We will not need to treat negative m when integrating the loop integral functions. As we will see below, $Z_p^{m,n}$ is tractable for all integer n and p , but m has to stay non-negative.

⁶⁹We suppress the constant of integration here and everywhere else.

⁷⁰It could also be adapted to other $n \neq \frac{1-p}{2}$, but in those cases eq. (C.22) is simpler.

In the case $p = 1$, $n > 0$ we integrate by parts:

$$Z_1^{m,n} = \frac{Z^{m,n}}{\beta} + \int \frac{Z^{m,n}}{\beta^2} d\beta. \tag{C.25}$$

Through eqs. (C.19) and (C.21), the integral will reduce to $p = 0$ terms that are easy to handle, plus various $Z_1^{m',n'}$ where $m' \leq m, n' \leq n$. At least one of the inequalities is strict, so we will eventually arrive at $Z_1^{0,n''}$ and $Z_1^{m'',0}$. The former can be run through the recurrence again, but the latter requires separate consideration.

C.3.3 Special cases for $n = 0$

The relations above are capable of reducing almost all Z 's to ζ 's, but they are unable to get rid of $Z_p^{m,0}$ for $p = 0, 1, 2$. If we could define a function $\Phi_\ell^{(p)}$ with derivative $-H' \Phi_{\ell-1}^{(p)}$ and base case $\frac{d}{d\beta} \Phi_0^{(p)} = \beta^{-p}$, then an elegant solution to this would be

$$Z_p^{m,0} = \sum_{\ell=0}^m \frac{m!}{(m-\ell)!} \Phi_\ell^{(p)} H^{m-\ell}, \quad m, p \geq 0. \tag{C.26}$$

Such a function can be constructed using polylogarithms, since

$$\frac{d}{d\beta} \text{Li}_\ell[f(\beta)] = \text{Li}_{\ell-1}[f(\beta)] \frac{f'(\beta)}{f(\beta)}. \tag{C.27}$$

The correct recurrence relation is obtained by solving a simple differential relation for $f(\beta)$, which gives

$$f(\beta) = K \frac{1+\beta}{1-\beta}, \quad \text{Li}_0[f(\beta)] = \frac{K(1+\beta)}{(1-K) - \beta(1+K)}. \tag{C.28}$$

We immediately see that $K = -1$ is suitable for $p = 0$, and $K = 1$ for $p = 2$.^{71,72}

$$\Phi_\ell^{(0)} = -2 \text{Li}_\ell\left(\frac{\beta+1}{\beta-1}\right), \quad \Phi_\ell^{(2)} = 2 \text{Li}_\ell\left(\frac{1+\beta}{1-\beta}\right). \tag{C.29}$$

We cannot handle $p = 1$ directly this way, but using the close connection between $\text{Li}_1(x)$ and the natural logarithm, we find that

$$\Phi_\ell^{(1)} = -\frac{\Phi_{\ell+1}^{(0)} + \Phi_{\ell+1}^{(2)}}{2} \tag{C.30}$$

gives the correct result.⁷³

⁷¹This process can in principle be continued to treat all $p \geq 2$, but it is more practical to rely on the recurrence relations to get rid of larger p .

⁷²Note that $\Phi_\ell^{(2)}$ is complex-valued for $\ell > 0$, since $\frac{1+\beta}{1-\beta} > 1$ for z above threshold.

⁷³Although $\frac{d}{d\beta} \Phi_0^{(1)} = 1/\beta$, we have $\Phi_0^{(1)} = \ln(-\beta)$, not $\ln \beta$, so one must be careful to place β on the correct side of the branch cut.

