

Systematics of aligned axions

Thomas C. Bachlechner,^a Kate Eckerle,^{a,b} Oliver Janssen^c and Matthew Kleban^c

^a*Department of Physics, Columbia University,
New York, U.S.A.*

^b*Dipartimento di Fisica, Università di Milano-Bicocca,
Milano, Italy*

^c*Center for Cosmology and Particle Physics, New York University,
New York, U.S.A.*

E-mail: t.bachlechner@columbia.edu, kate.eckerle@unimib.it,
opj202@nyu.edu, mk161@nyu.edu.

ABSTRACT: We describe a novel technique that renders theories of N axions tractable, and more generally can be used to efficiently analyze a large class of periodic potentials of arbitrary dimension. Such potentials are complex energy landscapes with a number of local minima that scales as $\sqrt{N!}$, and so for large N appear to be analytically and numerically intractable. Our method is based on uncovering a set of approximate symmetries that exist in addition to the N periods. These approximate symmetries, which are exponentially close to exact, allow us to locate the minima very efficiently and accurately and to analyze other characteristics of the potential. We apply our framework to evaluate the diameters of flat regions suitable for slow-roll inflation, which unifies, corrects and extends several forms of “axion alignment” previously observed in the literature. We find that in a broad class of random theories, the potential is smooth over diameters enhanced by $N^{3/2}$ compared to the typical scale of the potential. A Mathematica implementation of our framework is available online.

KEYWORDS: Effective Field Theories, Superstring Vacua

ARXIV EPRINT: [1709.01080](https://arxiv.org/abs/1709.01080)

Contents

1	Introduction	1
1.1	Systematic framework	2
1.2	Results for random ensembles	3
2	Exact and approximate axion symmetries	4
2.1	Auxiliary fields and a geometric picture	5
2.2	Exact periodicities and periodic domains	7
2.3	Approximate symmetries and well-aligned theories	8
2.4	Aligned coordinates and similar tiles	9
2.5	Phases	13
3	Minima and saddle points	15
3.1	Systematics of all minima	15
3.2	Neighboring minima	16
3.3	Minima to quadratic order	16
3.4	A uniform sample over all minima	19
3.5	Saddle points and maxima	20
3.6	Estimates in random ensembles	21
3.6.1	The quantity and energies of minima for $P = N + 1$	22
3.6.2	Hessian eigenvalues	23
3.6.3	Neighboring minima	25
3.6.4	Phases	25
3.7	Band structure of subleading terms	26
4	Aligned axion diameters	27
4.1	Diameters and field ranges in well-aligned theories	28
4.2	N-flation, lattice and kinetic alignment	30
4.2.1	N-flation	30
4.2.2	Lattice alignment	31
4.2.3	Kinetic alignment	32
4.3	Alignment in random ensembles	33
4.3.1	Diameter estimates	35
4.3.2	Unit metric	36
4.3.3	Wishart metric	37
4.3.4	Heavy-tailed metric	38
4.3.5	Variable couplings: dynamic alignment	38
5	Conclusions	39
A	Separating flat directions in V	41
B	Eliminating very massive axions	42

C	Extending a sublattice basis	43
D	An explicit example	43
E	Derivation of (3.12)	45
F	Derivation of (4.12)	46
G	Axion potentials and Gaussian random fields	46
	G.1 Gaussian fields	46
	G.2 Axion theories at large P	48

1 Introduction

Fields protected by exact or approximate shift symmetries play an integral role in a wide variety of physical phenomena, ranging from magnetic fluctuations of topological insulators in condensed matter physics to inflation in early universe cosmology. Axion fields are a particularly interesting example: to all orders in perturbation theory they have a continuous shift symmetry which is broken to a discrete one by non-perturbative effects. Axions were originally proposed to solve the strong CP problem of QCD [1], and many (~ 100) axions often arise in compactifications of string theory (see e.g. [2–12]). Axions can be dark matter [13–17], drive inflation [18–24] and (similar to quantized fluxes [25]) can account for the observed vacuum energy [26]. In order to analyze the interplay between these distinct phenomena, one requires a comprehensive framework for multi-axion theories. The purpose of this paper is to present such a framework, which can be employed to unify the cosmological mechanisms mentioned above.

Theories of $N \gg 1$ axion fields constitute an extremely complex “landscape” — that is, they have an exponentially large number of minima with different energies and a large diversity of regions of the potential. We will study general multi-axion theories, providing a systematic approach that renders even complex theories analytically tractable.¹ In this paper we focus on properties of the axion potential. We identify all exact and approximate shift symmetries, provide the location of local minima and characterize features of the potential through a natural partition of the axion field space. In a companion paper [27] we study the dynamics of these theories in the context of cosmology. A brief summary of our results can be found in [28]. In this paper and its companions we find that generic theories of $N \sim 100$ axions with a single energy scale close to the fundamental scale and with $\mathcal{O}(1)$ random coefficients can simultaneously account for the Big Bang (tunneling from a parent minimum), inflation (because such potentials generically have light directions with enhanced field ranges), “fuzzy” dark matter [13] with roughly the correct abundance [14, 15],

¹While we focus on axion field theories, our techniques carry over to the analysis of more general theories where the potential is a sum of terms with discrete shift symmetries.

and provide many minima with energies consistent with observation that can solve the cosmological constant problem anthropically [29–33].²

This paper is divided into three parts. In section 2 we identify the exact and approximate discrete shift symmetries of multi-axion theories by introducing a P -dimensional auxiliary field space. We discuss how in many cases the approximate shift symmetries can be used to eliminate all phases from the potential to good accuracy. Following the discussion of symmetries, the two subsequent sections can be read independently. In section 3 we apply the framework of symmetries to locate the critical points of the potential. We provide an algorithm that finds all minima in exponential time, while a representative sample of all minima can be obtained in polynomial time. More specifically we demonstrate how an exponentially large number of minima can be located analytically via a polynomial in N algorithm. We estimate the magnitude of the remaining phases, as well as the number of minima in certain ensembles of random axion theories. In section 4 we provide a general discussion of the geometry of the approximately quadratic regions of the potential surrounding minima. This discussion generalizes and corrects misleading prior results in the literature (including those of one of the authors [35]).

A Mathematica demonstration of our framework for multi-axion theories is available online [36].

1.1 Systematic framework

Consider the general two-derivative theory of N axions θ^i . The N continuous shift symmetries of the free theory are broken to discrete ones by P leading non-perturbative effects.³ The Lagrangian takes the form

$$\mathcal{L} = \frac{1}{2} \partial\theta^\top \mathbf{K} \partial\theta - \sum_{I=1}^P \Lambda_I^4 \left[1 - \cos(\mathbf{Q}\theta + \delta)^I \right] + \dots, \quad (1.1)$$

where \mathbf{K} is the metric on field space, \mathbf{Q} is the $P \times N$ integer matrix where the I th row contains the charge associated with the axions' coupling to the I th non-perturbative effect, and Λ_I is the energy scale of this effect. The dots denote subleading terms in the potential that we will generally neglect (but see section 3.7), as well as a possible additive constant (a bare cosmological constant) that will be irrelevant in this paper, as we do not consider coupling to gravity here (but see [27]). We denote matrices and vectors by bold font, with upper and lower indices identifying rows and columns, respectively. The lower case indices i, j run from 1 to N , a runs from 1 to $P - N$, and I runs from 1 to P .⁴ Throughout this work we assume that \mathbf{K} is independent of the axions θ .

Potentials of the form appearing in (1.1) are very complex when $N, P \gg 1$. However, because the cosine arguments consist of integer linear combinations of the axions, the

²See also [34] and references therein.

³In some cases the shift symmetries are entirely broken for some linear combinations of fields by couplings to classical objects, such as sources or fluxes. In this work we restrict our attention to fields that retain discrete shift symmetries. Appendix B discusses the appropriate coordinate transformations that eliminate axions that couple to classical sources from the theory.

⁴Without loss of generality we may assume that $P \geq N$ and \mathbf{Q} has maximal rank N . See appendix A for a discussion.

potential is manifestly invariant under the N discrete shifts $\theta^i \rightarrow \theta^i + 2\pi$.⁵ The existence of exact discrete symmetries is a fundamental characteristic of axion theories, and it allows us to restrict our attention to a finite region in field space — a single periodic domain defined by these symmetries. What is not so obvious is that the potential in (1.1) additionally possesses as many as $P - N$ *approximate* discrete shift symmetries, which can be used to eliminate all phases δ^I to good accuracy. We develop a framework that identifies these symmetries, and provides a natural division of the field space into domains over which none of the individual terms in the potential exceed their period. As we shall see, the identification of approximate symmetries consists of finding short lattice vectors of a P -dimensional rank $P - N$ lattice, which at fixed $P - N$ requires a number of evaluations that scales polynomially in N . The (approximate) symmetries then allow us to identify regions that are very similar. Furthermore, our framework allows us to identify a vast number of distinct minima by considering the approximate symmetry transformations away from a given minimum.

Axions are protected from perturbative corrections of the potential and therefore constitute prime candidates in the constructions of models for large field inflation and tests of quantum gravity more generally [37–46]. Large field inflation requires very flat potentials, therefore there has been much interest in the invariant distances over which axion potentials remain featureless [35, 47–51]. The potential certainly is featureless in a field space region within which none of the cosine terms traverses more than its period. The invariant size of these regions depends both on the kinetic matrix \mathbf{K} and the charge matrix \mathbf{Q} . Historically, two mechanisms have been proposed to construct theories with potentials that remain featureless over large invariant distances: lattice alignment [20], which relies upon an almost exact degeneracy between the axion charges, and kinetic alignment [47], which relies upon the delocalization of eigenvectors of the kinetic matrix. In this work we clarify the relation between these mechanisms and demonstrate that the diameter of featureless regions is bounded from above by

$$\mathcal{D} \leq 2\pi\sqrt{P} \frac{1}{\lambda_{\min}(|\mathbf{Q}\mathbf{K}^{-1/2}|)}, \tag{1.2}$$

where $\lambda_{\min}(\cdot)$ returns the smallest eigenvalue and we defined $|\mathbf{Q}| \equiv \sqrt{\mathbf{Q}^T \mathbf{Q}}$. The bound (1.2) is approximately saturated in large classes of axion theories. Note that the diameter, perhaps surprisingly, scales with $\sqrt{P} \geq \sqrt{N}$.

1.2 Results for random ensembles

To illustrate our framework we apply it to random ensembles of axion theories defined by a collection of integer charge matrices \mathbf{Q} , energy scales Λ_f^4 and axion-independent field space metrics \mathbf{K} that are loosely motivated by flux compactifications of string theory. While our techniques apply to all N, P , they are most powerful in the regime $N \gg 1$ and $P < 2N$. We will take \mathbf{Q} to be a $P \times N$ matrix of independent, identically distributed

⁵In general these shifts are linear combinations of “minimal” discrete symmetries of the potential, in a sense which we will clarify.

(i.i.d.) random integer entries with mean zero and standard deviation $\sigma_{\mathcal{Q}}$. As long as at least a small fraction $\gtrsim 3/N$ of the entries is non-vanishing — which at large N includes very sparse matrices — the universality of random matrix theory takes over and yields simple analytic results.⁶ As it turns out, in this regime the approximate shift symmetries become exponentially close to exact.

Even for the simplest case $P = N + 1$ we will find that the number of minima scales factorially with the number of terms in the potential (see also [26]),

$$\mathcal{N}_{\text{minima}} \propto \sigma_{\mathcal{Q}}^P \sqrt{P!}, \tag{1.3}$$

with a simple generalization to larger $P - N$. In these potentials there is a natural definition of neighboring minimum: those that are separated by no more than one traversal of the maxima of each cosine term of the potential. We will find that even when $P = N + 1$ the neighboring minima realize a wide variety of energy densities so long as $N \gg 1$. In other words these theories have extremely complex potentials that look random in the vicinity of any point or along a randomly chosen direction. However, they also possess nearly exact symmetries that make their analysis tractable. In particular, we can use the symmetries to identify minima with energy close to any desired value to exponential accuracy in polynomial time [28]. This kind of tractability in complex landscapes was recently discussed in [52].

We will consider both specific examples and ensembles of isotropic, positive definite kinetic matrices and parametrize the resulting field space diameters in terms of the largest eigenvalue f_{max}^2 of \mathbf{K} . Both the field space diameters and the distribution of energies in minima have previously been studied in such random axion theories. In this work we unify and generalize many of those results. We find that the field space distance suitable for inflation is typically as large as (see also [35])

$$\mathcal{D} \lesssim \begin{cases} N^{3/2} f_{\text{max}}, & \text{for } P - N = \text{constant}, \\ N^{1/2} f_{\text{max}}, & \text{for } P - N \propto N. \end{cases} \tag{1.4}$$

This result is robust even when large hierarchies are present between the energy scales Λ_I^4 .

The approximate shift symmetries are lost in the limit $P \gg N$. In this case the potential ceases to be analytically tractable and approaches a Gaussian random field instead. In appendix G we discuss a connection between multi-axion theories in this limit and Gaussian random fields with a Gaussian power spectrum.

2 Exact and approximate axion symmetries

The leading non-perturbative potential for the N axions $\boldsymbol{\theta}$ in (1.1) is

$$V(\boldsymbol{\theta}) = \sum_{I=1}^P \Lambda_I^4 \left[1 - \cos(\mathcal{Q}\boldsymbol{\theta})^I \right], \tag{2.1}$$

⁶Note that $P \times N$ matrices with a fraction of non-zero entries fewer than $1/P$ cannot be full rank, and can be dealt with using the techniques of appendix A. Hence our methods apply to all matrices except those with a fraction of non-zero entries between $1/P$ and $3/N$.

where we postpone the discussion of non-vanishing phases δ^I to section 2.5. This potential depends on the P energy scales Λ_I and the PN integers in the charge matrix \mathcal{Q} . At large N the field space volume becomes large, but given the limited number of parameters and the periodicity of the cosines, one might suspect that the entire structure of the potential is analytically tractable, at least so long as P is not too great. In the following we will make this expectation precise.

2.1 Auxiliary fields and a geometric picture

To analyze (2.1) it turns out to be useful to consider a set of P real scalar fields ϕ , subject to an auxiliary potential

$$V_{\text{aux}}(\phi) \equiv \sum_{I=1}^P \Lambda_I^4 [1 - \cos(\phi^I)] . \tag{2.2}$$

Comparing to (2.1) we observe that the argument of the I th cosine $(\mathcal{Q}\theta)^I$ has been replaced by an independent field ϕ^I . Hence the physical potential (2.1) is identical to (2.2) if $\phi^I = (\mathcal{Q}\theta)^I$, or more compactly

$$\phi|_{\Sigma} = \mathcal{Q}\theta , \tag{2.3}$$

where this equation defines a hyperplane Σ in the auxiliary field space \mathbb{R}^P (which we call ϕ -space). Specifically, note that

$$\mathcal{Q}\theta = \theta^1 \mathcal{Q}_1 + \theta^2 \mathcal{Q}_2 + \dots + \theta^N \mathcal{Q}_N \tag{2.4}$$

is a linear combination of the columns \mathcal{Q}_j of \mathcal{Q} . The surface Σ is the hyperplane spanned by these columns (the column space of $\mathcal{Q} \equiv \text{colsp}(\mathcal{Q})$), and (2.1) and (2.2) coincide when ϕ is constrained to Σ :

$$V(\theta) = V_{\text{aux}}(\phi|_{\Sigma}) . \tag{2.5}$$

For this reason we refer to Σ as the constraint surface (cf. figures 1 and 2). An efficient way to impose this constraint on ϕ is to introduce $P - N$ Lagrange multiplier fields into the action; we will do so explicitly in section 2.5. For $P \geq N$ the dimension of Σ is N if the columns of \mathcal{Q} are linearly independent, and so the map is one-to-one. In this case \mathcal{Q} is called full rank. We can assume this to be true without loss of generality and will do so from now on (cf. appendix A).

The utility of framing the problem in the extended P -dimensional space stems from the fact that the symmetries of V_{aux} are manifest: $\phi^I \rightarrow \phi^I + 2\pi n^I$ with $n^I \in \mathbb{Z}$, so that V_{aux} is identical in P -cubes of side-length 2π that we take to be centered on the sites of the scaled integer lattice $2\pi\mathbb{Z}^P$. Each cube can be labeled by an integer P -vector \mathbf{n} :

$$\{\phi : \|\phi - 2\pi\mathbf{n}\|_{\infty} \leq \pi\} , \quad \mathbf{n} \in \mathbb{Z}^P , \tag{2.6}$$

where the ℓ_{∞} -norm of a vector returns its largest absolute component.⁷ Within a single P -cube the potential V_{aux} is relatively featureless and every P -cube contains a single minimum

⁷Clearly the ℓ_{∞} -norm is basis dependent. We denote the basis of a vector by its symbol, i.e. ϕ is a vector in the standard basis for ϕ -space, while θ is a vector in the standard basis for θ -space (\mathbb{R}^N).




Figure 1. *Top:* constraint surface (red line) along with the lattice $2\pi\mathbb{Z}^P$ (gray dots). Arrows show the aligned basis vectors \mathbf{t}^{\parallel} and \mathbf{t}^{\perp} . *Bottom:* axion potential as a function of the aligned coordinate ω_{\parallel} (defined later in (2.18)). Distinct tiles $\mathcal{T}_{\mathbf{n}}$ are numbered and shaded.

at its center $\phi = 2\pi\mathbf{n}$ where $V_{\text{aux}} = 0$. Points where Σ passes through the center of a P -cube are therefore global minima of the physical potential. This set of points forms a sublattice $\mathcal{L}_{\Sigma} \equiv \Sigma \cap 2\pi\mathbb{Z}^P$. It is simple to see that this sublattice is rank N if \mathcal{Q} has integer entries and full rank.⁸ The auxiliary potential is manifestly invariant under shifts between any pair of such points, and therefore so is the physical potential V . In other words, this sublattice defines the N exact shift symmetries of (1.1).

The tiling of ϕ -space into P -cubes (2.6) provides a useful way to divide the physical field space into distinct regions. The constraint surface Σ slices across the cubes, and the regions of intersection of Σ with various P -cubes are an N -dimensional tiling of Σ (see figure 2). Within each tile the potential is relatively smooth because none of the individual cosine terms in (2.1) traverses its respective period. We can label each tile by the integer P -vector \mathbf{n} of the corresponding P -cube (2.6):

$$\mathcal{T}_{\mathbf{n}} = \{\boldsymbol{\theta} : \|\mathcal{Q}\boldsymbol{\theta} - 2\pi\mathbf{n}\|_{\infty} \leq \pi\}, \quad \mathbf{n} \in \mathbb{Z}^P. \quad (2.7)$$

Not every integer P -vector \mathbf{n} corresponds to a tile because some P -cubes do not intersect Σ .

When some or all of the angles of Σ with respect to the grid defined by (2.6) are small, one expects that at least some shifts from an initial tile to an adjacent one (adding 1 to one of the components of \mathbf{n}) will leave the physical potential close to invariant. This is the case, but we will see in section 2.3 that we can define a different set of shifts that are in general closer to exact symmetries than these, and have the desirable property that they form a complete (but not overcomplete) basis for the set of all distinct tiles (2.7).

⁸One might worry that $\text{GL}(N, \mathbb{R})$ transformations of the axion fields do not preserve the fact that \mathcal{Q} has integer entries. In fact, the necessary and sufficient condition on \mathcal{Q} such that $\mathcal{L}_{\Sigma} = \text{colsp}(\mathcal{Q}) \cap 2\pi\mathbb{Z}^P$ is rank N is that $\mathbf{P} = \mathcal{Q}(\mathcal{Q}^{\top}\mathcal{Q})^{-1}\mathcal{Q}^{\top}$, which is the orthogonal projector onto $\text{colsp}(\mathcal{Q})$, contains only rational entries. This property is preserved under $\text{GL}(N, \mathbb{R})$ transformations.

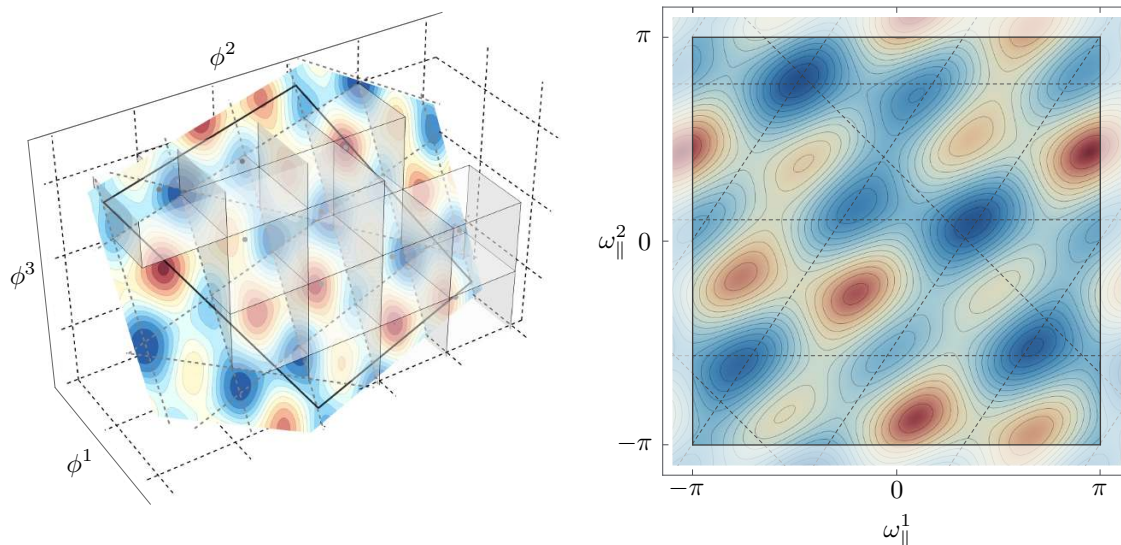


Figure 2. *Left:* contour plot of the auxiliary potential on the constraint surface Σ in ϕ -coordinates, for an example with $N = 2, P = 3$. The solid black lines denote one periodic domain of the potential, while the dashed gray lines denote the boundaries of the tiles defined in (2.7) (the intersections of the cubical tiling (2.6) of the auxiliary potential with the constraint surface). The depicted cubes constitute a full set of those with distinct intersections with Σ . *Right:* contour plot of the physical potential and its tiles in aligned coordinates ω_{\parallel} (defined later in (2.18)). Opposing boundaries of the periodic domain are identified.

2.2 Exact periodicities and periodic domains

As discussed above, displacements between points in the sublattice $\mathcal{L}_{\Sigma} \equiv \Sigma \cap 2\pi\mathbb{Z}^P$ leave V_{aux} unchanged and lie within the constraint surface Σ . These are the exact shift symmetries of the physical potential.

In general, a basis for a rank M lattice is a set of M linearly independent vectors with the property that *integer* linear combinations generate all lattice points. Consider an M -parallelepiped, with edges defined by the basis vectors of a lattice. This parallelepiped is a periodic domain for the lattice and contains exactly one lattice point. A simple example is a P -hypercube in (2.6) with e.g. $\mathbf{n} = \mathbf{0}$, which is a periodic domain for the P -dimensional lattice $2\pi\mathbb{Z}^P$.

