

The desert and the swampland

Cody Long,^a Miguel Montero,^a Cumrun Vafa^a and Irene Valenzuela^{a,b}

^a*Department of Physics, Harvard University,
Cambridge, MA 02138, U.S.A.*

^b*Instituto de Física Teórica UAM-CSIC and Departamento de Física Teórica,
Universidad Autónoma de Madrid,
Cantoblanco, 28049 Madrid, Spain*

E-mail: c.e.cody.long@gmail.com, mmontero@g.harvard.edu,
vafa@g.harvard.edu, irene.valenzuela@uam.es

ABSTRACT: The most natural expectation away from asymptotic limits in moduli space of supergravity theories is the *desert scenario*, where there are few states between massless fields and the quantum gravity cutoff. In this paper we initiate a systematic study of these regions deep in the moduli space, and use it to place a bound on the number of massless modes by relating it to the black hole species problem. There exists a consistent sub-Planckian UV cutoff (the species scale) which resolves the black hole species problem without bounding the number of light modes. We reevaluate this in the context of supersymmetric string vacua in the desert region and show that even though heuristically the species scale is compatible with expectations, the BPS states of the actual string vacua lead to a stronger dependence of the cutoff scale on the number of massless modes. We propose that this discrepancy, which can be captured by the “BPS desert conjecture”, resurrects the idea of a uniform bound on the number of light modes as a way to avoid the black hole species problem. This conjecture also implies a stronger form of the Tadpole Conjecture, which leads to an obstruction in stabilizing all moduli semi-classically for large number of moduli in flux compactifications.

KEYWORDS: Black Holes in String Theory, Superstring Vacua, Extended Supersymmetry, D-Branes

ARXIV EPRINT: [2112.11467](https://arxiv.org/abs/2112.11467)

Contents

1	Introduction	1
2	The general setup	3
3	BPS mass gap for theories with more than 8 supercharges	4
3.1	32 supercharges	5
3.2	16 supercharges	6
4	Species bound and mass gap in Calabi-Yau compactifications	9
5	BPS mass gap in 6d $\mathcal{N} = \text{EFTs}$	12
5.1	Review of 6d $\mathcal{N} = 1$ and CY_3 's with a toric base	12
5.2	Minimization algorithm and results	15
5.3	Field theory vs. black hole mass gap and species bound	18
5.4	Relation to the Tadpole Conjecture	21
6	Conclusions and finiteness of the CY landscape	22

1 Introduction

One of the principles of the Swampland program is that there are a finite number of quantum gravity backgrounds [1–4]. More precisely if we fix a cutoff we expect a finite number of string vacua. This in particular suggests that there is an upper bound on N the number of massless or light species modes allowed in a consistent quantum gravity theory. A heuristic explanation of this may have been that if there were no upper bound on N then the entropy of a fixed size black hole which should be at least N (as the black hole can be made up of such constituents) and this would lead to a contradiction if there were no bound on N . However, this reasoning is not quite correct.

In a quantum theory of gravity we expect that any effective theory breaks down for energies $E \geq M_P$, where quantum gravitational effects become strong and invalidate the low energy effective theory. Therefore in any effective theory including quantum gravity we expect to have a cutoff $\Lambda \leq M_P$. In some cases we expect $\Lambda \sim M_P$. This is the case for example for M-theory in 11 dimensions. On the other hand there are also cases where $\Lambda \ll M_P$. For example for 10d string theory the cutoff is given by the string scale $\Lambda \sim M_s \sim g_s^{1/4} M_P$ which for $g_s \ll 1$ is far below the Planck scale.

For a given cutoff, the radius of black holes which can be described reliably by the effective theory are bounded by the cutoff scale: $R_{BH} > 1/\Lambda$. In d dimensions, the entropy of such a black hole is thus bounded by $S > (M_P/\Lambda)^{d-2}$. Since the entropy of

such a black hole should be bigger than the number of light species N , we learn that if $N < (M_{\text{P}}/\Lambda)^{d-2}$ there would be no contradiction with the fact that the black hole entropy should be more than the number of light species. In other words if $\Lambda \leq \Lambda_{\text{sp}}$ with the species scale $\Lambda_{\text{sp}} = M_{\text{P}} \cdot N^{-\frac{1}{d-2}}$, there would be no contradiction between having an arbitrarily large number of light species N and the finiteness of the black hole entropy [5–7].

The main goal of this work is to evaluate this resolution of the black hole species problem from the perspective of string landscape. More specifically we focus on supersymmetric string landscape for which we have more analytic control. To test this proposal we need to investigate points on the landscape for which the number of light/massless modes can be large. This happens when we have 8 or less supercharges in the theory. The first such case is $\mathcal{N} = (1, 0)$ supersymmetric theories in 6d. We focus on the number of massless tensor multiplets N and attempt to compute the cutoff $\Lambda(N)$. First of all we need to be more precise what we mean by $\Lambda(N)$. As we change the moduli of the supersymmetric theory by changing the vev of scalar fields, we expect that mass of excited states change, and so in particular the first massive state which is not part of the EFT which defines Λ will depend on the massless moduli. The first question is whether we can have Λ become equal to the M_{P} ? In other words, can there be a *desert scenario* where there is no light states all the way till the Planck scale? This is interesting since the largest possible value of Λ provides the strongest constraint on the number of light species N via the species bound above. To address this issue, we show that no matter what moduli we choose, there is a $\Lambda < M_{\text{P}}$ for which there is a particle of that mass $m \sim \Lambda$ which is not part of the EFT. We do this by focusing on BPS states for which we can compute the mass exactly as a function of the moduli using supersymmetry. In particular we cannot quite have the desert all the way to the Planck scale. The next question would be whether the Λ we find scales with the number of massless modes as $N^{-1/(d-2)}$ as one would expect based on species scale. Quite surprisingly we find, using supersymmetry and some genericity assumptions, that for theories in $d = 4, 5, 6$ with 8 supercharges this behavior would naively be expected to hold. However, upon more careful scrutiny by focusing on actual examples (a mini-Landscape of 34864 6d F-theory models with toric base) and focusing on BPS states we show that this genericity assumption seems false and that $\Lambda \sim N^{-a}$ with $a > 1/(d-2)$. At first site this may sound contradictory with the species scale. However the Λ we have computed is based on actual supersymmetric states. However, there are potentially other lighter states in the theory which are not supersymmetric which can resurrect the dependence of species scale with N . This explanation can work as long as $N < N_0$ for some upper bound N_0 . If $N > N_0$ since the slope of the supersymmetric cutoff Λ with N is bigger than the species scale, the supersymmetric cutoff would not agree with the species scale. We propose that this may explain why there is an actual upper bound in the number of massless modes for theories with 8 supercharges.