C.3.4 The treatment of $(z - v)$

The factors that arise due to the dispersion relations are highly problematic when $v \neq 0$ (which corresponds to $s \neq 0$ or $u \neq 0$ in eq. (5.2)). We only need to consider $k = 0$, since

$$\int \frac{\zeta_p^{m,n}}{(z-v)^{k+1}} dz = \frac{1}{k!} \frac{d^k}{dv^k} \int \frac{\zeta_p^{m,n}}{z-v} dz; \quad (\text{C.31})$$

the derivative can be taken after evaluating the integral. Trying to get rid of the last power of $(z - v)$ is futile, so we instead change variables to β and find

$$\int \frac{\zeta_p^{m,n}}{z-v} dz = 2 \int \frac{4}{zv} \frac{\zeta_{p-1}^{m,n+2}}{\frac{4}{v} - \frac{4}{z}} d\beta = \frac{-4}{v} \int \frac{\zeta_{p-1}^{m,n+1}}{\beta^2 - \beta_v^2} d\beta, \quad (\text{C.32})$$

where $\beta_v \equiv \beta(v) = \sqrt{1 - 4/v}$.⁷⁴ This relation is singular when $v = 0$, but then $\frac{z^n}{(z-v)^{k+1}} = z^{n-(k+1)}$ so it is not needed.⁷⁵

It turns out that $\int H^m / (\beta^2 - \beta_v^2) d\beta$ is tractable, so the strategy is to separate that from the rest of the integral by repeatedly applying partial fractions:

$$\frac{\zeta_p^{m,n}}{\beta^2 - \beta_v^2} = \frac{v}{4} \left[\zeta_p^{m,n} - \frac{2\zeta_p^{m,n-1}}{\beta^2 - \beta_v^2} \right]. \quad (\text{C.33})$$

If n is negative, we instead use eq. (C.17). Once n has been reduced to zero in the $(\beta^2 - \beta_v^2)$ -containing term this way, we remove p similarly:

$$\frac{\zeta_p^{m,0}}{\beta^2 - \beta_v^2} = \zeta_{p+2}^{m,0} \left[\frac{\beta_v^2}{\beta^2 - \beta_v^2} + 1 \right] = \frac{1}{\beta_v^2} \left[\frac{\zeta_{p-2}^{m,0}}{\beta^2 - \beta_v^2} - \zeta_p^{m,0} \right]. \quad (\text{C.34})$$

This reduces p to 0 or -1 ; the latter can be handled with the expansion

$$\frac{\beta}{\beta^2 - \beta_v^2} = \frac{1 - \beta_v}{(1 - \beta)(\beta - \beta_v)} - \frac{\beta_v}{\beta^2 - \beta_0^2} + \frac{1}{2}(1 + \beta)H'. \quad (\text{C.35})$$

The second term on the right-hand side corresponds to $p = 0$, and the third does not involve β_v at all. In the spirit of eq. (C.26), the first is equal to $\frac{d}{d\beta} \Psi_0^-$ if we identify

$$\Psi_\ell^\pm = -\text{Li}_{\ell+1} \left[\frac{(1 + \beta)(1 \pm \beta_v)}{(1 - \beta)(1 \mp \beta_v)} \right] \quad (\text{C.36})$$

guided by eq. (C.28). Since $\frac{d}{d\beta} \Psi_\ell^\pm = -H' \Psi_{\ell-1}^\pm$,

$$\int \frac{\zeta_{-1}^{m,0}}{\beta^2 - \beta_v^2} d\beta = \frac{Z^{m,1} + Z_{-1}^{m,1}}{2} - \int \frac{\beta_v \zeta^{m,0}}{\beta^2 - \beta_v^2} d\beta + \sum_{\ell=0}^m \frac{m!}{(m-\ell)!} \Psi_\ell^- H^{m-\ell}. \quad (\text{C.37})$$

⁷⁴The same relation is useful if one wishes to explicitly evaluate a partial-wave expansion like eq. (5.3), since the Legendre polynomials consist of powers of $1/(z-4)$. It is further simplified by $\beta(4) = 0$.

⁷⁵Note that β_v is imaginary for $v \in (0, 4)$, which forces a more involved detour through the complex plane than above. For instance, the functions Ψ_ℓ^\pm below will be complex-valued, although in the end the integral will of course remain real.

