


Axion landscape cosmology

To cite this article: Thomas C. Bachlechner *et al* JCAP09(2019)062

Recent citations

- [Many-field inflation: universality or prior dependence?](#)
Perseas Christodoulidis *et al*

View the [article online](#) for updates and enhancements.



Credit: ESO/M. Kornmesser



AMERICAN
ASTRONOMICAL
SOCIETY

IOP | ebooksTM

Your first choice for astronomy, astrophysics,
solar physics, and planetary science ebooks.

Start exploring the collection—download the
first chapter of every title for free.

Axion landscape cosmology

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

^aDepartment of Physics, University of California San Diego,
Urey Hall Addition, 9500 Gilman Dr #1110-115, La Jolla, CA 92093, U.S.A.

^bDipartimento di Fisica, Università di Milano-Bicocca,
Piazza della Scienza 3, 20126 Milan, Italy

^cINFN, sezione di Milano-Bicocca,
Piazza della Scienza 3, 20126 Milan, Italy

^dCenter for Cosmology and Particle Physics, New York University,
726 Broadway, New York, NY 10003, U.S.A.

E-mail: tcb63@cornell.edu, kate.eckerle@unimib.it, opj202@nyu.edu,
mk161@nyu.edu

Received March 13, 2019

Accepted September 16, 2019

Published September 27, 2019

Abstract. We study the cosmology of complex multi-axion theories. With $\mathcal{O}(100)$ fields and GUT scale energies these theories contain a vast number of vacua, inflationary trajectories and a natural dark matter candidate. We demonstrate that the vacua are stable on cosmological timescales. In a single theory, both large- and small-field inflation are possible and yield a broad range of cosmological observables, and vacuum decay can be followed by a relatively large number (> 60) of efolds of inflation. Light axions stabilized by gravitational instantons may constitute a natural dark matter candidate that does not spoil an axion solution to the strong CP problem.

Keywords: axions, physics of the early universe, initial conditions and eternal universe, string theory and cosmology

ArXiv ePrint: [1810.02822](https://arxiv.org/abs/1810.02822)

Contents

1	Introduction	1
2	Vacuum transitions	3
2.1	Hawking-Moss decays	4
2.2	Coleman-de Luccia decays	4
2.2.1	Neighboring minima	5
2.2.2	Thin-wall tension	6
2.2.3	Numerical checks and the “gradient flow approximation”	7
3	Inflation	7
3.1	Sampling the theory	10
3.2	An explicit example	11
3.3	Inflation after tunneling	13
3.4	Reheating	15
4	Light axion phenomenology	16
4.1	Light and lightest axions	17
4.2	Fuzzy dark matter	18
4.3	QCD axion	19
A	Neighbors and saddle points	20
A.1	Degree- k saddle points	20
A.2	Hessian eigenvalues at large N	23
A.3	Face neighbor distances	24
B	Refined stability bound	28
B.1	Stability	28
B.2	Vacuum energy differences	29

1 Introduction

The landscape paradigm for solving the cosmological constant (CC) problem requires the existence of an enormous number of meta-stable phases (“vacua”) with differing vacuum energies [1, 2]. For theories where the fundamental scale is of order the Planck scale there must be $N_{\text{vac}} \gtrsim M_{\text{Pl}}^4/\rho_{\text{DE}} \sim 10^{120}$ such local minima of the potential, where ρ_{DE} is the observed dark energy density. Theories complex enough to contain such a large number of phases are in general extremely difficult to analyze at any level. Unfortunately, the dynamics of the theory are essential to the putative solution of the CC problem. The reason is that while small vacuum energy is necessary to allow structures to form, it is not sufficient. The cosmological histories may be such that other effects prohibit structure formation. The prototypical example of this is when the low-CC vacua are populated by tunneling, as one expects to be generic in the landscape. In this case the negative curvature of the initial universe after the tunneling inhibits structure formation even when the CC is small [3]. To avoid this, the tunneling must be followed by ~ 60 efolds of inflation (assuming a level of

initial density perturbations roughly commensurate with observation). Without inflation no structures form despite the small CC and the anthropic argument for the small CC fails. But landscapes with small CC minima may not contain such trajectories (for instance, the “double well to the power N ” toy landscape of [4]).

In [5, 6] we developed a powerful framework for analyzing general theories involving N axion fields θ^i coupled through a potential comprised of $P > N$ non-perturbative effects, and this is the third paper in this series. Our technique is based on identifying the set of exact and approximate shift symmetries of the axion potential. These symmetries are an extraordinarily powerful tool because the approximate symmetries are often extremely close to exact. For instance, this renders the task of locating the potential’s critical points tractable, even in field spaces with hundreds of dimensions. Once equipped with the symmetries one may apply repeated shifts to the global minimum by the approximate symmetries, and mod out by the exact ones, to enumerate the distinct local minima. In addition to the field space locations of extrema, we also showed how to retrieve important features of the potential like Hessian eigenvalues, and field ranges in the basins of attraction of minima from simple computations. The purpose of this paper is to apply that formalism to the context of cosmology and to determine whether axion landscapes can solve the CC problem.

The techniques developed in [6] apply to axion potentials of the form

$$V = V_0 + \sum_{I=1}^P \Lambda_I^4 [1 - \cos(\mathbf{Q}\boldsymbol{\theta} + \boldsymbol{\delta})^I]. \quad (1.1)$$

Our notation is that of [6]: bold represents a vector or matrix. Here V_0 is a constant, the Λ_I^4 are the couplings of the axions to the non-perturbative effects, $\boldsymbol{\theta}$ are the N axion fields, \mathbf{Q} is a $P \times N$ rank N matrix containing integer charges Q_j^I , and δ^I are constant phases. When $P > N$ and $P - N \ll N$ the δ^I can be set to zero to a very good approximation by a shift in field space [6]. The kinetic term is assumed to be field-independent:

$$\mathcal{L}_{\text{axion}} = \frac{1}{2} \partial \boldsymbol{\theta}^\top \mathbf{K} \partial \boldsymbol{\theta} - V, \quad (1.2)$$

with \mathbf{K} a positive definite $N \times N$ matrix.

The theory (1.2) is motivated by the study of compactifications of string theory, where there are often hundreds axion fields [7–15]. The shift symmetry of the axions is broken by non-perturbative effects, giving rise to a potential of the form (1.1). We model this by taking the charge matrix \mathbf{Q} as a random matrix with independent identically distributed integer entries Q_j^I , with variance σ_Q^2 . The simplest choice for the Λ_I and \mathbf{K} is $\Lambda_I = \Lambda$, $\forall I$, and $\mathbf{K} = f^2 \mathbb{1}$ with fixed f . In [6] we considered more general random (positive) ensembles. In most of this paper we will stick with the simplest choice, although in section 3 we discuss more general \mathbf{K} , and in section 4 we consider the large hierarchy in the Λ_I that can arise when some of the terms in (1.1) come from gravitationally suppressed instantons.

Because we are motivated by string theory, we will choose all the dimensionful parameters at around the same scale: $\Lambda \lesssim f \sim M_{\text{GUT}} \sim M_{\text{string}} \sim 10^{-2} M_{\text{Pl}}$ and $N \sim \text{few} \times 10^2$. In the later sections of the paper we will consider a minimal coupling to QCD or a U(1) with random, M_{GUT} -suppressed couplings, and the effects of gravitational instantons with actions of order M_{Pl}/f . With these parameters and the techniques of [6] we can test whether this landscape truly solves the CC problem — that is, whether those cosmological histories in which collapsed structures form (à la [1, 3]) do in fact resemble our universe.

Remarkably, without any model building and with only this simple requirement, typical cosmological histories have the following features our universe:

- An extremely small CC, of order $\rho_{\text{DE}} \propto e^{-\mathcal{O}(1) \times \frac{M_{\text{Pl}}}{M_{\text{GUT}}}}$.
- An age of over 10 billion years.
- Approximately 60 efolds of slow roll inflation with a primordial power spectrum $\delta\rho/\rho \sim 10^{-5}$.
- Reheating following inflation.
- Roughly the observed abundance of dark matter.

In the final work [16] of this series of papers we demonstrate how the above features arise naturally in multi-axion theories simply from restricting to cosmological histories in which gravitationally collapsed structures can form. For instance, the exponential suppression of the CC in the first bullet point originates from the small dark matter density relative to the density of radiation, which itself arises naturally from ultra-light axions that interact only with gravity. In this paper we lay the groundwork for that analysis by studying the cosmology of multi-axion theories more generally, including observables such as the inflationary power spectrum, the abundance of dark matter and the status of an axion solution to the strong CP problem.

The structure of this paper follows the order of our list. In section 2 we consider vacuum decay and demonstrate that a vast number of vacua are stable on cosmological timescales. In section 3 we discuss how the inflationary dynamics and observables in extremely complex multi-axion theories can be sampled efficiently, and we demonstrate that a single theory allows for a broad range of observables. We discuss fuzzy dark matter in section 4 and show that gravitational instantons do not spoil the axion solution of the strong CP problem.

2 Vacuum transitions

In the semiclassical approximation to the decay of de Sitter vacua, there are two mechanisms at work:¹ quantum tunneling through a barrier and thermal evaporation to the top of a barrier [17]. The decay proceeds either by a Coleman-de Luccia (CdL) instanton [18] that represents the least-action combination of these two mechanisms, or solely by thermal evaporation via a Hawking-Moss (HM) instanton [19].

In this section we examine decays for a certain class of vacua in our axion landscape, namely those with vacuum energy density sufficiently close to the global minimum that a quadratic approximation of the potential is applicable. Specifically, this means we will focus on vacua for which the arguments of the cosines, $(\mathcal{Q}\theta)^I$, in (1.1) are close to integer multiples of 2π (for all $1 \leq I \leq P$). Such vacua are under very good analytic control. We will avoid a precise definition of the quadratic domain, i.e. which vacua we consider to be well-described by a quadratic expansion of the potential, as our qualitative results are independent of the precise choice.

When $-V_0 = |V_0| \ll \Lambda^4$, all vacua with nearly zero or negative vacuum energy are guaranteed to fall into the quadratic domain, because if the argument of any cosine is substantially different from zero (mod 2π) its positive contribution to the potential (1.1) renders

¹TB does not concur with the results of [17–19] regarding vacuum transition rates in gravitational theories [20, 21].

the total potential energy of order $+\Lambda^4 \gg 0$. The size of the hierarchy required between $|V_0|$ and Λ^4 depends on the desired accuracy of the quadratic expansion. A factor of a few suffices to estimate the decay exponents (the instanton action) to $\mathcal{O}(1)$ accuracy. Therefore, at least if $0 < -V_0 \ll \Lambda^4$ this quadratic approximation suffices for purposes of studying vacua with small vacuum energy, and for studying the decay of such vacua to any lower energy minima.

We will see that in the parameter regime we are focusing on, CdL decays are the dominant channel and vacua with small CC are typically long-lived on cosmological timescales. However the quadratic approximation does not suffice for studying all decays *into* such vacua, as those can originate from higher regions of the potential that might not fall into the quadratic domain. For vacua in the quadratic domain, we find an upper bound on the decay rate, and conclude that no significant fine-tuning is necessary for the vacua to be meta-stable on cosmological timescales.

2.1 Hawking-Moss decays

HM instantons are configurations in Euclidean signature de Sitter space where the field is constant at a saddle point of the potential. These instantons are potentially relevant for decay only if the saddle has degree $k = 1$, meaning the Hessian evaluated there has exactly one negative eigenvalue $V_*'' < 0$, and if

$$M_{\text{Pl}}^2 \frac{|V_*''|}{V_*} \leq \frac{4}{3}, \quad (2.1)$$

where V_* is the potential energy at the saddle point supporting the instanton interpolating between the “parent” and “target” vacua (see e.g. [17]). If these conditions are satisfied the instanton has a single negative mode and contributes an imaginary part to the energy and hence to the decay of the state. For saddle points where (2.1) is not satisfied or there is more than one negative direction ($k > 1$), the HM instanton always has multiple negative modes and presumably does not contribute to the decay. In such cases a CdL instanton always exists [17].

It is simple to estimate whether (2.1) typically holds using the results of [6] (in particular section 3.6.2) and appendix A in this paper. The analysis there shows that the $k = 1$ saddles adjacent to minima in the quadratic domain have $V_* \approx 2\Lambda^4$ because one cosine reaches its maximum roughly halfway in between while the rest are constant. The negative direction at the saddle satisfies $\langle |V_*''| \rangle \sim N\sigma_Q^2 \Lambda^4 / 2f^2 \propto N$, with a standard deviation that scales only as \sqrt{N} . We are most interested in the rough parameter regime where $\sigma_Q \gtrsim \mathcal{O}(1/\sqrt{N})$, $N \approx P \gg 1$ and $f \ll M_{\text{Pl}}$, such that with (2.1) HM decays are typically suppressed,

$$\langle M_{\text{Pl}}^2 \frac{|V_*''|}{V_*} \rangle \propto N\sigma_Q^2 \left(\frac{M_{\text{Pl}}}{f} \right)^2 \gg 1. \quad (2.2)$$

We conclude HM instantons are irrelevant for the decay of quadratic domain minima.

2.2 Coleman-de Luccia decays

When the condition (2.1) for HM decays is not satisfied, decays will proceed via CdL transitions. We are mainly interested in studying the decay of vacua with small vacuum energy, those that can contain collapsed structures like galaxies. By Weinberg’s famous anthropic argument [1] these have $|V_{\text{vac}}| \lesssim 10^{-120} M_{\text{Pl}}^4$. This very narrow band is nevertheless densely populated in the landscapes we are studying [5, 6]. For $|V_0| \ll \Lambda^4$, all such vacua are in the quadratic domain. As mentioned above this ensures the validity of the approximations we

make to certain characteristics of the potential such as the location of critical points, their heights and their Hessian eigenvalues.

In general little is rigorously known about tunneling in high-dimensional landscapes, especially when the effects of gravity are included. For flat space tunneling more is known. For instance, we can assume the dominant instanton has maximal spherical symmetry [22, 23]. In the following, we will use the thin-wall approximation to bound the decay rates. Thin-wall should be justified when (2.1) is strongly violated, and the numerical checks described in section 2.2.3 support this conclusion.

2.2.1 Neighboring minima

A given (quadratic domain) vacuum can tunnel into any neighboring vacuum that has lower energy.² A method that accurately locates the neighboring minima can be deduced from some considerations regarding the potential (1.1). V is invariant under shifts of the arguments of the cosines by $2\pi\mathbf{v}_k$, where \mathbf{v}_k is an integer P -vector with k non-vanishing components of ± 1 . We call a degree- k neighbor a vacuum that is displaced from another vacuum by a shift in the N -dimensional axion field space, such that k cosine arguments shift by roughly 2π , i.e.

$$\mathcal{Q}\theta_{\text{neighbor}} = \mathcal{Q}\theta_{\text{vacuum}} + 2\pi\mathbf{v}_k. \quad (2.3)$$

Each vacuum has at most $3^P - 1$ neighboring vacua. However, since the argument of each cosine is a linear combination of the N axion fields, shifting the cosine arguments this way requires solving a set of P linear equations in N variables. When $P \leq N$ a solution always exists. When $P > N$ there are more equations than variables and such shifts (in general) do not exist. Nonetheless, when $P, N \gg P - N$ one can solve these equations approximately. This shows that low-lying vacua have roughly 3^P neighbors [6]. A priori, the decay could proceed by tunneling to any of this huge number of neighboring vacua. A special class are those separated from the decaying vacuum by shifting the argument of a *single* cosine by $\pm 2\pi$. There are $2P$ such $k = 1$ neighbors, some fraction of which have lower vacuum energy density than the decaying vacuum.³ We refer to these $k = 1$ neighbors as “face neighbors” (because they are separated from the decaying vacuum by a face of a cube in the auxiliary field space defined in [6]).