We will use the notation $\{\mathbf{t}_i^{\parallel}\}$ to denote the N integer vectors that generate the lattice $\Sigma \cap \mathbb{Z}^P$ (so that the vectors $\{2\pi\mathbf{t}_i^{\parallel}\}$ generate \mathcal{L}_{Σ}).⁹ A single cell of this lattice sublattice is a region in which all distinct features of the potential are captured — in other words, it is a periodic domain of the physical potential, and we can restrict our attention to one such cell.

⁹Given a hyperplane and a lattice, it is non-trivial to find a basis for the sublattice resulting from their intersection. For a rank M sublattice a set of M linearly independent lattice vectors that lie within in the hyperplane do not in general generate all points in the sublattice. For instance, the columns of \mathcal{Q} do not generally serve as a basis for \mathcal{L}_{Σ} . It is possible, however, to find the sublattice basis algorithmically, for instance with the extended LLL algorithm [53].

Any basis $\{2\pi\mathbf{t}_i^\parallel\}$ for \mathcal{L}_Σ forms a *primitive set* for the full auxiliary lattice $2\pi\mathbb{Z}^P$ (see appendix C for a proof). This means there exists a set of $P - N$ lattice vectors $\{\mathbf{t}_1^\#, \dots, \mathbf{t}_{P-N}^\#\}$ that are not parallel to Σ and that when combined with the N vectors $\{\mathbf{t}_i^\parallel\}$ form a basis for \mathbb{Z}^P . It will be important in a moment that the $P - N$ supplemental vectors are not unique. The only condition is that the matrix containing all P basis vectors

$$[\mathbf{t}_1^\parallel \dots \mathbf{t}_N^\parallel \mathbf{t}_1^\# \dots \mathbf{t}_{P-N}^\#] \tag{2.8}$$

is unimodular (determinant one with integer entries). Whenever the vectors $\mathbf{t}_a^\#$ are (at least somewhat) aligned with Σ , in a sense we shall make more precise below, we refer to this basis for \mathbb{Z}^P as the *aligned lattice basis*.

2.3 Approximate symmetries and well-aligned theories

In general the transverse lattice vectors $\mathbf{t}_a^\#$ will not be orthogonal to Σ . Their decomposition into Σ and its orthogonal complement Σ^\perp will be important. We label the orthogonal projectors onto these subspaces by \mathbf{P} and \mathbf{P}^\perp , respectively. Since $\Sigma = \text{colsp}(\mathbf{Q})$, it follows that $\Sigma^\perp = \ker(\mathbf{Q}^\top)$ and we have the following explicit form of the projectors in terms of the charge matrix,

$$\mathbf{P} = \mathbb{1}_P - \mathbf{P}^\perp = \mathbf{Q}(\mathbf{Q}^\top \mathbf{Q})^{-1} \mathbf{Q}^\top. \tag{2.9}$$

Now consider a shift of the fields generated by the projection of a non-parallel lattice vector onto Σ :

$$\mathbf{Q}\theta \rightarrow \mathbf{Q}\theta + 2\pi\mathbf{P}\mathbf{t}_a^\#. \tag{2.10}$$

This shift is projected onto Σ and hence corresponds to a physical shift of the potential $V = V_{\text{aux}}|_\Sigma$. However it is not an exact symmetry because $\mathbf{P}\mathbf{t}_a^\#$ is not an integer vector. The amount by which this shift breaks the symmetry is proportional the projection of $\mathbf{t}_a^\#$ onto Σ^\perp :

$$V_{\text{aux}}(\phi) \rightarrow V_{\text{aux}}(\phi + 2\pi\mathbf{P}\mathbf{t}_a^\#) \tag{2.11}$$

$$= V_{\text{aux}}(\phi + 2\pi\mathbf{t}_a^\# - 2\pi\mathbf{P}^\perp\mathbf{t}_a^\#) \tag{2.12}$$

$$= V_{\text{aux}}(\phi - 2\pi\mathbf{P}^\perp\mathbf{t}_a^\#), \tag{2.13}$$

where in the second step we used that the potential is invariant under $\phi^I \rightarrow \phi^I + 2\pi$. Therefore, if the integer vectors $\mathbf{t}_a^\#$ can be chosen so that each component of $\mathbf{P}^\perp\mathbf{t}_a^\#$ is much less than one — if $\|\mathbf{P}^\perp\mathbf{t}_a^\#\|_\infty \ll 1$ — the correction to each cosine term in (2.2) is small and the shift $\phi \rightarrow \phi + 2\pi\mathbf{P}\mathbf{t}_a^\#$ is an approximate symmetry.

To identify both the exact and approximate symmetries, we choose a basis (2.8) for \mathbb{Z}^P which is as aligned as possible with Σ . The first N vectors \mathbf{t}_i^\parallel lie in Σ (and are a basis for the lattice $\Sigma \cap \mathbb{Z}^P$), thus

$$\mathbf{P}^\perp\mathbf{t}_i^\parallel = \mathbf{0}, \quad \forall i \in \{1, \dots, N\}. \tag{2.14}$$

These vectors generate the N exact shift symmetries of the physical potential and any parallelepiped with the \mathbf{t}_i^\parallel as edges is a periodic domain of the potential. The remaining

$P - N$ vectors $\mathbf{t}_a^\#$ should satisfy

$$\left\| \mathbf{P}^\perp \mathbf{t}_a^\# \right\|_\infty \text{ are smallest possible, } \forall a \in \{1, \dots, P - N\}. \quad (2.15)$$

The vectors $2\pi \mathbf{P} \mathbf{t}_a^\#$ generate $P - N$ approximate symmetries of the potential.¹⁰ The aligned lattice basis for a simple axion potential, one with $P = 2$ and $N = 1$, is shown in figure 1. We refer to theories where the orthogonal projections of all elements of the aligned basis are small as *well-aligned*. That is, a well-aligned theory satisfies

$$\left\| \mathbf{P}^\perp \mathbf{t}_a^\# \right\|_\infty \ll \frac{1}{P - N}, \quad \forall a \in \{1, \dots, P - N\} \quad (2.16)$$

(cf. (3.28) for the origin of the $1/(P - N)$ on the right-hand side).

To illustrate the utility of these approximate symmetries, suppose ϕ is chosen to be a global minimum of the potential that lies on Σ (for instance $\phi = \mathbf{0}$). Repeated shifts by $2\pi \mathbf{P} \mathbf{t}_a^\#$ then identify the approximate location (and determine the energy, as we discuss in section 3.3) of many physically distinct minima with slightly different properties (see figure 1). We will see later that in the random ensembles we study $\left\| \mathbf{P}^\perp \mathbf{t}_a^\# \right\|_\infty$ can generically be chosen so that it is exponentially small in N , for all $1 \leq a \leq P - N$. More generally it can be small when $\det(\mathcal{Q}^\top \mathcal{Q})$ is large, since this determinant appears in the denominator of the projector. This allows us to locate and characterize many minima in an otherwise intractably complex landscape very easily and to good accuracy. Furthermore, in well-aligned theories all P phases δ^I can be eliminated to good accuracy (cf. section 2.5).

2.4 Aligned coordinates and similar tiles

The tiling (2.7) of Σ was introduced in section 2.1 as a means of delineating relatively featureless sections of the potential by using the basic infrastructure provided by the auxiliary lattice. Here we show how this tiling enables the identification of many similar regions within one periodic domain of the physical potential.

As a preliminary illustration of a more general method, consider the periodic domain surrounding a global minimum of the potential (say the origin $\theta = \mathbf{0}$), and a P -cube centered at this position on Σ . Now choose a specific non-parallel lattice vector $\mathbf{t}_a^\#$ and consider the set of P -cubes obtained by sequentially shifting the center of each cube by $2\pi \mathbf{t}_a^\#$, together with the tiles defined as their intersections with Σ . In well-aligned theories the auxiliary potential evaluated on successive intersections in the list (and hence the physical potential in the corresponding tiles) will be very similar. Now, regardless of whether or not a model is well-aligned, after a certain number of shifts the P -cube will have receded far enough from Σ that it no longer intersects it. If the number of shifts before this point

¹⁰Equation (2.15) defines what is known as a reduced basis for the lattice generated by \mathbf{P}^\perp . We are purposefully vague in the precise definition of “smallest possible” and “reduced”: there are multiple definitions, such as Minkowski, LLL, or Rankin reduction. For example, depending on the precise application, one may be interested in a basis that aligns only some of the $P - N$ transverse directions. These details are irrelevant for the present discussion and we refer to the literature [54–60]. A particularly simple approximation is given by the Mathematica package for the extended LLL algorithm [53]. We thank Liam McAllister and John Stout for discussion on this point.

is m , we have identified m distinct tiles that are labeled by consecutive integer multiples of \mathbf{t}_a^\parallel .¹¹ These tiles may be scattered across multiple periodic domains because accumulating shifts eventually push part or all of the intersection out of the original periodic domain. Shifts by integer linear combinations of the $2\pi\mathbf{t}_i^\parallel$ can then be used to uniquely return all portions of tiles into the periodic domain containing the origin. All such tiles are distinct because they originated from distinct intersections with Σ .

A complete tiling of the periodic domain can be found by following a generalization of the procedure outlined above — shifting the P -cube containing the global minimum by linear combinations of all non-parallel directions $2\pi\mathbf{t}_a^\parallel$, and scanning this space until all intersecting P -cubes are identified. A more convenient labeling of the regions that tile the periodic domain is achieved by defining *aligned coordinates* $\boldsymbol{\omega}$ for the auxiliary space:¹²

$$\boldsymbol{\phi} \equiv (\mathbf{t}^\parallel \mid \mathbf{t}^\perp) \boldsymbol{\omega}. \tag{2.17}$$

Recall that the matrix appearing in (2.17), which is identical to (2.8), is unimodular (determinant one) and thus has a unimodular inverse. Integer P -vectors in $\boldsymbol{\phi}$ -coordinates are in one-to-one correspondence with integer vectors in $\boldsymbol{\omega}$ -coordinates. The components of $\boldsymbol{\omega}$ separate into components parallel and not parallel to Σ :

$$\boldsymbol{\omega} = \begin{pmatrix} \boldsymbol{\omega}_\parallel \\ \boldsymbol{\omega}_\perp \end{pmatrix}. \tag{2.18}$$

A shift by 2π of any of the first N entries leaves the physical potential invariant, since this is a shift of $\boldsymbol{\phi}$ by a $2\pi\mathbf{t}_i^\parallel$. The periodic domain in $\boldsymbol{\omega}$ -coordinates is simply an N -cube of side-length 2π in the $\boldsymbol{\omega}_\parallel$ -plane. Since opposing sides of the periodic domain are identified, any point on Σ is readily identified with a corresponding point in the central periodic domain $\boldsymbol{\omega}_\parallel \pmod{2\pi\mathbb{Z}^N}$. It is likewise easy to recognize similar tiles in well-aligned theories: they correspond to P -cubes with similar $\boldsymbol{\omega}_\parallel$ -coordinates. The above is illustrated in figure 3. Note finally that an advantage of using aligned coordinates is that distinct tiles within one periodic domain are labeled by distinct integer $(P - N)$ -vectors \mathbf{m} ,¹³ as $P - N$ is the amount of \mathbf{t}_a^\parallel s and integral shifts along the $\boldsymbol{\omega}_\parallel$ -directions do not change the tile.

Notational intermezzo. To simplify equations like (2.17), from now on we adopt the notation $\mathbf{T}_{\alpha\beta}$ for a transformation between the components of a vector in two bases. The subscript pair is to be read left to right: “transform from $\boldsymbol{\alpha}$ -coordinates to $\boldsymbol{\beta}$ -coordinates”. In this notation (2.17) becomes

$$\boldsymbol{\phi} = \mathbf{T}_{\boldsymbol{\omega}\boldsymbol{\phi}} \boldsymbol{\omega}.$$

The inverse transformation is simply

$$\boldsymbol{\omega} = \mathbf{T}_{\boldsymbol{\phi}\boldsymbol{\omega}} \boldsymbol{\phi}.$$

¹¹Note that $m \sim 1/\|\mathbf{P}^\perp\mathbf{t}_a^\parallel\|_\infty$, which is large in a well-aligned theory.

¹²It may be useful to reduce the basis for the lattice spanned by the \mathbf{t}_i^\parallel .

¹³And vice-versa, modulo those related by $\mathbf{m} \leftrightarrow -\mathbf{m}$, as the P -cube grid and N -parallelepiped domains in Σ are symmetric under the reflection about any of the P Cartesian coordinate axes.




Figure 3. Illustration of the tiling of the periodic domain of the physical potential induced by the intersections of the constraint surface with P -cubes surrounding lattice sites in the auxiliary lattice. We used $P = N + 2 = 4$ and labeled each distinct intersection by a $P - N = 2$ -vector. This example is well-aligned only in the direction \mathbf{t}_1^\parallel corresponding to the upper label (if both directions are well-aligned, the plot becomes too dense to be legible).

When the transformation is between spaces of equal dimension as here (so that \mathbf{T} is square), $\mathbf{T}_{\phi\omega} = (\mathbf{T}_{\omega\phi})^{-1}$. When a transformation matrix's columns or rows separate in a useful way, as is the case here,

$$\mathbf{T}_{\omega\phi} = \left(\mathbf{t}^\parallel \mid \mathbf{t}^\sharp \right), \tag{2.19}$$

the (rectangular) submatrices are labeled in the natural way:

$$\left(\underbrace{\mathbf{t}^\parallel}_{\equiv \mathbf{T}_{\omega\parallel\phi}} \mid \underbrace{\mathbf{t}^\sharp}_{\equiv \mathbf{T}_{\omega\sharp\phi}} \right). \tag{2.20}$$

We can apply this notation to the transformation (2.3) between the N -vector $\boldsymbol{\theta}$ and the P -vector $\boldsymbol{\phi}$ constrained to Σ ,

$$\mathbf{T}_{\theta\phi} = \mathcal{Q}, \quad \mathbf{T}_{\phi\theta} = (\mathcal{Q}^\top \mathcal{Q})^{-1} \mathcal{Q}^\top. \tag{2.21}$$

Combining the transformations (2.20) and (2.21) we can also relate the coordinates $\boldsymbol{\omega}_\parallel$ and $\boldsymbol{\theta}$,

$$\mathbf{T}_{\omega_\parallel\theta} = \mathbf{T}_{\phi\theta} \mathbf{T}_{\omega_\parallel\phi}, \quad \mathbf{T}_{\theta\omega_\parallel} = (\mathbf{T}_{\omega_\parallel\theta})^{-1}. \tag{2.22}$$

Finally, we can express the exact and approximate shift symmetries of the theory in terms of the $\boldsymbol{\theta}$ -coordinates we started off with in (1.1) (in section 2.2 and section 2.3 these were only expressed in $\boldsymbol{\phi}$ -coordinates). The exact symmetries are given by

$$\boldsymbol{\theta} \rightarrow \boldsymbol{\theta} + 2\pi \mathbf{T}_{\phi\theta} \mathbf{t}_i^\parallel \tag{2.23}$$

$$= \boldsymbol{\theta} + 2\pi (\mathcal{Q}^\top \mathcal{Q})^{-1} \mathcal{Q}^\top \mathbf{t}_i^\parallel, \tag{2.24}$$

while the approximate symmetries are given by

$$\boldsymbol{\theta} \rightarrow \boldsymbol{\theta} + 2\pi \mathbf{T}_{\phi\theta} \mathbf{P} \mathbf{t}_a^\# \tag{2.25}$$

$$= \boldsymbol{\theta} + 2\pi (\mathbf{Q}^\top \mathbf{Q})^{-1} \mathbf{Q}^\top \mathbf{t}_a^\#. \tag{2.26}$$

Returning to the tiling of the periodic domain, it is not hard to see that there are only finitely many tiles: consider sliding the center of a P -cube along any real linear combination of the $\mathbf{t}_a^\#$ (which corresponds to a line emanating from the origin in the $(P-N)$ -dimensional $\boldsymbol{\omega}_\parallel$ -space) — a generalization of the illustration at the beginning of this section. As the perpendicular distance from the center of the cube to Σ is ever-increasing along the line, eventually the cube will no longer intersect Σ . Thus there are only finitely many distinct tiles, which are labeled by certain integer $(P-N)$ -vectors \mathbf{m} . We now argue that the allowed values for \mathbf{m} lie in a particular (compact) convex region \mathfrak{C} in $\boldsymbol{\omega}_\parallel$ -space. We define \mathfrak{C} by the set of *all* points in $\boldsymbol{\omega}_\parallel$ -space which correspond to centers of P -cubes in ϕ -space that have *some* intersection with Σ (note that this includes also vectors with irrational entries). This is the same as the collection of the $\boldsymbol{\omega}_\parallel$ -coordinates of all the points lying inside the P -cube of side-length 2π centered at the origin in ϕ -space, or yet in other terms, the orthogonal projection of that P -cube onto Σ^\perp (in $\boldsymbol{\omega}_\parallel$ -coordinates). It follows that distinct tiles are labeled by distinct integer vectors $\mathbf{m} \in \mathfrak{C}$ with¹⁴

$$\mathfrak{C} = \left\{ \mathbf{T}_{\phi\omega_\parallel} \boldsymbol{\phi} \mid \|\boldsymbol{\phi}\|_\infty \leq \pi \right\}. \tag{2.27}$$

Quite clearly the extreme points¹⁵ of \mathfrak{C} correspond to the projections of certain vertices of the P -cube. Since a compact convex set equals the (closed) convex hull¹⁶ of its extreme points by the Krein-Milman theorem [54, 61], we have the following alternate definition of \mathfrak{C} :

$$\mathfrak{C} = \text{Conv} \left(\left\{ \mathbf{T}_{\phi\omega_\parallel} \mathbf{e} \mid e^I = \pm\pi, I \in \{1 \dots P\} \right\} \right), \tag{2.28}$$

where $\text{Conv}(\cdot)$ denotes the convex hull of a set. \mathfrak{C} is a polytope in the $(P-N)$ -dimensional \mathbf{m} -space, illustrated in figure 5.¹⁷

To recap, the aligned coordinates $\boldsymbol{\omega}$ are very convenient to identify a complete set of tiles covering exactly one periodic domain of the potential of (1.1). First, it suffices to fix $\boldsymbol{\omega}_\parallel = \mathbf{0}$. This guarantees that only one periodic domain's worth of tiles will be counted, and the value of $\boldsymbol{\omega}_\parallel$ at a point is irrelevant to how or if a P -cube centered there intersects Σ . Therefore, all distinct tiles are labeled by the $\boldsymbol{\omega}_\parallel$ -coordinates of the centers of

¹⁴This is indeed a convex set: if $\mathbf{T}_{\phi\omega_\parallel} \boldsymbol{\phi}_1, \mathbf{T}_{\phi\omega_\parallel} \boldsymbol{\phi}_2 \in \mathfrak{C}$ then also $\lambda \mathbf{T}_{\phi\omega_\parallel} \boldsymbol{\phi}_1 + (1-\lambda) \mathbf{T}_{\phi\omega_\parallel} \boldsymbol{\phi}_2 \in \mathfrak{C}$ since it is of the form $\mathbf{T}_{\phi\omega_\parallel} [\lambda \boldsymbol{\phi}_1 + (1-\lambda) \boldsymbol{\phi}_2]$ with indeed $\|\lambda \boldsymbol{\phi}_1 + (1-\lambda) \boldsymbol{\phi}_2\|_\infty \leq \lambda \|\boldsymbol{\phi}_1\|_\infty + (1-\lambda) \|\boldsymbol{\phi}_2\|_\infty \leq \pi$.

¹⁵Extreme points of a convex set are points which are not interior points of any line segment belonging to the set.

¹⁶The convex hull of a set is the intersection of all convex sets containing that set. For the vertices of a polytope, the convex hull is the polytope.

¹⁷In principle one could eliminate all the redundant points in the set of which the convex hull is being taken in (2.28), which correspond to vertices of cubes that lie on the constraint surface but which are not the only intersection of the P -cube with Σ (cf. figure 5), and maintain the same polytope \mathfrak{C} . The amount of remaining (extreme) points of \mathfrak{C} is much less than the 2^P used in (2.28): we believe an upper bound scales only polynomially as P^{P-N-1} . However, we are not aware of a polynomial time algorithm that can find \mathfrak{C} .

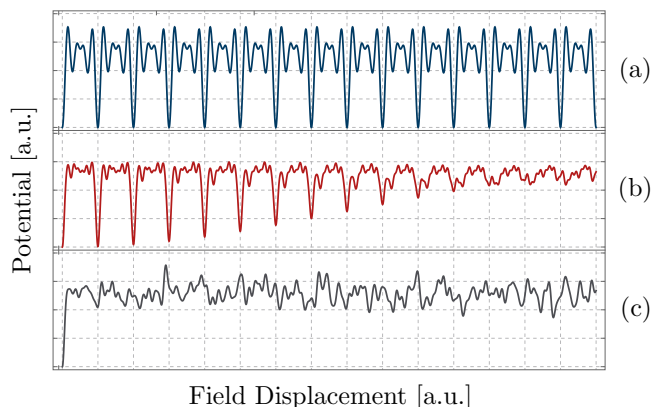


Figure 4. The potential plotted along three different rays through field space (starting at the global minimum), for an example of the potential in (1.1) with $N = 23, P = 40$. *Top pane:* a line oriented along an exact symmetry direction. *Middle pane:* a line oriented along an approximate symmetry direction. *Bottom pane:* a random direction.

their P -cubes, and there is some compact and convex region \mathfrak{C} of ω_{\parallel} -space that contains them all.

In general it is computationally very hard to identify the vertices that define the polytope \mathfrak{C} in (2.28). A simple sufficient condition for a lattice site $2\pi\mathbf{m}$ to lie within the polytope is that its projections onto the ϕ -coordinate axes do not exceed π ,

$$\|2\pi\mathbf{P}^{\perp}\mathbf{T}_{\omega_{\parallel}\phi}\mathbf{m}\|_{\infty} \leq \pi \quad \Rightarrow \quad \mathbf{m} \in \mathfrak{C}, \quad (2.29)$$

while the inverse is not true. This subregion of the polytope is illustrated in figure 5.

As mentioned above, another advantage of the ω -coordinates is that in well-aligned theories, regions of Σ corresponding to intersections with P -cubes that are close together in ω_{\parallel} -space will be nearly identical, because they differ by only a small number of shifts by approximate discrete symmetries. The corresponding regions may be very far apart even after modding to one periodic domain in ω_{\parallel} -space, because the shifts $2\pi\mathbf{P}\mathbf{t}_a^{\parallel}$ can be very long. Neighboring regions on Σ are not in general similar, while specific distant regions are. This is illustrated in figure 4: there is no clear structure in the potential along an arbitrary ray in field space, but when considering lines that intersect widely separated tiles related by exact or approximate shift symmetries the potential becomes structured.

2.5 Phases

We now return to the phases δ appearing in the original Lagrangian (1.1), with potential

$$V = \sum_{I=1}^P \Lambda_I^4 \left[1 - \cos(\mathcal{Q}\theta + \delta)^I \right]. \quad (2.30)$$

Just as before, we promote the cosine arguments to P independent fields ϕ^I that must be constrained to a hyperplane in order to reproduce the physical potential. That is, we require

$$\phi = \mathcal{Q}\theta + \delta, \quad (2.31)$$




Figure 5. Illustration of the ω_{\parallel} -coordinates for every distinct tile on Σ , in an example where $P = N + 2 = 8$. The central six-sided region denotes the area bounded by the simple sufficient condition (2.29), while the full shaded region \mathcal{C} contains all lattice sites whose P -cubes intersect the constraint surface. Blue crosses (large and small) denote the coordinates of centers of cubes which have a vertex that lies on Σ .

which defines a hyperplane parallel to $\Sigma = \text{colsp}(\mathcal{Q})$, such that the constraint surface on which the auxiliary potential reproduces the physical potential (2.30) is $\Sigma + \delta$.