The organization of this paper is as follows: in section 2 we discuss the general setup for the supersymmetric landscape. In section 3 we discuss the situation for theories with 16 or 32 supercharges. In section 4 we give a heuristic argument why in theories with 8 supercharges the supersymmetric cutoff Λ for theories with dimensions $d = 4, 5, 6$ agrees with the species bound. In section 5, however, we show this heuristic expectation based on

genericity is incorrect by actually looking at a large class of examples for $\mathcal{N} = (1, 0)$ supersymmetric theories in $d = 6$ and offer our resolution of this. We also discuss applications to the Tadpole Conjecture in this section. In section 6 we present our conclusions.

2 The general setup

Let us consider a supersymmetric theory coupled to gravity. Such theories typically have massless moduli fields, if there is enough supersymmetry. In such a case a natural question to ask is whether the low energy description involving only the massless modes is adequate all the way to Planck energy. In other words, is the first excited state of the order of Planck mass or higher? If there is a large number of massless modes in the supersymmetric theory, we would expect this not to be possible because otherwise the Planck mass black holes (or slightly above it) would be reliably described by such an effective theory and would yield an entropy too small to account for all the massless modes that the black hole can be made of, violating the species bound explained in the Introduction.

We know that the mass of excited states will depend on massless moduli $\Lambda(\phi)$. Note that we do not expect $\Lambda(\phi)$ to have a non-vanishing minimum, as there are always infinite distance points on the moduli where we expect exponentially light states in accord with the Distance Conjecture [8]. On the other hand it is natural to ask how close to the Planck scale can we make the cutoff Λ where the first excited particle appears by varying moduli. In other words we are interested in finding a point on the moduli space of massless fields where the cutoff is maximized and we call this the Λ of the supergravity theory.

$$\Lambda = \max_{\phi} \Lambda(\phi)$$

The difficulty with finding Λ is that we do not have a reliable method to compute the excited spectrum of the supergravity theory. However, there is a subset of states in supersymmetric theories that we control, namely the BPS states. Such states have masses or tensions which are determined exactly by supersymmetry. We can use this to find an upper bound on the value of Λ . Namely if we define

$$\Lambda_{BPS}(\phi) = \min_{BPS} \{m_{BPS}(\phi)\}$$

by scanning over the BPS states, then we can define the desert among all the BPS states by considering

$$\Lambda_{BPS} = \max_{\phi} \{\Lambda_{BPS}(\phi)\} = \max_{\phi} \{\min_{BPS} \{m_{BPS}(\phi)\}\}$$

where $\Lambda_{BPS}(\phi)$ denotes the first massive BPS state at a given point ϕ in moduli. It is clear from the definition that the actual Λ is bounded by Λ_{BPS} :

$$\Lambda \leq \Lambda_{BPS}.$$

Here we have to be a bit more careful about what we mean by the mass associated to a BPS object. For a BPS particle it is simply the mass of the state. However, for BPS p-branes whose tensions T_p^{BPS} is determined using the BPS algebra, we need to form localized

particle states (small brane “blobs”) which would necessarily be non-BPS. Although computing the mass of these states would involve a quantum description of their worldvolume, which is presently out of reach except for the fundamental string case, a natural guess (and in fact, the only possibility allowed by dimensional analysis in a Nambu-Goto like model, which only takes the brane tension into account) is that the mass scales of the order of $m_{BPS} \sim (T_p^{BPS})^{1/(p+1)}$. This is of course rather familiar in the context of fundamental strings.¹ Therefore in the above for each BPS state what we have in mind is actually

$$m_{BPS} \sim (T_p^{BPS})^{1/(p+1)}$$

In other words we have

$$\Lambda_{BPS} = \max_{\phi} \{ \min_{BPS_p} \{ (T_p^{BPS})^{1/(p+1)}(\phi) \} \}$$

Throughout this paper, we will compute Λ_{BPS} for different theories, organizing the discussion according to the level of supersymmetry. For theories with 8 supercharges, the Landscape is rich enough to study the dependence of Λ_{BPS} with the number of massless modes, i.e. $\Lambda_{BPS}(N_{\text{massless}})$. In that case, we can also compare the results with the species scale [5–7, 9],

$$\Lambda_{\text{sp}} \sim M_{\text{P}} \cdot N^{-\frac{1}{d-2}} \tag{2.1}$$

where N is the number of weakly-coupled species with a mass at or below Λ_{sp} . Notice that, in this definition, N are not only the massless modes but include also all the light modes that can be described within the EFT below this quantum gravity cut-off. In fact, near the infinite distance limits of the moduli space, this species scale is dominated by the infinite towers of states becoming light [8, 10, 11]. However, in the deep interior of the moduli space where the BPS mass gap gets maximized, we have a quasi-dessert scenario in which the asymptotic towers of states are very heavy and the species bound will instead be dominated by the massless modes. In other words, in the interior of the moduli space, we expect to have, in general, many more massless modes than light massive modes below Λ_{sp} . In that case, it should be a good approximation to replace N by N_{massless} in (2.1), as we will do when comparing our results with the species bound in section 5.

3 BPS mass gap for theories with more than 8 supercharges

We will now analyze the maximal gap predicted by the BPS sector in theories with 32 and 16 supercharges, where complete calculations are possible. Notice that we will use BPS masses and tensions as a proxy of the actual mass gap; therefore, what we obtain is actually an upper bound on how large the mass gap can be made in these higher-dimensional moduli spaces. Due to the simplicity of the setups involved, we also have some degree of control for some non-BPS states.

¹There could be other potential effects that invalidate the above simple conclusion. For instance, taking into account higher-derivative terms, the action of a D-brane deviates to pure Nambu-Goto and the string scale enters independently in the action. Although a proof is out of the scope of this paper, we believe these effects are subleading with respect to the leading Nambu-Goto piece for the above particle states, since the string scale is much smaller than the D-brane tension in perturbative string theory.

3.1 32 supercharges

The first and easiest example to analyze is the 11-dimensional vacuum of M-theory. The massless fields include no moduli; there is no optimization to perform, and the mass gap set by BPS states is the Planck scale. More precisely, if we estimate the gap from the BPS sector, there are mass scales associated to the M2 and M5 branes, which are

$$M_{M2} \sim T_{M2}^{1/3}, \quad M_{M5} \sim T_{M5}^{1/6}. \tag{3.1}$$

These two quantities are of order the 11d Planck mass. There is an intrinsic $\mathcal{O}(1)$ ambiguity in determining a mass scale out of the tension of an extended object, which explains the twiddles in (3.1). We will face similar issues in most of the lower-dimensional vacua we will analyze in this paper.