Each face neighbor minimum is separated from the decaying vacuum by a barrier of height approximately $2\Lambda^4$. The top of the barrier is generically a degree $k = 1$ saddle point. Neighbors where k cosines shift by 2π are typically separated from the decaying vacuum by a barrier with height approximately $2k\Lambda^4$, and by a degree- k saddle point. We will call these “degree- k neighbors”. The typical distance to a degree- k neighbor scales as \sqrt{k} (due to a famous result of Pythagoras). Hence in addition to being the set of minima separated from the parent vacuum by the lowest barriers, the face neighbors are also those typically located within the shortest distance. This makes it plausible that the dominant decay channel will be to a face neighbor.

²There may be decay channels to minima even further away, but these are presumably suppressed. We only consider tunneling to lower energy vacua in this paper. Upward transitions from de Sitter minima are not impossible, but are exponentially suppressed.

³The amount depends on the height of the decaying minimum above the global minimum. The number will be small for very low-lying minima, while for higher-lying minima (but still in the quadratic domain) it is well-approximated by P .

2.2.2 Thin-wall tension

We write the semiclassical bubble nucleation rate per unit four-volume as

$$\Gamma \sim A e^{-B}. \quad (2.4)$$

For a single scalar in flat space, and in the thin-wall approximation, we have [24]

$$B_{\text{flat}} \sim \frac{27\pi^2}{2} \frac{\sigma^4}{\epsilon^3}, \quad (2.5)$$

where $\epsilon = V_{\text{max}} - V_{\text{min}}$ is the difference in energy density between the two vacua and σ is the tension of the bubble wall,

$$\sigma = \int_{\varphi_{\text{min}}}^{\varphi_{\text{max}}} d\varphi \sqrt{2(V(\varphi) - V_{\text{min}})}. \quad (2.6)$$

($\varphi_{\text{min,max}}$ denote the locations of the lower-lying and higher-lying vacuum respectively.) By introducing a minimal wall tension,

$$\sigma_{\text{min}} \equiv \int_{\varphi_0}^{\varphi_{\text{max}}} d\varphi \sqrt{2(V(\varphi) - V_{\text{max}})}, \quad (2.7)$$

where φ_0 is defined by $V(\varphi_0) = V_{\text{max}}$, the thin-wall formula turns into a lower bound for B_{flat} [25]:

$$B_{\text{flat}} \geq \frac{27\pi^2}{2} \frac{\sigma_{\text{min}}^4}{\epsilon^3}. \quad (2.8)$$

This inequality holds for any (single) scalar field theory — even those for which the thin-wall approximation is not valid. Even if gravitational effects are not negligible, they only serve to increase B for thin-wall tunneling from flat (or nearly flat) spacetime to AdS, which is the case at hand. Finally, we have studied the instantons numerically and found that the thin-wall approximation does seem accurate in the regime we are focusing on (see also section 3.3).

One can now use (2.6) to estimate a lower bound on B , for a decay to a minimum separated from the parent by shifts of k cosines by approximately 2π . For the reason discussed above the height of the saddle point along such a direction is $V_* \sim k 2\Lambda^4$, with second derivative $|V''| \approx N\sigma_Q^2 \Lambda^4 / 2f^2$. The typical field space distance across the barrier is $\Delta\varphi \sim 2\sqrt{2V_*/|V''|} \times \pi/4 \sim \pi(f/\sigma_Q) \sqrt{2k/N}$ (where we've approximated the barrier as a parabola), which gives

$$\sigma_{\text{min}} \gtrsim \Delta\varphi \times \sqrt{2V_*} \sim 2\pi k \sqrt{\frac{2}{N}} \frac{f\Lambda^2}{\sigma_Q}. \quad (2.9)$$

The energy difference ϵ between a zero energy vacuum and one with negative energy cannot exceed V_0 (the energy of the global minimum); $\epsilon < |V_0| \ll \Lambda^4$.

Using these estimates in (2.8) gives

$$B \gtrsim \frac{27\pi^2}{2} \frac{\sigma_{\text{min}}^4}{\epsilon^3} \gtrsim \left(\frac{900}{N}\right)^2 \times \left(\frac{\Lambda^4}{V_0}\right)^3 \times \left(\frac{f}{\sigma_Q \Lambda}\right)^4 \times k^4, \quad (2.10)$$

which suggests that decays to face neighbors with $k = 1$ are dominant. Thus, even with $N \approx 10^3$, with $\Lambda < f$ and/or $V_0 < \Lambda^4$ one can achieve $B \gg 1$. Stability on the order of 10^{10} years requires $B \gtrsim 10^3$ (since $4 \log \frac{10^{10} \text{ years}}{t_{\text{Pl}}} \approx 10^3$). We have suppressed several steps in this analysis to give the reader the option of bypassing technical details if they wish. A thorough derivation can be found in appendices A and B. We also include an analysis of the distribution of vacuum energy differences across the sets of degree- k neighbors in section B.2.

2.2.3 Numerical checks and the “gradient flow approximation”

It is difficult to check these approximations numerically due to the high dimension of the field space. Even in field theories without gravity, to our knowledge the best current codes for studying vacuum decay can only handle roughly $N = 5$ field space dimensions [26]. However, the tools developed in [6] make a semi-analytic check available for axion theories.

As we described above, out of $\sim 3^P$ neighboring vacua the $2P$ “face neighbor” channels are likely to dominate the decay rate. For those decays (or any degree- k neighbor, in general) we can use our techniques to locate the lowest saddle point that separates the parent from the target. We then find the gradient line that connects the parent to the target and passes through this saddle point. The potential along this line can then be treated as if it were the potential for a single scalar, allowing us to numerically compute the CdL instanton. We refer to this as the “gradient flow approximation.” It is approximate because the exact instanton does not necessarily follow the gradient flow, but we expect this method to correctly compute the instanton action and trajectory up to $\mathcal{O}(1)$ corrections. The instanton computes the decay rate and provides the initial conditions for the cosmological dynamics after the tunneling, which we calculate using the full N -dimensional potential (see section 3.3).

We verified that the gradient flow indeed approximates the numerical results of [26] in the “sum of cosines” example considered there, and agrees with the results of the thin-wall analytic approximation we turn to next. At least for our class of potentials, it may be an improvement over the “straight line” approximation introduced in [27].

3 Inflation

We now turn the topic of inflationary dynamics. Much of the discussion in this section applies to general multi-axion theories, but we will focus particular attention on well-aligned theories (which are generic when $N \approx P \gg 1$) where it is a very good approximation to set the phases in the non-perturbative axion potential to zero (see [5] for details).

Let us begin by recalling the action relevant for the inflationary dynamics driven by N canonically normalized axions $\Theta \equiv \sqrt{K}\theta$,

$$S = \int d^4x \sqrt{-g} \left(\frac{M_{\text{Pl}}^2}{2} R - \frac{1}{2} \partial \Theta^\top \partial \Theta - V_{\text{axion}}(\Theta) - V_0 \right). \quad (3.1)$$

Here $g_{\mu\nu}$ is the flat FLRW metric with scale factor $a(t)$,

$$ds^2 = -dt^2 + a(t)^2 d\mathbf{x}^2. \quad (3.2)$$

As above we choose V_0 so that the axion contribution to the potential is non-negative, with its global minimum at zero:

$$V_{\text{axion}} = \sum_{I=1}^P \Lambda_I^4 \left[1 - \cos(\mathbf{Q}\Theta)^I \right], \quad (3.3)$$

where $\mathbf{Q} \equiv \mathcal{Q} K^{-1/2}$ is the charge matrix for the canonically normalized fields Θ .

The equations of motion for the scale factor and axions are

$$\begin{aligned} (\Theta^i)'' &= (\epsilon - 3)(\Theta^i)' - \frac{1}{H^2} \frac{\partial V_{\text{axion}}}{\partial \Theta^i}(\Theta), \\ 3M_{\text{Pl}}^2 H^2 &= \frac{V_{\text{axion}}(\Theta) + V_0}{1 - \epsilon/3}, \end{aligned} \quad (3.4)$$

where $' \equiv d/dN_e$, $N_e = \log a$ denotes the number of e-folds, $H = (da/dt)/a \equiv \dot{a}/a$ is the Hubble scale and the Hubble slow roll parameter ϵ is defined by

$$\epsilon = -\frac{\dot{H}}{H^2} = -\frac{H'}{H}. \quad (3.5)$$

In general the dynamics of this system are quite complicated. However, the evolution may effectively be that of a single field if no isocurvature perturbations are sourced. To make this manifest, following [28, 29], we decompose the fields into a basis defined by unit vectors $\{\mathbf{E}_i\}$ along the inflationary trajectory:

$$\mathbf{E}_i = \frac{\mathbf{P}_{i-1}^\perp \boldsymbol{\Theta}^{(i)}}{\|\mathbf{P}_{i-1}^\perp \boldsymbol{\Theta}^{(i)}\|_2}, \quad \mathbf{P}_i^\perp = \mathbb{1} - \sum_{j=1}^i \mathbf{E}_j \otimes \mathbf{E}_j, \quad (3.6)$$

where $i = 1, \dots, N$, $\mathbf{P}_0^\perp = \mathbb{1}$ and \mathbf{P}_i^\perp is a projection operator onto the subspace perpendicular to $\langle \mathbf{E}_1, \mathbf{E}_2, \dots, \mathbf{E}_i \rangle$. This decomposition of the field is very convenient: \mathbf{E}_1 corresponds to the instantaneous direction of the field velocity, while \mathbf{E}_2 indicates the direction of the acceleration transverse to the field velocity and signals multifield behavior. With this basis in mind we can decompose the second slow roll parameter $\boldsymbol{\eta}$ into components parallel and perpendicular to the field trajectory,

$$\boldsymbol{\eta} = \frac{1}{H} \frac{\ddot{\boldsymbol{\Theta}}}{\|\dot{\boldsymbol{\Theta}}\|_2} = \frac{\boldsymbol{\Theta}'' - \epsilon \boldsymbol{\Theta}'}{\|\boldsymbol{\Theta}'\|_2}, \quad \eta_{\parallel} = \boldsymbol{\eta} \cdot \mathbf{E}_1, \quad \eta_{\perp} = \boldsymbol{\eta} \cdot \mathbf{E}_2. \quad (3.7)$$

This decomposition is particularly well-suited to study perturbations. Curvature (adiabatic) perturbations are described by perturbations in the direction of \mathbf{E}_1 and are the only relevant perturbations for single field inflation. In the case of single field inflation the curvature perturbations can immediately be related to density perturbations. Isocurvature (entropy) perturbations describe the relative decomposition of the energy density into the different field components. In the basis we have chosen, one isocurvature mode is distinguished in that it is the only one that couples to the adiabatic perturbation. This coupling is proportional to η_{\perp} , such that single field behavior is recovered when the trajectory does not turn in field space, i.e. $\eta_{\perp} = 0$.

Given an initial condition, we will be concerned with solving the classical equations of motion and evaluating some basic observables, such as the spectral index and the tensor-to-scalar ratio, using the leading expressions in the slow roll regime. The transverse slow roll parameter η_{\perp} provides us with some information about the consistency of the single field, slow roll approximation.

We can define effective slow roll parameters by differentiating the potential along the inflationary trajectory. These “potential slow roll parameters” are related to those defined above by

$$\epsilon_V \equiv \frac{M_{\text{Pl}}^2}{2} \left(\frac{\partial_{\Theta} V}{V} \right)^2 \approx \epsilon, \quad \eta_V \equiv M_{\text{Pl}}^2 \frac{\partial_{\Theta}^2 V}{V} \approx \epsilon - \eta_{\parallel}, \quad (3.8)$$

where ∂_{Θ} denotes differentiation in the direction of the field space velocity \mathbf{E}_1 , and the approximation is valid in slow roll. In the single field, slow roll approximation where ϵ and $\sqrt{\epsilon} \eta_{\parallel}$ are small and η_{\perp} vanishes throughout the inflationary evolution, the spectral index and tensor-to-scalar ratio are given by

$$n_s \approx 1 - 2\eta_{\parallel} - 4\epsilon, \quad r \approx 16\epsilon, \quad (3.9)$$

which are evaluated at the time when the CMB modes exit the horizon. The amplitude of scalar temperature anisotropies is

$$A_s \equiv \frac{1}{24\pi^2} \frac{V}{\epsilon M_{\text{Pl}}^4}, \quad (3.10)$$

with observed values of $A_s \approx 2.1 \times 10^{-9}$, $n_s \approx .965$, and $r \lesssim .07$ [30]. Using this value for A_s , the scale of inflation is related to the tensor-to-scalar ratio by (see e.g. [31])

$$V_{\text{inf}} \approx 3.2 \times 10^{-8} r M_{\text{Pl}}^4. \quad (3.11)$$

Near low-lying minima the axion potential is approximately quadratic. For inflation, the first requirement is that the field range be long enough to produce a sufficient number of efolds of inflation. In [6], we estimated the “diameter” of the region surrounding a typical minimum to be

$$\mathcal{D} \approx \mathcal{O}(1) \times \frac{Nf}{\sigma_{\mathcal{Q}}}, \quad (3.12)$$

which is valid for $P - N \ll N$. To attain 60 efolds of quadratic inflation requires $\mathcal{D} \gtrsim 10 M_{\text{Pl}}$, which is possible with e.g. $f \sim 10^{-2} M_{\text{Pl}}$, $N \approx P \approx 500$, $\sigma_{\mathcal{Q}} \approx 0.15 \approx \sqrt{11/N}$.

In a quadratic potential, the inflaton is at a distance of about $15 M_{\text{Pl}}$ from the minimum when 60 efolds of inflation remain. The mass-squared of the lightest direction along which inflation can proceed is typically on the order of [6]

$$m^2 \sim \frac{\sigma_{\mathcal{Q}}^2 \Lambda^4}{N f^2}, \quad (3.13)$$

and thus the energy scale 60 efolds before the end of inflation is roughly

$$V_{60} \sim \frac{100}{N} \left(\frac{M_{\text{Pl}}}{f/\sigma_{\mathcal{Q}}} \right)^2 \Lambda^4. \quad (3.14)$$

The slow roll parameter $\epsilon = 1/2N_e$ for a quadratic potential, so the amplitude of the observed scalar perturbations will be approximately

$$A_s \approx \frac{60}{N} \left(\frac{\Lambda}{M_{\text{Pl}}} \right)^2 \left(\frac{\Lambda}{f/\sigma_{\mathcal{Q}}} \right)^2, \quad (3.15)$$

while the scalar spectral tilt and tensor-to-scalar ratio are

$$n_s = 1 - 4\epsilon \approx 0.97, \quad (3.16)$$

$$r = 16\epsilon \approx 0.13. \quad (3.17)$$

For e.g. $f \sim 10^{-2} M_{\text{Pl}}$, $N \approx P \approx 500$, $\sigma_{\mathcal{Q}} \approx 0.15 \approx \sqrt{11/N}$ and GUT scale $\Lambda = f/4$, this gives roughly the observed $A_s \approx 10^{-9}$. Tensors are above the observed upper bound for quadratic inflation. However, as we will see multi-axion theories contain many other types of inflationary trajectories, some of which are far from quadratic, or indeed from those of any standard inflationary potential.