To impose the constraint (2.31) in the action we introduce $P - N$ Lagrange multipliers ν_a :

$$V = \sum_{I=1}^P \Lambda_I^4 [1 - \cos(\phi^I)] + \sum_{a=1}^{P-N} \nu_a \mathcal{R}^a(\phi - \mathbf{P}^\perp \delta). \quad (2.32)$$

Here \mathcal{R} is any $(P - N) \times P$ matrix with the property that its row space is Σ^\perp , the orthogonal complement to Σ . For instance, for \mathcal{R} one could use any $P - N$ linearly independent rows of the matrix $\mathbf{P}^\perp = \mathbf{1}_P - \mathbf{P}$. The equations of motion for the ν_a constrain $\phi - \mathbf{P}^\perp \delta$ to be perpendicular to Σ^\perp ; that is, they constrain ϕ to lie in $\Sigma + \mathbf{P}^\perp \delta = \Sigma + \delta$. In checking this the identity $\mathcal{R}\mathbf{P}^\perp = \mathcal{R}$ is useful.

Since $\mathbf{P}^\perp \delta$ is a vector in the $(P - N)$ -dimensional subspace Σ^\perp , the projection in (2.32) has already removed all but $P - N$ of the original P phases (this reduction is simply the obvious freedom to continuously redefine the N fields θ in (2.30)). We will now demonstrate that in well-aligned theories, the remaining phases $\mathbf{P}^\perp \delta$ can be reduced to small values using the approximate shift symmetries of the theory. Consider the shift $\phi \rightarrow \phi + 2\pi \mathbf{T}_{\omega_{\parallel}\phi} \mathbf{n}_\delta$, where \mathbf{n}_δ is an arbitrary integer $(P - N)$ -vector. This shift is an exact symmetry of the cosines, but affects the constraint terms in (2.32):

$$V = \sum_I \Lambda_I^4 [1 - \cos(\phi^I)] + \sum_a \nu_a \mathcal{R}^a \left(\phi + 2\pi \mathbf{T}_{\omega_{\parallel}\phi} \mathbf{n}_\delta - \mathbf{P}^\perp \delta \right). \quad (2.33)$$

To identify the integers \mathbf{n}_δ that minimize the remaining phases in (2.33), recall the relation between the aligned coordinates ω and ϕ -coordinates

$$\omega = \mathbf{T}_{\phi\omega} \phi \equiv \begin{pmatrix} \mathbf{T}_{\phi\omega_{\parallel}} \\ \mathbf{T}_{\phi\omega_{\perp}} \end{pmatrix} \phi. \quad (2.34)$$

Using these definitions, the vector \mathbf{n}_δ that minimizes the remaining phases is

$$\mathbf{n}_\delta = \left[\frac{1}{2\pi} \mathbf{T}_{\phi\omega_\#} \mathbf{P}^\perp \boldsymbol{\delta} \right]_{\text{n.i.}}, \quad (2.35)$$

where $[\dots]_{\text{n.i.}}$ denotes the nearest integer vector. Using $\mathbf{P}^\perp = \mathbf{P}^\perp \mathbf{T}_{\omega_\# \phi} \mathbf{T}_{\phi\omega_\#} \mathbf{P}^\perp$, one can see that this choice of \mathbf{n}_δ reduces the phases to zero with an error bounded above by $\pi \|\mathbf{P}^\perp \mathbf{T}_{\omega_\# \phi}\|_\infty$, which is small when the theory is well-aligned.¹⁸

Explicitly, the field redefinition $\boldsymbol{\theta} \rightarrow \boldsymbol{\theta} + \boldsymbol{\theta}_{\text{shift}}$ that reduces the phases, and the remaining phases are given with (2.21) and (2.35) by

$$\boldsymbol{\theta}_{\text{shift}} = \text{mod}_{[-\pi, \pi]} \left[\mathbf{T}_{\phi\theta} (2\pi \mathbf{T}_{\omega_\# \phi} \mathbf{n}_\delta - \boldsymbol{\delta}) \right], \quad \boldsymbol{\delta}_r = \text{mod}_{[-\pi, \pi]} [\mathcal{Q}\boldsymbol{\theta}_{\text{shift}} + \boldsymbol{\delta}]. \quad (2.36)$$

From now on, to simplify the discussion we will focus on well-aligned theories where we can neglect the phases. As we shall see in section 3.6 this holds to good accuracy for large classes of axion theories with $P \lesssim 2N$. An explicit example illustrating the use of the aligned lattice basis and phase reduction can be found in appendix D.

3 Minima and saddle points

In the previous section we identified the most suitable basis to identify the symmetries, treating all P terms in the potential on equal footing. These symmetries can be employed to systematically explore all potential minimum locations, which in principle yields all minima to arbitrary accuracy. In typical applications, however, it may be more efficient to include other data as well, such as the scale of each of the non-perturbative terms. In order to find the stable minima, for example, non-perturbative terms that are entirely subleading will have essentially no effect on the location of the minimum, but may split the degeneracy between minima as discussed in [26] and reviewed in section 3.7. We will now discuss how the minima of a given axion theory can be determined to various degrees of accuracy.

3.1 Systematics of all minima

In section 2 we described how to decompose the field space into tiles, each of which corresponds to an intersection of Σ with a distinct P -cube in ϕ -space centered on a lattice point of $2\pi\mathbb{Z}^P$. The auxiliary potential $V_{\text{aux}}(\phi)$ is identical inside all P -cubes, but they can have a distinct (or empty) intersection with the constraint surface, and therefore contain a distinct region of the physical potential $V(\boldsymbol{\theta})$.

¹⁸ The ℓ_∞ -norm of a matrix \mathbf{A} can be defined as the maximum absolute row sum, $\|\mathbf{A}\|_\infty = \max_i \left\{ \sum_j |A^i_j| \right\}$. After the shift specified by (2.35), the remaining phase is $\boldsymbol{\delta}_r = 2\pi \mathbf{P}^\perp \mathbf{T}_{\omega_\# \phi} \boldsymbol{\alpha}$ for some $(P-N)$ -vector $\boldsymbol{\alpha}$ with $\|\boldsymbol{\alpha}\|_\infty \leq 1/2$. We may bound the magnitude of the largest component of this remaining phase by using the general inequality $\|\mathbf{A}\mathbf{v}\|_\infty \leq \|\mathbf{A}\|_\infty \|\mathbf{v}\|_\infty$ for any matrix \mathbf{A} and vector \mathbf{v} : $\|\boldsymbol{\delta}_r\|_\infty \leq \pi \|\mathbf{P}^\perp \mathbf{T}_{\omega_\# \phi}\|_\infty$. To make the connection with our definition of well-aligned theories (2.16), note $\|\mathbf{P}^\perp \mathbf{T}_{\omega_\# \phi}\|_\infty = \max_I \left\{ \sum_{a=1}^{P-N} \left| \left(\mathbf{P}^\perp \mathbf{T}_{\omega_\# \phi} \right)_a^I \right| \right\} \leq \sum_{a=1}^{P-N} \max_I \left\{ \left| \left(\mathbf{P}^\perp \mathbf{T}_{\omega_\# \phi} \right)_a^I \right| \right\} = \sum_{a=1}^{P-N} \|\mathbf{P}^\perp \mathbf{t}_a^\# \|_\infty$.

Just as already done in section 2.5 we introduce $P - N$ Lagrange multipliers ν_a that enforce the constraint of the auxiliary potential to Σ . To find extrema on Σ within a tile labeled by \mathbf{m} we then minimize the potential,

$$V = V_{\text{aux}}(\phi) + \sum_{a=1}^{P-N} \nu_a \mathcal{R}^a \phi, \quad (3.1)$$

within the corresponding P -cube $\phi = 2\pi \mathbf{T}_{\omega_{\parallel} \phi} \mathbf{m} + \delta\phi$, where $\|\delta\phi\| \leq \pi$. Requiring a vanishing gradient gives

$$\begin{aligned} \Lambda_I^4 \sin(\delta\phi^I) + (\nu \mathcal{R})^I &= 0, \quad \forall I \in \{1, \dots, P\}, \\ \mathcal{R}(2\pi \mathbf{T}_{\omega_{\parallel} \phi} \mathbf{m} + \delta\phi) &= \mathbf{0}. \end{aligned} \quad (3.2)$$

Since \mathcal{R} is a set of row vectors that span Σ^\perp , the first condition is the requirement that the gradient of V_{aux} is perpendicular to Σ — in other words, that the gradient projected onto Σ vanishes. The second condition ensures that the point is in Σ . Solving the optimization problem (3.2) within all P -cubes labeled by $\mathbf{m} \in \mathcal{C}$ yields all distinct extreme points, including all minima.

3.2 Neighboring minima

The aligned basis is ideally suited to identify similar regions of the axion potential. These tiles need not be close to one another in the physical field space, as we saw in section 2.4. Recall that this is because the \mathbf{t}_a^\parallel may contain large integers and therefore generate a large separation between the tiles' associated P -cubes (in ϕ -coordinates). This means that similar tiles are not generally immediate neighbors. For the purpose of this paper, we define immediately neighboring minima to be those whose P -cubes share a face or corner.

A minimum located at ϕ , which is within in the P -cube labeled by $\mathbf{n} = [\phi/(2\pi)]_{\text{n.i.}}$, has 3^P neighboring P -cubes labeled by

$$\mathbf{n}_{\text{neighbor}} = \mathbf{n} + \mathbf{e}, \quad \mathbf{e}^I \in \{0, \pm 1\}, \quad (3.3)$$

only some of which have a non-vanishing intersection with Σ . In order to identify all immediately neighboring minima we have to consider all neighboring P -cubes (3.3) that do intersect Σ . That means we need to consider all vectors \mathbf{e} that correspond to lattice sites which satisfy $2\pi \mathbf{T}_{\phi\omega_{\parallel}} \mathbf{P}^\perp(\mathbf{n} + \mathbf{e}) \in \mathcal{C}$ (cf. (2.28)). Again, a simple sufficient condition is given by

$$\|2\pi \mathbf{P}^\perp(\mathbf{n} + \mathbf{e})\|_\infty < \pi. \quad (3.4)$$

The precise locations of neighboring minima can be found by solving the optimization problem (3.2) with $\phi = 2\pi(\mathbf{n} + \mathbf{e}) + \delta\phi$ for each candidate \mathbf{e} .

3.3 Minima to quadratic order

In principle we can solve (3.2) and determine the location of all minima in the theory. However, although we can certainly solve (3.2) in any one P -cube to arbitrary accuracy, the very large number of domains make this impossible at large P . Instead, we will employ




Figure 6. Contour plot of the relative error made by expanding $\cos(\phi^1) + \cos(\phi^2)$ to quadratic order. The dashed lines indicate, from inner to outer, the levels of 0.25, .5, 0.75 and 1 relative error, respectively. The maximum relative error is 1.5. The quadratic domain is indicated by the gray square of side length π while the periodic domain is the whole box.

a quadratic expansion of the potential and the approximate symmetries to find an analytic expression for the approximate location of many minima.

The auxiliary potential has one single minimum located at the center of each P -cube, around which we can use a quadratic expansion. We will refer to the (somewhat arbitrary) region within which the quadratic expansion is a good approximation as the *quadratic domain*. Since the non-perturbative potential consists of simple cosines, we define the quadratic domain as

$$-\frac{\pi}{2} \leq \phi^I \leq \frac{\pi}{2}, \quad \forall I \in \{1, \dots, P\}, \quad (3.5)$$

such that the relative error made never exceeds 25% within that region. This choice might change if the underlying periodic function deviates from a cosine, but we chose it with some foresight in a way that this region will, in well-aligned theories, capture many minima of the full non-linear potential. The periodic and quadratic domains along with the relative error made by approximating cosines by a quadratic function are illustrated in figure 6.

The quadratic expansion dramatically simplifies the problem of finding minima. Consider a small displacement $\delta\phi$ from the auxiliary lattice point $2\pi \mathbf{T}_{\omega\phi} \mathbf{m}$. The potential in the corresponding quadratic domain evaluates to

$$V_{\delta\phi} = \frac{1}{2} \delta\phi^\top \text{diag}(\Lambda_I^4) \delta\phi + \boldsymbol{\nu}^\top \mathcal{R}(2\pi \mathbf{T}_{\omega\phi} \mathbf{m} + \delta\phi) + \mathcal{O}(\delta\phi^4), \quad (3.6)$$

so a vanishing gradient is implied by the conditions

$$\begin{aligned} \text{diag}(\Lambda_I^4) \delta\phi + \mathcal{R}^\top \boldsymbol{\nu} &= \mathbf{0}, \\ \mathcal{R}(2\pi \mathbf{T}_{\omega\phi} \mathbf{m} + \delta\phi) &= \mathbf{0}. \end{aligned} \quad (3.7)$$

Solving this system of equations for the location of a minimum $\phi_{\mathbf{m}}$ on the constraint




Figure 7. The blue shaded region illustrates the polytope \mathcal{C} in ω_{\parallel} -coordinates containing all lattice points corresponding to distinct tiles for an example with $P = N + 2 = 8$ (cf. figure 5 where the same region is shown). The central six-sided region is bounded by the $2P$ half-planes that determine the validity of the quadratic approximation (3.10). Tiles that contain a minimum as determined by numerical minimization are denoted by blue crosses, and nearly perfectly overlap with the quadratic region.

surface gives¹⁹

$$\phi_{\mathbf{m}} = 2\pi(\mathbb{1} - \Delta^{\perp}) \mathbf{T}_{\omega_{\parallel}\phi} \mathbf{m}, \quad (3.8)$$

where $\mathbb{1} - \Delta^{\perp}$ is a non-orthogonal projector onto the constraint surface:

$$(\Delta^{\perp})^2 = \Delta^{\perp} \equiv \text{diag}(\Lambda_I^{-4}) \mathcal{R}^{\top} \left[\mathcal{R} \text{diag}(\Lambda_I^{-4}) \mathcal{R}^{\top} \right]^{-1} \mathcal{R}. \quad (3.9)$$

Scanning over all \mathbf{m} , we can check which $\phi_{\mathbf{m}}$ lie within the quadratic domain of their respective lattice site $2\pi \mathbf{T}_{\omega_{\parallel}\phi} \mathbf{m}$; namely, within the intersection of $2P$ half-planes,

$$\|2\pi \Delta^{\perp} \mathbf{T}_{\omega_{\parallel}\phi} \mathbf{m}\|_{\infty} \leq \frac{\pi}{2}. \quad (3.10)$$

For these minima we will have succeeded at finding an approximate location of the constrained system. The energy density is given by

$$V(\phi_{\mathbf{m}}) \approx 2\pi^2 \mathbf{m}^{\top} \left[\mathbf{T}_{\phi\omega_{\parallel}} \text{diag}(\Lambda_I^{-4}) \mathbf{T}_{\phi\omega_{\parallel}}^{\top} \right]^{-1} \mathbf{m}, \quad (3.11)$$

where we used the specific choice $\mathcal{R} = \mathbf{T}_{\phi\omega_{\parallel}}$ to simplify the expression.

In figure 7 we illustrate the tiles in which there is a minimum located by numerically minimizing the potential, along with those for which the quadratic approximation predicts a minimum (i.e. predicts a minimum located within the quadratic domain where the approximation is self-consistent). In general these sets of points are not immediately related, but for explicit examples we typically found a substantial overlap.

¹⁹For non-vanishing phases $\delta_{\mathbf{r}}$ in (2.36) the minima are located at $\phi_{\mathbf{m}} + \delta_{\mathbf{r}}$, and correspondingly at potentials $(\delta\phi_{\mathbf{m}} - \delta_{\mathbf{r}})^{\top} \text{diag}(\Lambda_I^4)(\delta\phi_{\mathbf{m}} - \delta_{\mathbf{r}})/2$, where we defined $\delta\phi_{\mathbf{m}} = 2\pi \Delta^{\perp} \mathbf{T}_{\omega_{\parallel}\phi} \mathbf{m}$. Note the relation $\phi = \mathbf{T}_{\omega_{\parallel}\phi} \omega_{\parallel} + \delta_{\mathbf{r}}$, which implies that the locations of minima in ω_{\parallel} -coordinates remain unchanged in the presence of small phases.

Finally, let us comment on the special case of equal scales $\Lambda_I = \Lambda$ and $P = N + 1$. In this case the energies of the minima are

$$V(\phi_m) \approx 2\pi^2 \Lambda^4 \left(\frac{cm}{\sqrt{\det \mathbf{Q}^\top \mathbf{Q}}} \right)^2, \quad \forall |m| < \mathcal{N}_{\text{vac}}, \quad (3.12)$$

where c is an integer. The derivation of this result can be found in appendix E. The approximate signs in (3.12) denote the quadratic approximation which is valid for $\mathcal{N}_{\text{vac}} \sim \sqrt{\det \mathbf{Q}^\top \mathbf{Q}}$ minima. No other approximations are made in (3.12).

3.4 A uniform sample over all minima

The most direct approach to find all minima is to consider the $(P-N)$ -dimensional polytope \mathcal{C} defined in (2.28). This region contains all vectors $\mathbf{m} \in \mathbb{Z}^{P-N}$ that label the tiles of one periodic domain of the potential constrained to Σ (see figure 3). Minimizing the potential in all the tiles in \mathcal{C} covers one entire periodic domain of the potential, yielding the set of all minima. However, the number of tiles is typically exponential in N or P , so even at moderately large values of these parameters this comprehensive approach becomes intractable.

Instead, we can take advantage of the approximate symmetries of the potential that we identified. The auxiliary potential has minima only at the centers of each P -cube, which suggests that many tiles containing minima of the physical potential will occur in those P -cubes for which the constraint surface Σ passes through the quadratic domain near the center of the cube. Such tiles lie within a connected region in \mathbf{m} -space; that is, a subset $\mathfrak{R} \subset \mathcal{C}$. (This compactness property only applies for the specific tiling and labeling of the tiles defined by the aligned coordinates.) The region \mathfrak{R} contains exponentially fewer lattice sites than \mathcal{C} , so the problem of comprehensively sampling \mathfrak{R} is much less computationally intensive, but still requires a number of computations that is exponential in P .

We can further simplify the problem if we are interested only in statistical properties of the potential for which a relatively small, but representative sample of distinct tilings suffices. Such a representative sample can be obtained by uniformly sampling over lattice sites contained within the polytope \mathcal{C} , for example by performing a random walk that samples the polytope in a time polynomial in $P - N$ [62–64]. A simpler but much more computationally intensive mechanism to uniformly sample a polytope would be to define a simple region that fully contains the polytope and sample that, rejecting any sample that is not contained in the polytope. Again, in order to determine a statistical sample of most of the *minima* it suffices to sample only the region \mathfrak{R} . The sampling techniques above apply for both the non-linear optimization problem of section 3.1 and the analytic result for the energies of the minima in the quadratic approximation, (3.11).

We illustrate the distribution of energies at the minima for a specific theory obtained via three different approaches in figure 8. The probability distribution of the energy density obtained by sampling a small number of all minima agrees well with the exact distribution. Furthermore, as expected, the quadratic approximation works best for relatively low minima, and becomes increasingly inaccurate for higher minima.




Figure 8. Histogram showing the number of minima over the energy density, for an example with $N = 5$, $P = 6$. Shaded region: the exact distribution of all minima; red/dark: the distribution obtained via random sampling; blue/light: the distribution obtained via quadratic approximation.

It is possible for the physical potential to have minima in tiles for which Σ does not intersect the quadratic domain (that is, tiles that are outside \mathfrak{R}). However such minima are rare, at least in the well-aligned regime $N \gg 1$, $P - N \ll N$. This can be understood qualitatively as follows. Each P -cube of sidelength 2π can be decomposed into various regions — the quadratic domain at the center $\|\phi\|_\infty < \pi/2$ (which contains a fraction 2^{-P} of the volume of the cube), surrounded by rectilinear regions defined by allowing some of the components of ϕ to exceed $\pi/2$ in magnitude. The Hessian of the auxiliary potential (2.2) is

$$(H_{\text{aux}})^I_J = \delta^I_J \Lambda_I^4 \cos(\phi^I). \tag{3.13}$$

This is positive definite precisely in the quadratic domain around the center. Because the Hessian of the physical potential is a projection of \mathbf{H}_{aux} onto Σ , any critical point of the physical potential that occurs in the quadratic domain must be a minimum. For each component of ϕ that lies outside the quadratic domain, the auxiliary Hessian matrix (3.13) has an additional negative eigenvalue. Critical points of the physical potential in such regions can be minima (rather than saddle points) only if the negative eigenvalue(s) of the Hessian are projected out when the potential is constrained to Σ . For $P - N \ll N$, only a small fraction $P - N/P$ of the P directions are projected out and therefore it is unlikely (or impossible if $P - N$ is less than the number of negative modes) that an critical point outside the quadratic domain will be a minimum. We have verified this expectation numerically. Therefore, in this regime the great majority of the exact minima lie within \mathfrak{R} .

3.5 Saddle points and maxima

The cosine function changes sign under a half-period shift, $\cos(\phi) = -\cos(\phi + \pi)$, so the physical characteristics of maxima mirror those of minima. Due to the “1”s in (1.1), the global maximum has energy $V_{\text{max}} \approx \sum_I 2\Lambda_I^4$. (As explained below, this would be an equality if the phases were exactly $\delta^I = \pi$, and is a very good approximation in well-aligned

theories.) All the techniques we apply to minima carry over to maxima and other critical points almost unchanged. In particular, the locations and energies of maxima and saddles can be found efficiently by using the approximate symmetries. In fact, the property of “well-alignedness” that allowed us to reduce the phases to values very close to zero allows us to set the phases to anything we like, subject to errors of order those in our original procedure. In other words, we can set the phases $\boldsymbol{\delta} = \boldsymbol{\delta}_{\text{arb}} + \mathcal{O}(\|\mathbf{P}^\perp \mathbf{T}_{\omega_j \phi}\|_\infty)$, where the P -vector $\boldsymbol{\delta}_{\text{arb}}$ denotes any arbitrary phases. Geometrically, this is possible because in well-aligned theories the angles between Σ and the grid in ϕ -space are small, so that Σ approaches very close to every distinct point in ϕ -space.

For studying maxima it is convenient to set all phases as close as possible to π . Relating ϕ and $\boldsymbol{\theta}$ in this way corresponds to shifting the centers of the P -cube tiling of ϕ -space so that they fall on global maxima rather than global minima — that is, the new cubes are centered on the corners of the original ones. With this change nearly every equation in this paper carries over unchanged or with the obvious changes from minima to maxima. In particular maxima have a quadratic domain defined in the same way as for minima in (3.5), as the cube of side-length π surrounding a now maximum of V_{aux} , and they have identical statistics for their number, energies (except subtracted from V_{max} rather than added to V_{min} , etc.