Compactifying M-theory on a circle we reach IIA string theory. There is now a modulus, the dilaton, which controls the ratio between string and Planck scales. The tensions of D-branes, NS5-brane, and the fundamental string itself, are controlled by different powers of g_s . Recall the formula for the tension of a Dp-brane in Planck units [12],

$$T_{Dp} = g_s^{\frac{p-3}{2}}. \tag{3.2}$$

The point of largest gap is then reached when $g_s \sim 1$, and all scales coincide with the Planck scale. Not surprisingly, we find that the point of largest gap is as far away from the infinite distance limits of moduli space as one can be; as we will see in later sections, it is also as far as it can be from any finite-distance singularity where charged states become light. From (3.2), we also note that since the tensions of BPS objects are determined by their couplings, pairs of dual D-branes satisfy

$$T_{Dp} \cdot T_{D(6-p)} \sim \text{independent of } g_s. \tag{3.3}$$

What this means is that, as one object gets heavy, the electromagnetic dual becomes light, so the natural point where both electric and magnetic objects are heavy is in the middle of moduli space where $g_s \sim 1$. This corresponds to compactifying M-theory on a circle with radius of Planck length. Notice that at the extremum there are more than one object with the same mass scale. This pattern will be repeated in all the examples below.

Lower-dimensional toroidal compactifications of M-theory follow the same pattern, with the BPS spectrum obtained from gravitational solitons such as KK monopoles, and M2 or M5 branes wrapped on the same cycles. The point of maximal gap is then a Planckian torus, with all sides similar to each other.

The one remaining higher-dimensional vacuum is 10-dimensional IIB. There is a complex modulus, the axiodilaton τ , and the spectrum of extended objects includes (p, q) -strings and five branes, D3-branes, as well as (p, q) 7-branes. The tension of the D3 brane in Planck units is $\mathcal{O}(1)$ and independent of τ , so it does not play a role in our discussion. We will also ignore (p, q) 7-branes, since, due to their long-range interactions, the mass scale obtained from their tension is not a good proxy for the mass gap of the theory. This

leaves (p, q) strings and 5-branes as the objects setting the gap. The formula for both of them is very similar, taking the form [12]

$$T_{(p,q) \text{ string}} = \frac{|p + q\tau|}{\sqrt{\tau_2}} \tag{3.4}$$

in Planck units for strings, and a similar expression but with τ replaced by $-1/\tau$ for the (p, q) 5-branes. Similarly to the discussion around (3.1), assuming the gap is at the Planck scale or below, it will be set by the (p, q) -strings and not the 5-branes, since in obtaining a mass scale one takes a fifth root rather than a square root for the latter versus the former. We are led to the problem of maximizing the minimum of (3.4) over $(p, q) \in \mathbb{Z}^2$ as a function of τ . This problem has a simple geometric interpretation: τ_2 is the volume of the lattice in \mathcal{C} generated by $\{1, \tau\}$, and so $|p + q\tau|$ is the length of a lattice vector. By rescaling the lattice vectors by a real number λ , we can make sure that the shortest lattice vector has norm 1. The problem then becomes finding the maximal value of the volume $\lambda^2 \tau_2$ of the fundamental cell of the lattice, subject to the constraint that the minimal length vector has volume one. By drawing a sphere of diameter 1 around each lattice point in such a lattice, one obtains a sphere packing. Thus, the problem we have is equivalent to the classic problem of sphere packing [13], which in two dimensions is solved by an hexagonal lattice. We conclude that the optimal point, with the maximal gap for (p, q) -strings, corresponds to the hexagonal lattice with τ a third root of unity. Although this corresponds to one of the special ‘‘cusps’’ in the torus complex structure moduli space, it is interesting that it does not correspond to the cusp at $\tau = i$, as one might have naively expected. In any case, the gap we obtain is again $\mathcal{O}(1)$ in Planck units, and corresponds to any of the fundamental string, the D1 brane, or the $(1, 1)$ string, all of which have the same tension.

3.2 16 supercharges

We now analyze theories with 16 supercharges. This class of theories includes minimally supersymmetric models in seven, eight, and nine dimensions. Unlike theories with 32 supercharges, for which only one massless multiplet (the gravity multiplet) is possible, theories with 16 supercharges can have both gravity and vector multiplets [14]. Low-energy interactions are fully determined by supersymmetry, to the point that the one free parameter in the theory is the number of vector multiplets at a generic point in moduli space, known as the rank of the theory [14, 15]. Theories with 16 supercharges have been the subject of much recent work from the Swampland perspective [16–21], to the point that Swampland principles are enough to ‘‘bootstrap’’ the known string theory compactifications, realizing the String Lamppost Principle. In particular, there is an upper bound on the rank of the theory, which is $26 - d$ for a d -dimensional theory [22].

We will now partly address the question of the maximal gap in the BPS spectrum for toroidal compactifications of heterotic string theory, which will illustrate the general features of the problem we will study later. The reason we will not do an in-depth study of these compactifications is that the Landscape of these models is finite, and not very large; for instance, in nine dimensions, the moduli space has four connected components only [14]. This means that we cannot increase the number of massless fields arbitrarily,

and therefore, we do not have enough statistics to meaningfully study the dependence of the gap with the number of fields. This problem will be remedied in section 5, by reducing the number of supercharges.

Since the theory contains vector multiplets, a natural set of charged objects to consider are particles, which can in fact be BPS [14]. The advantage of a perturbative heterotic setup is that we have control over all perturbative states, both BPS and non-BPS, so that we can also investigate (in perturbative corners) to what extent the BPS states give a good estimate of the actual mass gap. Of course, deep in moduli space, only the masses of BPS states are reliable.

The well-known perturbative mass formula in string units is [23]

$$m^2 = \frac{1}{\alpha'} \left(p_L^2 + p_R^2 + 2(N + \bar{N} - 1) \right) \tag{3.5}$$

where $p_{L,R}$ are the left-(right-)handed worldsheet momenta and $N(\bar{N})$ the left(right)-handed oscillation number. Upon imposing the level-matching condition

$$p_L^2 - p_R^2 = 2 - 2(N - \bar{N}), \tag{3.6}$$

the mass spectra reads

$$\alpha' m^2 = 2p_R^2 + 4\bar{N} = 2p_L^2 + 4N - 4. \tag{3.7}$$

BPS states, preserving half the supercharges, form a cone, and in heterotic string theory, correspond to states having $\bar{N} = 0$ [24]. This means that

$$\alpha' m_{BPS}^2 = 2p_R^2, \tag{3.8}$$

and must satisfy $p_L^2 - p_R^2 \leq 2$ due to (3.6). We can distinguish three types of BPS states:

- $p_L^2 - p_R^2 = 2$: field theory states. They can become massless at finite distance and induce an enhancement of the gauge group.
- $p_L^2 - p_R^2 = 0$: these towers of states can only become massless at infinite distance.
- $p_L^2 - p_R^2 < 0$: black hole states. Their mass in string units remains bounded below by an $\mathcal{O}(1)$ number everywhere in moduli space, as follows from (3.7) and the fact that $N > 1$ for these states (from (3.6)).