In section 3.1 we discuss inflation generally in axion landscapes. In section 3.2 we numerically analyze inflation in a specific example, where we sample the landscape by choosing the starting point uniformly randomly in the field space (with zero initial velocity). In section 3.3 we discuss inflation where the initial conditions are set by tunneling from a higher minimum. In section 3.4 we briefly sketch reheating when the axions are coupled to a gauge field.

3.1 Sampling the theory

One could study the inflationary dynamics in multi-axion theories, or random ensembles of such theories, in some generality by sampling over the parameters of the Lagrangian (1.2). These consist of the metric on moduli space \mathbf{K} , the charge matrix \mathbf{Q} and the global minimum of the vacuum energy density V_0 . The gravitational contributions to the axion potential will be irrelevant for the dynamics, but they may vastly increase the possible discrete vacuum energy densities in the theory, so we will assume that V_0 can be tuned to arbitrary accuracy. A choice for the parameter ensembles that is loosely inspired by explicit compactifications of string theory [7–15] is as follows

1. The metric \mathbf{K} is a positive definite random matrix (for instance a Wishart or inverse Wishart matrix) with largest eigenvalue $f_N^2 \lesssim M_{\text{Pl}}^2$.
2. The axion charge matrix \mathbf{Q} is a sparse matrix of i.i.d. random integers with a fraction $\gtrsim 3/N$ of non-vanishing entries.
3. The background vacuum energy density V_0 , uniformly distributed between $\pm M_{\text{Pl}}^4$.

Even after fixing an ensemble of effective theories or even a unique theory, significant uncertainty remains due to the unknown weight with which different cosmological histories contribute to the distributions of observables — i.e., the measure problem of inflationary cosmology. For example, there may exist a significant selection bias towards small final vacuum energy densities and sufficient inflation. To at least partially account for these selection biases we only retain inflationary trajectories that satisfy the following:

1. The vacuum energy density in the minimum the trajectory ends in is not substantially larger than the observed dark energy density in our universe [1].
2. Inflation lasts long enough to solve the horizon and flatness problems, which here for simplicity we take to mean that $N_e \geq 60$.

In the related work [16] we demonstrate how these assumptions follow from the single requirement of structure formation.

Since we assume that the background vacuum energy density is roughly uniformly distributed, any of the vacua of the axion theory may correspond to the cosmologically relevant late time vacuum. We will reject any dynamics that do not terminate in a vacuum with vanishingly small vacuum energy density, so we can obtain a representative sample of inflationary observables by picking a representative sample of vacua at energy densities $V_{\text{vac},i}$ and then setting $V_0 = -V_{\text{vac},i}$. For each i we then choose initial conditions that are uniformly distributed over the periodic domain of the axion potential, and reject any trajectory that gives rise to less than 60 e-folds of inflation, or does not terminate at a vacuum with vanishing energy density, consistent with our assumptions about selection bias. In fact, it is not necessary to sample all vacua, nor to consider initial conditions uniformly distributed over the entire periodic domain of the potential. Merely considering the attractor regions in the vicinity of a representative sample of potential late time vacua provides a representative sample of the inflationary dynamics. This allows for a systematic study of potentials with exponentially many distinct vacua. We illustrate some possible inflationary trajectories for the specific potential discussed in the next section in figure 1.

Now that we discussed how to systematically sample ensembles of axion theories we might embark on a detailed study of the distribution of inflationary observables. However,




Figure 1. Complete sample of inflationary trajectories with more than 60 efolds terminating at low-CC vacua for theory defined in (3.18). The full trajectories are shown in gray, while the last 60 efolds are colored red. The contour plot represents the axion potential. The periodic domain is highlighted and surrounded by gray solid lines. The boundaries of all tiles are denoted by dashed lines.




Figure 2. Sample of tensor-to-scalar ratios r and spectral indices n_s for the particular axion theory discussed in section 3.2. The red crosses denote trajectories with 60 or more efolds of inflation that terminate in the set of minima with $V_{\text{vac}} \approx 0.64\Lambda^4$, while the blue crosses denote > 60 efold trajectories that terminate in the vacua with all other vacuum energies in (3.19). As discussed in the text, for a trajectory terminating in vacuum i , V_0 is set equal to $-V_{\text{vac},i}$ so that the total energy density vanishes at the endpoint. For comparison we show the observables corresponding to chaotic, natural and hilltop inflation. Some data points are outside the range displayed.

as long as the measure-dependent weight of each inflationary trajectory is unknown such a study is tentative to some extent, as we cannot make definite predictions for cosmological observables. Still, it may be instructive to sample the inflationary dynamics. A comprehensive study of this kind is beyond the scope of this work, so in the following section we simply consider one particular axion theory.

3.2 An explicit example

Let us consider a particularly simple axion potential to illustrate our technique for the systematic sampling of inflationary trajectories. The theory is discussed at length in the

appendix D of [6], but here we have set the phases to zero. The relevant parameters are

$$N = 2, \quad P = 3, \quad \mathbf{K} \approx M_{\text{Pl}}^2 \begin{pmatrix} 107 & 41.4 \\ 41.4 & 48.2 \end{pmatrix}, \quad \mathbf{Q} = \begin{pmatrix} 1 & 1 \\ 2 & -3 \\ -3 & 0 \end{pmatrix}, \quad \Lambda_I = \Lambda, \quad (3.18)$$

where we chose the largest eigenvalue of the kinetic matrix $f_2^2 = (11M_{\text{Pl}})^2$ to allow for inflation, as axion alignment is inefficient at $N = 2$. As discussed above, V_0 is chosen to successively set each of the vacua to vanishing vacuum energy density. We chose initial positions that are uniformly distributed over the periodic domain of the potential, and vanishing initial velocities. Solving the equations of motion (3.4) for the classical trajectories, and selecting all trajectories that terminate at vanishing energy density after more than 60 efolds of inflation, we obtain a representative sample of the inflationary dynamics, as illustrated in figure 1. There are nine stable vacua, four of which are doubly degenerate, at vacuum energy densities

$$\frac{V_{\text{vac},i}}{\Lambda^4} \approx \{0, 0.17, 0.64, 1.3, 1.9\}. \quad (3.19)$$

The field ranges (defined by the distance to the edge of the “tile” surrounding the minimum, cf. [6]) along the lightest direction around each of the vacua are given by

$$\begin{aligned} \mathcal{R}_{\text{light},+,i}/M_{\text{Pl}} &= \{16.4, 17.8, 13.8, 19.6, 18.8\}, \\ \mathcal{R}_{\text{light},-,i}/M_{\text{Pl}} &= \{16.4, 15, 18.9, 11.4, 4.4\}. \end{aligned} \quad (3.20)$$

These field ranges can be read off from figure 1.

Since there are multiple fields active during inflation it is not easy to obtain the correlation functions of perturbations. Still, we can evaluate the spectral index and the tensor-to-scalar ratio assuming single field, slow roll inflation. The resulting observables are shown in figure 2. Clearly a very wide range of observables is possible, even within this extremely simple theory.⁴ This is a very direct example of how large the theoretical uncertainties remain, even if we were able to uniquely identify the effective theory governing our landscape.

It is important to check whether the single field, slow roll approximation is valid. In figure 3 we display the effective potential and the slow roll parameters ϵ , η_{\parallel} and η_{\perp} for the last 60 efolds of inflation of three specific inflationary runs. The examples were chosen to illustrate the wide variety of dynamics: the slow roll parameter η_{\perp} can be large, or small, compared to η_{\parallel} . A more sophisticated multifield analysis of the perturbations is required to study the possible non-Gaussianity signatures of these dynamics. Some of the trajectories exhibit turns, and the tensor-to-scalar ratio can vary between 5×10^{-4} and 1. Considering the effective potential along the trajectory we see that inflation does not proceed in a quadratic potential. In part this is because we chose the scale of the kinetic matrix such that Planckian displacements are possible, but the quadratic domain alone is not large enough to support 60 efolds of inflation. We expect this qualitative finding to hold much more generally than in this particular example: inflation in axion theories allows for an extremely wide range of observables and inflationary energy scales.

⁴That the inflationary phenomenology of multiple-axion theories is richer than that of single field natural inflation was shown in [32] for two-axion theories (with $N = P = 2$).




Figure 3. Effective potential and slow roll parameters during the last 60 efolds of three particular trajectories. $N_e - N_{60}$ denotes the number of efolds before the end of inflation. The observable 60 efolds from the end of inflation in the single field, slow roll approximation are: (a) $n_s = 0.73$, $r = 5 \times 10^{-4}$; (b) $n_s = 0.96$, $r = 0.03$; (c) $n_s = 0.93$, $r = 0.03$. Ticks along the potential mark $\Delta N_e = 10$ intervals.

3.3 Inflation after tunneling

We now turn to the specific scenario where the inflationary initial conditions originate from the decay of a meta-stable vacuum. Generally speaking, inflation after barrier tunneling seems to require fine-tuning. The condition for thin-wall tunneling is that $M_{\text{Pl}}^2 |V_*''|/V_* \gg 1$ (cf. (2.1)), while a necessary condition for slow roll is the opposite, that $\eta_V \equiv M_{\text{Pl}}^2 |V''|/V \ll 1$. These conditions are not logically incompatible because the former applies at the maximum of the barrier V_* while the latter applies to the potential slope after the barrier, but there is nevertheless a clear tension [33, 34].

One of the interesting features of random multi-axion theories is the existence of a hierarchy of Hessian eigenvalues — the fact that at large N different directions in field space can have very different second derivatives. Since the least-action path for tunneling tends to coincide with directions in which the barrier is thinnest and the height is lowest, it is plausible that tunneling will proceed in directions where $|V_*''|/V_*$ is large. Tunneling in such a direction can leave the field displaced from the minimum along a direction or directions with much smaller $|V''|/V$ (see also [35]). Therefore inflation after tunneling does not necessarily require tuning beyond the large number of fields $N \gg 1$.

If the tunneling creates a region with a field value that is in or near the “quadratic domain” of a low-lying minimum, the potential will be roughly quadratic. For a given choice of axion parameters this makes sharp inflationary predictions, as essentially all low-lying minima in the class of theories we are considering are very similar. For instance, the amplitude of density perturbations will be (3.15).

As discussed in section 2, our analytic control over the potential is strongest for minima that are “low-lying”; that is, those with vacuum energy less than Λ^4 above the global minimum of the potential. When $V_0 \ll \Lambda^4$ this includes all small-CC vacua in which structure can form. However, it does not include all parent vacua from which the universe might have tunneled to a given small-CC target. Inflation following tunneling from a high minimum could in principle take place on some feature of the potential outside the quadratic region

surrounding the target minimum where the inflationary trajectory should end. (In section 3.1 we analyzed general inflationary trajectories.)

We cannot rule out the possibility that inflationary histories might be dominated by such non-quadratic potentials, but for the rest of this section we will focus on inflation following tunneling that takes place in the quadratic region. In the quadratic region a necessary condition for at least 60 efolds of inflation is that the typical field range in (3.12) should satisfy

$$\mathcal{R} \sim \frac{\mathcal{D}}{2} > 15M_{\text{Pl}}, \quad (3.21)$$

where \mathcal{D} was estimated in (3.12). Assuming the parameters are such that \mathcal{R} satisfies this condition, tunneling from high minima should sometimes produce $N_e > 60$ efolds, with a power spectrum set by (3.15).

We can say much more about tunneling between low-lying minima. Clearly, achieving large amounts of inflation following a tunneling from a low-lying parent is more difficult, because the starting point on the potential is lower. Nevertheless we will see that it is possible, albeit with more restrictive conditions on the parameters, and we will exhibit an explicit numerical example.

As discussed in section 2, the dominant tunneling trajectories are generally those between neighboring minima that are separated by a 2π shift in one or a few cosines ($k = 1$ or $k = \text{few}$, respectively). The mean separation between face neighbor ($k = 1$) vacua is⁵

$$\langle \|\Theta_{\text{parent}} - \Theta_{\text{target}}\|_2 \rangle \equiv \langle \|\Delta\Theta\|_2 \rangle \approx \frac{2\pi f}{\sigma_Q} \times \mathcal{O}(1). \quad (3.22)$$

Importantly, the distribution of distances has a polynomial tail (see appendix section A.3). Vacuum separations significantly larger than (3.22) are much more frequent than separations significantly smaller than (3.22) due to this tail.

When the separation is greater than M_{Pl} (as required for inflation in the quadratic regime), we expect roughly half of the vacuum separation to be relevant for a possible period of inflation after the tunneling event. There are two reasons for this. First, the saddle point is located roughly halfway in between the parent and the target vacuum. Second, the field space distance traversed by a Coleman-de Luccia instanton is in general sub-Planckian. To see this, note that if the instanton enters the regime of slow roll, dimensional analysis suggests that

$$\Delta\phi = \dot{\phi} \Delta\tau \approx \frac{V'}{3H} \Delta\tau \leq \frac{V'}{3H^2} \approx \sqrt{2\epsilon_V} M_{\text{Pl}}, \quad (3.23)$$

where we have used the slow roll equations and the fact that the instanton exists only for a Euclidean time of order $1/H$ (the radius of the four-sphere). If $\Delta\phi \gtrsim M_{\text{Pl}}$, (3.23) implies that a slow roll condition would be violated, making a significant support of the instanton in the slow roll regime inconsistent.

If (as just argued for above) roughly half the field space separation between the minima is available for inflation following tunneling, the probability to find at least 60 efolds of inflation scales (for $N \gg 1$ and $P - N \ll N$) as

$$\begin{aligned} \text{Prob}(N_e > 60) &\sim \text{Prob}(\|\Delta\Theta\|_2 > 30M_{\text{Pl}}) \\ &\sim \left(\frac{2\pi f/\sigma_Q}{30M_{\text{Pl}}} \right)^{P-N+1}, \end{aligned} \quad (3.24)$$

⁵For a neighbor of degree k , the mean vacuum separation scales $\propto \sqrt{k}$, see section 2.2.1.




Figure 4. The potential between two “face neighbor” minima in an example with randomly chosen charges \mathcal{Q} , plotted along the gradient flow line that passes over the saddle in between. The red dots indicate the two extreme values of the field for the approximate Coleman-de Luccia instanton found numerically using this one-dimensional potential. There are approximately 109 efolds of inflation following tunneling, and (as expected for inflation post-tunneling) the inflationary parameters 60 efolds from the end are close to those of standard quadratic inflation. The axion model parameters for this example are $N = 100$, $P = 101$, $\sigma_{\mathcal{Q}}^2 = 1/20$, $f = 10^{-1} M_{\text{Pl}}$ and $\Lambda = 10^{-2} M_{\text{Pl}}$. Most randomly drawn \mathcal{Q} matrices with these parameters will not yield this much inflation after tunneling between face neighbors (cf. (3.24)); this example was found after $\mathcal{O}(10)$ draws.

where the final scaling is estimated from numerical observations of the tail of the distribution, as discussed in section A.3.