This trick of setting the phases to a desired value is also useful for studying saddle points of any given degree. For instance, to study saddles of degree one (critical points where the Hessian has 1 negative eigenvalue and $N - 1$ positive eigenvalues) we should set one phase equal to π and the rest as close as possible to zero. These points are the centers of the faces of the original cubes, and are points where the auxiliary Hessian (3.13) has precisely one negative mode (and the auxiliary potential has a degree one saddle). Tiles where Σ passes through the quadratic domain of these points often contain degree one saddles of the physical potential, and tiles that do not may not, for the same reason described in the previous subsection for the case of minima. Such degree one saddles occur between tiles that correspond to P -cubes that are neighbors along a face, and play a crucial role in the analysis of tunneling transitions (cf. [28]).

3.6 Estimates in random ensembles

In the previous section we discussed how to systematically enumerate and locate all minima of a given axion theory. We found an analytic expression for their energy densities in the quadratic approximation. We now turn to a discussion of the expected number and distribution of minima in ensembles of random axion theories.

We define these theories by ensembles of random integer charge matrices \mathcal{Q} and energy scales Λ_J^4 , as discussed in section 1. To repeat our assumptions, \mathcal{Q} is a $P \times N$ matrix of independent, identically distributed random integer entries with vanishing mean and standard deviation $\sigma_{\mathcal{Q}}$. We assume the universal limit of random matrix theory such that the precise distribution (including the fact that the entries of \mathcal{Q} are integer) becomes irrelevant and all expectation values only depend on $\sigma_{\mathcal{Q}}$. This assumption roughly holds whenever $\gtrsim 3/N$ of the entries of the charge matrix are non-vanishing, and the distribution

of the entries is not heavy-tailed. The field space metric is irrelevant for the discussion in this section.

In general it is a very difficult task to analytically obtain the distribution of minima, or even the number of stable minima. Even in the quadratic approximation this problem amounts to determining the number of lattice sites within a non-trivial high-dimensional polytope defined by (3.10). In this section we therefore mostly restrict our attention to the simplest case of one single auxiliary field, $P = N + 1$ and equal scales $\Lambda_I = \Lambda$, unless otherwise noted. We will find that the energy density (3.12) is valid for super-exponentially many minima.

3.6.1 The quantity and energies of minima for $P = N + 1$

We now determine the number of distinct minima \mathcal{N}_{vac} that are well-approximated by the quadratic expansion in (3.12). This count is simply given by the number of sites of the P -dimensional, rank $P - N = 1$ sublattice $\delta\phi_m = 2\pi\mathbf{P}^\perp\mathbf{T}_{\omega_{\#}\phi}m$ that are contained within the quadratic domain $\|\delta\phi_m\|_\infty \leq \pi/2$. When all phases in the original Lagrangian exactly vanish there exists a two-fold degeneracy of all minima. If the phases do not precisely vanish, they can be absorbed up to a finite remainder that is typically of order the change of $\delta\phi$ between similar minima, see section 2.5. To further simplify the problem we assume identical scales $\Lambda_I = \Lambda$. The number of distinct minima in the quadratic domain is then simply

$$\mathcal{N}_{\text{vac}} = \frac{1}{2\|\mathbf{P}^\perp\mathbf{T}_{\omega_{\#}\phi}\|_\infty}. \quad (3.14)$$

Note that since \mathcal{Q} has independent, identically distributed random entries, \mathbf{P}^\perp projects onto a random direction that is isotropically distributed, hence the vector $\mathbf{P}^\perp\mathbf{T}_{\omega_{\#}\phi}$ is isotropically distributed, with (E.2) its two-norm is given by

$$\|\mathbf{P}^\perp\mathbf{T}_{\omega_{\#}\phi}\|_2^2 = \frac{c^2}{\det\mathcal{Q}^\top\mathcal{Q}}, \quad (3.15)$$

and we defined the positive integer c as in (E.2). A vector that is distributed isotropically on the sphere consists of independent, normally distributed entries. Matching the expected norm of that vector to (3.15) then determines the distribution of the entries,

$$(\mathbf{P}^\perp\mathbf{T}_{\omega_{\#}\phi})^I \in \mathcal{N}\left(0, \frac{c}{\sqrt{P\det\mathcal{Q}^\top\mathcal{Q}}}\right), \quad \forall I \in \{1, \dots, P\}, \quad (3.16)$$

where $\mathcal{N}(0, \sigma)$ denotes a normal distribution of mean zero and standard deviation σ . It is now straightforward to evaluate the median of the largest absolute entry of $\mathbf{P}^\perp\mathbf{T}_{\omega_{\#}\phi}$, which yields the number of minima as

$$\mathcal{N}_{\text{vac}} \approx \frac{\sqrt{P}}{2\ell(P)}\sqrt{\det\mathcal{Q}^\top\mathcal{Q}}, \quad (3.17)$$

where $\ell(N) \equiv \sqrt{2}\text{erf}^{-1}(2^{-1/N})$ is the median largest absolute entry of an N -vector with entries that are unit normal distributed. In (3.17) we used that c is an order one integer, which we confirmed in extensive simulations for the ensembles under consideration.

The matrix $\mathcal{Q}^\top \mathcal{Q}$ is a real Wishart matrix, the determinant of which is distributed as the product of $P - N$ chi-squared random variables with $P, P - 1, \dots, P - N + 1$ degrees of freedom, respectively [65], which gives for our case

$$\langle \det \mathcal{Q}^\top \mathcal{Q} \rangle = \sigma_{\mathcal{Q}}^{2(P-1)} \frac{P!}{1!}. \quad (3.18)$$

Finally, we find a simple expression for the expected number of minima,

$$\mathcal{N}_{\text{vac}} \approx \frac{\sigma_{\mathcal{Q}}^{P-1}}{2\ell(P)} \sqrt{PP!}, \quad (3.19)$$

which is exponentially large in N in the universal regime, where at least a fraction $3/N$ of the entries in \mathcal{Q} are non-vanishing. To give a sense of these numbers, with $P = N + 1 = 150$, and $\sigma_{\mathcal{Q}} = 1$, one obtains $\mathcal{N}_{\text{vac}} \approx 10^{131}$. Note that this result was only derived for $P - N = 1$, but as we discuss below we expect similar results to hold more generally (see (3.30)). The scaling with N is identical to that observed in [26] for a specific case where $P \gg N$.

Finally, let us estimate the energy levels at which the minima arise. To that end, we will assume that the distribution of minima can be well-approximated by all lattice sites that lie within the quadratic domain, i.e. $|\delta\phi^I| \leq \pi/2$. Since the displacements $\delta\phi$ are proportional to $\mathbf{P}^\perp \mathbf{T}_{\omega\phi}$, which by (3.16) is roughly normal distributed, we can easily estimate the typical magnitude of the entries of the displacement vector, when the largest component is $\pi/2$,

$$|\delta\phi^I| \approx \frac{\pi}{2\ell(P)}. \quad (3.20)$$

The maximum energy density in a minimum is therefore well-approximated by²⁰

$$\max (V_{\text{vac}}) \approx \frac{1}{2} \left(\frac{\pi}{2\ell(P)} \right)^2 \langle V \rangle \approx 0.14 \langle V \rangle, \quad (3.21)$$

where we used $\ell(N) \approx 3$ for $N \sim 100$ in the last approximation, and used the mean of the potential $\langle V \rangle = \sum_{I=1}^P \Lambda_I^4$. Since the potential is simply quadratic the median energy density is given by

$$\text{median} (V_{\text{vac}}) \approx \frac{1}{4} \max (V_{\text{vac}}) \approx 0.034 \langle V \rangle. \quad (3.22)$$

3.6.2 Hessian eigenvalues

Beyond their energies, another interesting characteristic of critical points is the spectrum of eigenvalues of the Hessian. If the kinetic matrix is $K_{ij} = f^2 \delta_{ij}$, the canonically normalized fields are $\Theta \equiv f\theta$. Defining $\mathbf{Q} \equiv \mathbf{K}^{-1/2} \mathcal{Q} = \mathcal{Q}/f$, the potential in canonically normalized coordinates is

$$V(\Theta) = \sum_{I=1}^P \Lambda_I^4 \left[1 - \cos(\mathbf{Q}\Theta)^I \right]. \quad (3.23)$$

The eigenvalues of the Hessian of this potential at a critical point are then the masses of the canonical fields at that point. In section 4 we will perform a more detailed analysis

²⁰Note that this expression applies for general P .

of the size of the tiles surrounding minima and the masses of the canonically normalized fields for various less trivial choices of kinetic matrix.

Expanded around a minimum labeled by \mathbf{m} , the Hessian of (3.23) is

$$\mathbf{H} = \mathbf{Q}^\top \text{diag}(\Lambda_I^4) \mathbf{Q} + \mathcal{O}[(\mathbf{P}^\perp \mathbf{m})^2]. \quad (3.24)$$

For simplicity let us take all $\Lambda_I = \Lambda$. In that case $\mathbf{Q}^\top \mathbf{Q}$ is a Wishart matrix, and at large N the empirical density (i.e. the amount of eigenvalues in a small interval) follows the Marchenko-Pastur distribution [66]. For a Wishart matrix with standard deviation 1, the mean smallest eigenvalue is $\mathcal{O}(1/N)$ while the largest is $\mathcal{O}(N)$ (if $P = N + 1$ the precise values are $1/4N$ and $4N$, respectively). The Marchenko-Pastur distribution has a sharp peak near the minimum and a long tail to larger values. The mean and median are both of order N . Putting the dimensions back in, this means the masses will range from

$$\frac{\sigma_{\mathbf{Q}}^2}{N} \left(\frac{\Lambda^2}{f}\right)^2 \lesssim m^2 \lesssim N \sigma_{\mathbf{Q}}^2 \left(\frac{\Lambda^2}{f}\right)^2.$$

As mentioned previously, the Hessian on a degree one (one negative mode) saddle is of interest for questions involving tunneling from one minimum to another [28]. Such saddles are most easily analyzed by setting one phase to π and the rest to zero. This corresponds to considering points where Σ passes close to the center of one face of the P -cube (note that such points are degree one saddles of V_{aux}). Since \mathbf{Q} is isotropic it does not matter which phase we set to π . Choosing the first one, it is easy to see that (3.24) becomes

$$\mathbf{H} \approx \Lambda^4 \mathbf{Q}^\top \text{diag}(-1, 1, 1, \dots) \mathbf{Q}. \quad (3.25)$$

This is not a Wishart matrix and we are unaware of any analytic results for its eigenvalue spectrum. It has at most one negative eigenvalue. If the smallest eigenvalue λ_{\min} turns out to be positive, this critical point is in fact a minimum rather than a saddle. Numerically we established that the mean and standard deviation of the minimum eigenvalue are

$$\langle \lambda_{\min}(\mathbf{H}) \rangle \approx -\frac{N}{2} \sigma_{\mathbf{Q}}^2 \left(\frac{\Lambda^2}{f}\right)^2, \quad \sigma_{\lambda_{\min}(\mathbf{H})} \approx \sqrt{\frac{3}{2N}} |\langle \lambda_{\min}(\mathbf{H}) \rangle|.$$

Hence, at large N is extremely unlikely that there is no negative eigenvalue and the would-be degree one saddle is actually a minimum. This at least partially confirms the expectation explained in section 3.4, that most local minima occur in tiles where the constraint surface intersects the quadratic domain of the minimum of V_{aux} , rather than in neighboring regions such as these. Similarly most saddles of degree k will occur in regions where Σ intersects the quadratic domain of a degree k saddle of V_{aux} . The relation between the number of saddles $\mathcal{N}(k)$ of degree k and the number of minima can then be estimated from (3.25):

$$\mathcal{N}(k) \approx \binom{P}{k} \mathcal{N}_{\text{vac}}.$$

3.6.3 Neighboring minima

In the previous section we estimated the total number of distinct, non-degenerate minima in the entire potential. For some questions one might however only be interested in the immediate neighborhood of one particular minimum. We therefore turn to determining the expected number of immediate neighboring minima, including degenerate ones. To allow for a simple estimate, consider all sites neighboring the origin, $\mathbf{n} = \mathbf{0} + \mathbf{e}$, with unit or zero entries for \mathbf{e} as in (3.3). Let us count the immediate neighbors that lead to a minimum within the quadratic domain (3.5),

$$\|2\pi \mathbf{P}^\perp \mathbf{e}\|_\infty \leq \frac{\pi}{2}. \quad (3.26)$$

We verified numerically that in the universal limit the entries of the orthogonal projector matrix \mathbf{P}^\perp have variance $(P - N)/P^2$. Using the central limit theorem and denoting the number of non-vanishing entries of \mathbf{e} by n_e , we can approximate

$$\left(\mathbf{P}^\perp \mathbf{e}\right)^I \in \mathcal{N}\left(0, \sqrt{n_e(P - N)}/P\right). \quad (3.27)$$

The median largest entry of $\mathbf{P}^\perp \mathbf{e}$ evaluates to $\ell(P)\sqrt{n_e(P - N)}/P$, such that for $P \lesssim 2N$ a large fraction of the 3^P neighbors are stable minima, i.e. $\|\mathbf{P}^\perp \mathbf{e}\|_\infty \ll \pi$. Therefore, not only is the total number of minima extremely large, but each minimum has a vast number of immediately neighboring minima.

3.6.4 Phases

As discussed in section 2.5, in well-aligned axion theories the N exact and $P - N$ approximate shift symmetries allows one to set N phases to precisely to zero and make the remaining $P - N$ phases very small. The accuracy to which all phases can be eliminated, as measured by the largest remaining phase δ_{\max} , depends on how aligned the basis is:

$$\delta_{\max} \leq \pi \left\| \mathbf{P}^\perp \mathbf{T}_{\omega \parallel \phi} \right\|_\infty \leq \pi(P - N) \max_a \left\{ \|\mathbf{P}^\perp \mathbf{t}_a^\parallel\|_\infty \right\}. \quad (3.28)$$

In order to get some analytical intuition for how well-aligned theories in our ensemble tend to be, let us assume that the vectors $\mathbf{P}^\perp \mathbf{t}_a^\parallel$ are orthogonal and are the shortest they could possibly be, as in (3.15), and that the volume of the cubic, but arbitrarily oriented periodic domain of the lattice generated by \mathbf{P}^\perp is given by $(\det \mathbf{Q}^\top \mathbf{Q})^{-1/2}$. These assumptions yield the distribution for the components of all projections,

$$\left(\mathbf{P}^\perp \mathbf{t}_a^\parallel\right)^I \in \mathcal{N}\left(0, \frac{(\det \mathbf{Q}^\top \mathbf{Q})^{-\frac{1}{2(P-N)}}}{\sqrt{P}}\right), \quad \forall I \in \{1, \dots, P\}, \quad a \in \{1, \dots, P - N\}. \quad (3.29)$$

This gives an upper bound on the number of minima in the quadratic domain,

$$\mathcal{N}_{\text{vac}} \lesssim \left(\frac{\sqrt{P}}{2\ell(P)}\right)^{P-N} \sqrt{\det \mathbf{Q}^\top \mathbf{Q}} \approx \left(\frac{\sqrt{P}}{2\ell(P)}\right)^{P-N} \sigma_{\mathbf{Q}}^N \sqrt{\frac{P!}{(P - N)!}}, \quad (3.30)$$

reproducing (3.19) in the special case $P = N + 1$.




Figure 9. Ensemble average of largest phase (3.28) as a function of P , with $N = 20$ and $\sigma_{\mathcal{Q}}^2 = 2/3$. The solid line is the estimate (3.31).

We can furthermore easily obtain the largest components of the orthogonal projections of the aligned lattice basis,

$$\max_a \left\{ \|\mathbf{P}^\perp \mathbf{t}_a^\# \|_\infty \right\} \sim \frac{\ell(P[P-N])}{\sqrt{P}(\det \mathcal{Q}^\top \mathcal{Q})^{\frac{1}{2(P-N)}}} \approx \frac{\ell(P[P-N])}{\sqrt{P}} \left(\frac{1}{\sigma_{\mathcal{Q}}^{2N}} \frac{(P-N)!}{P!} \right)^{\frac{1}{2(P-N)}}. \quad (3.31)$$

Using Stirling's approximation we observe that when random matrix universality applies the basis is well-aligned for $P \approx N$, and for order unity $\sigma_{\mathcal{Q}}$ the basis ceases to be well-aligned with growing P at $P \lesssim 2N$. We illustrate how the largest phase increases with the number of non-perturbative terms along with the analytic estimate (3.31) in figure 9.

3.7 Band structure of subleading terms

Finally, let us address the last feature of the axion Lagrangian that we ignored so far; the subleading terms in the non-perturbative potential, denoted only by ellipses in (1.1). Explicitly, we have the full axion potential

$$V = \sum_{I=1}^P \Lambda_I \left[1 - \cos(\mathcal{Q}\theta)^I \right] + V_{\text{sl}}(\theta), \quad (3.32)$$

where we introduced a subleading potential $-\Lambda_{\text{sl}}^4 \leq V_{\text{sl}} \leq \Lambda_{\text{sl}}^4$ of scale Λ_{sl}^4 that is negligible compared to the leading P terms in the non-perturbative potential. Remember that we chose coordinates such that $\theta^i \rightarrow \theta^i + 2\pi$ are the discrete shift symmetries respected by the full theory, such that also $V_{\text{sl}}(\theta)$ breaks any larger symmetries respected by the P leading terms to those fundamental symmetries. If there are any shift symmetries respected by the P leading terms that are broken by the subleading potential this will result in a multiplicative increase in the number of distinct minima, as discussed in [26]. This effect is related to, but distinct from the mechanism discussed thus far.

The leading potential is invariant under the P shifts $(\mathcal{Q}\theta)^I \rightarrow (\mathcal{Q}\theta)^I + 2\pi$, which generates an N -dimensional lattice denoting the shift symmetries in terms of the θ -coordinates.




Figure 10. *Top:* constraint surface (red line) along with the lattice $2\pi\mathbb{Z}^P$ (gray dots). Distinct tiles of the leading cosine terms are numbered and shaded dark, while the periodicity of the subleading terms is indicated by the light shading. *Bottom:* axion potential. The subleading terms further split the energies in the minima of the leading terms.

In the notation of section 2 a basis for this lattice is given with (2.22) by $\mathcal{B} = \mathbf{T}_{\omega_{\parallel}\theta}$, i.e. the leading potential is (minimally) invariant under the N shifts $\theta \rightarrow \theta + 2\pi\mathcal{B}_i$. The subleading potential, however, is only invariant under shifts on the integer lattice $2\pi\mathbb{Z}^N$. This means that the periodic domain of the full potential contains $\mathcal{N}_{\text{sl}} = 1/\sqrt{\det \mathcal{B}^T \mathcal{B}}$ periodic domains of the leading potential (as this is the inverse volume of that domain). If the leading P terms in the non-perturbative potential contain $\mathcal{N}_{\mathcal{Q}}$ minima, each of these minima degenerates into \mathcal{N}_{sl} distinct minima due to the further symmetry breaking in the subleading potential. The total number of minima therefore becomes

$$\mathcal{N}_{\text{vac}} = \mathcal{N}_{\text{sl}} \times \mathcal{N}_{\mathcal{Q}}. \tag{3.33}$$

We illustrate this energy level splitting in figure 10.

Of course we could have included the charges of the subleading terms in the P rows of the leading potential and found the corresponding aligned basis that includes all possible charges in the theory. However, the subleading potential is irrelevant for all practical purposes when identifying approximate shift symmetries of the potential and therefore would only introduce a spurious complication of the computational problem by increasing the dimensionality P of the auxiliary lattice. When identifying the shift symmetries according to section 2 it is therefore important to identify which terms in (1.1) can safely be ignored for a problem at hand.

4 Aligned axion diameters

In this section we provide a systematic discussion of the theory in the vicinity of local minima; that is, within the tiles \mathcal{T}_n . Recall that the tiles are defined as regions within

which none of the individual terms in the potential exceeds its maximum (see section 1), and therefore define the characteristic scale on which the potential changes. Within each of these tiles the potential is relatively flat and hence provides for a natural environment to study large field inflation. Several specific cases were previously studied in the literature [19, 20, 35, 47–50, 67–76]. In this section we describe a systematic approach to determine the size of an arbitrary tile. We restrict our discussion to well-aligned theories (cf. section 2) where we can set all P phases in (1.1) to zero to good accuracy by a shift of θ .

Recall the Lagrangian (1.1) of a well-aligned axion theory,

$$\mathcal{L} = \frac{1}{2} \partial \theta^\top \mathbf{K} \partial \theta - \sum_{I=1}^P \Lambda_I^4 \left[1 - \cos(\mathbf{Q}\theta)^I \right], \quad (4.1)$$

where we retain only $P \geq N$ leading terms in the potential. In this section we are interested in invariant field space distances, so it is convenient to introduce canonically normalized fields Θ ,

$$\Theta \equiv \sqrt{\mathbf{K}} \theta, \quad (4.2)$$

where $\sqrt{\mathbf{K}}$ is the positive matrix square root.²¹ The Lagrangian in canonically normalized coordinates reads

$$\mathcal{L} = \frac{1}{2} \partial \Theta^\top \partial \Theta - \sum_{I=1}^P \Lambda_I^4 \left[1 - \cos(\mathbf{Q}\Theta)^I \right], \quad (4.3)$$

where the canonical charges are related to the integer charges by

$$\mathbf{Q} \equiv \mathcal{Q} \mathbf{K}^{-1/2}. \quad (4.4)$$

In canonically normalized coordinates the tiles (2.7) are given by

$$\mathcal{T}_n = \{ \Theta : \|\mathbf{Q}\Theta - 2\pi \mathbf{n}\|_\infty \leq \pi \}. \quad (4.5)$$

The tiles are polytopes in N dimensions defined by the intersection of $2P$ half-planes,²² and spherical shells determine the surfaces of constant invariant distance to the center of the sphere. We illustrate this polytope in figure 11.

4.1 Diameters and field ranges in well-aligned theories

To characterize the scale of these domains we will consider two distinct measures of size. One is the *diameter* \mathcal{D}_n , that is, the length of the longest straight line contained in the tile \mathcal{T}_n . It is clear that this line will run between two vertices of the polytope, and that the tile with the largest diameter is the one at the origin $\mathbf{n} = \mathbf{0}$. If we denote the vertices of the polytope by $\mathbf{d}_{n,l}$, we have therefore the corresponding diameter

$$\mathcal{D}_n = \max_{l,k} \|\mathbf{d}_{n,l} - \mathbf{d}_{n,k}\|_2 \leq \mathcal{D}_0 = 2 \max_l \|\mathbf{d}_{0,l}\|_2. \quad (4.6)$$

²¹The matrix square root satisfies $\sqrt{\mathbf{K}}\sqrt{\mathbf{K}} = \mathbf{K}$, and is related to the matrix $\mathbf{S}_\mathbf{K}$ containing the (column) eigenvectors of \mathbf{K} and its eigenvalues f_i^2 by $\sqrt{\mathbf{K}} = \mathbf{S}_\mathbf{K} \text{diag}(f_i) \mathbf{S}_\mathbf{K}^\top$.

²²Note that \mathcal{T}_n is indeed a polytope, i.e. a finite volume subset of \mathbb{R}^N bounded by hyperplanes of codimension one, since \mathbf{Q} is full rank.




Figure 11. Illustration of the tile $\mathcal{T}_{(30)\tau}$ of the potential shown in figure 3 in canonical fields Θ . The solid arrows denote the field ranges \mathcal{R}_{\pm} along the two lightest directions in the vicinity of the minimum, while the dashed arrow denotes the diameter \mathcal{D} of the tile.