Intuitively, we expect that the maximum gap will occur at a regular point in the interior of the moduli space, and as far away as possible from all singularities including the gauge field theory enhancements. In this sense, it is the *deepest* we can be in the moduli space. Only states with $p_L^2 - p_R^2 \geq 0$ can become lighter than the string scale, so they will be the relevant states in setting the gap Λ_{BPS} . Contrarily, the black hole states will always be more massive and play no role in determining this mass gap.

To give an illustrative example, we will now discuss the case of a circle compactification of heterotic string theory (nine dimensions) without Wilson lines. We can view this as




Figure 1. Mass behavior (in string units) of different BPS states in terms of the circle radius for a 9d theory with 16 supercharges (heterotic on S^1 with Wilson lines switched off). The blue, red, yellow and green curves denote masses of field theory, black hole, and the two infinite distance light towers respectively.

keeping $E_8 \times E_8$ as part of the massless modes, leaving only the radius (in addition to string coupling constant) as the moduli to vary. The momenta are then given by [23]

$$p_{L,R} = \frac{1}{\sqrt{2}R} (n \pm R^2 w) \tag{3.9}$$

where n, w are the KK and winding number, while R is the circle radius in string units. Moreover, BPS states satisfy $p_L^2 - p_R^2 = 2nw \leq 2$. Indeed, the asymptotic towers of states (either purely KK or purely winding modes) have $p_L^2 - p_R^2 = 0$, while the only field theory state here has $n = w = 1$ and becomes massless at the self-dual point $R = 1$.

The mass formulae above are given in string units. For Swampland purposes, it is more natural to write masses in nine-dimensional Planck units, which are related to the above via a rescaling by a factor depending on the 9d dilaton. This implies that the masses also depend on ϕ , the heterotic dilaton, in a way that is independent of the charges. The BPS spectrum masses can grow without bound; at strong coupling, where the system is described by type I' string theory [25], BPS states correspond to D0-D8 bound states, and can become arbitrarily massive, being far above the actual gap of the theory, which is set by non-BPS fundamental type I' string states. For this reason, the spectrum of BPS states is only a good proxy for the actual gap of the theory at heterotic coupling constant of order 1 and below. Fortunately, it is precisely at $g_s \sim \mathcal{O}(1)$ where we expect to find the maximum gap of the theory, by a similar discussion as in the previous section.

The mass behavior of the different states is plotted in figure 1. The maximum mass gap occurs at the point at which the mass for field theory state crosses the mass of the towers. As one can read from the plot, this occurs at $R = \frac{1}{\sqrt{2}}$ and yields $\Lambda_{BPS} = 0.25M_s$.

It is also interesting to check the mass gap that arises when including the field theory states of $SU(2)$ as part of the EFT. This will then freeze the radius at $R = 1$, and the only moduli left is the string coupling constant which we expect to lead to $g_s \sim 1$ for maximal gap of KK (or dual winding) towers.

Due to (3.7), non-BPS states ($\bar{N} \neq 0$) have a mass lower-bounded by 4 in string units. Hence, the gap will be set by BPS states, at least perturbatively, as long as there is always a BPS state with mass less than 4, or equivalently, with $p_R^2 \leq 2$. This is automatically true in the restricted sections of the moduli space of toroidal compactifications of the heterotic string that we just analyzed.

We will see that, also in theories with 8 supercharges, Λ_{BPS} will be determined by field theory states, and the maximum gap will occur at the point which is furthest away from all finite and infinite distance singularities. The asymptotic towers play the leading role only when field theory states are not included.

As emphasized above, we did not carry out a systematic analysis of theories with 16 supercharges because the corresponding Landscape is too small to provide us with any interesting dependence on the number of moduli. We will nevertheless finish with some general comments on how the analysis would be carried out. Just like in heterotic, the natural BPS objects to consider are particles. For a rank r theory, these charged states live naturally in a charge lattice $\Lambda_{10-d,10-d+r}$, which is equipped with an inner product of signature $(10-d, 10-d+r)$. All these theories can be described perturbatively in corners of their moduli space, and in these corners BPS states are naturally Dabholkar-Harvey states [23] just as in heterotic. Then the analysis above can be repeated.

4 Species bound and mass gap in Calabi-Yau compactifications

Let us next reduce the level of supersymmetry and consider compactifications with eight supercharges. This includes 6d $\mathcal{N} = 1$, 5d $\mathcal{N} = 1$ and 4d $\mathcal{N} = 2$ theories arising from Calabi-Yau (CY) threefold compactifications. From now on, we are going to focus only on the maximum mass gap generated by BPS states, since the full non-BPS spectrum is not well enough known to allow for a systematic analysis.

In general, the mass scale of a $(p+1)$ -dimensional BPS state coming from a brane wrapping a cycle Σ is given by

$$m_{BPS} = M_D (vol_\Sigma)^{1/(p+1)} = \frac{M_d}{\mathcal{V}^{1/(d-2)}} (vol_\Sigma)^{1/(p+1)} \tag{4.1}$$

where vol_Σ is the volume of the wrapped cycle and M_D is the higher dimensional Planck mass for M-theory compactifications or the string scale in the case of string compactifications. Furthermore, M_d is the lower dimensional Planck mass and the overall volume \mathcal{V} is written in M_D units. Depending on the space-time dimension, the lightest BPS object will be a particle or a string.

The goal is to find CY moduli where for all BPS states we can push up their mass to the maximum value. Hence, we want to make the volume of all cycles Σ as large as possible while keeping the overall volume \mathcal{V} fixed to one. In principle one may think we can get this as large as we wish by sending $\mathcal{V} \rightarrow 0$, however, it turns out that $(vol_\Sigma)^{1/(p+1)}$ scales the same way as $\mathcal{V}^{1/(d-2)}$ for all the cases of interest here, so we can set $\mathcal{V} = 1$ with no loss of generality. The independence of the mass in Planck units from the volume follows from the decoupling of hypermultiplets (which includes the volume moduli) from the vector/tensor

multiplets due to having 8 supercharges. Unlike in the previous sections, the landscape of EFTs with 8 supercharges is much richer and the BPS mass gap highly depends on the topological data characterizing each CY and entering through the value of the volumes. In the next section we will calculate explicitly the BPS mass gap in the case of 6d $\mathcal{N}=1$ theories by plugging this topological data from a list of known CY's. However, we can first provide some preliminary estimation of Λ_{BPS} by assuming the simplest possible distribution of this data, in particular a diagonal topological intersection matrix as will be discussed next. Intriguingly, this naive estimation of Λ_{BPS} reproduces the results predicted by the species bound discussed in section 2, although it differs from the actual results computed numerically using explicit CY data in section 5. We show this coincidence with the species bound in the following, which seems to be rooted in the independence of the BPS states with the overall volume, and leave the comparison with the numerical results for section 5.