We have numerically tested this paradigm by constructing an effective 1D scalar potential which captures the essential properties of the axion potential between the two vacua. Using the numerical technique described in section 2.2.3, the 1D effective potential was constructed by following the gradient of the full N -dimensional potential starting at the parent minimum and ending at the target minimum, *and* passing directly through the (degree one) saddle in between. This 1D potential indeed has the expected shape: a super-Planckian range slow roll regime starting at one minimum that connects to a sharp, sub-Planckian range barrier with height $\approx 2\Lambda^4$, connecting back to another super-Planckian range slow roll regime that ends at the other minimum. Using this 1D potential we solved the Euclidean Einstein and scalar field equations in the inverted potential to find the instanton, and then solved the Lorentzian Einstein equations in the full N -dimensional field space, using the extreme value of the instanton as an initial condition, to find the evolution after tunneling (including any inflation). We plot an example in figure 4.

3.4 Reheating

Axions may interact with gauge fields F via the coupling

$$\frac{\alpha}{8\pi f_{\text{inf}F}} \Theta_{\text{inf}} F \tilde{F}, \quad (3.25)$$

where α is a dimensionless coupling constant and $f_{\text{inf}F}$ is the effective axion decay constant for the inflaton Θ_{inf} . The interaction (3.25) is topological when the inflaton evolves slowly, but becomes important at the end of slow roll inflation. This allows an efficient energy transfer from the inflationary to the gauge field sector that can drive reheating. Reheating proceeds

through a combination of perturbative and non-perturbative processes. More details appear in [16].

4 Light axion phenomenology

We now turn towards a brief discussion of light axion phenomenology [36–38]. In particular we will be interested in whether a coupling between QCD and light axions can resolve the strong CP problem and if fuzzy dark matter can be accommodated in a multi-axion theory without fine-tuning.

Thus far we assumed that the leading non-perturbative contributions to the potential (1.1) stabilize all axions, i.e. the charge matrix \mathbf{Q} , or equivalently the canonically normalized charge matrix $\mathbf{Q} \equiv \mathbf{Q} \mathbf{K}^{-1/2}$, is full rank. This is consistent, as very light axions would be frozen by Hubble friction in the early cosmological evolution. In order to study extremely light states, however, let us now assume that the P leading and relatively strong non-perturbative effects stabilize all but $L \geq 1$ of the N axions, that is, $\text{Rank}(\mathbf{Q}) = N - L < N$ (cf. appendix A of [6]). It is generally believed that theories of quantum gravity do not permit global continuous symmetries [39–48]. Therefore, at least the gravitational axion potential V_{gr} breaks the L remaining shift symmetries. Since we are interested in light axions, we assume that the energy scale associated to the gravitational potential V_{gr} is vastly smaller than that of the leading potential, so when the light fields are relevant we can ignore the dynamics of fields stabilized by the P leading non-perturbative effects. The L -dimensional subspace relevant for low energy dynamics is then the null space of the leading charge matrix \mathbf{Q} . We define an orthonormal basis $\{\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_L\}$ of $\ker(\mathbf{Q})$,

$$\mathbf{Q} \mathbf{t}_l = \mathbf{0}, \quad \forall l \in \{1, \dots, L\}, \quad (4.1)$$

which we extend to an orthonormal basis of the full field space \mathbb{R}^N by $N - L$ vectors $\mathbf{t}_{L+1}, \mathbf{t}_{L+2}, \dots, \mathbf{t}_N$. We further decompose the field into light and heavy components via

$$\boldsymbol{\Theta} = \mathbf{T} \begin{pmatrix} \boldsymbol{\Theta}_{\text{light}} \\ \boldsymbol{\Theta}_{\text{heavy}} \end{pmatrix} = (\mathbf{T}_{\text{light}} | \mathbf{T}_{\text{heavy}}) \begin{pmatrix} \boldsymbol{\Theta}_{\text{light}} \\ \boldsymbol{\Theta}_{\text{heavy}} \end{pmatrix}, \quad (4.2)$$

where we have split $\boldsymbol{\Theta}$ into a piece $\boldsymbol{\Theta}_{\text{light}}$ of length L and a piece $\boldsymbol{\Theta}_{\text{heavy}}$ of length $N - L$, and the matrices $\mathbf{T}_{\text{light,heavy}}$ are composed by placing the $\mathbf{t}_{1,\dots,L}$ respectively the $\mathbf{t}_{L+1,\dots,N}$ on consecutive columns, see also [6].

It will be instructive to consider two distinct sectors that contribute to the subleading non-perturbative potential: the axions couple to QCD instantons, as well as gravitational wormhole instantons [49]. The former coupling gives rise to a Peccei-Quinn (PQ) axion that can solve the strong CP problem [50], while the latter coupling is motivated by popular beliefs about quantum gravity. As pointed out in [38], gravitational strength breaking of axion shift symmetries might result in a viable candidate for ultra-light axion dark matter. However, the two mechanisms might seem mutually exclusive. The subleading contributions couple to all axions, including the QCD axion and there is no reason to expect that the relative phase between gravitational and gauge instantons is correlated. This imposes an upper bound on the strength of gravitational instantons in order not to spoil the PQ solution to the strong CP problem [51]. On the other hand, the QCD axion could potentially give rise to an overwhelming amount of phenomenologically unacceptable dark matter, which in turn strongly constrains the effective decay constant or the initial misalignment of the axion acting as a PQ symmetry (see e.g. [52]).

4.1 Light and lightest axions

We write the potential relevant for the L canonically normalized light axions Θ_{light} as

$$V_{\text{gr}}(\Theta_{\text{light}}) = \Lambda_{\text{QCD}}^4 [1 - \cos(\mathbf{Q}_{\text{QCD}} \mathbf{T}_{\text{light}} \Theta_{\text{light}} + \delta_{\text{QCD}})] + \sum_{\alpha} M_{\text{Pl}}^4 e^{-S^{\alpha}} [1 - \cos(\mathbf{Q}_{\text{gr}} \mathbf{T}_{\text{light}} \Theta_{\text{light}} + \delta_{\text{gr}})^{\alpha}] , \quad (4.3)$$

where α is an index running over the subleading contributions to the axion potential and \mathbf{Q}_{gr} and \mathbf{Q}_{QCD} are the integer charges of the gravitational and QCD instantons, respectively. We've also defined the canonical charge matrices

$$\mathbf{Q}_{\text{gr}} = \mathbf{Q}_{\text{gr}} \mathbf{K}^{-1/2}, \quad \mathbf{Q}_{\text{QCD}} = \mathbf{Q}_{\text{QCD}} \mathbf{K}^{-1/2}, \quad (4.4)$$

where \mathbf{Q}_{gr} is the integer charge matrix of the gravitational instantons. In defining the coordinates θ in (1.2) we assumed that any instanton has integer charges with respect to all axions. Gravity indiscriminately couples to all axions, so we expect that \mathbf{Q}_{gr} contains the entire integer lattice \mathbb{Z}^N . Since the instanton actions depend on the instanton charges, for large instanton actions only a small number of gravitational instantons will provide a relevant contribution. Let us consider instantons whose action can be parametrized as

$$S^{\alpha} = \mathcal{S} M_{\text{Pl}} \|\mathbf{Q}_{\text{gr}}^{\alpha}\|_2 + \delta S^{\alpha} = \mathcal{S} M_{\text{Pl}} \sqrt{\mathbf{Q}_{\text{gr}}^{\alpha} (\mathbf{Q}_{\text{gr}}^{\alpha})^{\top}} + \dots, \quad (4.5)$$

where the ellipses denote corrections to the classical instanton action that we will ignore and, for clarity, $\mathbf{Q}_{\text{gr}}^{\alpha} = (\mathbf{Q}_{\text{gr}})_{i=1,\dots,N}^{\alpha}$ is the α^{th} row of \mathbf{Q}_{gr} which contains the couplings of the α^{th} gravitational instanton to the N axions in the Θ -basis. For the example of Euclidean wormholes the prefactor of the classical action is given by $\mathcal{S} = \sqrt{6}\pi/8 \approx 1$ [44, 53, 54]. We note that while we can compute the action for Euclidean wormholes, our argument holds more generally as long as the parametrization (4.5) is valid and $\mathcal{S} \approx \mathcal{O}(1)$.

While it is generally believed that quantum gravity breaks all continuous global symmetries, it is not immediately obvious to what extent these symmetries are broken. One natural guess is to assume that all gauge interactions mediate forces that are no weaker than the gravitational interactions. This assumption is known as the weak gravity conjecture [40]. The conjecture has been generalized to multi-axion theories in [42], where it becomes the requirement that the convex hull of the vectors

$$\mathbf{z}^a = \frac{\mathbf{Q}^a}{S^a} M_{\text{Pl}}, \quad (4.6)$$

contains the unit ball.⁶ Note that for the actions (4.5) all vectors \mathbf{z}^a have length $1/\mathcal{S}$, so that whenever $\mathcal{S} < 1$ the convex hull condition of the weak gravity conjecture is satisfied. Perhaps not surprisingly, gravitational instantons roughly saturate this bound.

Consider now the case of $L = 2$ light axions. Whenever the QCD instantons constitute an important contribution to the potential of the light axions we can simply write the relevant potential as

$$V_{\text{gr}}(\tilde{\Theta}_{\text{light}}) = \Lambda_{\text{QCD}}^4 \left[1 - \cos\left(\frac{\Theta_{\text{QCD}}}{f_{\text{QCD}}} + \delta_{\text{QCD}}\right) \right] + M_{\text{Pl}}^4 e^{-S^1} \left[1 - \cos\left(\frac{\Theta_{\text{DM}}}{f_{\text{DM}}} + \delta_{\text{gr}}^1\right) \right], \quad (4.7)$$

⁶The index a runs over *all* the non-perturbative contributions to the potential, and \mathbf{Q} here is the full canonically normalized charge matrix (i.e. it contains the rows of what we have previously called \mathbf{Q} , namely the canonically normalized charge matrix of leading instantons, \mathbf{Q}_{gr} and \mathbf{Q}_{QCD}).

where S^1 denotes the action corresponding to the most important gravitational instanton contribution and we defined the fields $\tilde{\Theta}_{\text{light}} = (\Theta_{\text{QCD}}, \Theta_{\text{DM}})$ by the linear combinations

$$\frac{\Theta_{\text{QCD}}}{f_{\text{QCD}}} \equiv \mathbf{Q}_{\text{QCD}} \mathbf{T}_{\text{light}} \Theta_{\text{light}}, \quad \frac{\Theta_{\text{DM}}}{f_{\text{DM}}} \equiv \mathbf{Q}_{\text{gr}}^1 \mathbf{T}_{\text{light}} \Theta_{\text{light}}, \quad (4.8)$$

and we assumed $\Lambda_{\text{QCD}}^4 \gg M_{\text{Pl}}^4 e^{-S^1}$ so we can drop all other gravitational contributions. The relevant axion decay constants are given by

$$f_{\text{QCD}} = \|\mathbf{Q}_{\text{QCD}} \mathbf{T}_{\text{light}}\|_2^{-1}, \quad f_{\text{DM}} = \|\mathbf{Q}_{\text{gr}}^1 \mathbf{T}_{\text{light}}\|_2^{-1}. \quad (4.9)$$

Note that the transformations leading to the light axions Θ_{QCD} and Θ_{DM} are not orthogonal, so kinetic couplings can remain to the heavy fields. However, in the parameter regime we are interested in there is a vast hierarchy in the axion masses, such that the heavy fields will be stationary when the light fields are dynamical and we can ignore these kinetic couplings.

4.2 Fuzzy dark matter

Let us discuss the impact of the lightest axion Θ_{DM} on the cosmological evolution. This axion is frozen during most of the cosmological history, but begins to oscillate when the Hubble scale drops below the axion mass. After this time, the axion will oscillate in its approximately quadratic potential and act as dark matter. In this section we merely comment on some generic features, and defer a more detailed discussion to [16].

In principle, the gravitational sector contributes an infinite number of terms to the non-perturbative potential: the charges \mathbf{Q}_{gr} contain all sites of the lattice \mathbb{Z}^N . However, in the regime of perturbative control only a small number of those terms will be relevant for low-energy physics. Given the instanton actions (4.5), the leading non-perturbative effects correspond to those sites in the charge lattice with smallest two-length $\|\mathbf{Q}_{\text{gr}}^\alpha\|_2$. In the simple ensemble where $\mathbf{K} = f^2 \mathbb{1}$ is fixed,⁷

$$\min_{\alpha} (\|\mathbf{Q}_{\text{gr}}^\alpha\|_2) = 1/f. \quad (4.10)$$

With this we can estimate the most important contribution to the lightest axion Θ_{DM} . To that end, let us approximate the light field space directions $\mathbf{T}_{\text{light}}$ as isotropic and independent of the charges corresponding to the most important gravitational instanton. In general this approximation is violated: both $\mathbf{T}_{\text{light}}$ and \mathbf{Q}_{gr} depend on the kinetic matrix \mathbf{K} , and therefore can be correlated to some extent. Assuming this correlation can be neglected, we expect $f_{\text{DM}} \approx \sqrt{N}/\|\mathbf{Q}_{\text{gr}}^1\|_2$. With (4.5), (4.9) and (4.10) we then have for the dark matter decay constant and corresponding instanton action at large N ,

$$f_{\text{DM}} \approx \sqrt{N} f, \quad S^1 \approx S \frac{M_{\text{Pl}}}{f}. \quad (4.11)$$

The mass of the light axion is given by

$$m_{\text{DM}} = \frac{M_{\text{Pl}}^2}{f_{\text{DM}}} e^{-S^1/2}. \quad (4.12)$$

⁷Since the dark matter mass is exponentially sensitive to $\|\mathbf{Q}_{\text{gr}}^\alpha\|_2 = \|(\mathbf{Q}_{\text{gr}} \mathbf{K}^{-1/2})^\alpha\|_2$, this analysis should carefully be redone for different ensembles of \mathbf{K} . Generally an expectation value will have to be taken on the left-hand side of eq. (4.10) and a non-trivial N -dependence will be introduced on the right-hand side (see for example [44] which discusses the cases where \mathbf{K} is a matrix drawn from the Wishart or inverse-Wishart ensemble).