Note that the diameters of the tiles depend only on the charge and kinetic matrices of the theory (not on the couplings Λ_J^4).

The expression (4.6) defines a unique size for every tile, but in practice there are exponentially many vertices, making it hard to evaluate. Furthermore, the low energy physics in the vicinity of a minimum is generally dominated by the lightest degrees of freedom, which do not necessarily coincide with the axis of largest diameter. This motivates our second characterization of the scale (which does depend on the couplings), namely the *field range* $\mathcal{R}_{n\pm}$ within \mathcal{T}_n along the lightest direction — the line defined by the eigenvector $\hat{\Psi}_H$ of smallest eigenvalue of the Hessian H at the minimum in the tile. The Hessian is given by

$$H = \mathbf{Q}^\top \text{diag}(\Lambda_J^4) \mathbf{Q} + \mathcal{O}[(\mathbf{P}^\perp \mathbf{n})^2]. \quad (4.7)$$

More precisely we define two field ranges $\mathcal{R}_n(\pm \hat{\Psi}_H) \equiv \mathcal{R}_{n\pm}$ as the canonically normalized field space distance between a minimum at Θ_n and the boundary of the corresponding tile in the least massive directions $\pm \hat{\Psi}_H$. Solving the equation defining the boundary of the tile,

$$\left\| \pm \mathcal{R}_{n\pm} \mathbf{Q} \hat{\Psi}_H + \text{mod}_{[-\pi, \pi]} \mathbf{Q} \Theta_n \right\|_\infty = \pi, \quad (4.8)$$

yields the field ranges

$$\mathcal{R}_n(\pm \hat{\Psi}_H) = \min_I \left\{ \frac{\pi}{|(\mathbf{Q} \hat{\Psi}_H)^I|} \mp \frac{\text{mod}_{[-\pi, \pi]} (\mathbf{Q} \Theta_n)^I}{(\mathbf{Q} \hat{\Psi}_H)^I} \right\}. \quad (4.9)$$

Note that (4.9) holds for the field range $\mathcal{R}_n(\hat{\Theta})$ along an arbitrary direction $\hat{\Theta}$.

In well-aligned theories the sizes of exponentially many tiles \mathcal{T}_n are well-approximated by the size of the tile \mathcal{T}_0 containing the origin $\Theta = \mathbf{0}$ (or, when $P = N$, this is the only tile), so we will focus our attention on this last tile in the remainder of this section. In \mathcal{T}_0 the expressions for the diameter and field ranges simplify. For any unit N -vector $\hat{\Theta}$ (in

particular the lightest directions $\pm \hat{\Psi}_{\mathbf{H}}$) we have

$$\mathcal{R}_0(\hat{\Theta}) = \mathcal{R}_0(-\hat{\Theta}) = \frac{\pi}{\|\mathbf{Q}\hat{\Theta}\|_\infty}, \quad (4.10)$$

which indeed follows from the general expression (4.9). The diameter can alternatively be expressed as

$$\mathcal{D}_0 = \max \left\{ 2\mathcal{R}_0(\hat{\Theta}) \mid \hat{\Theta} \in S^{N-1} \right\} = \max \left\{ \frac{2\pi}{\|\mathbf{Q}\hat{\Theta}\|_\infty} \mid \hat{\Theta} \in S^{N-1} \right\}. \quad (4.11)$$

This last expression for the diameter can be used to derive bounds on it in an arbitrary theory, namely²³

$$\frac{2\pi}{\lambda_{\min}(|\mathbf{Q}|)} < \mathcal{D}_0 \leq \frac{2\pi\sqrt{P}}{\lambda_{\min}(|\mathbf{Q}|)}, \quad (4.12)$$

where $\lambda_{\min}(|\mathbf{Q}|)$ denotes the smallest eigenvalue of the matrix $|\mathbf{Q}| \equiv \sqrt{\mathbf{Q}^\top \mathbf{Q}}$, i.e. it is the smallest singular value of \mathbf{Q} .

4.2 N-flation, lattice and kinetic alignment

In the previous sections we defined two notions of size of the axion field space in the vicinity of minima. To illustrate these definitions we now apply them to three special cases, focussing in particular on how the diameter in multi-axion theories may be enhanced compared to the single-axion diameter $2\pi f$. We will consider N-flation [19], lattice²⁴ (or KNP) alignment [20], and finally kinetic alignment [47]. Each of these models was originally restricted to $P = N$ non-perturbative terms, but we will generalize the main ideas behind lattice and kinetic alignment to well-aligned theories with $P \geq N$. As mentioned in section 4.1 in well-aligned theories it suffices to consider the representative tile \mathcal{T}_0 around the origin, which we will do in the following.

4.2.1 N-flation

As our first example we consider the N -axion theory with diagonal kinetic matrix $\mathbf{K} = \text{diag}(f_I^2)$ and $P = N$ trivial charges $\mathcal{Q} = \mathbb{1}_N$ [19]. In terms of canonical coordinates the Lagrangian is given by

$$\mathcal{L} = \frac{1}{2} \partial \Theta^\top \partial \Theta - \sum_{I=1}^N \Lambda_I^4 \left[1 - \cos \left(\frac{\Theta^I}{f_I} \right) \right]. \quad (4.13)$$

This theory has only one distinct tile and has one minimum at the origin $\Theta = \mathbf{0}$, as illustrated in figure 12. The mass matrix is diagonal, $\mathbf{H} = \text{diag}(\Lambda_I^4/f_I^2) \propto \mathbb{1}_N$, where for simplicity we selected the scales Λ_I such that all masses are equal.

²³See appendix F for a short derivation.

²⁴The term “lattice alignment” is used synonymous with “KNP alignment” and should not be confused with the “aligned lattice basis” introduced in section 2.2. The two terms refer to unrelated mechanisms.




Figure 12. Illustration of the tile \mathcal{T} for N -flation. The diameter is the Pythagorean sum of the f_I .

The tile consists of an N -dimensional hyperrectangle with side-lengths $2\pi f_I$, which yields the diameter as the Pythagorean sum

$$\mathcal{D} = 2\pi \sqrt{\sum_I f_I^2}. \tag{4.14}$$

For fixed $f_{\max} \equiv \max_I \{f_I\}$, the largest possible diameter is obtained when all metric eigenvalues are equal, $f_I = f_{\max}$: $\mathcal{D} = 2\pi\sqrt{N}f_{\max}$. By contrast, if there are large hierarchies in the f_I the diameter is $\mathcal{D} \gtrsim 2\pi f_{\max}$. Since all directions are equally massive, the lightest direction is degenerate and the field ranges accessible from the minimum at the origin are just half of the diameter, $\mathcal{R}_{\pm} = \mathcal{D}/2$. In the cosmological context this scenario is known as *N-flation*, a particular realization of assisted inflation [77]: while none of the individual fields Θ^I traverse a displacement larger than f_{\max} , the simultaneous displacement of N fields realizes an invariant field range parametrically as large as $\sqrt{N}f_{\max}$.

4.2.2 Lattice alignment

We now consider the *lattice alignment* (or KNP) mechanism, first discussed by Kim, Nilles and Peloso [20] for the special case $N = P$. Lattice alignment relies on a small singular value of the charge matrix. For simplicity, we assume a kinetic matrix proportional to the identity, $\mathbf{K} = f^2 \mathbf{1}_N$, while the $P \geq N$ integer charges are left general,

$$\mathcal{L} = \frac{1}{2} \partial \Theta^\top \partial \Theta - \sum_{I=1}^P \Lambda_I^4 \left[1 - \cos \frac{(\mathcal{Q}\Theta)^I}{f} \right]. \tag{4.15}$$

Recall the general bound (4.12) on the diameter of the tile \mathcal{T}_0 ,

$$\mathcal{D}_0 > \frac{2\pi f}{\lambda_{\min}(|\mathcal{Q}|)}. \tag{4.16}$$

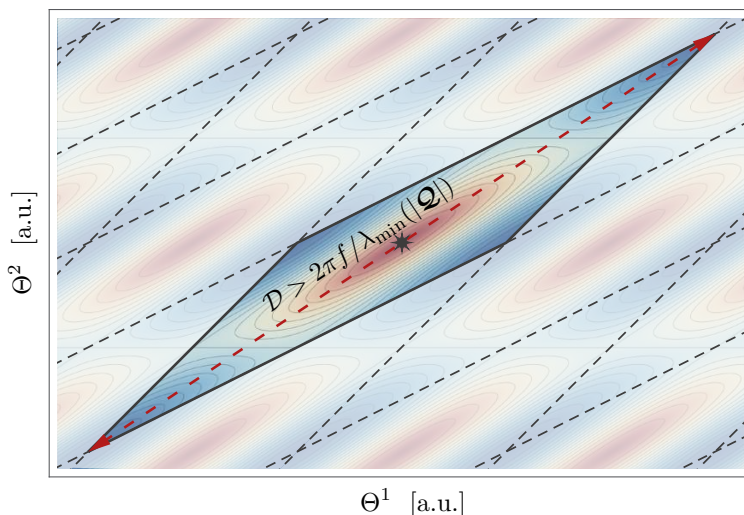


Figure 13. Illustration of the tile \mathcal{T} for lattice alignment. The diameter is enhanced by the inverse of the smallest singular value of the canonical charge matrix.

Lattice alignment is the observation that one can arbitrarily enhance the field range in these theories compared to the single-field $2\pi f$ by decreasing the smallest eigenvalue λ_{\min} of $|\mathcal{Q}| = \sqrt{\mathcal{Q}^\top \mathcal{Q}}$.²⁵ As λ_{\min} decreases, the lower bound on the diameter increases. We illustrate this phenomenon in figure 13.

4.2.3 Kinetic alignment

Finally let us discuss models with *kinetic alignment* [47]. In the original discussion one assumed $P = N$, a trivial charge matrix $\mathcal{Q} = \mathbb{1}_N$ and a general kinetic matrix \mathbf{K} , but the definition of kinetic alignment can just as easily be given in the more general context of well-aligned theories with $P \geq N$ and \mathbf{K}, \mathcal{Q} unspecified. Note that the upper bound in (4.12) is saturated if

$$\|\mathbf{Q}\hat{\Psi}_{|\mathcal{Q}|\infty} = \frac{1}{\sqrt{P}}\|\mathbf{Q}\hat{\Psi}_{|\mathcal{Q}|}\|_2, \tag{4.17}$$

in other words, if the direction defined by $\mathbf{Q}\hat{\Psi}_{|\mathcal{Q}|}$ aligns with a diagonal of the P -cube, and $\hat{\Psi}_{|\mathcal{Q}|}$ denotes the eigenvector of $|\mathcal{Q}|$ with smallest eigenvalue. For this to be possible it is in particular necessary that the constraint surface Σ contains a diagonal of the P -cube. The sufficient condition (4.17) to saturate the upper bound in (4.12) is the extension of the original kinetic alignment proposal to arbitrary well-aligned theories with $P \geq N$. In models with (perfect) kinetic alignment we therefore have a diameter

$$\mathcal{D} = \frac{2\pi}{\|\mathbf{Q}\hat{\Psi}_{|\mathcal{Q}|\infty}} = \frac{2\pi\sqrt{P}}{\lambda_{\min}(|\mathcal{Q}|)}. \tag{4.18}$$

One might naively expect that alignment with a diagonal requires a great amount of fine-tuning in the canonical charge matrix \mathcal{Q} . In large dimensions $N, P \gg 1$, however, the

²⁵Practically this can be achieved by having some columns of \mathcal{Q} be nearly degenerate, meaning that their normalized variants have an overlap nearly equal to one (recall that \mathcal{Q} is integer-valued).




Figure 14. Illustration of the tile \mathcal{T} for kinetic alignment in the original model where $P = N$, $\mathbf{Q} = \mathbb{1}_N$. The diameter is enhanced by \sqrt{N} relative to f_{\max} .

converse is true: a P -cube has many more vertices (2^P) than faces ($2P$). Therefore, an isotropically oriented vector within an isotropically oriented constraint surface Σ is much more likely be pointing towards a vertex of a P -hypercube than towards a face. In section 4.3 we will make this expectation more precise and demonstrate that in broad classes of random axion theories the relation (4.17) is indeed approximately satisfied, and kinetic alignment is generic.

For completeness let us consider the original model where $P = N$, $\mathbf{Q} = \mathbb{1}_N$. Here $\mathbf{Q} = \mathbf{K}^{-1/2}$, and the eigenvector $\hat{\Psi}_{\mathbf{K}^{-1/2}}$ is equal to the eigenvector corresponding to the largest eigenvalue f_{\max}^2 of \mathbf{K} . The vector $\mathbf{Q}\hat{\Psi}_{\mathbf{K}^{-1/2}} = \hat{\Psi}_{\mathbf{K}^{-1/2}}/f_{\max}$ points towards a diagonal when $\hat{\Psi}_{\mathbf{K}^{-1/2}}$ does, which is the condition for (perfect) kinetic alignment as given in [47]. In this case

$$\mathcal{D} = 2\pi\sqrt{N}f_{\max}. \tag{4.19}$$

Note that here the diameter only depends on the largest metric eigenvalue f_{\max}^2 , and is independent of all other $f_I \leq f_{\max}$. When there are large hierarchies in the metric eigenvalues, the diameter (4.19) in kinetically aligned theories is larger by a factor of \sqrt{N} relative to the diameter (4.14) of the N-flation scenario. The enhancement of the diameter by \sqrt{N} relative to f_{\max} was originally referred to as kinetic alignment, which we illustrate in figure 14. More generally we have (4.18): the enhancement of the diameter by \sqrt{P} relative to the inverse of the smallest singular value of \mathbf{Q} .

4.3 Alignment in random ensembles

We now discuss diameters and field ranges in ensembles of random axion theories.²⁶ The Lagrangian is given by

$$\mathcal{L} = \frac{1}{2}\partial\theta^\top \mathbf{K}\partial\theta - \sum_{I=1}^P \Lambda_I^4 \left[1 - \cos(\mathbf{Q}\theta)^I \right], \tag{4.20}$$

²⁶This was previously considered in [35], but important aspects were missed that we discuss here.

and we study random ensembles of kinetic matrices \mathbf{K} , integer charges \mathbf{Q} , and dynamical scales Λ_f^4 introduced in section 1 and used already in section 3.6. We work at $2N \gtrsim P \geq N \gg 1$, where the theory (1.1) is generically very well-aligned so that it is consistent set the phases to zero in (4.20). Furthermore as in the previous section we restrict our attention to the tile around the global minimum at $\Theta = \mathbf{0}$, because at least for minima in the quadratic regime (see section 3.4) the diameters and field ranges along particular directions are similar up to $\mathcal{O}(1)$ factors.

Before discussing the details we first briefly review the main results for diameters in random axion theories. With (4.18) the diameter of a tile in a well-aligned theory is given by

$$\mathcal{D} \approx \frac{2\pi\sqrt{P}}{\lambda_{\min}(|\mathbf{Q}|)}, \quad (4.21)$$

where, again, $|\mathbf{Q}| = \sqrt{\mathbf{Q}^\top \mathbf{Q}}$. The diameter is enhanced by \sqrt{P} due to the fact that a random P -vector is very likely to be aligned with a vertex of the P -cube periodic domain of the auxiliary lattice rather than with one of its faces (kinetic alignment, see section 4.2.3), as well as by $1/\lambda_{\min}(|\mathbf{Q}|)$ (lattice alignment, cf. section 4.2.2).

In the universal limit the matrix $\mathbf{Q}^\top \mathbf{Q}$ resembles a Wishart matrix so we expect its eigenvalue distribution to depend only on P , N and the scale of the charges. In the simple case of $\mathbf{K} = f^2 \mathbf{1}_N$ and random \mathbf{Q} this scale is $(\sigma_{\mathbf{Q}}/f)^2$. We can substitute a naive random matrix theory expectation [78] for the smallest eigenvalue in (4.21) and obtain

$$\mathcal{D} \approx 2\pi\sqrt{P} \frac{f}{\sigma_{\mathbf{Q}}} \frac{1}{\sqrt{P}(1 - \sqrt{N/P})} \approx 2\pi\sqrt{P} \frac{f}{\sigma_{\mathbf{Q}}} \frac{2\sqrt{N}}{P - N}, \quad (4.22)$$

where the last approximate equality is valid when $P - N \ll N$.

For aligned theories where $P \approx N$, there are three parametric enhancements that each can scale as $\sim \sqrt{N}$. The first factor of $\sqrt{P} \approx \sqrt{N}$ in (4.22) is due to kinetic alignment and depends on the fact that the canonical charge matrix is isotropic. The second factor can arise from the sparsity of the charge matrix, encoded in $\sigma_{\mathbf{Q}}$, which may be as small as $\approx 3/\sqrt{N}$, while retaining universality. The last factor is due to the eigenvalue distribution of a Wishart matrix and leads to two different parametric scalings: when $P - N = \text{constant}$ and N is large, we have a third parametric enhancement of \sqrt{N} , while for $P - N \propto N$ and N large the third term decreases the diameter parametrically as $1/\sqrt{N}$. Using the least possible entries in the integer charge matrix we therefore have the following scaling with N :

$$\mathcal{D} \lesssim \begin{cases} N^{3/2} f, & \text{for } P - N = \text{constant}, \\ N^{1/2} f, & \text{for } P - N \propto N, \end{cases} \quad (4.23)$$

both valid when N is sufficiently large.

Even though these naive expectations are very crude, they turn out to accurately represent the mean diameter in a broad class of random models as we show below. Furthermore, we will find that the field range $\mathcal{R}_{\mathbf{0}\pm}$ along the lightest direction scales with N in a manner very similar to the diameter, and that this scaling is robust even when there are large hierarchies present in the dynamical scales Λ_f^4 . The simple expectation (4.23) from

random matrix theory can then be compared to fundamental theories with axions, such as compactifications of string theory [10].

Finally, via (4.21) the results (4.22) and (4.23) can also be applied to determine the scaling of the smallest eigenvalue m^2 of the Hessian matrix (4.7) — that is, the mass-squared of the lightest field around a minimum. Up to $\mathcal{O}(1)$ factors and with all $\Lambda_I = \Lambda$ equal,

$$m \approx 2\pi\sqrt{P} \frac{\Lambda^2}{\mathcal{D}}. \tag{4.24}$$

4.3.1 Diameter estimates

To estimate the diameter in random axion theories, recall from section 4.2.3 that perfect kinetic alignment implies that the diameter \mathcal{D}_0 of the tile is given by twice the field range along the eigenvector $\hat{\Psi}_{|\mathbf{Q}|}$. In the random theories we introduced (modulo a caveat on the kinetic matrices \mathbf{K} that we will discuss) we claim that kinetic alignment is well-satisfied at large P . More precisely

$$\frac{2\pi}{\|\mathbf{Q}\hat{\Psi}_{|\mathbf{Q}|}\|_\infty} \approx \frac{2\pi\sqrt{P}}{\ell(P)} \frac{1}{\|\mathbf{Q}\hat{\Psi}_{|\mathbf{Q}|}\|_2} \tag{4.25}$$

is satisfied with ever-increasing probability as $P \rightarrow \infty$.²⁷ This hinges on the following fact: the images of eigenvectors of $\mathbf{Q}^\top \mathbf{Q}$ under \mathbf{Q} are uniformly distributed on the unit P -sphere and therefore their entries are approximately normally distributed. In other words they are *delocalized*.²⁸ The asymptotic exactness of the relation (4.25) as $P \rightarrow \infty$ provides us with a reliable lower bound on the diameter \mathcal{D}_0 at large P ,

$$\mathcal{D}_0 \geq 2\mathcal{R}_0(\hat{\Psi}_{|\mathbf{Q}|}) \rightarrow \frac{2\pi\sqrt{P}}{\ell(P)\sigma_{\mathbf{Q}} \lambda_{\min}(|\hat{\mathbf{Q}}|)}. \tag{4.26}$$

Here we have extracted a scale $\sigma_{\mathbf{Q}}$ from the entries in $\mathbf{Q} = \mathbf{Q} \mathbf{K}^{-1/2}$ via the definition $\hat{\mathbf{Q}} = \mathbf{Q} \mathbf{K}^{-1/2}$, where the entries of $\hat{\mathbf{Q}}$ are distributed according to $\mathcal{N}(0, 1)$ in the universal regime. This separates a trivial scaling factor $\sigma_{\mathbf{Q}}$ appearing in $\lambda_{\min}(|\mathbf{Q}|)$ from its more intrinsic scaling properties with N, P . The lower bound (4.26) is significant because it differs from an upper bound on \mathcal{D}_0 (cf. section 4.1) only by the logarithmic factor $\ell(P)$:

$$\mathcal{D}_0 \leq \frac{2\pi\sqrt{P}}{\sigma_{\mathbf{Q}} \lambda_{\min}(|\hat{\mathbf{Q}}|)}. \tag{4.27}$$

An intuitive understanding of (4.25) was given in section 4.2.3: in a large-dimensional P -cube the number of vertices vastly outnumbers the number of faces, thus it is much more likely for a vector to (approximately) point towards a vertex than towards a face. More quantitatively, the matrix $\mathbf{Q}^\top \mathbf{Q}$ is rotationally invariant (i.e. its form is preserved under $\mathbf{Q} \rightarrow \mathbf{O}\mathbf{Q}$ with \mathbf{O} a $P \times P$ orthogonal matrix) so from general considerations in random matrix ensembles [79] we expect the eigenvectors of $\mathbf{Q}^\top \mathbf{Q}$ (including $\hat{\Psi}_{|\mathbf{Q}|}$) to

²⁷Recall the definition of $\ell(P)$ in (3.17), $\ell(P) = \sqrt{2} \operatorname{erf}^{-1}(2^{-1/P}) \approx \sqrt{2 \log P}$ at large P .

²⁸In fact, the eigenvectors of $\mathbf{Q}^\top \mathbf{Q}$ themselves are delocalized (in particular $\hat{\Psi}_{|\mathbf{Q}|}$), but this is of subordinate relevance.

be delocalized (see also [80]). Furthermore, provided \mathbf{K} does not introduce significant anisotropy, the vector $\mathbf{Q}\hat{\Psi}_{|\mathbf{Q}|}$ will be delocalized as well.

In the following three sections we will verify the delocalization of $\mathbf{Q}\hat{\Psi}_{|\mathbf{Q}|}$ in various ensembles of kinetic matrices. Having established this, we will use (4.26) to analytically obtain a reliable lower bound on the diameter of \mathcal{T}_0 in the different ensembles. We will examine two distinct regimes:

- “hard edge”: as $N \rightarrow \infty$, $P - N$ is held fixed,
- “soft edge”: as $N \rightarrow \infty$, N/P is held fixed.