We can organize the discussion according to the space-time dimension of the EFT as follows:

- 6d $\mathcal{N}=1$ theories (F-theory on elliptic CY_3)

The relevant BPS objects are strings arising from wrapping D3-branes on 2-cycles of the base of the elliptically fibered CY. Denoting J as the Kahler form of the base, the tension of the strings read

$$T = \int_{\Sigma_2} J = q^I t_I \quad (4.2)$$

where q_I are the quantized charges and t_I are the Kahler moduli defined as periods of the Kahler form over a suitable basis of 2-cycles in $H_2(B, \mathbb{Z})$. More details can be found in section 5. In order to get the tension in Planck units, we need to divide over the volume of the base or, equivalently, set the base volume to one. Hence, the above formula provides the string tension in Planck units subject to the constraint

$$\mathcal{V} = \frac{1}{2} t_I C^{IJ} t_J = 1 \quad (4.3)$$

where C^{IJ} is the inverse of the intersection form matrix in $H_2(B, \mathbb{Z})$. If one assumes that the intersection matrix is diagonal (so the trace grows linearly in N , the dimension of this space, i.e. $h^{1,1}(B) = N$), the constraint (4.3) implies

$$N t^2 \sim 1 \rightarrow t \sim N^{-1/2} \quad (4.4)$$

so that the BPS tension scales as $T \sim N^{-1/2}$. Hence, the energy scale associated with this naive estimation of the tension is

$$\Lambda_{BPS} \simeq T^{1/2} \sim N^{-1/4} \quad (4.5)$$

which coincides with the species bound in 6d, as given in (2.1).

- 5d $\mathcal{N}=1$ theories (M-theory on CY_3)

The field content is organized into a gravity multiplet, vector multiplets and hypermultiplets. In total there are $h^{1,1}$ vectors but since one combination of them

controlling the volume belongs to the hypermultiplets, there are only $h^{1,1} - 1$ vector multiplets, whose scalar components parametrize the BPS masses. The vector multiplet moduli space can then be parametrised by $h^{1,1}$ real fields subject to the constraint that the overall volume is constant. In this case, the lightest BPS states are particles arising from wrapping M2-branes on real 2-cycles and BPS strings which come from M5 branes wrapping 4-cycles. Let us first focus on the M2 branes. Their mass in Planck units is given by

$$m_{BPS} = \int_{\Sigma_2} J = q_I t^I \tag{4.6}$$

where in addition we need to set the overall volume to one,

$$\mathcal{V} = \frac{1}{6} C_{IJK} t^I t^J t^K = 1 \tag{4.7}$$

Here, the t^I are the Kahler moduli, parameterizing the Kahler form as $J = t^I \omega_I$, $\omega_I \in H^{1,1}(CY)$, the C_{IJK} are the triple intersection numbers of the divisors D_I dual to the ω_I and the number of moduli is $N = h^{1,1}(CY_3)$. Assuming again that the intersection matrix is diagonal, the constraint (4.7) implies

$$N t^3 \sim 1 \rightarrow t \sim N^{-1/3} \tag{4.8}$$

By plugging this into (4.6), the BPS mass scales as

$$\Lambda_{BPS} \sim N^{-1/3} \tag{4.9}$$

which coincides again with the species bound in 5d.

As far as M5 branes wrapping 4-cycles, the tension of the resulting string is given by

$$T_{BPS} = \int_{\Sigma_4} J \wedge J = \frac{1}{2} C_{IJK} t^I t^J q^K \tag{4.10}$$

where q^K captures the charge class of the 4-cycle. We thus have

$$T_{BPS} \sim t^2 \sim N^{-2/3} \rightarrow \Lambda_{BPS} \sim N^{-1/3} \tag{4.11}$$

We thus learn that both the M2 branes and M5 branes lead to the same naive N dependence for the cutoff for large N and both agree with the species bound in 5d.

- 4d $\mathcal{N} = 2$ theories (Type II on CY_3)

The field content includes the gravity multiplet, vector multiplets and hypermultiplets. Again, the BPS central charges are determined by the scalars on the vector multiplets, which correspond to Kahler moduli of Type IIA on a CY_3 or to the complex structure moduli of Type IIB on the mirror CY_3 . A similar classical analysis to the above for Kahler moduli sector would naively give a result that does not match with the species bound. However, the naive Kahler moduli space of Type IIA on CY_3 receives quantum corrections and it is better to compute the BPS masses in the complex

structure sector of the mirror Type IIB theory which receives no corrections and where the following exact BPS mass formula holds for D3 branes wrapped around 3-cycles:

$$m_{BPS} = \frac{1}{[\int \Omega \wedge \bar{\Omega}]^{\frac{1}{2}}} \int_{\Sigma_I} \Omega = \frac{q^I \Pi_I}{[\Pi_I \bar{\Pi}_{\bar{J}} - \bar{\Pi}_I \Pi_{\bar{J}}]^{\frac{1}{2}}} \quad (4.12)$$

Here, Π_I are the periods of the $(3,0)$ -form of the CY over a symplectic basis of 3-cycles, and are holomorphic functions on the complex structure moduli. Note that m_{BPS} is scale invariant and we can set the scale of Ω to one without loss of generality:

$$\int \Omega \wedge \bar{\Omega} = \Pi_I \bar{\Pi}_{\bar{J}} - \bar{\Pi}_I \Pi_{\bar{J}} = 1 \quad (4.13)$$

Assuming that there are no significant cancellations and the number of terms is again given by $N = h^{2,1}$, this suggest that $N \Pi^2 \sim 1$, implying that the BPS mass scales as

$$m_{BPS} \sim N^{-1/2} \quad (4.14)$$

which coincides with the species bound in 4d.

To sum up, by assuming a diagonal intersection matrix, where all nonzero entries have the same value, we get an estimation for the maximum BPS mass gap which coincides with the species scale in each dimension. This is very intriguing as the black hole arguments behind the species bound do not have direct information about the BPS spectrum of the theory. We will check this with explicit CY data in the next section, finding that the actual growth of Λ_{BPS} with $1/N$ is even bigger.

5 BPS mass gap in 6d $\mathcal{N} = \text{EFTs}$

Having explored the desert scenario in compactifications with sixteen supercharges and after giving a heuristic derivation of species bound for the case of 8 supercharges, we now move to make more precise computations for the latter case.

Although the philosophy and many of the details are very similar to the case of sixteen supercharges, with eight supercharges we encounter, for the first time, a large landscape of solutions, related to the existence of a large number of Calabi-Yau manifolds. By contrast, minimally supersymmetric compactifications in $d > 6$ are all related to M theory on $K3$ [21], possibly with frozen singularities, which makes the number of possibilities very limited.