It may be instructive to consider the case of $N \approx 100$ axions, a GUT scale decay constant $f \approx 4 \times 10^{-3} M_{\text{Pl}}$ and $\mathcal{S} = 2.3$. These values give a light axion mass $m_{\text{DM}} \approx 10^{-22} \text{ eV}$, consistent with the mass required to realize fuzzy dark matter. The expected axion decay constant $f_{\text{DM}} \approx 4 \times 10^{-2} M_{\text{Pl}}$ coincides with the value giving the desired dark matter abundance [37, 38],

$$\Omega_{\text{axion}} \sim 0.2 \left(\frac{f_{\text{DM}}}{.04 M_{\text{Pl}}} \right)^2 \left(\frac{m_{\text{DM}}}{10^{-22} \text{ eV}} \right)^{1/2}. \quad (4.13)$$

4.3 QCD axion

In order to estimate the decay constant relevant to the QCD axion, we again first estimate the two-norm $\|\mathbf{Q}_{\text{QCD}}\|_2$. We have, in the $\mathbf{K} = f^2 \mathbb{1}$ ensemble,

$$\|\mathbf{Q}_{\text{QCD}}\|_2 = \|\mathbf{Q}_{\text{QCD}} \mathbf{K}^{-1/2}\|_2 = \frac{\sigma_{\text{QCD}}}{f}, \quad (4.14)$$

where σ_{QCD} denotes the root-mean-squared value of the entries in \mathbf{Q}_{QCD} . Assuming again that the light directions $\mathbf{T}_{\text{light}}$ are uncorrelated with the canonically normalized charges of the QCD axion we roughly expect $\|\mathbf{Q}_{\text{QCD}} \mathbf{T}_{\text{light}}\|_2 \sim \|\mathbf{Q}_{\text{QCD}}\|_2 / \sqrt{N}$, or equivalently

$$f_{\text{QCD}} \sim \sqrt{N} \frac{f}{\sigma_{\text{QCD}}}. \quad (4.15)$$

Finally, the non-perturbative contribution due to QCD stabilizes the corresponding axion at $-\delta_{\text{QCD}} f_{\text{QCD}}$, eliminating the CP-violating phase. However, there exist a large number of gravitational instantons that also couple to the same fields and potentially introduce a large amount of CP violation. Very roughly, we can estimate the shift of the phase due to the gravitational instantons as

$$\frac{\Delta \Theta_{\text{QCD}}}{f_{\text{QCD}}} \lesssim \frac{M_{\text{Pl}}^4 e^{-S^1}}{\Lambda_{\text{QCD}}^4}. \quad (4.16)$$

Using $\Lambda_{\text{QCD}}^4 \approx 10^{-78} M_{\text{Pl}}^4$ and the estimate (4.11) we have

$$\log_{10} \left(\frac{\Delta \Theta_{\text{QCD}}}{f_{\text{QCD}}} \right) \sim 78 - \frac{\mathcal{S}}{\log 10} \frac{M_{\text{Pl}}}{f}. \quad (4.17)$$

Using the values from the end of section 4.2, we see that the gravitational instantons contribute a CP-violating phase of roughly

$$\frac{\Delta \Theta_{\text{QCD}}}{f_{\text{QCD}}} \sim 10^{-172} \ll 10^{-10}, \quad (4.18)$$

where the last inequality denotes the comparison to the experimental bound. Hence gravitational instantons do not spoil the PQ solution to the strong CP problem.

Note that the QCD axion might also act as dark matter and potentially lead to an overwhelming amount of phenomenologically unacceptable dark matter, depending on the initial misalignment angle [52]. The severity of this tuning depends on the parameter choice, and any accidental alignment between the axions. We return to this issue in the more comprehensive discussion in [16].

Acknowledgments

We thank Ben Freivogel, Mark Hertzberg, Liam McAllister, Alex Vilenkin and Jeremy Wachter for useful discussions at various stages of the development of this paper. The work of TB is supported in part by DOE under grants no. DE-SC0011941 and DE-SC0009919 and by the Simons Foundation SFARI 560536. The work of KE is in part supported by the INFN. OJ is supported by a James Arthur Graduate Fellowship. The work of MK is supported in part by the NSF through grants PHY-1214302 and PHY-1820814.

A Neighbors and saddle points

In this appendix we provide a derivation of two results used in section 2 and section 3.3.

The first result is the large N behavior of the Hessian eigenvalues of degree- k saddle points of the potential. These eigenvalues are important because they determine the sharpness of the barrier separating adjacent minima, and hence the characteristics of the instanton that mediates the decay. This is particularly important to determine the lifetime of a low-CC minimum (so that the “parent” minimum has nearly zero CC, and the “target” has lower, probably negative, vacuum energy). Because of the large number of possible decay channels the *distribution* of these eigenvalues is important, not just the mean.

The second result is the large N behavior of the canonically normalized field space distance between two face neighboring, low-lying vacua. This is important for the question of slow roll inflation following tunneling between two such neighbors (although we emphasize again that decay from a higher minimum or other point on the potential might give rise to inflation even when decays between face neighbors do not). Of particular importance is the fact that the distribution of the distances between face neighbors is heavy tailed, so that the extreme cases can be much larger than the mean. For this question the target minimum is the one with nearly zero CC, while the parent has larger, positive CC (although small enough for the quadratic approximation to be valid).

To avoid confusion we use subscripts A and B for the vacua, only specifying which has higher energy when necessary. Familiarity with the geometric picture developed in [6], where the field space is identified with an N -dimensional hyperplane Σ slicing through a uniform lattice in P -dimensional Euclidean space, is assumed.

This appendix is organized as follows. In section A.1 we review the necessary technology from [6], in section A.2 we study the distribution of Hessian eigenvalues in detail, and in section A.3 we study the distribution of face neighbor distances.

A.1 Degree- k saddle points

To study tunneling to/from a quadratic domain vacuum Θ_A in a tile labeled by \mathbf{n}_A , we first identify its nearby tiles. We expect that the dominant decay channels will be between Θ_A and its neighboring vacua. What constitutes a “neighboring” vacuum is to some degree arbitrary, but a reasonable definition adopted in [6] is to deem vacua in tiles whose P -cubes are separated by at most one step in each of the ϕ^J -coordinate directions as “neighbors”, i.e. those at

$$2\pi\mathbf{n}_B = 2\pi\mathbf{n}_A + 2\pi\mathbf{v}, \quad (\text{A.1})$$

where \mathbf{v} is any P -vector with entries in $\{-1, 0, 1\}$, and \mathbf{n}_A and \mathbf{n}_B are the integer P -vectors for the vacuum and neighbor’s lattice points, respectively. There are a total of 3^P candidates

tiles, but only the subset of P -cubes that intersect Σ may actually correspond to neighboring vacua of Θ_A .

We can label the displacement vectors $2\pi\mathbf{v}$ by the number of non-zero entries they contain, i.e. the number of mutually orthogonal steps taken from $2\pi\mathbf{n}_A$ in the ambient space to reach the center of the neighbor cube. We indicate this by a natural number subscript k , writing \mathbf{v}_k from now on. For a given k there are a total of $2^k \binom{P}{k}$ distinct vectors \mathbf{v}_k . A degree- k saddle point of V_{aux} ⁸ lies exactly half-way between the pair of auxiliary minima $2\pi\mathbf{n}_A$ and $2\pi(\mathbf{n}_A + \mathbf{v}_k)$, at

$$\pi\mathbf{n}_{S,k} \equiv 2\pi\mathbf{n}_A + \pi\mathbf{v}_k. \quad (\text{A.3})$$

This is because $\mathbf{n}_{S,k}$ contains $P - k$ even integers and k odd integers, resulting in k of the cosine terms of V_{aux} being maximized at $\phi = \pi\mathbf{n}_{S,k}$, while the rest are minimized. In the vicinity of the auxiliary saddle all P cosines are well-approximated by quadratics. In other words, there are quadratic domains situated not only at the points of the auxiliary lattice $2\pi\mathbb{Z}^P$ (the auxiliary minima), but more generally at the points $\pi\mathbb{Z}^P$ (auxiliary critical points).

For very low values of k , *any* of the $2^k \binom{P}{k}$ integer displacement vectors \mathbf{v}_k will typically identify a nonempty neighboring tile of Θ_A . This follows from the fact that Σ is randomly oriented with respect to the standard basis elements $\{\mathbf{e}^{(J)}\}$ of \mathbb{R}^P , and so the size of $\mathbf{P}^\perp \mathbf{e}^{(J)}$ is $\sim \sqrt{\nu/P} \ll 1$, for all J .⁹ The expected length of $\mathbf{v}_k^\perp \equiv \mathbf{P}^\perp \mathbf{v}_k$ is

$$\|\mathbf{v}_k^\perp\|_2 \equiv \|\mathbf{P}^\perp \mathbf{v}_k\|_2 = \left\| \sum_{J=1}^P v_k^J \mathbf{P}^\perp \mathbf{e}^{(J)} \right\|_2 \approx \sqrt{k\nu/P}, \quad (\text{A.4})$$

where the \approx follows since this is a sum of k non-zero vectors $\mathbf{P}^\perp \mathbf{e}^{(J)}$ of random orientation and length $\sim \sqrt{\nu/P}$.

The largest distance one can move away from Σ by shifting from lattice point $2\pi\mathbf{n}_A$ by $2\pi\mathbf{v}_k$ is $2\pi\|\mathbf{v}_k^\perp\|_2$. This is just a statement of the triangle inequality: the two-norm distance of a candidate lattice point is bounded by

$$d_B \lesssim d_A + 2\pi\sqrt{k\nu/P}, \quad (\text{A.5})$$

where d_A is the ℓ_2 -distance between the original vacuum in Σ and its lattice point in \mathbb{R}^P ,

$$d_A \equiv \|\mathbf{Q}\Theta_A - 2\pi\mathbf{n}_A\|_2. \quad (\text{A.6})$$

Since the vacuum Θ_A admits a quadratic description, d_A is small (the vacuum energy in the $V_0 = 0$ theory is proportional to the two-norm distance in the case of equal couplings).

⁸The auxiliary potential is defined as the following function of P scalars,

$$V(\phi) \equiv \sum_{I=1}^P \Lambda_I^4 \left(1 - \cos(\phi^I) \right). \quad (\text{A.2})$$

The axion potential can be identified as the auxiliary potential evaluated at $\phi = \mathbf{Q}\theta$. See [6] for further details.

⁹Note that this is a much weaker notion of alignment than that of the $\mathbf{t}_a^\#$ of the aligned lattice basis, which have *exponentially* small perpendicular components.

Concretely,

$$\frac{1}{2}d_A^2 = 2\pi^2 \|\mathbf{P}^\perp \mathbf{n}_A\|_2^2 < \mu_{\max} P, \quad \mu_{\max} \ll 1. \quad (A.7)$$

So, for $k\nu/P \ll 1$ we can conclude d_B remains small for all the neighbor candidates defined by (A.1).

As k increases the bound on d_B from the triangle inequality weakens, but for any k one can expect some fraction of the $2^k \binom{P}{k}$ candidates to remain close to Σ . For such \mathbf{v}_k (those which either have particularly short \mathbf{v}_k^\perp themselves, or those whose \mathbf{v}_k^\perp are largely canceled by $\mathbf{P}^\perp \mathbf{n}_A$) it certainly follows that the auxiliary saddle between the neighboring lattice points is *also* in close proximity to Σ . The reason is because the auxiliary saddle's distance from Σ is bounded above and below by d_A and d_B , respectively, or vice-versa.

In other words, for general k not all neighboring P -cubes defined by (A.1) necessarily give rise to quadratic domain vacuum neighbors, but the fraction that *do* are separated from $2\pi\mathbf{n}_A$ by a degree- k saddle of the auxiliary potential which itself has a quadratic domain that is intersected by Σ . Consequently, the physical potential in the region between the \mathbf{n}_A and \mathbf{n}_B tiles is well-described by the orthogonal projection of the specific auxiliary saddle domain. Thus, the physical potential has a degree- k saddle point with ambient coordinates $\mathbf{Q}\Theta_{S,k} \approx \pi\mathbf{n}_{S,k}$, or in canonical coordinates,

$$\Theta_B \approx \Theta_A + 2\pi(\mathbf{Q}^\top \mathbf{Q})^{-1} \mathbf{Q}^\top \mathbf{v}_k \quad (A.8)$$

$$\Theta_{S,k} \approx \Theta_A + \pi(\mathbf{Q}^\top \mathbf{Q})^{-1} \mathbf{Q}^\top \mathbf{v}_k. \quad (A.9)$$

where we've used $\mathbf{Q}\Theta_A \approx 2\pi\mathbf{n}_A$.

Furthermore, the Hessian of the physical potential V evaluated at (physical) saddle points $\Theta_{S,k}$ can be approximated in a simple manner [6]. This was used in section 2 in estimating the stability of quadratic domain vacua Θ_A . First, note the chain rule implies a simple expression for the Hessian of the physical potential $H_{ij} \equiv \partial_i \partial_j V$ in terms of the auxiliary one, $H_{IJ} \equiv \partial_I \partial_J V_{\text{aux}}$. In canonical coordinates,

$$\mathbf{H}_{\text{phys}}(\Theta) = \mathbf{Q}^\top \mathbf{H}_{\text{aux}}|_{\phi=\mathbf{Q}\Theta} \mathbf{Q}, \quad (A.10)$$

where

$$\mathbf{H}_{\text{aux}}(\phi) = \Lambda^4 \text{diag}\{\cos(\phi^1), \dots, \cos(\phi^P)\}. \quad (A.11)$$

At the saddle points (A.3), where

$$\mathbf{Q}\Theta_{S,k} \approx 2\pi\mathbf{n}_A + \pi\mathbf{v}_k, \quad (A.12)$$

we therefore have

$$\mathbf{H}_{\text{phys}}(\Theta_{S,k}) \approx \Lambda^4 \mathbf{Q}^\top \mathbf{D} \mathbf{Q}, \quad (A.13)$$

where \mathbf{D} is a diagonal matrix containing k negative ones and $P - k$ positive ones along the diagonal because

$$\mathbf{H}_{\text{aux}}(\mathbf{Q}\Theta_{S,k}) \approx \mathbf{H}_{\text{aux}}(2\pi\mathbf{n}_A + \pi\mathbf{v}_k) = \Lambda^4 \mathbf{D}. \quad (A.14)$$

¹⁰ μ_{\max} is defined as the maximum quadratic domain vacuum energy, divided by the average height of the potential, here $P\Lambda^4$. For example, $\mu_{\max} \equiv \frac{V_{\max}}{P\Lambda^4}$ resulted in a value of about $\mu_{\max} = 0.014$ in an $N = 100$ and $\nu = 1$ theory [6]. The value is determined by the parameters and one's choice for the threshold value on the ℓ_∞ -displacement in their definition of the quadratic domain. The quantity $\mu \equiv \frac{V_{\text{vac}}}{P\Lambda^4}$ is useful in that it enables a comparison between quadratic domain vacua of different landscapes by removing the routine factor of P arising simply from the total number of cosine terms, which may differ in the two landscapes.

The most negative Hessian eigenvalue at saddle points of V is important in our estimate of the CdL decay rate. Our analytical tools are only applicable to physical saddles well-described by the orthogonal projection of auxiliary ones. As noted, this will in general include higher values of k as well. So, we proceed by studying the eigenvalue distributions of ensembles of matrices $\mathbf{Q}^\top \mathbf{D} \mathbf{Q}$ for the full range of k , from 1 to P .