4.3.2 Unit metric

In random axion theories where $\mathbf{K} = f^2 \mathbb{1}_N$ with f a fixed scale, we can formulate the most precise results. In this case $\mathbf{Q}^\top \mathbf{Q} = (\sigma_{\mathcal{Q}}/f)^2 \hat{\mathcal{Q}}^\top \hat{\mathcal{Q}}$, is well-described by a Wishart matrix. At large N , the eigenvector $\hat{\Psi}_{|\mathbf{Q}|}$ is delocalized and we further verified numerically that the P -vectors $\mathbf{Q}\hat{\Psi}_{|\mathbf{Q}|}$ are delocalized as well. The field range along $\hat{\Psi}_{|\mathbf{Q}|}$ is therefore given via (4.26) by

$$\mathcal{R}_0(\hat{\Psi}_{|\mathbf{Q}|}) \approx \frac{\pi\sqrt{P}f}{\ell(P)\sigma_{\mathcal{Q}}\lambda_{\min}(|\hat{\mathcal{Q}}|)}. \quad (4.28)$$

For any specific N, P the probability distribution of the smallest eigenvalue in the Wishart ensemble can be calculated (either recursively [81, 82] or directly [83]). When $P - N$ is held fixed as N tends to infinity, the asymptotics of its mean satisfy $\langle \lambda_{\min}(\hat{\mathcal{Q}}^\top \hat{\mathcal{Q}}) \rangle \sim 1/N$ as $N \rightarrow \infty$, hence the name “hard edge” statistics, as the smallest eigenvalue approaches the constraint that the matrix is positive definite. The knowledge of the distribution of the smallest eigenvalue can be translated to calculate the probability distribution of a lower bound on the diameter \mathcal{D}_0 via (4.28). In general only the first $P - N$ moments of the probability distribution of the diameter along $\hat{\Psi}_{|\mathbf{Q}|}$ are finite, while higher moments diverge.²⁹ More specifically, at large N , we find for the z th moment ($z \leq P - N$):

$$\langle \mathcal{R}_0(\hat{\Psi}_{|\mathbf{Q}|})^z \rangle \sim c(z, P - N) \left(\frac{\pi N f}{\ell(P)\sigma_{\mathcal{Q}}} \right)^z \quad \text{as } N \rightarrow \infty \text{ with } P - N = \text{constant}, \quad (4.29)$$

for some constants $c(z, P - N)$.³⁰ As $\sigma_{\mathcal{Q}} \sim 1/\sqrt{N}$ in sparse models, we see the mean field range along $\hat{\Psi}_{|\mathbf{Q}|}$ scales as $N^{3/2}$ for all $P - N > 0$, up to a logarithmic correction factor.³¹ The standard deviation also exhibits this $N^{3/2}$ scaling with N .³² It may be instructive

²⁹In particular for $P = N$ the distribution of the diameter along $\hat{\Psi}_{|\mathbf{Q}|}$ is heavy-tailed. In that case one can show that the median diameter behaves as $c(1, 0)(2\pi f N/\ell(P)\sigma_{\mathcal{Q}})$ with $c(1, 0) = (\sqrt{\log 4 + 1} - 1)^{-1} \approx 1.84$.

³⁰We found $c(1, 1) = \sqrt{\pi/2} \approx 1.25$, $c(1, 3) = \sqrt{\pi}e[3I_1(1) - I_0(1)]/3\sqrt{2} \approx 0.49$, $c(2, 3) = (e^2 - 5)/8 \approx 0.30$ and $c(3, 3) = \sqrt{\pi}e[5I_0(1) - 9I_1(1)]/15\sqrt{2} \approx 0.28$. To our knowledge there is no closed-form expression for $c(z, P - N)$ — generic values must be determined numerically. This can be done algorithmically [81].

³¹This parametrically improves the lower bound on the diameter of the tile obtained in [80] for this random model, where for $P - N > 0$ one found a lower bound that scales only linearly with N .

³²However, the probability distribution on $\mathcal{R}_0(\hat{\Psi}_{|\mathbf{Q}|})$ is super-exponentially suppressed at small values, and only power-like suppressed at large values. Thus the field range along $\hat{\Psi}_{|\mathbf{Q}|}$ may easily become larger than the mean, but not smaller. This holds in the hard edge limit for all $P \geq N$. For the $P = N$ case see also the discussion in appendix A of [80].

to recapitulate how the $N^{3/2}$ scaling arises, namely as the product of three factors $N^{1/2}$ with different origins: one comes from the alignment of $\mathbf{Q}\hat{\Psi}_{|\mathbf{Q}|}$ along a diagonal of the P -hypercube (kinetic alignment), another from the square root of the smallest eigenvalue of a Wishart matrix scaling like $1/\sqrt{N}$ in the large N -limit where $P - N$ is held fixed (lattice alignment), and finally a \sqrt{N} arising from the assumed sparsity of the charge matrices.

For soft edge statistics where N/P is held fixed, we may use the result of [78], which implies $\langle \lambda_{\min}(|\hat{\mathbf{Q}}|) \rangle \rightarrow (\sqrt{P/N} - 1)\sqrt{N}$. We only discuss the mean field range along $\hat{\Psi}_{|\mathbf{Q}|}$ in this case. We find

$$\langle \mathcal{R}_0(\hat{\Psi}_{|\mathbf{Q}|}) \rangle \sim \frac{\pi f}{\ell(P)\sigma_{\mathcal{Q}}(1 - \sqrt{N/P})} \quad \text{as } N \rightarrow \infty \text{ with } N/P \text{ constant.} \quad (4.30)$$

So in sparse models where the amount of non-perturbative effects scales linearly with N , the mean field range along $\hat{\Psi}_{|\mathbf{Q}|}$ (and hence, typically, the diameter \mathcal{D}_0) is enhanced only by the minimal amount $N^{1/2}$ compared to the single-axion f .

4.3.3 Wishart metric

Here we discuss the random ensemble where \mathbf{K} is a Wishart matrix, constrained to have largest eigenvalue equal to a fixed scale f_{\max}^2 . Specifically, we draw the entries of a matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ from a normal distribution with zero mean and unit variance, and form the combination $\mathbf{A}^T \mathbf{A}$. After, we rescale this matrix to have largest eigenvalue f_{\max}^2 .

As in the case of unit kinetic matrix, we find the eigenvector $\hat{\Psi}_{|\mathbf{Q}|}$ to be delocalized to good accuracy for all $P \geq N$. For constant $P - N$ as $N \rightarrow \infty$, we numerically established that

$$\lambda_{\min}(|\mathbf{Q}\mathbf{K}^{-1/2}|) \approx \frac{2}{f_{\max}} \lambda_{\min}(|\mathbf{Q}|) \quad (4.31)$$

to good approximation (in the distributional sense). For fixed N/P , we found

$$\left\langle \frac{1}{\lambda_{\min}(|\mathbf{Q}\mathbf{K}^{-1/2}|)} \right\rangle = g(N/P) \left\langle \frac{1}{\lambda_{\min}(|\mathbf{Q}|)} \right\rangle f_{\max} \quad (4.32)$$

for some profile $g(N/P)$ which decreases monotonically from 1 at $N/P = 0$ to $1/2$ at $N/P = 1$. This allows us use the results of the $\mathbf{K} \propto \mathbb{1}$ ensemble discussed in the previous section. Thus, in the hard edge limit, the field range along $\hat{\Psi}_{|\mathbf{Q}|}$ satisfies

$$\langle \mathcal{R}_0(\hat{\Psi}_{|\mathbf{Q}|})^z \rangle \sim c(z, P - N) \left(\frac{\pi N f_{\max}}{2\ell(P)\sigma_{\mathcal{Q}}} \right)^z \quad \text{as } N \rightarrow \infty \text{ with } P - N \text{ constant,} \quad (4.33)$$

while for soft edge asymptotics we find

$$\langle \mathcal{R}_0(\hat{\Psi}_{|\mathbf{Q}|}) \rangle \sim \frac{\pi g(N/P) f_{\max}}{\ell(P)\sigma_{\mathcal{Q}}(1 - \sqrt{N/P})} \quad \text{as } N \rightarrow \infty \text{ with } N/P \text{ constant.} \quad (4.34)$$

As in the ensemble with $\mathbf{K} \propto \mathbb{1}$, the mean field range along $\hat{\Psi}_{|\mathbf{Q}|}$ scales as $N^{3/2}$ in the hard edge limit and as $N^{1/2}$ for soft edge statistics, up to a logarithmic correction factor.

4.3.4 Heavy-tailed metric

Finally we consider an example of a heavy-tailed ensemble of kinetic matrices \mathbf{K} , i.e., large fluctuations of its eigenvalues are polynomially suppressed (as opposed to exponentially, as in the Wishart ensemble). Specifically, we consider inverse-Wishart matrices \mathbf{K} , with largest eigenvalue rescaled to f_{\max}^2 . (We rescale the combination $(\mathbf{A}^\top \mathbf{A})^{-1}$ of the previous section.)³³

Once again $\hat{\Psi}_{|\mathbf{Q}|}$ is delocalized to good approximation, and we established numerically that the mean field range along $\hat{\Psi}_{|\mathbf{Q}|}$ qualitatively behaves as

$$\langle \mathcal{R}_0(\hat{\Psi}_{|\mathbf{Q}|}) \rangle \propto \frac{\sqrt{N} f_{\max}}{\ell(P) \sigma_{\mathcal{Q}}} \quad \text{as } N \rightarrow \infty \text{ with } P - N \text{ constant,} \quad (4.35)$$

and as

$$\langle \mathcal{R}_0(\hat{\Psi}_{|\mathbf{Q}|}) \rangle \propto \frac{f_{\max}}{\ell(P) \sigma_{\mathcal{Q}}} \quad \text{as } N \rightarrow \infty \text{ with } N/P \text{ constant.} \quad (4.36)$$

In the hard edge case, a factor of \sqrt{N} is lost compared to the unit and Wishart kinetic matrix ensembles because the distribution of the smallest singular value of the canonical charge matrix is qualitatively different. In particular, we found

$$\frac{f_{\max}}{\sigma_{\mathcal{Q}}} \left\langle \lambda_{\min}(|\mathcal{Q} \mathbf{K}^{-1/2}|) \right\rangle = \mathcal{O}(1) \quad \text{as } N \rightarrow \infty \text{ with } P - N \text{ fixed.} \quad (4.37)$$

An intuitive explanation of this goes as follows: in the ensemble where the kinetic matrix is a rescaled Wishart matrix, $\mathbf{K} = \mathbf{A}^\top \mathbf{A} / \lambda_{\max}(\mathbf{A}^\top \mathbf{A})$, the largest eigenvalue λ_{\max} is not too different from a typical eigenvalue; the ratio $\lambda_{\max} / \text{median}_i(\lambda_i)$ is of order 1. This is because large eigenvalues occur with exponentially small probability. So a typical eigenvalue of \mathbf{K} is expected to be broadly distributed on the interval $[0, f_{\max}^2]$. In other words,

$$f_{\max}^2 \lambda_{\min}(\mathbf{K}^{-1/2} \mathcal{Q}^\top \mathcal{Q} \mathbf{K}^{-1/2}) \approx \lambda_{\min}(\mathcal{Q}^\top \mathcal{Q}) \approx \sigma_{\mathcal{Q}}^2 / N. \quad (4.38)$$

In the ensemble $\mathbf{K} = (\mathbf{A}^\top \mathbf{A})^{-1} / \lambda_{\max}[(\mathbf{A}^\top \mathbf{A})^{-1}]$, on the other hand, the largest eigenvalue is on average much larger than a typical eigenvalue. So eigenvalues of \mathbf{K} will be small ($\ll f_{\max}^2$) with high probability. With this one can appreciate how $f_{\max} \langle \lambda_{\min} \rangle / \sigma_{\mathcal{Q}} = \mathcal{O}(1)$ in this ensemble.

For soft edge statistics the reason for the reduction of the expected field range down to $N^{1/2}$ is the same as in the other ensembles: $\langle \lambda_{\min} \rangle \propto N$. We summarize these results together with the mean field range behaviour along $\hat{\Psi}_{|\mathbf{Q}|}$ in the unit and Wishart kinetic matrix ensembles (sections 4.3.2 and 4.3.3) in figure 15.

4.3.5 Variable couplings: dynamic alignment

To investigate the dynamics we consider the field range along the lightest direction $\hat{\Psi}_{\mathbf{H}}$ emanating from $\Theta = \mathbf{0}$, as discussed in section 4.1. From $\mathbf{H} = \mathbf{Q}^\top \text{diag}(\Lambda_I^4) \mathbf{Q}$ we observe

³³With this rescaling the eigenvalues of \mathbf{K} no longer follow a heavy-tailed distribution. However, we will see that the field range distribution along $\hat{\Psi}_{|\mathbf{Q}|}$ is qualitatively different in this ensemble compared to the others we have discussed, and we will argue that this is precisely due to the heavy-tailed character of the non-rescaled eigenvalues of \mathbf{K} .




Figure 15. Mean field range along $\hat{\Psi}_{|\mathbf{Q}|}$ for three ensembles of kinetic matrices \mathbf{K} . The field range along $\hat{\Psi}_{|\mathbf{Q}|}$ is a significant lower bound on the diameter \mathcal{D}_0 as explained in section 4.3.1. *Left:* the hard edge, $P = N + 5$. *Right:* soft edge, $P = 2N$. Top, middle and bottom (blue, green and red) lines denote the unit kinetic matrix, Wishart and inverse Wishart ensembles with largest eigenvalue f_{\max}^2 (set to 1 here). We chose $\sigma_{\mathcal{Q}} = 7/N$. Dashed lines show the analytic scaling.

that if the couplings Λ_I^4 are not all equal, this is not the same direction as the previously considered $\hat{\Psi}_{|\mathbf{Q}|}$ (which we showed was well-aligned with the direction providing the actual diameter \mathcal{D}_0 of the tile \mathcal{T}_0 due to kinetic alignment). However if the two directions are sufficiently aligned the available field range within the tile along each will be similar. Below we illustrate in two specific examples how much the couplings may deviate from overall equality before this alignment fails and the expected field ranges along $\hat{\Psi}_{|\mathbf{Q}|}$ and $\hat{\Psi}_{\mathbf{H}}$ become parametrically different in N . In these examples we find that $\hat{\Psi}_{|\mathbf{Q}|}$ and $\hat{\Psi}_{\mathbf{H}}$ remain well-aligned although the couplings may differ from one another to a certain degree. While a general analysis of the alignment between the lightest and the kinematic direction lies beyond the scope of this work,³⁴ the insensitivity of this alignment to the distribution of couplings has been called *dynamic alignment* in a previous discussion [35].

For simplicity we consider ensembles with trivial kinetic matrix, $\mathbf{K} = f^2 \mathbb{1}$, and consider two qualitatively different hierarchies in the couplings. These are illustrated in figure 16. In a first example, assume the couplings Λ_I^4 are uniformly distributed on the interval $[\Lambda_{\min}^4, 1]$. As we dial down Λ_{\min}^4 from one to zero, we expect the field range along $\hat{\Psi}_{\mathbf{H}}$ to diminish with respect to the field range along $\hat{\Psi}_{|\mathbf{Q}|}$. In a second example, consider the case where the couplings are log-uniformly distributed on the interval $[10^{-10}, 1]$. Although the couplings may wildly differ from one another, the expected field range along the lightest direction is very robust.

5 Conclusions

In this paper we presented the details of a novel formalism that allows us to analyze a class of periodic functions of many variables, focusing on those of the form 2. Our technique

³⁴An interesting question is whether there exists a simple criterion on the distribution of the couplings Λ_I^4 that assures the vectors $\hat{\Psi}_{|\mathbf{Q}|}$ and $\hat{\Psi}_{\mathbf{H}}$ are aligned. A useful definition of “aligned” would relate the field ranges $\mathcal{R}(\hat{\Psi}_{|\mathbf{Q}|})$ and $\mathcal{R}(\hat{\Psi}_{\mathbf{H}})$ as $N, P \rightarrow \infty$.




Figure 16. Scaling of mean field range along the lightest direction $\hat{\Psi}_{\mathbf{H}}$ with N , $\mathcal{R}(\hat{\Psi}_{\mathbf{H}})/f \propto N^\alpha/\ell(P)$, for the hard edge case $P = N + 5$. *Left:* Λ_I^4 are uniformly distributed over $[\Lambda_{\min}^4, \Lambda_{\max}^4]$. *Right:* $\log_{10}(\Lambda_{\max}^4/\Lambda_I^4)$ are uniformly distributed over $[0, 10]$. The scaling is extracted from sampling the range $100 \leq N \leq 200$. Observe that even when large hierarchies are present in the couplings, the scaling behavior with N of the mean field range along the lightest direction may remain very similar to the scaling behavior of the mean diameter.

identifies the exact periods, breaks up a unit cell of the resulting lattice into conveniently labelled tiles, and makes it possible to identify *approximate* shift symmetries. This last feature is very powerful, because (at least in the large N random ensembles we consider) these approximate symmetries are extremely close to exact. As a result if we analyze one region of the function, the results can be translated to exponentially many other regions with exponential accuracy. In particular the number of critical points is exponentially large, and the spacing of their energy levels is exponentially fine.

We employ this technology to determine the number and characteristics of critical points of the potential 2, and to analyze the vicinity of a typical minimum. For $N \sim 100$ there are generically (in our ensembles) an enormous number of distinct minima, each with a unique vacuum energy. This number can be larger than 10^{120} and the distribution of energies is smooth, so this theory provides values for the vacuum energy consistent with observation even if the energy scales in the potential are close to the Planck scale [26, 28]. Furthermore we find that there is a range of masses for the canonically normalized fields that is enhanced by powers of N . The lightest of these provide long gentle slopes that may turn out to be suitable for large-field inflation if the energy scales are high (or small field inflation if they are lower) [28]. Lastly, the characteristics of the critical points are such that the barriers between basins of attractions of adjacent minima tend to be quite thin. We will explore the consequences of this for tunneling transitions in [27], where we will also discuss the natural candidate for dark matter that arises in these theories.

Acknowledgments

We thank Nick Cook, Frederik Denef, Mario Kieburg, Liam McAllister, Arman Mimar, Ruben Monten, Tom Rudelius, Geoff Ryan, Marjorie Schillo and John Stout for valuable

discussions. OJ is supported by a James Arthur fellowship. The work of MK is supported in part by the NSF through grant PHY-1214302, and he acknowledges membership at the NYU-ECNU Joint Physics Research Institute in Shanghai. The work of TB and KE was supported by DOE under grant no. DE-SC0011941.

A Separating flat directions in V

In this appendix we discuss how to separate flat directions in V (that is, directions in field space with exactly zero potential) from the non-flat directions. Such directions are present when the charge matrix \mathcal{Q} has rank $R < N$. This can happen either when $P < N$ (in which case necessarily $R \leq P < N$) or because \mathcal{Q} is not full rank. After the separation procedure described in this appendix, we are left with $N - R$ massless fields decoupled from a reduced theory of R axions with a full rank charge matrix, to which the techniques in the bulk of our paper apply.

We start with the N -axion theory (1.1),

$$\mathcal{L} = \frac{1}{2} \partial \boldsymbol{\theta}^\top \mathbf{K} \partial \boldsymbol{\theta} - \sum_{I=1}^P \Lambda_I^4 \left[1 - \cos(\boldsymbol{\mathcal{Q}} \boldsymbol{\theta} + \boldsymbol{\delta})^I \right] + \dots, \quad (\text{A.1})$$

and change coordinates to canonically normalized fields $\boldsymbol{\Theta} = \sqrt{\mathbf{K}} \boldsymbol{\theta}$, in which the charge matrix takes the form $\mathbf{Q} = \boldsymbol{\mathcal{Q}} \mathbf{K}^{-1/2}$. Suppose the rank of $\boldsymbol{\mathcal{Q}}$ is $R < N$. This happens either when $P < N$, or when $P \geq N$ but not all columns of $\boldsymbol{\mathcal{Q}}$ are linearly independent. Then there are $L = N - R$ flat directions; moving along these directions does not change V . In other words, the null space of \mathbf{Q} , $\ker(\mathbf{Q})$, is L -dimensional, which is the same as the dimension of $\ker(\boldsymbol{\mathcal{Q}})$. Find an orthonormal basis $\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_L$ of $\ker(\mathbf{Q})$, and extend it to a basis of \mathbb{R}^N by the adherence of R vectors $\mathbf{t}_{L+1}, \mathbf{t}_{L+2}, \dots, \mathbf{t}_N$ (note these are generally not integer-valued vectors). Now define new coordinates $\boldsymbol{\Omega}$ via the rule

$$\boldsymbol{\Theta} = \mathbf{T}_{\Omega\Theta} \boldsymbol{\Omega} = (\mathbf{T}_{\Omega_L\Theta} \mid \mathbf{T}_{\Omega_R\Theta}) \begin{pmatrix} \boldsymbol{\Omega}_L \\ \boldsymbol{\Omega}_R \end{pmatrix}, \quad (\text{A.2})$$

where we have split the N -vector $\boldsymbol{\Omega}$ into a piece of length L and a piece of length R , and the matrices $\mathbf{T}_{\Omega_L\Theta}$ and $\mathbf{T}_{\Omega_R\Theta}$ are composed by placing the $\mathbf{t}_{1,\dots,L}$ respectively the $\mathbf{t}_{L+1,\dots,N}$ on consecutive columns. Note that $\mathbf{T}_{\Omega\Theta}$ is an orthogonal matrix. In these coordinates the flat directions are manifestly separated from the non-flat ones. Indeed, since

$$\mathbf{Q} \boldsymbol{\Theta} = \mathbf{Q} \mathbf{T}_{\Omega_R\Theta} \boldsymbol{\Omega}_R \equiv \mathbf{Q}_R \boldsymbol{\Omega}_R, \quad (\text{A.3})$$

only the R fields $\boldsymbol{\Omega}_R$ appear in the potential. Furthermore the kinetic term reads

$$\frac{1}{2} \partial \boldsymbol{\Omega}_L^\top \partial \boldsymbol{\Omega}_L + \frac{1}{2} \partial \boldsymbol{\Omega}_R^\top \partial \boldsymbol{\Omega}_R, \quad (\text{A.4})$$

such that the massless fields $\boldsymbol{\Omega}_L$ decouple. In (A.3) \mathbf{Q}_R is a full rank $P \times R$ matrix (but it is not integer-valued, in general). For the final step, note that there exists an

invertible $R \times R$ matrix \mathbf{R} such that $\mathbf{Q}_R \mathbf{R}$ is an integer-valued matrix.³⁵ Transforming to coordinates $\boldsymbol{\Omega}_R = \mathbf{R} \boldsymbol{\Xi}_R$, the Lagrangian in $\boldsymbol{\Xi}_R$ -coordinates has a kinetic matrix $\mathbf{R}^\top \mathbf{R}$ and an integer-valued, full rank charge matrix $\mathbf{Q}_R \mathbf{R}$. As $P \geq R$, there are effectively more non-perturbative contributions to the potential than axions. So one can apply the techniques developed in the body of this work to the reduced system of R axions $\boldsymbol{\Xi}_R$.