This leaves $p = 0$, for which a similar solution is

$$\int \frac{\beta_v \zeta^{m,0}}{\beta^2 - \beta_v^2} d\beta = \sum_{\ell=0}^m \frac{m!}{(m-\ell)!} \frac{\Psi_\ell^- - \Psi_\ell^+}{2} H^{m-\ell}. \quad (\text{C.38})$$

The very last piece in the puzzle of evaluating $D_k^J(\lambda, v, t)$ is the conceptually simple derivative in eq. (C.31):

$$\frac{d\beta_v}{dv} = \frac{2}{\beta_v v^2}, \quad \frac{d}{dv} \Psi_\ell^\pm = \pm \frac{v}{\beta_v} \Psi_{\ell-1}^\pm, \quad \frac{d\beta}{dv} = \frac{d}{dv} Z_p^{m,n} = 0. \quad (\text{C.39})$$

C.3.5 The completed integral

The above recurrence relations allow for the integration of all terms that appear in $D_k^J(\lambda, v, t)$. $Z_p^{m,n}$ diverges at $z \rightarrow 4$ for some values of p, m, n (in particular those that contain $\Phi_1^{(2)}$ or negative powers of β), so obtaining a finite lower limit of the overall integral requires careful (albeit straightforward) extraction and cancellation of those divergences. The resulting expressions are very lengthy in most cases, so we do not reproduce them here.⁷⁶

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] J. Gasser and H. Leutwyler, *Chiral Perturbation Theory to One Loop*, *Annals Phys.* **158** (1984) 142 [[INSPIRE](#)].
- [2] J. Gasser and H. Leutwyler, *Chiral Perturbation Theory: Expansions in the Mass of the Strange Quark*, *Nucl. Phys. B* **250** (1985) 465 [[INSPIRE](#)].
- [3] S. Weinberg, *Phenomenological Lagrangians*, *Physica A* **96** (1979) 327 [[INSPIRE](#)].
- [4] S. Scherer and M.R. Schindler, *A Primer for Chiral Perturbation Theory*, Springer-Verlag, 1 edition (2012) [[DOI](#)] [[INSPIRE](#)].
- [5] A. Pich, *Effective Field Theory with Nambu-Goldstone Modes*, in *Les Houches summer school: EFT in Particle Physics and Cosmology*, (2018) [[arXiv:1804.05664](#)] [[INSPIRE](#)].
- [6] J. Bijnens and G. Ecker, *Mesonic low-energy constants*, *Ann. Rev. Nucl. Part. Sci.* **64** (2014) 149 [[arXiv:1405.6488](#)] [[INSPIRE](#)].
- [7] A. Martin, *Scattering Theory: Unitarity, Analyticity and Crossing*, Springer-Verlag, 1 edition (1969) [[DOI](#)] [[INSPIRE](#)].
- [8] T.N. Pham and T.N. Truong, *Evaluation of the Derivative Quartic Terms of the Meson Chiral Lagrangian From Forward Dispersion Relation*, *Phys. Rev. D* **31** (1985) 3027 [[INSPIRE](#)].
- [9] B. Ananthanarayan, D. Toublan and G. Wanders, *Consistency of the chiral pion pion scattering amplitudes with axiomatic constraints*, *Phys. Rev. D* **51** (1995) 1093 [[hep-ph/9410302](#)] [[INSPIRE](#)].

⁷⁶The FORM implementation of the relations, and the expressions produced by it, are available from Mattias Sjö upon request.