A.2 Hessian eigenvalues at large N

It is conducive to pull out the overall scales introduced trivially by the random charge matrix \mathbf{Q} and kinetic matrix \mathbf{K} . To that end, define $\hat{\mathbf{Q}}$ such that

$$\mathbf{Q} = \sigma_{\mathbf{Q}} \hat{\mathbf{Q}}. \quad (\text{A.15})$$

Then \mathbf{Q} , the charge matrix in canonical coordinates, is expressed in terms of a Wigner matrix $\hat{\mathbf{Q}}$ whose entries are normally distributed with variance 1, as follows

$$\mathbf{Q} = \frac{\sigma_{\mathbf{Q}}}{f} \hat{\mathbf{Q}}. \quad (\text{A.16})$$

Then the (physical) Hessian is

$$\mathbf{H} = \Lambda^4 \left(\frac{\sigma_{\mathbf{Q}}}{f} \right)^2 \hat{\mathbf{H}}, \quad (\text{A.17})$$

where

$$\hat{\mathbf{H}} \equiv \hat{\mathbf{Q}}^\top \mathbf{D} \hat{\mathbf{Q}}. \quad (\text{A.18})$$

Thus it suffices to study the ensemble (A.18) generated by $\hat{Q}_j^i \sim \mathcal{N}(0, 1)$ at large N (provided $\sigma_{\mathbf{Q}}$ is not too small, $\gtrsim \sqrt{3/N}$ [6, 11]).

The spectrum of $\hat{\mathbf{H}}$ for $k < P$ negative signs in \mathbf{D} must in some sense be bounded by a pair of Marchenko-Pastur distributions; one located on the positive axis and one on the negative axis. This is simply because in the two extremes, $k = 0$ or $k = P$, the matrix \mathbf{D} is either plus or minus the identity, making $\hat{\mathbf{H}}$ exactly plus or minus the Wishart matrix $\hat{\mathbf{Q}}^\top \hat{\mathbf{Q}}$.

In fact, it suffices to study just the evolution of these distributions over half of the k values. For even P , the range $k = 0$ to $k = P/2$ suffices, and for odd P , $k = 0$ to $k = \frac{P-1}{2}$ does, because $\mathbf{D}_{P-k} = -\mathbf{D}_k$. This means the series of $\hat{\mathbf{H}}$ eigenvalue distributions for the second half of k values is given by reflecting each of the distributions in the first half-series onto the negative axis (and then reordering the series of distributions in reverse).

To simplify discussion, let us take P to be even. The numerics confirm the naive expectation of a gradual transition between the Marchenko-Pastur distribution with support on the interval $\sim [1/N, \alpha N]$ for $k = 0$, to the — necessarily symmetric — distribution for $k = P/2$, whilst at each k value always keeping the (possibly disjoint) support interval contained within $[-\alpha N, \alpha N]$. The Marchenko-Pastur distribution for the Wishart ensemble (for all fixed $\nu = P - N$) has a right-edge at $\alpha = 4$, and this holds approximately at finite but large N .

The evolution in $\hat{\mathbf{H}}$ eigenvalue distributions is shown in figure 5. Each time k increases the shape of the PDF changes as follows: a layer of the once-Marchenko-Pastur distribution on the positive axis (corresponding to $k = 0$) is shaved off in a uniform fashion, and its area is deposited onto the negative axis, beginning with a small bump at $-N/2$ ($k = 1$), which ultimately grows into the reflected Marchenko-Pastur shape (on the negative axis). The symmetric distribution that results from flipping exactly half the signs in \mathbf{D} is shown in the bottom-right panel of figure 5.

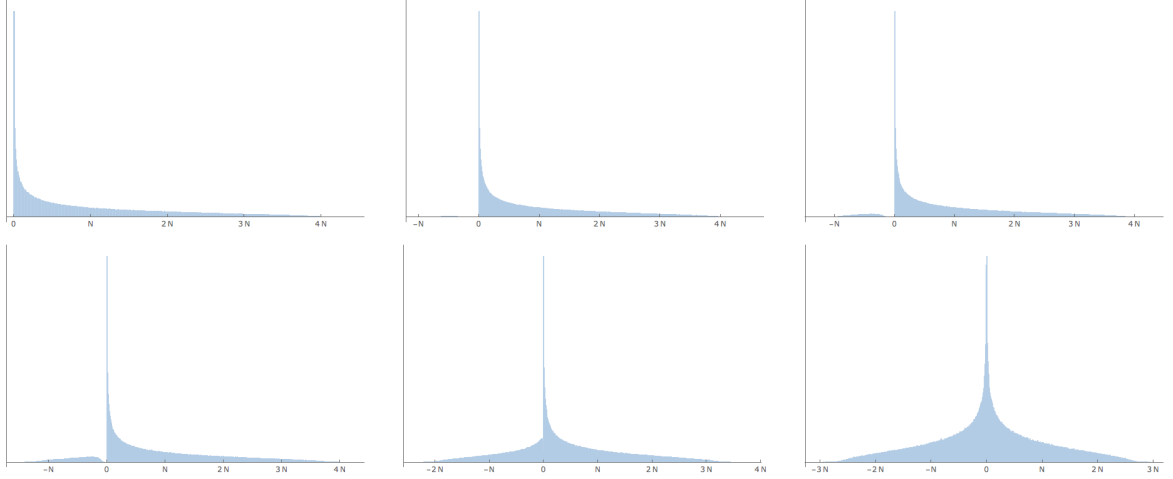


Figure 5. Spectrum of the “sign-flipped Wishart ensemble” defined by $\hat{\mathbf{Q}}^\top \mathbf{D} \hat{\mathbf{Q}}$ where \mathbf{Q} is an i.i.d. $P \times N$ matrix with real, normally distributed entries and \mathbf{D} is a $P \times P$ diagonal matrix with entries ± 1 , as the number of negative signs in \mathbf{D} increases from $k = 1$ to $k = P/2$ from the upper left to lower right.

Turning to the behavior of the minimum eigenvalue of $\hat{\mathbf{H}}$, which we’ll denote by λ_- , consider again the extreme case $k = P$. Here the set of most negative eigenvalues of an ensemble of $\hat{\mathbf{H}}$ is given by the set of *maximum* eigenvalues of the Wishart ensemble $\hat{\mathbf{Q}}^\top \hat{\mathbf{Q}}$, just multiplied by negative one. For $N \gg 1$ and $P - N \ll N$ we have

$$\langle \lambda_- \rangle \approx -\alpha N, \quad (\text{A.19})$$

with $\alpha \approx 4$, with standard deviation of order $N^{1/3}$ (e.g. [55]).

This generalizes to $k < P$. The left edge of the full eigenvalue distribution for an ensemble of $\hat{\mathbf{H}}$ generated from random $\hat{\mathbf{Q}}$ using \mathbf{D} with fixed $k < P$ negative signs converges to $\langle \lambda_- \rangle = -c(k)N$, where $c(k)$ is an order one number less than 4. The scaling of the width of the distribution of λ_- with N is also suppressed with respect to the mean $\langle \lambda_- \rangle$, just as for the Tracy-Widom distribution describing the $k = 0$ and $k = P$ cases. For example, when $k = 1$, $\langle |\lambda_-| \rangle = N/2$ and the standard deviation is $\sqrt{3N/2}$. A sample of the λ_- distributions for varying k is shown in figure 6.

Therefore we see that for all k and at large N , the standard deviation in the smallest (most negative) eigenvalue is much smaller than its mean.

A.3 Face neighbor distances

Each face neighbor is specified by the signed unit-normal vector to the face it shares with vacuum A ’s cube,

$$\pm \mathbf{e}^{(J)} = (0, 0, \dots, 0, \pm 1, 0, \dots, 0). \quad (\text{A.20})$$

In canonical coordinates the face neighbor vacuum’s location is well-approximated by

$$\boldsymbol{\Theta}_B \approx \boldsymbol{\Theta}_A \pm 2\pi(\mathbf{Q}^\top \mathbf{Q})^{-1} \mathbf{Q}^\top \mathbf{e}^{(J)}. \quad (\text{A.21})$$

With the hatted notation introduced in section A.2, the mean of the distribution of the face neighbor distances $\|\boldsymbol{\Theta}_A - \boldsymbol{\Theta}_B\|_2$ is

$$\langle \|\boldsymbol{\Theta}_A - \boldsymbol{\Theta}_B\|_2 \rangle = \frac{2\pi f}{\sigma_Q} \times \langle \|(\hat{\mathbf{Q}}^\top \hat{\mathbf{Q}})^{-1} \hat{\mathbf{Q}}^\top \mathbf{e}^{(J)}\|_2 \rangle. \quad (\text{A.22})$$




Figure 6. Minimum eigenvalue distributions of \hat{H} ensembles defined as $\hat{\mathbf{Q}}^\top \mathbf{D} \hat{\mathbf{Q}}$ for $N = 100$ and $\nu = 3$. Moving from right to left are results for $k = 1, 5, 20, 40, 75$, and 100 negative signs in the diagonal of \mathbf{D} out of $P = N + \nu$. A Gaussian with standard deviation $\sqrt{3N/2}$ centered at $-N/2$ is plotted for the sake of comparison with the $k = 1$ data.

A non-obvious fact is that the hatted quantity on the lefthand side of (A.22) turns out to be order one at large N ,

$$\langle \|(\hat{\mathbf{Q}}^\top \hat{\mathbf{Q}})^{-1} \hat{\mathbf{Q}}^\top \mathbf{e}^{(J)}\|_2 \rangle = \mathcal{O}(1). \quad (\text{A.23})$$

The resulting canonical field distances between face neighbors is

$$\langle \|\Theta_A - \Theta_B\|_2 \rangle \approx \frac{2\pi f}{\sigma_Q} \times \mathcal{O}(1). \quad (\text{A.24})$$

Note the face neighbor distance is suppressed by a factor N relative to the tile diameters, which go like N when $\sigma_Q \sim \mathcal{O}(1)$, and $N^{3/2}$ in the sparse charge matrix case where $\sigma_Q = \mathcal{O}(1)/\sqrt{N}$.

An important aspect of the distribution of $\|\Theta_A - \Theta_B\|_2$, and of the distribution of maximum distances, is that it appears to be heavy-tailed for $\nu \ll N$. Taking the asymptotic form $\|\Theta_A - \Theta_B\|_2^{-f(\nu)}$ for the tail, numerics suggest

$$f(\nu) = 2 + \nu, \quad (\text{A.25})$$

although we have no analytic argument for this relation. A numerical PDF of the maximum distance is shown in figure 7. Though lacking a derivation of the fall-off exponent $f(\nu)$, the basic fact that the distribution is heavy-tailed at all can be understood heuristically. Moreover, this reasoning simultaneously explains the $\mathcal{O}(1)$ expectation for the hatted distance between face neighbors (the quantity on the left-hand side of (A.23)).

Consider first the case when $P = N$, so the charge matrix is square and invertible. Then $(\hat{\mathbf{Q}}^\top \hat{\mathbf{Q}})^{-1} \hat{\mathbf{Q}}^\top$ reduces to $\hat{\mathbf{Q}}^{-1}$, so we recognize the hatted distance between two face neighboring minima

$$\|(\hat{\mathbf{Q}}^\top \hat{\mathbf{Q}})^{-1} \hat{\mathbf{Q}}^\top \mathbf{e}^{(J)}\|_2 \quad (\text{A.26})$$

as the length of the J th column of $\hat{\mathbf{Q}}^{-1}$.




Figure 7. Numerical probability density of the maximum 2-norm length of the columns of the matrix $(\hat{\mathbf{Q}}^\top \hat{\mathbf{Q}})^{-1} \hat{\mathbf{Q}}^\top$, where the entries of $\hat{\mathbf{Q}}$ are i.i.d. according to $\mathcal{N}(0, 1)$. This quantity appears in an estimate for the field space distance between face neighbor minima. In this example $N = 200$, $P = 201$, and we have fit the tail of the distribution to the form $(\text{length})^{-f(1)}$, with $f(1) \approx 3$.

When $P > N$ the matrix $\hat{\mathbf{Q}}$ of course is no longer invertible, but the rectangular matrix $(\hat{\mathbf{Q}}^\top \hat{\mathbf{Q}})^{-1} \hat{\mathbf{Q}}^\top$ is — in a precise sense — the closest thing to the inverse of $\hat{\mathbf{Q}}$. Start with a singular value decomposition of $\hat{\mathbf{Q}}$,

$$\hat{\mathbf{Q}} = \mathbf{V} \mathbf{\Sigma} \mathbf{U}^\top. \quad (\text{A.27})$$

Here \mathbf{U} and \mathbf{V} are orthogonal $N \times N$ and $P \times P$ matrices respectively, and the $P \times N$ matrix $\mathbf{\Sigma}$ contains the singular values of $\hat{\mathbf{Q}}$ along the diagonal of its top $N \times N$ diagonal subblock, followed by $P - N$ zero rows. In other words $\mathbf{\Sigma}$ has the form,

$$\mathbf{\Sigma} = \begin{bmatrix} \sigma_1 & 0 & \dots & 0 \\ 0 & \sigma_2 & \dots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \dots & \sigma_N \\ 0 & 0 & \dots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \dots & 0 \end{bmatrix}. \quad (\text{A.28})$$

The first N columns of \mathbf{V} form an orthonormal basis for the constraint surface Σ when it is embedded in \mathbb{R}^P , while \mathbf{V} 's remaining $P - N$ columns do so for the orthogonal complement.

To simplify notation we label the inverse matrix as $\mathbf{A} \equiv (\hat{\mathbf{Q}}^\top \hat{\mathbf{Q}})^{-1}$. Note that the eigenvectors of \mathbf{A} are the right-singular vectors of $\hat{\mathbf{Q}}$ (the columns of \mathbf{U} , which we'll label as \mathbf{u}_i , and similarly for \mathbf{V}). The eigenvalues of \mathbf{A} are $a_i = 1/\sigma_i^2$. While $\hat{\mathbf{Q}}$ maps the unit vector \mathbf{u}_i to the (generally not normalized) P -vector $\sigma_i \mathbf{v}_i$, the matrix $\mathbf{A} \hat{\mathbf{Q}}^\top$ maps \mathbf{v}_i to $1/\sigma_i \mathbf{u}_i$. This is the sense in which the pair $\mathbf{A} \hat{\mathbf{Q}}^\top$ and $\hat{\mathbf{Q}}$ can be thought of as the rectangular analog of a matrix and its inverse. The rectangularity is reflected in the fact that \mathbf{v}_i plus any P -vector in Σ^\perp still maps to $1/\sigma_i \mathbf{u}_i$.

Going forward we also suppress the (J) superscript on the face vector. To develop an expectation for the magnitude of the lengths $\|\mathbf{A} \hat{\mathbf{Q}}^\top \mathbf{e}\|_2$, start by expanding $\hat{\mathbf{Q}}^\top \mathbf{e}$ in the

right-singular vectors of $\hat{\mathbf{Q}}$,

$$\hat{\mathbf{Q}}^\top \mathbf{e} = \hat{\mathbf{Q}}^\top \mathbf{e}_\parallel = \sum_{i=1}^N b_i \mathbf{u}_i. \quad (\text{A.29})$$

Then

$$\mathbf{A} \hat{\mathbf{Q}}^\top \mathbf{e} = \sum_{i=1}^N \frac{b_i}{\sigma_i^2} \mathbf{u}_i. \quad (\text{A.30})$$

Since \mathbf{U} is orthogonal, the length of $\mathbf{A} \hat{\mathbf{Q}}^\top \mathbf{e}$ is the square root of the sum of the squares of the expansion coefficients,

$$\|\mathbf{A} \hat{\mathbf{Q}}^\top \mathbf{e}\|_2 = \sqrt{\sum_{i=1}^N \frac{b_i^2}{\sigma_i^4}}. \quad (\text{A.31})$$

The spectrum of \mathbf{A} is known because it is an inverse Wishart matrix ($\hat{\mathbf{Q}}$ is a Wigner matrix). The task then amounts to correctly estimating the expansion coefficients b_i . It might be tempting to think these would be roughly equal, or, perhaps a bit more carefully, that they be normally distributed with approximately equal variances. This is not the case, however.