B Eliminating very massive axions

In this appendix we discuss how to eliminate axions that receive large masses, e.g. due to their coupling to classical sources. At low energies these axions are effectively fixed to a certain value. Specifically let us assume the N -axion theory (1.1),

$$\mathcal{L} = \frac{1}{2} \partial \boldsymbol{\theta}^\top \mathbf{K} \partial \boldsymbol{\theta} - \sum_{I=1}^P \Lambda_I^4 \left[1 - \cos(\boldsymbol{\mathcal{Q}} \boldsymbol{\theta} + \boldsymbol{\delta})^I \right] + \dots, \quad (\text{B.1})$$

is supplemented with $L < N$ such classical sources, where the rows of a full rank $L \times N$ matrix \mathbf{C} specify which axion combinations couple to each source,³⁶ and that these directions are fixed according to

$$\mathbf{C} \boldsymbol{\theta} = \boldsymbol{\delta}_C, \quad (\text{B.2})$$

where $\boldsymbol{\delta}_C$ is a certain L -vector. We would like to perform a (linear) coordinate transformation $\boldsymbol{\theta} \rightarrow \boldsymbol{\xi}$ that disentangles the massive directions from the others in (B.1). In order to retain the same discrete shift symmetries in the $\boldsymbol{\xi}$ -basis as in the $\boldsymbol{\theta}$ -basis, such a transformation must be unimodular. To construct it, note that the directions unaffected by the classical sources lie in $\ker(\mathbf{C})$, which has dimension $N - L \equiv R$. The intersection $\ker(\mathbf{C}) \cap \mathbb{Z}^N$ is thus a lattice of rank R .³⁷ Extend a basis $\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_R$ of this lattice to a basis for \mathbb{Z}^N by adding L integer vectors $\mathbf{t}_{R+1}, \mathbf{t}_{R+2}, \dots, \mathbf{t}_N$ (see appendix C for a proof that this can always be done). The $N \times N$ matrix $\mathbf{T}_{\xi\theta} = (\mathbf{T}_{\xi_R\theta} \mid \mathbf{T}_{\xi_L\theta})$ is unimodular, where $\mathbf{T}_{\xi_R\theta}$ ($\mathbf{T}_{\xi_L\theta}$) is formed by placing the $\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_R$ ($\mathbf{t}_{R+1}, \mathbf{t}_{R+2}, \dots, \mathbf{t}_N$) on consecutive columns. If we denote the first R components of the N -vector $\boldsymbol{\xi}$ by $\boldsymbol{\xi}_R$ and the final $L = N - R$ by $\boldsymbol{\xi}_L$, and relate the coordinates $\boldsymbol{\xi}$ to $\boldsymbol{\theta}$ by $\boldsymbol{\theta} = \mathbf{T}_{\xi\theta} \boldsymbol{\xi}$, we have $\mathbf{C} \boldsymbol{\theta} = \mathbf{C} \mathbf{T}_{\xi_L\theta} \boldsymbol{\xi}_L$, and thus via (B.2)

$$\boldsymbol{\xi}_L = (\mathbf{C} \mathbf{T}_{\xi_L\theta})^{-1} \boldsymbol{\delta}_C. \quad (\text{B.3})$$

The Lagrangian (1.1) then effectively becomes

$$\mathcal{L} = \frac{1}{2} \partial \boldsymbol{\xi}_R^\top \mathbf{K}_R \partial \boldsymbol{\xi}_R - \sum_{I=1}^P \Lambda_I^4 \left[1 - \cos(\boldsymbol{\mathcal{Q}}_R \boldsymbol{\xi}_R + \boldsymbol{\delta}_R)^I \right] + \dots, \quad (\text{B.4})$$

³⁵This is because $\text{colsp}(\mathbf{Q}_R) = \text{colsp}(\boldsymbol{\mathcal{Q}})$, and the projector onto a linear subspace is basis-independent. So we know the projector onto $\text{colsp}(\mathbf{Q}_R)$ has rational entries, implying the existence of \mathbf{R} (see also footnote 8).

³⁶In general identical combinations may couple to more than one source, implying that \mathbf{C} would not be full rank, or L may be greater than N . These cases are easily dealt with.

³⁷We assume the projector onto $\ker(\mathbf{C})$ contains only rational entries.

where

$$\mathbf{K}_R = \mathbf{T}_{\xi_R\theta}^\top \mathbf{K} \mathbf{T}_{\xi_R\theta}, \tag{B.5}$$

$$\mathbf{Q}_R = \mathbf{Q} \mathbf{T}_{\xi_R\theta}, \tag{B.6}$$

$$\delta_R = \delta + \mathbf{Q} \mathbf{T}_{\xi_L\theta} (\mathbf{C} \mathbf{T}_{\xi_L\theta})^{-1} \delta_C. \tag{B.7}$$

The axions that couple to classical sources have been eliminated while preserving the original form of the theory.

C Extending a sublattice basis

In this appendix we prove that any basis for the rank N sublattice defined by the intersection of an N -dimensional linear subspace Σ with the integer lattice \mathbb{Z}^P can be extended to a basis for the full integer lattice. In particular, one can always supplement the N P -vectors \mathbf{t}_i^\parallel with $P - N$ additional vectors \mathbf{t}_a^\parallel to form a basis for \mathbb{Z}^P (cf. section 2).

Before giving the proof, it is perhaps worth giving an example of a sublattice that *cannot* be extended this way. First, recall that since we are discussing lattices, one should consider only linear combinations of the basis vectors with integer coefficients. Now consider the rank one sublattice of \mathbb{Z}^2 that is the even integers along the x -axis; that is, the sublattice generated by the vector $(2, 0)$. It is clear that this cannot be extended to a basis for \mathbb{Z}^2 by the addition of any vector. However, note that this sublattice is *not* the intersection of any linear subspace with \mathbb{Z}^2 — the intersection of the x -axis with \mathbb{Z}^2 is generated by the vector $(1, 0)$.

The proof is as follows:³⁸ every rank P lattice \mathcal{L} can be thought of as a finitely generated free Abelian group (under addition of the lattice vectors). Any rank $N \leq P$ sublattice \mathcal{L}' of \mathcal{L} is then a subgroup. The structure theorem for finitely generated Abelian groups implies that there always exists a special basis $\mathbf{B} \equiv \{\mathbf{b}_1, \dots, \mathbf{b}_P\}$ for \mathcal{L} with the property that $\{a_1 \mathbf{b}_1, \dots, a_N \mathbf{b}_N\}$ is a basis for \mathcal{L}' , where $\{a_1, \dots, a_N\}$ are a set of integers with the property that each divides the next. However if \mathcal{L}' is the intersection of a linear subspace Σ with the lattice \mathcal{L} , then $a_i \mathbf{b}_i \in \mathcal{L}'$ implies $\mathbf{b}_i \in \mathcal{L}'$. Therefore $\{\mathbf{b}_1, \dots, \mathbf{b}_N\}$ must in fact be a basis for \mathcal{L}' (because it generates $\{a_1 \mathbf{b}_1, \dots, a_N \mathbf{b}_N\}$). But this basis can trivially be extended to the basis \mathbf{B} for \mathcal{L} by appending $\{\mathbf{b}_{N+1}, \dots, \mathbf{b}_P\}$.

To see that *any* basis for \mathcal{L}' can be extended to a basis for \mathcal{L} , note that any basis for \mathcal{L}' is related to any other basis (for instance, $\{\mathbf{b}_1, \dots, \mathbf{b}_N\}$) by some $N \times N$ unimodular matrix. But any such $N \times N$ unimodular matrix can obviously be extended to a block-diagonal $P \times P$ unimodular matrix. Acting with this matrix on \mathbf{B} gives the extended basis.

D An explicit example

In this appendix we illustrate the construction of the aligned lattice basis and the reduction of the relative phases in an explicit example with $P = N + 1 = 3$. In particular, we consider

³⁸MK would like to thank Arman Mimar for explaining this to him.

a theory with charges and phases

$$\mathbf{Q} = \begin{pmatrix} 1 & 1 \\ 2 & -3 \\ -3 & 0 \end{pmatrix}, \quad \boldsymbol{\delta} = \begin{pmatrix} 2.04 \\ 6.20 \\ 4.16 \end{pmatrix}, \quad (\text{D.1})$$

and take the non-perturbative scales Λ_I^4 to be identical for simplicity. The potential is therefore

$$V(\boldsymbol{\theta}) = \Lambda^4 [3 - \cos(\theta^1 + \theta^2 + 2.04) - \cos(2\theta^1 - 3\theta^2 + 6.20) - \cos(-3\theta^1 + 4.16)]. \quad (\text{D.2})$$

The auxiliary coordinates $\boldsymbol{\phi}$ are constrained by the condition $\mathbf{P}^\perp \boldsymbol{\phi} = \boldsymbol{\delta}$ to reproduce (D.2) on-shell, where the orthogonal projector onto the orthogonal complement of the constraint surface Σ is given by

$$\mathbf{P}^\perp = \mathbb{1} - \mathbf{Q}(\mathbf{Q}^\top \mathbf{Q})^{-1} \mathbf{Q}^\top = \frac{1}{115} \begin{pmatrix} 81 & 27 & 45 \\ 27 & 9 & 15 \\ 45 & 15 & 25 \end{pmatrix}. \quad (\text{D.3})$$

As expected, the rank of the orthogonal projector is $P - N = 1$ in this example. Employing the LLL lattice reduction algorithm [53, 55], we find the aligned basis vectors,

$$\mathbf{T}_{\omega\phi} = \begin{pmatrix} -1 & -1 & 1 \\ 3 & -2 & -1 \\ 0 & 3 & -1 \end{pmatrix}, \quad \mathbf{T}_{\omega_\parallel\phi} = \begin{pmatrix} -1 & -1 \\ 3 & -2 \\ 0 & 3 \end{pmatrix}, \quad \mathbf{T}_{\omega_\# \phi} = \begin{pmatrix} 1 \\ -1 \\ -1 \end{pmatrix}. \quad (\text{D.4})$$

It is easy to verify that the first two basis vectors, $\mathbf{T}_{\omega_\parallel\phi}$ are parallel to Σ , while the last basis vector $\mathbf{T}_{\omega_\# \phi}$ has a very small projection onto the orthogonal complement of Σ ,

$$\mathbf{P}^\perp \mathbf{T}_{\omega_\# \phi} = \frac{1}{115} \begin{pmatrix} 9 \\ 3 \\ 5 \end{pmatrix}. \quad (\text{D.5})$$

Note that the length of the shortest lattice vector of the lattice generated by \mathbf{P}^\perp agrees with $1/\sqrt{\det(\mathbf{Q}^\top \mathbf{Q})} = 1/\sqrt{115}$, as expected. The inverse transformation is given by

$$\mathbf{T}_{\phi\omega} = \begin{pmatrix} 5 & 2 & 3 \\ 3 & 1 & 2 \\ 9 & 3 & 5 \end{pmatrix}, \quad \mathbf{T}_{\phi\omega_\parallel} = \begin{pmatrix} 5 & 2 & 3 \\ 3 & 1 & 2 \end{pmatrix}, \quad \mathbf{T}_{\phi\omega_\#} = \begin{pmatrix} 9 & 3 & 5 \end{pmatrix}. \quad (\text{D.6})$$

We now can express the potential in terms of the aligned coordinates $\boldsymbol{\omega}_\parallel = \mathbf{T}_{\phi\omega_\parallel} \boldsymbol{\phi}$,

$$V(\boldsymbol{\omega}_\parallel) = \Lambda^4 \left[3 - \cos(2.04 - \omega_\parallel^1 - \omega_\parallel^2) - \cos(3\omega_\parallel^1 - 2\omega_\parallel^2 + 6.20) - \cos(-3\omega_\parallel^2 + 4.16) \right]. \quad (\text{D.7})$$

Finally, we note that considering the shift $\boldsymbol{\phi} \rightarrow \boldsymbol{\phi} + 2\pi \mathbf{T}_{\omega_\# \phi} n_\delta$ the constraint equation can be rewritten as

$$\mathbf{P}^\perp(\boldsymbol{\phi} + 2\pi \mathbf{T}_{\omega_\# \phi} n_\delta - \mathbf{P}^\perp \boldsymbol{\delta}) = 0, \quad (\text{D.8})$$

which allows to reduce the phase via (2.35),

$$\mathbf{n}_\delta = \left[\frac{1}{2\pi} \mathbf{T}_{\phi\omega_{\parallel}} \mathbf{P}^\perp \boldsymbol{\delta} \right]_{\text{n.i.}} = 21, \quad (\text{D.9})$$

which gives the potential

$$V(\omega_{\parallel}) = \Lambda^4 \left[3 - \cos(-\omega_{\parallel}^1 - \omega_{\parallel}^2 - 0.09) - \cos(3\omega_{\parallel}^1 - 2\omega_{\parallel}^2 - 0.03) - \cos(-3\omega_{\parallel}^2 - 0.05) \right]. \quad (\text{D.10})$$

As expected, the phases are significantly reduced by employing the approximate shift symmetry.

E Derivation of (3.12)

In this appendix we derive (3.12), starting from (3.6). For the case of equal Λ_I , one has $\boldsymbol{\Delta}^\perp = \mathbf{P}^\perp = \mathbf{R}^\top (\mathbf{R}\mathbf{R}^\top)^{-1} \mathbf{R}$. Using (3.6), it is easy to see that energies at the minima can be written

$$V(\phi_{n_\omega}) \approx 2\pi^2 \Lambda^4 \left[\mathbf{T}_{\omega_{\parallel}\phi}^\top \mathbf{P}^\perp \mathbf{T}_{\omega_{\parallel}\phi} \right] n_\omega^2 \approx 2\pi^2 \Lambda^4 \left(\frac{c n_\omega}{\sqrt{\det \mathcal{Q}^\top \mathcal{Q}}} \right)^2, \quad \forall |m| < \mathcal{N}_{\text{vac}}, \quad (\text{E.1})$$

where the positive integer c is

$$c \equiv \sqrt{\det \left[\left(\mathbf{P}^\perp \mathbf{T}_{\omega_{\parallel}\phi} \right)^\top \left(\mathbf{P}^\perp \mathbf{T}_{\omega_{\parallel}\phi} \right) \right] \det \mathcal{Q}^\top \mathcal{Q}} = \sqrt{\frac{\det \mathcal{Q}^\top \mathcal{Q}}{\det \mathbf{T}_{\omega_{\parallel}\phi}^\top \mathbf{T}_{\omega_{\parallel}\phi}}}. \quad (\text{E.2})$$

To derive (E.2) we first used that $\mathbf{T}_{\omega_{\parallel}\phi}^\top \mathbf{P}^\perp \mathbf{T}_{\omega_{\parallel}\phi} = \mathbf{T}_{\omega_{\parallel}\phi}^\top (\mathbf{P}^\perp)^2 \mathbf{T}_{\omega_{\parallel}\phi} = \left(\mathbf{P}^\perp \mathbf{T}_{\omega_{\parallel}\phi} \right)^\top \left(\mathbf{P}^\perp \mathbf{T}_{\omega_{\parallel}\phi} \right)$ is a number, because $P - N = 1$, and thus equal to the determinant of the matrices that form it. Then we use the identity

$$\det \left[\left(\mathbf{P}^\perp \mathbf{T}_{\omega_{\parallel}\phi} \right)^\top \left(\mathbf{P}^\perp \mathbf{T}_{\omega_{\parallel}\phi} \right) \right] = \left(\det \mathbf{T}_{\omega_{\parallel}\phi}^\top \mathbf{T}_{\omega_{\parallel}\phi} \right)^{-1}. \quad (\text{E.3})$$

To see this note that

$$\begin{aligned} \det \mathbf{T}_{\omega\phi} &= 1 \\ &= \det \begin{pmatrix} \mathbf{T}_{\omega_{\parallel}\phi} & \mathbf{T}_{\omega_{\parallel}\phi} \end{pmatrix} \\ &= \det \begin{pmatrix} \mathbf{T}_{\omega_{\parallel}\phi} & \mathbf{P}^\perp \mathbf{T}_{\omega_{\parallel}\phi} \end{pmatrix} \equiv \det \mathbf{T}'_{\omega\phi}, \end{aligned} \quad (\text{E.4})$$

where we used the invariance of the determinant under adding linear combinations of some columns to other columns. Then, by computing $\det \left(\mathbf{T}'_{\omega\phi} \right)^\top \mathbf{T}'_{\omega\phi}$ and using that the columns of $\mathbf{T}_{\omega_{\parallel}\phi}$ and $\mathbf{P}^\perp \mathbf{T}_{\omega_{\parallel}\phi}$ are orthogonal, one obtains (E.3).

F Derivation of (4.12)

To derive the bounds (4.12) on the diameter of the tile containing the origin, note first the general inequality for P -vectors \mathbf{v} : $\|\mathbf{v}\|_2/\sqrt{P} \leq \|\mathbf{v}\|_\infty \leq \|\mathbf{v}\|_2$, or

$$\frac{1}{\|\mathbf{v}\|_2} \leq \frac{1}{\|\mathbf{v}\|_\infty} \leq \frac{\sqrt{P}}{\|\mathbf{v}\|_2}. \quad (\text{F.1})$$

The field range inside \mathcal{T}_0 along any specific direction is a lower bound for (half of) the diameter (cf. (4.11)). This holds in particular for the field range along the eigenvector $\hat{\Psi}_{|\mathbf{Q}|}$ corresponding to the smallest eigenvalue of $|\mathbf{Q}| \equiv \sqrt{\mathbf{Q}^\top \mathbf{Q}}$. Combining this with the left-most inequality in (F.1) we have therefore

$$\mathcal{D}_0 \geq \frac{2\pi}{\|\mathbf{Q}\hat{\Psi}_{|\mathbf{Q}|}\|_\infty} \geq \frac{2\pi}{\|\mathbf{Q}\hat{\Psi}_{|\mathbf{Q}|}\|_2} = \frac{2\pi}{\lambda_{\min}(|\mathbf{Q}|)}. \quad (\text{F.2})$$

On the other hand, each P -vector $\mathbf{Q}\hat{\Theta}$ is subject to the right-most inequality in (F.1). Therefore the same inequality holds between the maxima of both sides over all $\hat{\Theta} \in S^{N-1}$. Using this relation in the expression (4.11) for the diameter, we arrive at

$$\mathcal{D}_0 \leq \max \left\{ \frac{2\pi\sqrt{P}}{\|\mathbf{Q}\hat{\Theta}\|_2} \mid \hat{\Theta} \in S^{N-1} \right\} = \frac{2\pi\sqrt{P}}{\lambda_{\min}(|\mathbf{Q}|)}, \quad (\text{F.3})$$

where we used $\lambda_{\min}(|\mathbf{Q}|) = \|\mathbf{Q}\hat{\Psi}_{|\mathbf{Q}|}\|_2 \leq \|\mathbf{Q}\hat{\Theta}\|_2$ for all $\hat{\Theta} \in S^{N-1}$.

G Axion potentials and Gaussian random fields

In the body of this work we developed tools that allow for a systematic approach to general (multi-)axion theories. This analytic approach is most powerful for well-aligned axion theories. Unfortunately, when the number P of non-trivial non-perturbative terms becomes very large, alignment typically fails and all approximate shift symmetries are broken. We now turn to a complimentary description of the axion potential, in terms of a Gaussian random field, that is valid precisely when the theory ceases to be well-aligned, and again allows for a simple statistical description of the theory. In particular, we find that at large $P \gg N$ the potential statistics are typically well approximated by an isotropic Gaussian random field with Gaussian covariance function, henceforth referred to by the shorthand *Gaussian field*. The statistical properties and the distribution of minima in Gaussian fields is very well understood, and efficient numerical algorithms exist to numerically sample such fields [80]. Therefore, by providing an effective description of the axion potential in terms of Gaussian fields a host of tools become available to study axion theories.

G.1 Gaussian fields

Before discussing the connection to axion potentials, let us review some of the basic properties of a stationary, isotropic Gaussian field $V_{\mathcal{G}}(\boldsymbol{\chi})$ in N dimensions χ^i . We will assume

³⁹It is not hard to see that this lower bound can never be saturated.

a Gaussian covariance function for simplicity, although more general results exist. The ensemble is specified fully by the mean and the two-point function,

$$\langle V_{\mathcal{G}}(\boldsymbol{\chi}) \rangle = \bar{V}_{\mathcal{G}}, \tag{G.1}$$

$$\langle (V_{\mathcal{G}}(\boldsymbol{\chi}) - \bar{V}_{\mathcal{G}})(V_{\mathcal{G}}(\boldsymbol{\chi}') - \bar{V}_{\mathcal{G}}) \rangle = \Lambda_{\mathcal{G}}^8 e^{-\|\boldsymbol{\chi} - \boldsymbol{\chi}'\|_2^2 / 2\Delta_{\mathcal{G}}^2}, \tag{G.2}$$

where the typical length scale over which the potential varies significantly is called $\Delta_{\mathcal{G}}$ and the overall scale is set by $\Lambda_{\mathcal{G}}$. $\bar{V}_{\mathcal{G}}$ denotes the mean of the random function. In order to understand the distribution of minima in the potential defined above, we will be interested in the correlations between the potential, its gradient and the Hessian matrix $H_{ij} = \partial_i \partial_j V_{\mathcal{G}}$. The correlations of the Hessian are given by [80] (see also [84])

$$\langle H_{ab}(\boldsymbol{\chi}) H_{cd}(\boldsymbol{\chi}) \rangle = (\delta_{ab} \delta_{cd} + \delta_{ad} \delta_{bc} + \delta_{ac} \delta_{bd}) \frac{\Lambda_{\mathcal{G}}^8}{\Delta_{\mathcal{G}}^4}. \tag{G.3}$$

Let us consider the ensemble of points at which the random function takes on a particular value, V , and denote the corresponding ensemble average as $\langle \dots \rangle_V$. The only non-vanishing correlations between the field and its derivatives are given by

$$\langle H_{ab}(\boldsymbol{\chi}) \rangle_V = -\frac{V - \bar{V}_{\mathcal{G}}}{\Delta_{\mathcal{G}}^2} \delta_{ab}, \tag{G.4}$$

$$\langle H_{ab}(\boldsymbol{\chi}) H_{cd}(\boldsymbol{\chi}) \rangle_V = \left(\frac{(V - \bar{V}_{\mathcal{G}})^2}{\Lambda_{\mathcal{G}}^8} \delta_{ab} \delta_{cd} + \delta_{ad} \delta_{bc} + \delta_{ac} \delta_{bd} \right) \frac{\Lambda_{\mathcal{G}}^8}{\Delta_{\mathcal{G}}^4}. \tag{G.5}$$

Note that crucially the gradient is uncorrelated with the potential and the Hessian. These correlations can be cast into a simple random matrix model,

$$H = M - \frac{V - \bar{V}_{\mathcal{G}}}{\Delta_{\mathcal{G}}^2} \mathbb{1}, \tag{G.6}$$

where the matrix M is a real, symmetric random matrix in the Gaussian orthogonal ensemble (GOE), i.e. it can be written as

$$M = \frac{1}{\sqrt{2}}(A + A^{\top}), \quad A_{ij} \sim \mathcal{N}(0, \sigma_M), \tag{G.7}$$

where $\sigma_M = \Lambda_{\mathcal{G}}^4 / \Delta_{\mathcal{G}}^2$. In the large N -limit the eigenvalue spectrum of GOE matrices is given by the famous Wigner semicircle,

$$\rho(\Lambda) = \frac{1}{2\pi N \sigma_M^2} \sqrt{4N \sigma_M^2 - \Lambda^2}. \tag{G.8}$$

By considering the case where the smallest eigenvalue is no longer negative, $\lambda_{\min} > 0$, we can therefore easily solve for the mean value of stable minima in the large N -limit. We find

$$\langle V \rangle_{\text{minima}} \approx \bar{V}_{\mathcal{G}} - 2\sqrt{N} \Lambda_{\mathcal{G}}^4. \tag{G.9}$$

The probability distribution function of energies at minima with barely positive definite Hessian matrix is simply obtained by considering the smallest eigenvalue of the Hessian in (G.6) and solving for V : it is approximated by the convolution of the Tracy-Widom distribution with a normal distribution.