- [10] M.R. Pennington and J. Portoles, *The Chiral Lagrangian parameters, l_1 , l_2 , are determined by the rho resonance*, *Phys. Lett. B* **344** (1995) 399 [[hep-ph/9409426](#)] [[INSPIRE](#)].
- [11] P. Dita, *Positivity constraints on chiral perturbation theory pion pion scattering amplitudes*, *Phys. Rev. D* **59** (1999) 094007 [[hep-ph/9809568](#)] [[INSPIRE](#)].
- [12] J. Distler, B. Grinstein, R.A. Porto and I.Z. Rothstein, *Falsifying Models of New Physics via WW Scattering*, *Phys. Rev. Lett.* **98** (2007) 041601 [[hep-ph/0604255](#)] [[INSPIRE](#)].
- [13] A.V. Manohar and V. Mateu, *Dispersion Relation Bounds for $\pi\pi$ Scattering*, *Phys. Rev. D* **77** (2008) 094019 [[arXiv:0801.3222](#)] [[INSPIRE](#)].
- [14] V. Mateu, *Universal Bounds for SU(3) Low Energy Constants*, *Phys. Rev. D* **77** (2008) 094020 [[arXiv:0801.3627](#)] [[INSPIRE](#)].
- [15] B. Bellazzini, J. Elias Miró, R. Rattazzi, M. Riembau and F. Riva, *Positive moments for scattering amplitudes*, *Phys. Rev. D* **104** (2021) 036006 [[arXiv:2011.00037](#)] [[INSPIRE](#)].
- [16] S. Caron-Huot and V. Van Duong, *Extremal Effective Field Theories*, *JHEP* **05** (2021) 280 [[arXiv:2011.02957](#)] [[INSPIRE](#)].
- [17] A. Sinha and A. Zahed, *Crossing Symmetric Dispersion Relations in Quantum Field Theories*, *Phys. Rev. Lett.* **126** (2021) 181601 [[arXiv:2012.04877](#)] [[INSPIRE](#)].
- [18] A. Zahed, *Positivity and geometric function theory constraints on pion scattering*, *JHEP* **12** (2021) 036 [[arXiv:2108.10355](#)] [[INSPIRE](#)].
- [19] Y.-J. Wang, F.-K. Guo, C. Zhang and S.-Y. Zhou, *Generalized positivity bounds on chiral perturbation theory*, *JHEP* **07** (2020) 214 [[arXiv:2004.03992](#)] [[INSPIRE](#)].
- [20] A.J. Tolley, Z.-Y. Wang and S.-Y. Zhou, *New positivity bounds from full crossing symmetry*, *JHEP* **05** (2021) 255 [[arXiv:2011.02400](#)] [[INSPIRE](#)].
- [21] B. Alvarez, *Updating bounds on the low-energy constants of Chiral Perturbation Theory from exact bounds on amplitudes*, MSc Thesis LU TP 16-22, Lund University, June 2016.
- [22] O. Catà and V. Mateu, *Chiral perturbation theory with tensor sources*, *JHEP* **09** (2007) 078 [[arXiv:0705.2948](#)] [[INSPIRE](#)].
- [23] J.F. Donoghue and H. Leutwyler, *Energy and momentum in chiral theories*, *Z. Phys. C* **52** (1991) 343 [[INSPIRE](#)].
- [24] S.R. Coleman, J. Wess and B. Zumino, *Structure of phenomenological Lagrangians. 1*, *Phys. Rev.* **177** (1969) 2239 [[INSPIRE](#)].
- [25] C.G. Callan Jr., S.R. Coleman, J. Wess and B. Zumino, *Structure of phenomenological Lagrangians. 2*, *Phys. Rev.* **177** (1969) 2247 [[INSPIRE](#)].
- [26] G. Ecker, J. Gasser, A. Pich and E. de Rafael, *The Role of Resonances in Chiral Perturbation Theory*, *Nucl. Phys. B* **321** (1989) 311 [[INSPIRE](#)].
- [27] J. Bijnens, G. Colangelo and G. Ecker, *The Mesonic chiral Lagrangian of order p^6* , *JHEP* **02** (1999) 020 [[hep-ph/9902437](#)] [[INSPIRE](#)].
- [28] J. Bijnens, N. Hermansson-Truedsson and S. Wang, *The order p^8 mesonic chiral Lagrangian*, *JHEP* **01** (2019) 102 [[arXiv:1810.06834](#)] [[INSPIRE](#)].
- [29] J. Bijnens, G. Colangelo and G. Ecker, *Renormalization of chiral perturbation theory to order p^6* , *Annals Phys.* **280** (2000) 100 [[hep-ph/9907333](#)] [[INSPIRE](#)].