The b_i are strongly correlated with the a_i due to the fact that \mathbf{e} is normalized and $\hat{\mathbf{Q}} \mathbf{A} \hat{\mathbf{Q}}^\top$ is an orthogonal projector. Qualitatively, the components of $\mathbf{A} \hat{\mathbf{Q}}^\top \mathbf{e}$ in large-eigenvalued eigendirections must be small enough that their image under $\hat{\mathbf{Q}}$ have norm ≤ 1 . The correlation can be understood precisely by relating the b_i to the expansion coefficients of \mathbf{e}_\parallel in the *left*-singular vectors (\mathbf{v}_i), which we'll denote by β_i ,

$$\mathbf{e}_\parallel = \sum_{i=1}^N \beta_i \mathbf{v}_i. \quad (\text{A.32})$$

In light of the fact that $\mathbf{A} \hat{\mathbf{Q}}^\top : \mathbf{v}_i \mapsto 1/\sigma_i \mathbf{u}_i$ the expansion coefficients are related by $\beta_i = b_i/\sigma_i$. The β_i^2 must sum to a value ≤ 1 because \mathbf{e} is normalized. Recall that face vectors generically have $\|\mathbf{e}_\parallel\|_2 \approx 1$ and $\mathbf{e}_\perp \ll 1$ for $P - N \ll N$ because Σ is randomly oriented. Likewise, the projection onto Σ itself, \mathbf{e}_\parallel , bears no special relation to the orthogonal directions \mathbf{v}_i . Consequently, we expect values for *these* expansion coefficients — the β_i — to be comparable to one-another. They are the components of a normalized vector so the magnitude to expect is $\beta_i \sim 1/\sqrt{N}$ since,

$$\sum_{i=1}^N \beta_i^2 \approx 1 \quad \Rightarrow \quad N \beta_i^2 \sim 1. \quad (\text{A.33})$$

A rough estimate of the mean/median of $\|\mathbf{A} \hat{\mathbf{Q}}^\top \mathbf{e}\|_2$ can be obtained setting all $\beta_i^2 = 1/N$. One then concludes a value of $\langle \|\mathbf{A} \hat{\mathbf{Q}}^\top \mathbf{e}\|_2 \rangle \approx \langle \sqrt{\lambda_A} \rangle$, where λ_A represents a blindly-drawn eigenvalue of the matrices in an inverse Wishart ensemble.¹¹ Numerically we find the median of this is $\mathcal{O}(1)$. This is (literally) the observation (A.23).

Furthermore, the eigenvalue distribution of the inverse Wishart ensemble is heavy-tailed. This explains the heavy tail of the PDF in figure 7, albeit without an analytic estimate of

¹¹It should be noted that while $1/\sqrt{N}$ is the magnitude to expect for the β_i , the upper bound on any given $|\beta_i|$ is of course still 1. In terms of the $|b_i|$ this translates to a strict boundary at $|b_i| = 1/\sqrt{a_i}$. We have confirmed numerically that a marked decrease in the $|b_i|$ is observed at $3/\sqrt{N a_i}$, which is consistent with having gaussian distributed β_i with variance $1/N$.

the numerically observed fall-off rate $f(\nu)$. An important conclusion is that parent-to-target distances much larger than the median (A.22) are much more frequent than distances much smaller than (A.22) (which appear to be exponentially suppressed).

B Refined stability bound

The purpose of this appendix is provide a more detailed version of the stability analysis performed in section 2 for tunneling from minima with nearly zero vacuum energy. We also account for the variations among the neighbors of a given degree- k set, which is currently ignored in the estimate of an upper bound on thin-wall tunneling exponent B in section 2. The result will be a more precise bound. We also study the distribution of vacuum energy differences and show that this may improve the bound on B for tunneling to face neighbors.

B.1 Stability

The set of degree- k neighbors of a given parent are not identical. They vary both in the widths of the barriers separating them from the parent, and in their difference in vacuum energy density compared to the parent. The worst case scenario from the perspective of a parent's stability, is when the degree- k neighbor with thin-wall tension (σ) fluctuated the most toward small values is *also* that with largest vacuum energy difference, ϵ . Any other correlation between the two quantities would serve to increase B , and render the vacuum more stable.

Recall that

$$\sigma \approx \frac{2\pi k \Lambda^4}{\sqrt{|V_*''|}}, \quad (\text{B.1})$$

with $|V_*''|$ well-approximated by

$$|V_*''| = |\lambda_-| \frac{\sigma_Q^2 \Lambda^4}{f^2}, \quad (\text{B.2})$$

and where λ_- is the most negative eigenvalue of $\hat{\mathbf{H}}$. As shown in appendix section A.2, the mean of the distribution of the most negative eigenvalues in an $\hat{\mathbf{H}}$ -ensemble behaves at large N as

$$\langle \lambda_- \rangle = -c(k)N, \quad (\text{B.3})$$

where $c(k)$ is an $\mathcal{O}(1)$ constant, ranging from $c(1) \approx 1/2$, to $c(P) \approx 4$. Although using $\langle \lambda_- \rangle$ in (B.1) captures the contribution to the tension from the near-saddle region for a *typical* degree- k neighbor, what matters is the extremal neighbor; the leftmost outlier λ_- among the set of about $2^k \binom{P}{k}$ neighbors. To account for this we'll write $\lambda_{\text{edge}} = C_1 \langle \lambda_- \rangle$. The standard deviation of the distribution of λ_- goes like $\sqrt{3N/2}$, which is suppressed by a power of $N^{1/2}$ relative to the mean, implying C_1 will tend to 1 as $N \rightarrow \infty$. However, for $N = 500$, a value of about $C_1 = 1.5$ is appropriate. Putting things together, we have

$$B_{\text{degree-}k} \gtrsim \frac{6^3 \pi^6}{C_1^2} \frac{k^4}{c^2(k)} \frac{\Lambda^8 f^4}{\sigma_Q^4 N^2 \epsilon^3} \quad (\text{B.4})$$

An *upper* bound on $\epsilon = |\Delta V_{\text{lowest outlier}}|$ completes the calculation. For tunneling from a minimum with zero vacuum energy one has $\epsilon < |V_0|$, since the global minimum has energy $V_{\text{global min}} \geq -|V_0|$. (We have stated this as an inequality to include the case of non-vanishing phases δ in (1.1). The inequality is saturated for $\delta = \mathbf{0}$.) Plugging in $\epsilon = |V_0|$ reproduces the bound in section 2, (2.10) provided one sets $C_1 = 1$ and $c(k) = c(1) = 1/2$ (since the estimate

there ignores fluctuations in λ_- and the sublinear k -dependence of the random matrix theory coefficient $c(k)$).

It turns out that face neighbors have smaller ΔV outliers than $-|V_0|$ in certain parameter regimes. This is discussed in detail in the following section. For example, for $V_0 = \Lambda^4$ and $P = 500$ the net effect is about a factor of a half, $\epsilon_{k=1} < |V_0|/2$. If the decay rate is dominated by tunneling to these neighbors (as we expect), this factor of 2 contributes roughly an order of magnitude increase to B , since it enters via $1/\epsilon^3$.

B.2 Vacuum energy differences

The vacuum energy difference between two neighbors can be approximated by a simple Taylor expansion of

$$V_B = V_{\text{aux}}(\mathbf{Q}\Theta_B) \approx V_{\text{aux}}(\mathbf{Q}\Theta_A + 2\pi\mathbf{P}\mathbf{v}_k). \quad (\text{B.5})$$

The result to second order in $\|\mathbf{v}_k^\perp\|_2 \equiv \|\mathbf{P}^\perp\mathbf{v}_k\|_2$ is

$$\Delta V = V_B - V_A = \Lambda^4 \left(2\pi\sqrt{2\mu P}\|\mathbf{v}_k^\perp\|_2 \cos(\psi) + 2\pi^2\|\mathbf{v}_k^\perp\|_2^2 \right), \quad (\text{B.6})$$

where ψ is the angle between \mathbf{n}_A^\perp and \mathbf{v}_k^\perp , and $\mu = 2\pi^2\|\mathbf{n}_A^\perp\|_2^2/P$. The derivation of (B.6) is included at the end of this section for completeness.

The quadratic contribution to (B.6) is positive definite. Since the \mathbf{v}_k come in pairs that are equal in magnitude but exactly opposite in direction,¹² half of the degree- k neighbors have $\cos(\psi) > 0$, and half have $\cos(\psi) < 0$. In any given model, parent vacua with sufficiently high μ have first order contributions to the ΔV 's that tend to dominate the quadratic ones for most of the neighbors in a given degree- k set, resulting in the parent's set of ΔV 's being approximately symmetric about zero, for each k . As a parent's μ value decreases though, this balance is thrown off because an increasing number of negative first order contributions — which are proportional to $\sqrt{\mu}$ — will be partially canceled, if not entirely overwhelmed, by the positive quadratic term.

Whether positive or negative, the ΔV to a parent's lowest energy neighbor can be estimated with relative ease because of the observation that the two-norms of the \mathbf{v}_k^\perp and their orientations with respect to the \mathbf{n}_A^\perp appear to be uncorrelated in random axion landscapes. These are the only two sources of variability in a parent's ΔV 's for fixed k .¹³ When $\nu \gtrsim 10$ the $\cos(\psi)$ behave like the dot products between a set delocalized unit vectors in ν dimensions, implying that the $\cos(\psi)$ are approximately normal distributed with mean zero and standard deviation $1/\sqrt{\nu}$. The fall-off eventually deviates from the Gaussian, becoming sharper due to the fact that $|\cos(\psi)| \leq 1$. For a sample of different low CC parents, the distribution of maximum dot product factors — one $\text{Max}(\cos(\psi))$ from each parent taken over its fixed k neighbors — is peaked at $x^*/\sqrt{\nu}$ with x^* ranging from about 2.2 to 2.9 for $\nu = 10$ to $\nu = 15$ and $P = 400$.

Now for the two-norms; the mean of the distribution of the set of $\|\mathbf{v}_k^\perp\|_2$ for fixed k is in accordance with the naive prediction based on assuming P independent normal distributed entries with standard deviation $\sqrt{\nu k}/P$, namely the mean of a χ -distribution in P variables

¹²Because for every nontrivial entry there is a choice of ± 1 .

¹³This applies to all neighbors for which the second order Taylor expansion (B.6) is accurate. Since we are only interested in tunneling to lower energy neighbors, approximating the vacuum energy of the neighbor by $\Lambda^4 2\pi^2 d_B^2$ is even more accurate than the analogous expression is for the parent. The difference between these two quadratic approximations to the vacuum energy directly gives (B.6) without further assumption.

scaled by the standard deviation of the individual entries, i.e.

$$\langle \|\mathbf{v}_k^\perp\|_2 \rangle = \sqrt{\nu k/P} \quad (\text{B.7})$$

(cf. (A.4)). Though the means agree, the actual $\{\|\mathbf{v}_k^\perp\|_2\}$ are more widely distributed.¹⁴

With these considerations, the ΔV for a general degree- k neighbor can be expressed in terms of the natural scales by setting $|\cos(\psi)| = x/\sqrt{\nu}$ and $\|\mathbf{v}_k^\perp\|_2 = y\sqrt{\nu k/P}$ in equation (B.6). We also take the parent's μ to be $\mu_{\text{low CC}} \approx V_0/(P\Lambda^4)$ from now on. The result, expressed as a function of the dimensionless order 1 quantities x and y that entirely capture the variation in ΔV across neighbors of fixed k , is

$$\frac{\Delta V}{\Lambda^4} = \pm 2\pi\sqrt{2V_0/\Lambda^4} \, xy\sqrt{\frac{k}{P}} + 2\pi^2 \, y^2 \frac{\nu k}{P}, \quad (\text{B.8})$$

where the $+$ applies to the half with $\cos(\psi) > 0$ and $-$ to those with $\cos(\psi) < 0$, and we bear in mind that $0 \leq x \leq \sqrt{\nu}$ and $0 \leq y \leq \sqrt{P/(\nu k)}$.

The optimal conditions for making ΔV as negative as possible come from $x = \sqrt{\nu}$ (and obviously selecting the $-$ sign), and y given by

$$y_{\text{optimal}} = \frac{\sqrt{PV_0/\Lambda^4}}{\pi\sqrt{2k\nu}}. \quad (\text{B.9})$$

Evaluating $\Delta V(x, y)$ at the optimal values turns out to simply give $-V_0$, the energy gap to the global minimum. At higher k one can expect to saturate the inequality $\Delta V > -V_0$ because there are exponentially many neighbors, but for low k the $2^k \binom{P}{k}$ draws of (x, y) may be insufficient for ensuring that at least one pair of the (x, y) is close to $(x_{\text{optimal}}, y_{\text{optimal}})$.

Without an analytic expression for the PDF governing the $\{x_i\}$ and $\{y_i\}$ though, the only way to determine the outlier (x_i, y_i) at low k is numerically. For $\nu = 10$ to about 15 and $P = 250$ –500 we find that the aforementioned value of x^* together with $y^* = 1$ yields a good approximation to the $\Delta V_{\text{outlier}}$ for the special case of $k = 1$. Essentially, this is because the wider of the two distributions controls the outlier value if there are not enough draws to densely sample the allowed ΔV 's, coupled to the fact that the optimal y happens to be relatively close to 1 (unlike for high k , where it approaches zero). Hence for $k = 1$ we may apply the estimate

$$\frac{\Delta V_{\text{outlier}, k=1}}{\Lambda^4} = -2\pi\sqrt{2V_0/\Lambda^4} \frac{x^*(\nu, P)}{\sqrt{P}} + 2\pi^2 \frac{\nu}{P} \quad (\text{B.10})$$

and define

$$C_3(V_0, \nu, P) = 2\pi \left(-\sqrt{2V_0/\Lambda^4} \, x^* + \frac{\pi\nu}{\sqrt{P}} \right), \quad (\text{B.11})$$

¹⁴The spread of an actual set of $\|\mathbf{v}_k^\perp\|_2$ is wider than the naive expectation by a factor of 5 or so,

$$\begin{aligned} \text{Std. Dev. (scaled } \chi) &= \sqrt{\frac{\nu k}{P}} \sqrt{P-2 \left(\frac{\Gamma((P+1)/2)}{\Gamma(P/2)} \right)^2} \\ &\sim \sqrt{\frac{2\nu k}{P}}, \\ \text{Std. Dev. (actual } \|\mathbf{v}_k^\perp\|_2) &\approx 5\sqrt{\frac{2\nu k}{P}}. \end{aligned}$$

The reason for this discrepancy is that the entires of \mathbf{P}^\perp are correlated.

so that

$$\Delta V_{\text{outlier}, k=1} = \frac{C_3 \Lambda^4}{\sqrt{P}}. \quad (\text{B.12})$$

For $V_0 = \Lambda^4$ and $P = 500$ this results in a value of about -12 for C_3 .