As already noted in the previous section, compactifications with 8 supercharges include 4d $\mathcal{N} = 2$, 5d $\mathcal{N} = 1$ and 6d $\mathcal{N} = 1$ models. Their stringy descriptions are type II on a CY_3 , M-theory on a CY_3 , or F-theory on an elliptically fibered CY_3 , respectively. In this section we will address the question of gap only for 6d $\mathcal{N} = 1$ models. The reason for this restriction is a technical limitation, which we describe in more detail below.

5.1 Review of 6d $\mathcal{N} = 1$ and CY_3 's with a toric base

A 6d $\mathcal{N} = 1$ supergravity includes a number of tensor, vector, and hypermultiplets [12]. The F-theory model is an elliptic fibration over a complex surface base B , which is the

actual compactification manifold of IIB string theory, and the degenerations of the fiber describe the location and types of 7-branes. Dimensional reduction of the self-dual 4-form C_4 of IIB on the 2-cycles of the base B gives rise to the 6d tensor fields [26]; because the signature of the intersection pairing of B is $(1, T)$ for some number T (this is fixed by demanding that we have a $\mathcal{N} = 1$ model [27]), this reduction produces one anti-self dual tensor, which is part of the gravity multiplet, and $T = h^{11}(B) - 1$ tensor multiplets. Vectors arise from the worldvolume degrees of freedom of 7-branes, and neutral hypermultiplets arise from the 7-brane moduli as well as the volume of the base, totaling $h^{21}(\text{CY}_3) + 1$ hypermultiplets [26]. Additional charged hypers can arise at self-intersections of 7-branes.

We want to explore the desert scenario in this class of models — that is, we want to figure out how large can the gap be for the first heavy state. Just as before, there are many non-BPS states whose masses we cannot compute reliably: towers of KK modes, massive string excitations between branes, even the mass of charged states under vectors are unreliable because they cannot be BPS. There is a central charge in the 6d $\mathcal{N} = 1$ algebra corresponding to a string charge, representing the tension of the BPS strings charged under the tensor multiplets. Given the tension of a BPS string T , we can estimate the mass of a closed string loop as

$$m \sim \sqrt{T}. \tag{5.1}$$

This will be our estimate for the mass gap of the theory at a generic point in moduli space. The actual mass gap may be lower, but certainly not higher than this estimate. Thus, the BPS strings are an upper bound for the actual mass gap. We should note that, just like in the 16 supercharges case, some of the BPS strings that we consider actually become tensionless at finite-distance singularities in the tensor moduli space, corresponding to SCFT points [28]. By contrast, there are other strings, called supergravity strings in [28], where this does not happen. Including or excluding these SCFT strings from our considerations will lead to two different notions of maximum gap: when strings are included, the gap describes the regime of validity for the low-energy free EFT that describes the dynamics at a generic point in moduli space; when they are excluded, the gap is set by purely gravitational objects as we keep the SCFT strings modes as part of the EFT.

From the point of view of F-theory, the strings arise from wrapping $D3$ branes on 2-cycles (divisors) of the base B . The lattice of quantized string lattices $H_2(B, \mathbb{Z})$ is self-dual (but not necessarily even [29]). B is a complex manifold with vanishing odd Betti numbers, and admits a Kahler metric, and the volumes of calibrated 2-cycles are controlled by the Kahler form of the base $J_B \in H^2(B, \mathbb{R})$. Picking a two-cycle $\mathcal{C}_2 \in H_2(B, \mathbb{Z})$ the tension of the corresponding string is

$$T(\mathcal{C}_2) = \int_{\mathcal{C}_2} J_B. \tag{5.2}$$

Picking an arbitrary basis of $H_2(B, \mathbb{Z})$, which we denote $\{\alpha_I\}$, the Kahler form may be parametrized by the real Kahler parameters

$$t_I \equiv \int_{\alpha_I} J_B. \tag{5.3}$$

Since the generic charge vector in the string lattice is given by an integer linear combination of the generators, $q^I \alpha_I$ with $q^I \in \mathbb{Z}$, the BPS mass formula (5.2) can be rewritten as

$$T = q^I t_I. \tag{5.4}$$

Not every string charge admits a BPS representative, but any two BPS strings (which preserve the same supercharges) are mutually BPS, and so the set of charge sites admitting a BPS representative naturally forms a cone. Thus, there is a cone over the integers, generated by some charge vectors $\{q^{\text{BPS},I}\}$, and since the tension of any BPS state must be nonnegative, we have

$$q^{\text{BPS},I} t_I > 0, \quad \text{for all } q^{\text{BPS},I}. \tag{5.5}$$

Since BPS states correspond to linear combinations of irreducible complex curves in the base, the cone of BPS states is actually the *Mori cone* of the complex surface B . The conditions (5.5) define a dual cone in $J_B \in H^2(B, \mathbb{R})$, the Kahler cone. At boundaries of the Kahler cone, it can happen that a string becomes tensionless, signifying the appearance of an interacting SCFT.

The Kahler parameters t_I are dimensionful, but to study the mass gap in BPS particles, one typically works in six-dimensional Planck units. The 6d and 10d Planck masses are related by the volume of the surface B , which sits in one of the hypermultiplets in the F-theory description [26]. So to compute the tension in Planck units we can just divide (5.4) by the square root of the volume of the base,

$$\mathcal{V} = \frac{1}{2} \int_B J \wedge J = \frac{1}{2} C^{IJ} t_I t_J, \tag{5.6}$$

where C^{IJ} is the inverse of the intersection form matrix C_{IJ} in $H_2(B, \mathbb{Z})$. Due to the large amount of supersymmetry, hypermultiplet and tensor multiplets do not mix, and can be tuned independently. Thus, equivalently, we can just work with the constraint that the internal volume is set to 1,

$$\frac{1}{2} C^{IJ} t_I t_J = 1. \tag{5.7}$$

While working with a Planck-sized compactification space does not make much sense physically, and indeed unprotected quantities are completely unreliable, the BPS spectrum is insensitive to these details.

To address the question of the BPS maximal gap, we need to know the Mori cone generated by the $\{q_a^{\text{BPS},I}\}$. This is a difficult question in general, and we do not know the answer for generic threefolds, or even generic elliptically fibered ones. This is why we chose not to look at general 5d $\mathcal{N} = 1$ or 4d $\mathcal{N} = 2$ models. However, for those elliptic threefolds where the base B is toric, the full answer was provided in [30]. This reference enumerated all possible toric bases for 6d F-theory models, and there is roughly 60000 of them; so toric bases provide a relatively large “sub-landscape” of 6d $\mathcal{N} = 1$ compactifications where we can answer the question.