G.2 Axion theories at large P

In the previous section we reviewed the statistical properties of Gaussian random fields with Gaussian covariance function. We are now in a position to consider the statistics of the axion potential in the large N and $P \gg N$ limit, and compare those results to a Gaussian random field.

The non-perturbative potential for the axions with unbroken discrete shift symmetry is given in (2.1),

$$V = \sum_{I=1}^P \Lambda_I^4 \left[1 - \cos(\mathcal{Q}\boldsymbol{\theta})^I \right], \quad (\text{G.10})$$

where as above \mathcal{Q} is a $P \times N$ integer charge matrix. In the following we will assume that the couplings Λ_I^4 are of similar magnitude and independent of the charges \mathcal{Q}^I . In the large P -limit the potential approaches a Gaussian random field. However, the potential (G.10) clearly is not isotropic (it is periodic under $v_{\parallel}^i \rightarrow v_{\parallel}^i + 2\pi$ only for some directions $\boldsymbol{v}_{\parallel}$), nor does it have a Gaussian covariance function. Curiously, however, when sampling over random one-dimensional slices through the N -dimensional potential, the mean power spectrum is very well approximated by a Gaussian. It is therefore not very surprising that there are some similarities between the distribution of minima in Gaussian random landscapes and random axion landscapes, as we make precise below.

When sampling over ensembles of potentials defined by random \mathcal{Q} , containing i.i.d. random integers distributed uniformly in the interval $[-s, s]$ and random phases, the mean \bar{V}_{np} and variance Λ_{np}^8 of the potential V respectively are given by

$$\begin{aligned} \bar{V}_{\text{np}} &\equiv \langle V(\boldsymbol{\theta}) \rangle = \sum_{I=1}^P \Lambda_I^4, \\ \Lambda_{\text{np}}^8 &\equiv \langle [V(\boldsymbol{\theta}) - \langle V(\boldsymbol{\theta}) \rangle]^2 \rangle = \frac{1}{2} \sum_{I=1}^P \Lambda_I^8. \end{aligned} \quad (\text{G.11})$$

Just as in the case of a Gaussian field, the gradient is not correlated with the potential, or the Hessian. The correlations of the Hessian $H_{ab} \equiv \partial_a \partial_b V(\boldsymbol{\theta})$ of the non-perturbative potential are given by

$$\langle H_{ab}(\boldsymbol{\theta}) H_{cd}(\boldsymbol{\theta}) \rangle = \left(\delta_{ab} \delta_{cd} + \delta_{ad} \delta_{bc} + \delta_{ac} \delta_{bd} - \frac{6}{5} \delta_{ac} \delta_{bc} \delta_{cd} \right) \frac{\Lambda_{\text{np}}^8}{\Delta_{\text{np}}^4}, \quad (\text{G.12})$$

where we defined an effective correlation length of the axion potential,

$$\Delta_{\text{np}} = \frac{\sqrt{3}}{s}. \quad (\text{G.13})$$

The correlations (G.12) are very similar to the Hessian correlations of a Gaussian fields in (G.3), and only deviate for the diagonal elements of the Hessian. Furthermore, we have for the correlation between the Hessian and the potential

$$\langle H_{ab}(\boldsymbol{\theta}) \rangle_V = -\frac{V - \bar{V}_{\text{np}}}{\Delta_{\text{np}}^2} \delta_{ab}. \quad (\text{G.14})$$

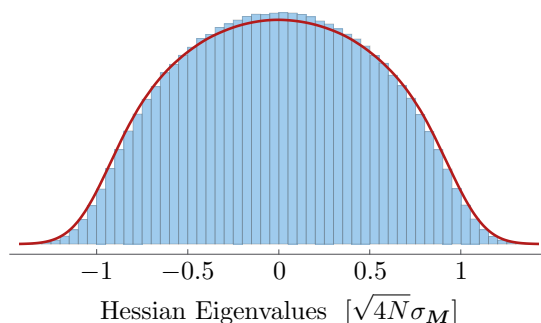


Figure 17. Normalized eigenvalue spectrum of Hessian of potential along with random matrix model for $N = 20$, $P = 2000$.

We can therefore propose an approximate random matrix model for the Hessian, that reproduces the correct correlations, except for the variance of the diagonal terms,

$$\mathbf{H} \approx \mathbf{M} - \frac{V - \bar{V}_{\text{np}}}{\Delta_{\text{np}}^2} \mathbb{1}, \quad (\text{G.15})$$

where \mathbf{M} is a GOE matrix with standard deviation $\sigma_{\mathbf{M}} = \Lambda_{\text{np}}^4 / \Delta_{\text{np}}^2$. We display the eigenvalue spectrum of the full Hessian matrix along with this simple random matrix model in the left part of figure 17. Using the main result of [85], this implies that in the large N -limit we expect most minima at

$$\langle V \rangle_{\text{minima}} \approx V - 2\sqrt{N}\Lambda_{\text{np}}^4 \left(1 - \frac{0.6}{N^{2/3}}\right) = \bar{V}_{\text{np}} - \sqrt{2N}\langle \Lambda_I^4 \rangle_{\text{r.m.s}} \left(1 - \frac{0.6}{N^{2/3}}\right), \quad (\text{G.16})$$

so that the leading term behaves just like in the case of Gaussian fields. It is extremely hard to accurately sample the distribution of minima of the axion potential at large P . However, we can obtain a (not necessarily representative) sample of minima by numerically solving for local minima. We find a good agreement between the numerical results and the random matrix theory expectation.

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] R.D. Peccei and H.R. Quinn, *CP Conservation in the Presence of Instantons*, *Phys. Rev. Lett.* **38** (1977) 1440 [[INSPIRE](#)].
- [2] X.G. Wen and E. Witten, *World Sheet Instantons and the Peccei-Quinn Symmetry*, *Phys. Lett.* **B 166** (1986) 397 [[INSPIRE](#)].
- [3] M. Dine and N. Seiberg, *Nonrenormalization Theorems in Superstring Theory*, *Phys. Rev. Lett.* **57** (1986) 2625 [[INSPIRE](#)].
- [4] M. Dine, N. Seiberg, X.G. Wen and E. Witten, *Nonperturbative Effects on the String World Sheet*, *Nucl. Phys.* **B 278** (1986) 769 [[INSPIRE](#)].

- [5] F. Denef and M.R. Douglas, *Distributions of nonsupersymmetric flux vacua*, *JHEP* **03** (2005) 061 [[hep-th/0411183](#)] [[INSPIRE](#)].
- [6] P. Svrček and E. Witten, *Axions In String Theory*, *JHEP* **06** (2006) 051 [[hep-th/0605206](#)] [[INSPIRE](#)].
- [7] F. Denef, M.R. Douglas and S. Kachru, *Physics of String Flux Compactifications*, *Ann. Rev. Nucl. Part. Sci.* **57** (2007) 119 [[hep-th/0701050](#)] [[INSPIRE](#)].
- [8] F. Denef, *Les Houches Lectures on Constructing String Vacua*, in *String theory and the real world: From particle physics to astrophysics. Proceedings, Summer School in Theoretical Physics, 87th Session, Les Houches, France, July 2–27, 2007*, pp. 483–610, (2008), [arXiv:0803.1194](#) [[INSPIRE](#)].
- [9] D. Baumann and L. McAllister, *Inflation and String Theory*, Cambridge University Press, (2015), [[arXiv:1404.2601](#)].
- [10] C. Long, L. McAllister and J. Stout, *Systematics of Axion Inflation in Calabi-Yau Hypersurfaces*, *JHEP* **02** (2017) 014 [[arXiv:1603.01259](#)] [[INSPIRE](#)].
- [11] J. Halverson, C. Long and P. Nath, *Ultralight axion in supersymmetry and strings and cosmology at small scales*, *Phys. Rev. D* **96** (2017) 056025 [[arXiv:1703.07779](#)] [[INSPIRE](#)].
- [12] J. Halverson, C. Long and B. Sung, *On Algorithmic Universality in F-theory Compactifications*, [arXiv:1706.02299](#) [[INSPIRE](#)].
- [13] W. Hu, R. Barkana and A. Gruzinov, *Cold and fuzzy dark matter*, *Phys. Rev. Lett.* **85** (2000) 1158 [[astro-ph/0003365](#)] [[INSPIRE](#)].
- [14] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper and J. March-Russell, *String Axiverse*, *Phys. Rev. D* **81** (2010) 123530 [[arXiv:0905.4720](#)] [[INSPIRE](#)].
- [15] L. Hui, J.P. Ostriker, S. Tremaine and E. Witten, *Ultralight scalars as cosmological dark matter*, *Phys. Rev. D* **95** (2017) 043541 [[arXiv:1610.08297](#)] [[INSPIRE](#)].
- [16] D.J.E. Marsh, *Axion Cosmology*, *Phys. Rept.* **643** (2016) 1 [[arXiv:1510.07633](#)] [[INSPIRE](#)].
- [17] A. Diez-Tejedor and D.J.E. Marsh, *Cosmological production of ultralight dark matter axions*, [arXiv:1702.02116](#) [[INSPIRE](#)].
- [18] K. Freese, J.A. Frieman and A.V. Olinto, *Natural inflation with pseudo - Nambu-Goldstone bosons*, *Phys. Rev. Lett.* **65** (1990) 3233 [[INSPIRE](#)].
- [19] S. Dimopoulos, S. Kachru, J. McGreevy and J.G. Wacker, *N-flation*, *JCAP* **08** (2008) 003 [[hep-th/0507205](#)] [[INSPIRE](#)].
- [20] J.E. Kim, H.P. Nilles and M. Peloso, *Completing natural inflation*, *JCAP* **01** (2005) 005 [[hep-ph/0409138](#)] [[INSPIRE](#)].
- [21] E. Silverstein and A. Westphal, *Monodromy in the CMB: Gravity Waves and String Inflation*, *Phys. Rev. D* **78** (2008) 106003 [[arXiv:0803.3085](#)] [[INSPIRE](#)].
- [22] L. McAllister, E. Silverstein and A. Westphal, *Gravity Waves and Linear Inflation from Axion Monodromy*, *Phys. Rev. D* **82** (2010) 046003 [[arXiv:0808.0706](#)] [[INSPIRE](#)].
- [23] N. Kaloper and L. Sorbo, *A Natural Framework for Chaotic Inflation*, *Phys. Rev. Lett.* **102** (2009) 121301 [[arXiv:0811.1989](#)] [[INSPIRE](#)].
- [24] N. Kaloper, A. Lawrence and L. Sorbo, *An Ignoble Approach to Large Field Inflation*, *JCAP* **03** (2011) 023 [[arXiv:1101.0026](#)] [[INSPIRE](#)].
- [25] R. Bousso and J. Polchinski, *Quantization of four form fluxes and dynamical neutralization of the cosmological constant*, *JHEP* **06** (2000) 006 [[hep-th/0004134](#)] [[INSPIRE](#)].

- [26] T.C. Bachlechner, *Axionic Band Structure of the Cosmological Constant*, *Phys. Rev. D* **93** (2016) 023522 [[arXiv:1510.06388](#)] [[INSPIRE](#)].
- [27] T. Bachlechner, O. Janssen, K. Eckerle and M. Kleban, to appear.
- [28] T.C. Bachlechner, K. Eckerle, O. Janssen and M. Kleban, *Axions of Evil*, [arXiv:1703.00453](#) [[INSPIRE](#)].
- [29] P.C.W. Davies and S.D. Unwin, *Why is the cosmological constant so small*, *Proc. Roy. Soc. Lond. A* **377** (1981) 147.
- [30] A.D. Sakharov, *Cosmological Transitions With a Change in Metric Signature*, *Sov. Phys. JETP* **60** (1984) 214 [[INSPIRE](#)].
- [31] T. Banks, *T C P, Quantum Gravity, the Cosmological Constant and All That...*, *Nucl. Phys. B* **249** (1985) 332 [[INSPIRE](#)].
- [32] A.D. Linde, *Inflation and Quantum Cosmology*, submitted to: *Newton Centenary Volume*, Cambridge University Press, (1986).
- [33] S. Weinberg, *Anthropic Bound on the Cosmological Constant*, *Phys. Rev. Lett.* **59** (1987) 2607 [[INSPIRE](#)].
- [34] A. Linde, *A brief history of the multiverse*, *Rept. Prog. Phys.* **80** (2017) 022001 [[arXiv:1512.01203](#)] [[INSPIRE](#)].
- [35] T.C. Bachlechner, C. Long and L. McAllister, *Planckian Axions in String Theory*, *JHEP* **12** (2015) 042 [[arXiv:1412.1093](#)] [[INSPIRE](#)].
- [36] T.C. Bachlechner, K. Eckerle, O. Janssen and M. Kleban, *A Systematic Framework for Axion Theories*, <http://cosmo.nyu.edu/kleban/>, (2017).
- [37] C. Vafa, *The string landscape and the swampland*, [hep-th/0509212](#) [[INSPIRE](#)].
- [38] N. Arkani-Hamed, L. Motl, A. Nicolis and C. Vafa, *The string landscape, black holes and gravity as the weakest force*, *JHEP* **06** (2007) 060 [[hep-th/0601001](#)] [[INSPIRE](#)].
- [39] T. Rudelius, *On the Possibility of Large Axion Moduli Spaces*, *JCAP* **04** (2015) 049 [[arXiv:1409.5793](#)] [[INSPIRE](#)].
- [40] C. Cheung and G.N. Remmen, *Naturalness and the Weak Gravity Conjecture*, *Phys. Rev. Lett.* **113** (2014) 051601 [[arXiv:1402.2287](#)] [[INSPIRE](#)].
- [41] B. Heidenreich, M. Reece and T. Rudelius, *Weak Gravity Strongly Constrains Large-Field Axion Inflation*, *JHEP* **12** (2015) 108 [[arXiv:1506.03447](#)] [[INSPIRE](#)].
- [42] T.C. Bachlechner, C. Long and L. McAllister, *Planckian Axions and the Weak Gravity Conjecture*, *JHEP* **01** (2016) 091 [[arXiv:1503.07853](#)] [[INSPIRE](#)].
- [43] T. Rudelius, *Constraints on Axion Inflation from the Weak Gravity Conjecture*, *JCAP* **09** (2015) 020 [[arXiv:1503.00795](#)] [[INSPIRE](#)].
- [44] A. Hebecker, P. Mangat, F. Rompineve and L.T. Witkowski, *Winding out of the Swamp: Evading the Weak Gravity Conjecture with F-term Winding Inflation?*, *Phys. Lett. B* **748** (2015) 455 [[arXiv:1503.07912](#)] [[INSPIRE](#)].
- [45] L.E. Ibáñez, M. Montero, A. Uranga and I. Valenzuela, *Relaxion Monodromy and the Weak Gravity Conjecture*, *JHEP* **04** (2016) 020 [[arXiv:1512.00025](#)] [[INSPIRE](#)].
- [46] B. Heidenreich, M. Reece and T. Rudelius, *Evidence for a sublattice weak gravity conjecture*, *JHEP* **08** (2017) 025 [[arXiv:1606.08437](#)] [[INSPIRE](#)].
- [47] T.C. Bachlechner, M. Dias, J. Frazer and L. McAllister, *Chaotic inflation with kinetic alignment of axion fields*, *Phys. Rev. D* **91** (2015) 023520 [[arXiv:1404.7496](#)] [[INSPIRE](#)].

- [48] K. Choi, H. Kim and S. Yun, *Natural inflation with multiple sub-Planckian axions*, *Phys. Rev. D* **90** (2014) 023545 [[arXiv:1404.6209](#)] [[INSPIRE](#)].
- [49] T. Higaki and F. Takahashi, *Natural and Multi-Natural Inflation in Axion Landscape*, *JHEP* **07** (2014) 074 [[arXiv:1404.6923](#)] [[INSPIRE](#)].
- [50] T. Higaki and F. Takahashi, *Axion Landscape and Natural Inflation*, *Phys. Lett. B* **744** (2015) 153 [[arXiv:1409.8409](#)] [[INSPIRE](#)].
- [51] G. Shiu, W. Staessens and F. Ye, *Widening the Axion Window via Kinetic and Stückelberg Mixings*, *Phys. Rev. Lett.* **115** (2015) 181601 [[arXiv:1503.01015](#)] [[INSPIRE](#)].
- [52] N. Bao, R. Bousso, S. Jordan and B. Lackey, *Fast optimization algorithms and the cosmological constant*, [arXiv:1706.08503](#) [[INSPIRE](#)].
- [53] W. van der Kallen, *Implementations of extended LLL*, <http://www.staff.science.uu.nl/~kalle101/lllimplementations.html>, (1998).
- [54] H. Minkowski, *Gesammelte Abhandlungen von Hermann Minkowski: Vol.: 2*, B.G. Teubner, (1911), <http://books.google.com/books?id=yRiiQwAACAAJ>.
- [55] A. Lenstra, H. Lenstra and L. Lovász, *Factoring polynomials with rational coefficients*, *Math. Ann.* **261** (1982) 515.
- [56] P.Q. Nguyen and D. Stehlé, *Low-Dimensional Lattice Basis Reduction Revisited*, Springer, (2004), pp. 338–357, [https://doi.org/10.1007/978-3-540-24847-7_26].
- [57] N. Gama, N. Howgrave-Graham, H. Koy, and P.Q. Nguyen, *Rankin's Constant and Blockwise Lattice Reduction*, Springer Berlin Heidelberg, Berlin, Heidelberg, (2006), pp. 112–130, [https://doi.org/10.1007/11818175_7].
- [58] Y. Chen and P.Q. Nguyen, *BKZ 2.0: Better Lattice Security Estimates*, pp. 1–20, Springer Berlin Heidelberg, Berlin, Heidelberg, (2011), [https://doi.org/10.1007/978-3-642-25385-0_1].
- [59] D. Dadush and D. Micciancio, *Algorithms for the densest sub-lattice problem*, in *Proceedings of the Twenty-fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '13*, pp. 1103–1122, Society for Industrial and Applied Mathematics, Philadelphia, PA, U.S.A., (2013), <http://dl.acm.org/citation.cfm?id=2627817.2627896>.
- [60] J. Li and P.Q. Nguyen, *Approximating the densest sublattice from Rankin's inequality*, *LMS Journal of Computation and Mathematics* **A 17** (2014) 92.
- [61] M. Krein and D. Milman, *On extreme points of regular convex sets*, *Studia Math.* **9** (1940) 133.
- [62] L. Lovasz and M. Simonovits, *Random walks in a convex body and an improved volume algorithm*, *Random Struct. Algorithm.* **4** (1993) 359.
- [63] B. Cousins and S. Vempala, *Bypassing kls: Gaussian cooling and an O^{n^3} volume algorithm*, in *Proceedings of the Forty-seventh Annual ACM Symposium on Theory of Computing, STOC '15*, ACM, New York, NY, U.S.A., (2015), pp. 539–548, [<https://doi.org/10.1145/2746539.2746563>].
- [64] B. Cousins, *Volume computation of convex bodies*, (2015), <https://www.mathworks.com/matlabcentral/fileexchange/43596-volume-computation-of-convex-bodies?requestedDomain=www.mathworks.com>.
- [65] N.R. Goodman, *The Distribution of the Determinant of a Complex Wishart Distributed Matrix*, *Ann. Math. Statist.* **34** (1963) 178.
- [66] V.A. Marčenko and L.A. Pastur, *Distribution of eigenvalues for some sets of random matrices*, *Math.USSR Sb.* **1** (1967) 457.

- [67] C. Long, L. McAllister and P. McGuirk, *Aligned Natural Inflation in String Theory*, *Phys. Rev. D* **90** (2014) 023501 [[arXiv:1404.7852](#)] [[INSPIRE](#)].
- [68] M. Czerny, T. Higaki and F. Takahashi, *Multi-Natural Inflation in Supergravity*, *JHEP* **05** (2014) 144 [[arXiv:1403.0410](#)] [[INSPIRE](#)].
- [69] S.H.H. Tye and S.S.C. Wong, *Helical Inflation and Cosmic Strings*, [arXiv:1404.6988](#) [[INSPIRE](#)].
- [70] T. Ali, S.S. Haque and V. Jejjala, *Natural Inflation from Near Alignment in Heterotic String Theory*, *Phys. Rev. D* **91** (2015) 083516 [[arXiv:1410.4660](#)] [[INSPIRE](#)].
- [71] C. Burgess and D. Roest, *Inflation by Alignment*, *JCAP* **06** (2015) 012 [[arXiv:1412.1614](#)] [[INSPIRE](#)].
- [72] J. Brown, W. Cottrell, G. Shiu and P. Soler, *Fencing in the Swampland: Quantum Gravity Constraints on Large Field Inflation*, *JHEP* **10** (2015) 023 [[arXiv:1503.04783](#)] [[INSPIRE](#)].
- [73] M. Montero, A.M. Uranga and I. Valenzuela, *Transplanckian axions!?*, *JHEP* **08** (2015) 032 [[arXiv:1503.03886](#)] [[INSPIRE](#)].
- [74] G. Shiu, W. Staessens and F. Ye, *Large Field Inflation from Axion Mixing*, *JHEP* **06** (2015) 026 [[arXiv:1503.02965](#)] [[INSPIRE](#)].
- [75] E. Palti, *On Natural Inflation and Moduli Stabilisation in String Theory*, *JHEP* **10** (2015) 188 [[arXiv:1508.00009](#)] [[INSPIRE](#)].
- [76] R. Kappl, H.P. Nilles and M.W. Winkler, *Modulated Natural Inflation*, *Phys. Lett. B* **753** (2016) 653 [[arXiv:1511.05560](#)] [[INSPIRE](#)].
- [77] A.R. Liddle, A. Mazumdar and F.E. Schunck, *Assisted inflation*, *Phys. Rev. D* **58** (1998) 061301 [[astro-ph/9804177](#)] [[INSPIRE](#)].
- [78] J.W. Silverstein, *The Smallest Eigenvalue of a Large Dimensional Wishart Matrix*, *Ann. Probab.* **13** (1985) 1364.
- [79] M. Rudelson and R. Vershynin, *Delocalization of eigenvectors of random matrices with independent entries*, *Duke Math. J.* **164** (2015) 2507 [[arXiv:1306.2887](#)].
- [80] T.C. Bachlechner, *On Gaussian Random Supergravity*, *JHEP* **04** (2014) 054 [[arXiv:1401.6187](#)] [[INSPIRE](#)].
- [81] A. Edelman, *Eigenvalues and condition numbers of random matrices*, *SIAM J. Matrix Anal. Appl.* **9** (1988) 543.
- [82] A. Edelman, *The distribution and moments of the smallest eigenvalue of a random matrix of wishart type*, *Linear Algebra Appl.* **159** (1991) 55.
- [83] T. Wirtz, G. Akemann, T. Guhr, M. Kieburg and R. Wegner, *The smallest eigenvalue distribution in the real Wishart-Laguerre ensemble with even topology*, *J. Phys. A* **48** (2015) 245202 [[arXiv:1502.03685](#)] [[INSPIRE](#)].
- [84] R. Easther, A.H. Guth and A. Masoumi, *Counting Vacua in Random Landscapes*, [arXiv:1612.05224](#) [[INSPIRE](#)].
- [85] C.A. Tracy and H. Widom, *On orthogonal and symplectic matrix ensembles*, *Comm. Math. Phys.* **177** (1996) 727.