As k increases there is a transition from an outlier energy difference of $\epsilon = |C_3| \Lambda^4 \sqrt{k/P}$ for $k = 1$ to an outlier with $\epsilon = V_0$. In this parameter regime it turns out that the transition occurs immediately at $k = 2$. The particular values here result in an outlier difference for $k = 1$ that is about half as large as the global bound. Accounting for such order one factors at $V_0 = \mathcal{O}(1) \Lambda^4$ can be significant because they enter with a third power in the denominator of B , and so are capable of increasing the numerical factor in the bound by orders of magnitude.

Expressing the outlier ΔV as $\sim 1/\sqrt{P}$, as we've done in (B.12), is useful in that it accurately reflects the N dependence for larger μ parents (where the linear term in ΔV dominates). However, writing the outlier ΔV in this manner obscures the fact that the absolute minimum value it ever takes on is $-|V_0|$. This fact is encoded in the definition of C_3 in (B.11), but the property that the refined bound on ΔV is always stronger than the global bound can be made manifest as follows. Evaluating (B.8) at the outlier values for x and y expressed in terms of the optimal ones: $x^* = r_x x_{\text{optimal}}$ and $y^* = r_y y_{\text{optimal}}$, results in the expression,

$$\Delta V_{\text{outlier}} = (r_y^2 - 2r_x r_y) V_0. \quad (\text{B.13})$$

By writing the outlier energy gap between quadratic domain vacua as a fraction times V_0 , one can immediately see the degree to which using the refined bound on ΔV over the global one improves the lower bound on the tunneling exponent B . The benefit amounts to an enhancement of B_{min} by the inverse of the fraction, cubed. An expansive numerical study the typical values of r_x and r_y in different parameter regimes would enable one to determine the typical benefit of employing the refined bound for very low k neighbors (possibly only $k = 1$, as the numerical survey may reveal that the transition to saturation never occurs above $k = 2$).

Finally, we include the derivation of the expression we began with, (B.6). Start by evaluating the potential at the leading order locations for the two vacua:

$$\begin{aligned} \Delta V &= V(\Theta_B) - V(\Theta_A) \\ &\approx V_{\text{aux}}(\mathbf{Q}\Theta_A + 2\pi \mathbf{P}\mathbf{v}_k) - V_{\text{aux}}(\mathbf{Q}\Theta_A) \\ &= \Lambda^4 \sum_{J=1}^P 1 - \cos \left[2\pi(\mathbf{n}_A + \mathbf{v}_k - \mathbf{n}_A^\perp - \mathbf{v}_k^\perp)^J \right] - \left(1 - \cos \left[2\pi(\mathbf{n}_A - \mathbf{n}_A^\perp)^J \right] \right) \\ &= \Lambda^4 \sum_{J=1}^P 1 - \cos \left[2\pi(\mathbf{n}_A - \mathbf{n}_A^\perp - \mathbf{v}_k^\perp)^J \right] - \left(1 - \cos \left[2\pi(\mathbf{n}_A - \mathbf{n}_A^\perp)^J \right] \right). \end{aligned}$$

Now define $\epsilon^J \equiv 2\pi(\mathbf{v}_k^\perp)^J$ and expand the first cosine term in the sum about the argument of the second, i.e. expand in ϵ^J . The result is

$$\Delta V = \Lambda^4 \sum_{J=1}^P \left(-\sin[2\pi(\mathbf{n}_A - \mathbf{n}_A^\perp)^J] \epsilon^J + \frac{|\epsilon^J|^2}{2} \cos[2\pi(\mathbf{n}_A - \mathbf{n}_A^\perp)^J] \right) + \mathcal{O}[(\epsilon^J)^3]. \quad (\text{B.14})$$

Now, further expand the sine and cosine about $(2\pi \mathbf{n}_A)^J$,

$$\Delta V = \Lambda^4 \sum_{J=1}^P \left(2\pi(\mathbf{n}_A^\perp)^J \epsilon^J + (1 - (\mathbf{n}_A^\perp)^2) \frac{|\epsilon^J|^2}{2} \right) + \mathcal{O}[(\epsilon^J)^3]. \quad (\text{B.15})$$

Using $(n_A^\perp)^J \ll 1$, we arrive at the following simple expression for the energy difference between quadratic domain neighbors:

$$\Delta V = 2\pi^2 \Lambda^4 \left(2\mathbf{n}_A^\perp \cdot \mathbf{v}_k^\perp + \|\mathbf{v}_k^\perp\|_2^2 \right), \quad (\text{B.16})$$

which we can recognize as the difference in the parent and neighbor lattice point's two-norm distances to the constraint surface, times $2\pi^2$ (as it had to be). Plugging in $\|\mathbf{n}_A^\perp\|_2 = \sqrt{\mu_A P}/2\pi^2$ and defining ψ as the angle between \mathbf{n}_A^\perp and \mathbf{v}_k^\perp gives equation (B.6).

References

- [1] S. Weinberg, *Anthropic bound on the cosmological constant*, *Phys. Rev. Lett.* **59** (1987) 2607 [[INSPIRE](#)].
- [2] 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](#)].
- [3] B. Freivogel, M. Kleban, M. Rodriguez Martinez and L. Susskind, *Observational consequences of a landscape*, *JHEP* **03** (2006) 039 [[hep-th/0505232](#)] [[INSPIRE](#)].
- [4] N. Arkani-Hamed, S. Dimopoulos and S. Kachru, *Predictive landscapes and new physics at a TeV*, [hep-th/0501082](#) [[INSPIRE](#)].
- [5] T.C. Bachlechner, K. Eckerle, O. Janssen and M. Kleban, *Multiple-axion framework*, *Phys. Rev. D* **98** (2018) 061301 [[arXiv:1703.00453](#)] [[INSPIRE](#)].
- [6] T.C. Bachlechner, K. Eckerle, O. Janssen and M. Kleban, *Systematics of aligned axions*, *JHEP* **11** (2017) 036 [[arXiv:1709.01080](#)] [[INSPIRE](#)].
- [7] F. Denef and M.R. Douglas, *Distributions of nonsupersymmetric flux vacua*, *JHEP* **03** (2005) 061 [[hep-th/0411183](#)] [[INSPIRE](#)].
- [8] M.R. Douglas and S. Kachru, *Flux compactification*, *Rev. Mod. Phys.* **79** (2007) 733 [[hep-th/0610102](#)] [[INSPIRE](#)].
- [9] D. Marsh, L. McAllister and T. Wrase, *The wasteland of random supergravities*, *JHEP* **03** (2012) 102 [[arXiv:1112.3034](#)] [[INSPIRE](#)].
- [10] C. Long, L. McAllister and P. McGuirk, *Heavy tails in Calabi-Yau moduli spaces*, *JHEP* **10** (2014) 187 [[arXiv:1407.0709](#)] [[INSPIRE](#)].
- [11] T.C. Bachlechner, C. Long and L. McAllister, *Planckian axions in string theory*, *JHEP* **12** (2015) 042 [[arXiv:1412.1093](#)] [[INSPIRE](#)].
- [12] C. Long, L. McAllister and J. Stout, *Systematics of axion inflation in Calabi-Yau hypersurfaces*, *JHEP* **02** (2017) 014 [[arXiv:1603.01259](#)] [[INSPIRE](#)].
- [13] J. Halverson, C. Long and B. Sung, *Algorithmic universality in F-theory compactifications*, *Phys. Rev. D* **96** (2017) 126006 [[arXiv:1706.02299](#)] [[INSPIRE](#)].
- [14] 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](#)].
- [15] M. Demirtas, C. Long, L. McAllister and M. Stillman, *The Kreuzer-Skarke axiverse*, [arXiv:1808.01282](#) [[INSPIRE](#)].
- [16] T.C. Bachlechner, K. Eckerle, O. Janssen and M. Kleban, *The axidental universe*, to appear.
- [17] P. Batra and M. Kleban, *Transitions between de Sitter minima*, *Phys. Rev. D* **76** (2007) 103510 [[hep-th/0612083](#)] [[INSPIRE](#)].
- [18] S.R. Coleman and F. De Luccia, *Gravitational effects on and of vacuum decay*, *Phys. Rev. D* **21** (1980) 3305 [[INSPIRE](#)].

- [19] S.W. Hawking and I.G. Moss, *Fluctuations in the inflationary universe*, *Nucl. Phys. B* **224** (1983) 180 [INSPIRE].
- [20] T.C. Bachlechner, *Action principle for isotropic general relativity*, [arXiv:1808.10448](#) [INSPIRE].
- [21] T.C. Bachlechner, *Vacuum decay in general relativity*, [arXiv:1809.05174](#) [INSPIRE].
- [22] S.R. Coleman, V. Glaser and A. Martin, *Action minima among solutions to a class of Euclidean scalar field equations*, *Commun. Math. Phys.* **58** (1978) 211 [INSPIRE].
- [23] K. Blum, M. Honda, R. Sato, M. Takimoto and K. Tobioka, *$O(N)$ invariance of the multi-field bounce*, *JHEP* **05** (2017) 109 [Erratum *ibid.* **06** (2017) 060] [[arXiv:1611.04570](#)] [INSPIRE].
- [24] S.R. Coleman, *The fate of the false vacuum. 1. Semiclassical theory*, *Phys. Rev. D* **15** (1977) 2929 [Erratum *ibid.* **D 16** (1977) 1248] [INSPIRE].
- [25] A.R. Brown, *Thin-wall approximation in vacuum decay: a lemma*, *Phys. Rev. D* **97** (2018) 105002 [[arXiv:1711.07712](#)] [INSPIRE].
- [26] A. Masoumi, K.D. Olum and B. Shlaer, *Efficient numerical solution to vacuum decay with many fields*, *JCAP* **01** (2017) 051 [[arXiv:1610.06594](#)] [INSPIRE].
- [27] A. Masoumi, K.D. Olum and J.M. Wachter, *Approximating tunneling rates in multi-dimensional field spaces*, *JCAP* **10** (2017) 022 [[arXiv:1702.00356](#)] [INSPIRE].
- [28] S. Groot Nibbelink and B.J.W. van Tent, *Scalar perturbations during multiple field slow-roll inflation*, *Class. Quant. Grav.* **19** (2002) 613 [[hep-ph/0107272](#)] [INSPIRE].
- [29] L. McAllister, S. Renaux-Petel and G. Xu, *A statistical approach to multifield inflation: many-field perturbations beyond slow roll*, *JCAP* **10** (2012) 046 [[arXiv:1207.0317](#)] [INSPIRE].
- [30] PLANCK collaboration, *Planck 2018 results. VI. Cosmological parameters*, [arXiv:1807.06209](#) [INSPIRE].
- [31] D. Baumann and L. McAllister, *Inflation and string theory*, Cambridge Monographs on Mathematical Physics, Cambridge University Press, Cambridge, U.K. (2015) [[arXiv:1404.2601](#)] [INSPIRE].
- [32] M. Peloso and C. Unal, *Trajectories with suppressed tensor-to-scalar ratio in aligned natural inflation*, *JCAP* **06** (2015) 040 [[arXiv:1504.02784](#)] [INSPIRE].
- [33] J. García-Bellido, *Single bubble open inflation: an overview*, in *2nd International Workshop on Birth of the Universe and Fundamental Physics*, Rome, Italy, 19–24 May 1997 [[hep-ph/9803270](#)] [INSPIRE].
- [34] A.D. Linde, *A toy model for open inflation*, *Phys. Rev. D* **59** (1999) 023503 [[hep-ph/9807493](#)] [INSPIRE].
- [35] J. García-Bellido and A.D. Linde, *Open hybrid inflation*, *Phys. Rev. D* **55** (1997) 7480 [[astro-ph/9701173](#)] [INSPIRE].
- [36] W. Hu, R. Barkana and A. Gruzinov, *Cold and fuzzy dark matter*, *Phys. Rev. Lett.* **85** (2000) 1158 [[astro-ph/0003365](#)] [INSPIRE].
- [37] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper and J. March-Russell, *String axiverse*, *Phys. Rev. D* **81** (2010) 123530 [[arXiv:0905.4720](#)] [INSPIRE].
- [38] 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].
- [39] C. Vafa, *The string landscape and the swampland*, [hep-th/0509212](#) [INSPIRE].
- [40] 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].

- [41] T. Rudelius, *On the possibility of large axion moduli spaces*, *JCAP* **04** (2015) 049 [[arXiv:1409.5793](#)] [[INSPIRE](#)].
- [42] C. Cheung and G.N. Remmen, *Naturalness and the weak gravity conjecture*, *Phys. Rev. Lett.* **113** (2014) 051601 [[arXiv:1402.2287](#)] [[INSPIRE](#)].
- [43] B. Heidenreich, M. Reece and T. Rudelius, *Weak gravity strongly constrains large-field axion inflation*, *JHEP* **12** (2015) 108 [[arXiv:1506.03447](#)] [[INSPIRE](#)].
- [44] T.C. Bachlechner, C. Long and L. McAllister, *Planckian axions and the weak gravity conjecture*, *JHEP* **01** (2016) 091 [[arXiv:1503.07853](#)] [[INSPIRE](#)].
- [45] T. Rudelius, *Constraints on axion inflation from the weak gravity conjecture*, *JCAP* **09** (2015) 020 [[arXiv:1503.00795](#)] [[INSPIRE](#)].
- [46] 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](#)].
- [47] 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](#)].
- [48] B. Heidenreich, M. Reece and T. Rudelius, *Evidence for a sublattice weak gravity conjecture*, *JHEP* **08** (2017) 025 [[arXiv:1606.08437](#)] [[INSPIRE](#)].
- [49] S.B. Giddings and A. Strominger, *String wormholes*, *Phys. Lett. B* **230** (1989) 46 [[INSPIRE](#)].
- [50] R.D. Peccei and H.R. Quinn, *CP conservation in the presence of instantons*, *Phys. Rev. Lett.* **38** (1977) 1440 [[INSPIRE](#)].
- [51] P. Svrček and E. Witten, *Axions in string theory*, *JHEP* **06** (2006) 051 [[hep-th/0605206](#)] [[INSPIRE](#)].
- [52] D.J.E. Marsh, *Axion cosmology*, *Phys. Rept.* **643** (2016) 1 [[arXiv:1510.07633](#)] [[INSPIRE](#)].
- [53] N. Arkani-Hamed, J. Orgera and J. Polchinski, *Euclidean wormholes in string theory*, *JHEP* **12** (2007) 018 [[arXiv:0705.2768](#)] [[INSPIRE](#)].
- [54] M. Montero, A.M. Uranga and I. Valenzuela, *Transplanckian axions!?*, *JHEP* **08** (2015) 032 [[arXiv:1503.03886](#)] [[INSPIRE](#)].
- [55] P. Vivo, S.N. Majumdar and O. Bohigas, *Large deviations of the maximum eigenvalue in Wishart random matrices*, *J. Phys. A* **40** (2007) 4317 [[cond-mat/0701371](#)].