As explained in [30], the intersection structure and Mori cone of a toric surface with a given nonvanishing Hodge number $h^{1,1}(B)$ is fully described in terms of an overcomplete

set of $h^{11}(B) + 2$ divisors β_K , with intersection form $C_{KL} = \beta_K \cdot \beta_L$, where

$$\beta_K \cdot \beta_{K+1} = \beta_K \cdot \beta_{K-1} = 1, \quad \beta_K \cdot \beta_L = 0 \quad \text{otherwise.} \quad (5.8)$$

The entries $\beta_K \cdot \beta_K$ are unconstrained and can have either sign, although at large h_{11} we will find that the entries are mostly negative. The Mori cone is generated by the dual cone spanned by the $\{\beta_K\}$. We can now choose $h^{11}(B)$ elements of this list, say the first $h^{11}(B)$ entries, to construct the dual basis $\{\alpha_I\}$ above, and the associated Kahler parameters t_I . Using duality, we can also write the Kahler form as

$$J(q) = t^I \alpha_I \cdot (q^K \beta_K) = q^K C_{KI} t^I, \quad (5.9)$$

where we have defined the dual Kahler parameters as

$$t^I = C^{IJ} t_J. \quad (5.10)$$

In terms of the t^I , the volume of the $h^{11}(B) + 2$ curves that generate the Mori cone is simply

$$\mathcal{V}_K = C_{KI} t^I, \quad (5.11)$$

and the condition (5.5) is simply that $\mathcal{V}_K \geq 0$.

5.2 Minimization algorithm and results

We are now ready to explain the algorithm we will follow to find the maximal BPS gap. The starting point is equation (5.11) for the volumes. We wish to maximize the minimum of the \mathcal{V}_K over all generators of the Mori cone, as a function of the t^I , subject to the constraint that the total volume is 1, (5.7). Because the curve volumes depend homogeneously on the Kahler parameters t^I , we can always rescale them such that e.g. the smallest volume curve has any volume that we wish, say 1. The total volume of the base \mathcal{V} is no longer one and the volume of the minimal curve before rescaling is related to the base volume \mathcal{V} after rescaling as

$$\min(\mathcal{V}_K)|_{\mathcal{V}=1} = \frac{1}{\sqrt{\mathcal{V}|_{\min(\mathcal{V}_K=1)}}}. \quad (5.12)$$

Our original problem, which is maximizing the left hand side of (5.12), can then be traded by the dual problem of minimizing the right hand side. This is constrained optimization of a single function, as opposed to minmax optimization of $h^{11}(B) + 2$ functions in the original problem.

In turn, this simplified problem can be solved exactly. The condition that the volume of all curves is ≥ 1 defines a subcone of the Kahler cone, called the *stretched Kahler cone* [31], and so our problem is minimizing the quadratic function

$$\mathcal{V} = \frac{1}{2} C_{IJ} t^I t^J \quad (5.13)$$

over a convex cone. The minima are located either at where $\nabla \mathcal{V}$ vanishes, on the boundary of the cone or at infinity. Since

$$\nabla \mathcal{V} \cdot q = C_{IJ} t^I q^J \quad (5.14)$$

is the volume of a curve, and all the curve volumes are ≥ 1 , the gradient never vanishes. Furthermore, since the volume of the curves are positive (i.e. they are effective), $\nabla\mathcal{V} > 0$, implying that the base volume decreases as the volume of the curves decreases. Hence, the minimum of \mathcal{V} is located on the boundary of the stretched Kahler cone, where the curves take their minimum allowed value. Since the cone is defined by the condition

$$C_{KI}t^I \geq 1 \quad \forall K \tag{5.15}$$

the boundary is piecewise linear, and each face is defined by setting the volume of one of the $h^{11}(B) + 2$ defining curves to 1. Within this face, we can run the argument again, to conclude that the minimum is located at another point where a second curve is set to 1, and proceeding iteratively, we conclude that the minimum is found at a point where the volumes of $h^{11}(B)$ curves are set exactly to one.

Once the volumes of $h^{11}(B)$ curves are set, the Kahler parameters t^I are fully determined, and so are the volumes of the remaining 2 curves. This leads to the following simple algorithm to find the minimum:

1. For each of the

$$\begin{pmatrix} h^{11}(B) + 2 \\ h^{11}(B) \end{pmatrix} = \frac{(h^{11}(B) + 2)(h^{11}(B) + 1)}{2} \tag{5.16}$$

possible choices of $h^{11}(B)$ curves, find the value of the parameters t^I that set all of their volumes to 1.

2. Find the volumes of the additional 2 curves. If either of these volumes is below 1, the corresponding t^I lies outside of the stretched Kahler cone, and we ignore it. Otherwise, the point is kept, and we compute the corresponding \mathcal{V} .
3. The global minimal volume corresponds to the minima of all the values of \mathcal{V} found in the previous point.

We applied this algorithm to the list of 61539 toric bases constructed in [30]. We excluded examples including curves of self-intersection -9,-10 and -11, since as explained in [30] the corresponding bases are not toric after blow-up, and the cycle structure described above does not apply to these examples. This still leaves us with 34868 different explicit bases, which constitute the “mini-Landscape” where we performed our numerical experiments.

We now illustrate the algorithm above with a particular case, corresponding to the entry labelled with number 10 in the database [30]. This is a toric Calabi-Yau with intersection matrix of ambient space cycles given by

$$q_{\text{big}} = \begin{pmatrix} 0 & 1 & 0 & 1 \\ 1 & -8 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 8 \end{pmatrix}. \tag{5.17}$$

As explained in [30], the corresponding toric base has $h_{11} = \text{rank}(q_{\text{big}}) - 2 = 2$. We pick a basis of 2 - cycles consisting of the first two rows in the matrix q_{big} , and define Kahler




Figure 2. Logarithmic plot of the minimal gap obtained from equation (5.1) versus $h^{11}(B)$, for the 34868 toric bases in the toric mini-Landscape described in [30]. There is some spread corresponding to different models with the same $h^{11}(B)$, but there is an overall linear trend, with the best fit shown, corresponding to the function $h^{11}(B)^{-n}$ with $n \sim 0.7$.

parameters t_1 and t_2 as described above. With these, we can write down explicit expressions for the volumes (equivalently, BPS masses) of all curves. In the case at hand, these are

$$\{t_1, t_1 - 8t_2, t_2, t_1\}. \tag{5.18}$$

The first step of the algorithm now requires us to consider all $\binom{4}{2} = 6$ pairs of curves above and, for each, find the values of (t_1, t_2) such that the corresponding volumes are one. We then compute the two remaining curve volumes and keep only those cases where these are bigger than one, as instructed in the second step. In this case, there are only two possibilities, corresponding to setting the volumes of either curves 1 and 4 or 3 and 4 to one, and both having $t_1 = 9, t_2 = 1$. The corresponding base volume is $\mathcal{V} = 5$. Per (5.12), the BPS gap is then $1/\sqrt{5}$ in Planck units in this case.