- [30] J. Bijnens, G. Colangelo, G. Ecker, J. Gasser and M.E. Sainio, *Pion-pion scattering at low energy*, *Nucl. Phys. B* **508** (1997) 263 [Erratum *ibid.* **517** (1998) 639] [[hep-ph/9707291](#)] [[INSPIRE](#)].
- [31] J. Bijnens, G. Colangelo, G. Ecker, J. Gasser and M.E. Sainio, *Elastic $\pi\pi$ scattering to two loops*, *Phys. Lett. B* **374** (1996) 210 [[hep-ph/9511397](#)] [[INSPIRE](#)].
- [32] J. Bijnens, P. Dhonte and P. Talavera, *$\pi\pi$ scattering in three flavor ChPT*, *JHEP* **01** (2004) 050 [[hep-ph/0401039](#)] [[INSPIRE](#)].
- [33] J. Bijnens, P. Dhonte and P. Talavera, *πK scattering in three flavor ChPT*, *JHEP* **05** (2004) 036 [[hep-ph/0404150](#)] [[INSPIRE](#)].
- [34] J. Bijnens and J. Lu, *Meson-meson Scattering in QCD-like Theories*, *JHEP* **03** (2011) 028 [[arXiv:1102.0172](#)] [[INSPIRE](#)].
- [35] R.S. Chivukula, M.J. Dugan and M. Golden, *Analyticity, crossing symmetry and the limits of chiral perturbation theory*, *Phys. Rev. D* **47** (1993) 2930 [[hep-ph/9206222](#)] [[INSPIRE](#)].
- [36] S. Weinberg, *Pion scattering lengths*, *Phys. Rev. Lett.* **17** (1966) 616 [[INSPIRE](#)].
- [37] J. Gasser and M.E. Sainio, *Two loop integrals in chiral perturbation theory*, *Eur. Phys. J. C* **6** (1999) 297 [[hep-ph/9803251](#)] [[INSPIRE](#)].
- [38] D.E. Neville, *Elastic Scattering of Pseudoscalar Mesons and $SU(N)$ Symmetry*, *Phys. Rev.* **132** (1963) 844 [[INSPIRE](#)].
- [39] F.P. Preparata and D.E. Muller, *Finding the intersection of n half-spaces in time $O(n \log n)$* , *Theor. Comput. Sci.* **8** (1979) 45.
- [40] M. Froissart, *Asymptotic behavior and subtractions in the Mandelstam representation*, *Phys. Rev.* **123** (1961) 1053 [[INSPIRE](#)].
- [41] A.V. Manohar, *Effective field theories*, *Lect. Notes Phys.* **479** (1997) 311 [[hep-ph/9606222](#)] [[INSPIRE](#)].
- [42] A. Gomez Nicola and J.R. Pelaez, *Meson meson scattering within one loop chiral perturbation theory and its unitarization*, *Phys. Rev. D* **65** (2002) 054009 [[hep-ph/0109056](#)] [[INSPIRE](#)].
- [43] G. Cacciapaglia, C. Pica and F. Sannino, *Fundamental Composite Dynamics: A Review*, *Phys. Rept.* **877** (2020) 1 [[arXiv:2002.04914](#)] [[INSPIRE](#)].
- [44] E. Farhi and L. Susskind, *A Technicolored $G.U.T.$* , *Phys. Rev. D* **20** (1979) 3404 [[INSPIRE](#)].
- [45] J. Bijnens and T. Husek, *Six-pion amplitude*, *Phys. Rev. D* **104** (2021) 054046 [[arXiv:2107.06291](#)] [[INSPIRE](#)].
- [46] P. Tol, *Colour schemes*, <https://personal.sron.nl/~pault/> (2018).
- [47] C.A. Brewer, *Colorbrewer, a web tool for selecting colors for maps*, <http://colorbrewer2.org> (2009).
- [48] C.B. Barber, D.P. Dobkin and H. Huhdanpaa, *The quickhull algorithm for convex hulls*, *ACM Trans. Math. Software* **22** (1996) 469.