The results for all the cases are shown in figure 2, where we plot $h^{11}(B)$ versus the minimal gap in Planck units (equivalently, the cutoff Λ_{BPS} set by the BPS states), given by (5.1) applied to the solution of the minmax problem in the left hand side of (5.12). We see some spread in the data, and a general downwards trend, which is roughly linear in the logarithmic plot we provide; the best fit is the curve

$$m_{\text{gap}} \sim \Lambda_{\text{BPS}} \sim \frac{1}{h^{11}(B)^n}. \tag{5.19}$$

with $n \sim 0.7$. The total volume of the base scales with $h^{11}(B)$ as $h^{11}(B)^{2n} = h^{11}(B)^{1.4}$, which will be relevant later.

Unfortunately we cannot provide a sharp upper bound on $h^{11}(B)$, since this would also depend on the $\mathcal{O}(1)$ coefficients responsible for the twiddles in (5.19) and (5.22) which we do not control. We should also point out that it has been shown that the Landscape of elliptically fibered Calabi-Yau threefolds is finite, see [32]. This is consistent with our results, since the mini-Landscape we look at is contained within this class.

At this point we would like to recall the general argument provided in section 3, where with some genericity assumption, we concluded that the cutoff coming from the strings should saturate the species bound. The reason for the mismatch is that the genericity assumption we made, that C_{IJ} is close to the identity matrix, is a bad one in our mini-Landscape. As shown in eq. (5.8), the actual intersection matrix we obtain has off-diagonal entries and diagonal values which can be far from 1. Writing the volume as

$$\mathcal{V} = \frac{1}{2} C^{IJ} t_I t_J, \tag{5.20}$$

the vector t_I of volumes has components all close to 1 at the minimum (and in fact, at least all components except for two are exactly one). The partial inverse matrix C^{IJ} has many nonzero entries of all sizes, and is in fact far from the identity for a typical base. As a result, $C_{IJ} t^I t^J$ grows faster than $N t^2$, resulting in a faster decrease of t with N .

The failure of this argument should be viewed as a cautionary tale: it is dangerous to make general arguments about string compactifications, involving topological data such as C_{IJ} , or triple intersections C_{IJK} using “genericity” assumptions without actually having concrete examples to back them up.

5.3 Field theory vs. black hole mass gap and species bound

It is also very important to understand the nature of the BPS strings providing the gap above. Roughly speaking, the strings come in three different classes [28], according to the sign of the self-intersection $q \cdot q = q^I C_{IJ} q^J$:

- Strings with $q \cdot q < 0$ can become tensionless at finite distance singularities in moduli space, where they give rise to SCFT’s. These are the analog of the field theory states described in subsection 3.2; in a sense, they are captured by (conformal) field theory description, and at a generic point in moduli space they are heavy only due to a field-theoretic Higgsing mechanism. They are also strongly interacting and self-repulsive: they do not form bound states at any point in moduli space, as the corresponding large charge states would be described by a CFT of negative central charge.
- Strings with $q \cdot q > 0$ are intrinsically gravitational objects, called *supergravity strings* in [28]. Unlike their field-theory counterparts, they can form bound states, and for large charge are described by black strings of nonzero horizon area, with near-horizon geometry $AdS_3 \times S^3$. They never become tensionless, and in fact, their tension is lower-bounded by 1 in Planck units.
- Finally, there are states with $q \cdot q = 0$, such as the fundamental heterotic string, which sit at the divide between field theory and gravity strings. For large values of the charge, these may be described by singular dilatonic zero-area solutions, such as the small black holes in [4]. Hence we may call them *small supergravity strings*. Unlike honest supergravity strings, whose tension remains superplanckian throughout moduli space, small supergravity strings become tensionless at infinite distance limits in moduli space. The equivalent version in 4d theories were called *EFT strings* [33] and play an essential role when testing the Distance Conjecture [8].

For $h_{11} > 2$, the minimal volume curves have negative self-intersection, and so the limits presented in figure 2 for h_{11} above this value all correspond to deformations of SCFT strings where the strings have acquired maximal tension. The physical interpretation of the cutoff is then the scale at which the IR free EFT at generic points in moduli space breaks down (similar to the GUT scale in the GUT scenario), and is replaced by an interacting, but still non-gravitational description. Above this cutoff, the effective description involves several interacting light strings that are not mutually local (each of them becomes tensionless at different points in moduli space), and so we are unlikely to be able to describe them via a local field theory.

We can now compare this BPS gap with the prediction of the species bound. To do this, we compute the total number of massless modes (tensors, vectors and hypermultiplets) for each CY, by identifying the existing non-higgsable clusters and using anomaly cancellation conditions, and using the results of [30] to count the number of vector and hyper multiplets. Each of these multiplets has eight bosonic degrees of freedom. Taking this into account, the result is plotted in figure 3 where the best fit reads

$$\Lambda_{\text{BPS}} \sim \frac{1}{N^{0.62}} \tag{5.21}$$

Intriguingly, the gap associated to the string states decreases with N faster than what the six-dimensional species bound would suggest,

$$\Lambda_{\text{sp}} \sim \frac{1}{N^{\frac{1}{4}}}. \tag{5.22}$$

As we argued in section 2, we can have points which are desert-like in which case the species bound cutoff should depend mainly on the number of exactly massless fields. Assuming this we will now check if this leads to any contradictions between (5.21) and (5.22).

We will also make the reasonable assumption that below the species scale Λ_{sp} there exists a consistent EFT coupled to gravity including all the states with mass less than this scale. On the other hand we expect that in the string examples we studied the points which give the maximal gap cannot be described in the context of a consistent EFT as in that case we have no bounds in how large we can make the BPS tensions. By contrast, in an actual field theory description, with gravity decoupled, the vevs of tensor multiplets (and hence, the string tensions) can grow without bound. Therefore this implies that $\Lambda_{\text{BPS}} \geq \Lambda_{\text{sp}}$.

Since the slope of Λ_{BPS} and Λ_{sp} are different, they are going to cross at some point. Whenever Λ_{BPS} is above the species bound, there is no conflict, as the mismatch can be accounted for in principle by lighter non-BPS states we cannot control, such as Kaluza-Klein modes, etc. Having the species bound above Λ_{BPS} would not only go against the desert scenario, but would also imply that we can have a field theory description above Λ_{BPS} leading to a contradiction with $\Lambda_{\text{BPS}} \geq \Lambda_{\text{sp}}$. Hence, the dependence of Λ_{BPS} we found experimentally from the mini-Landscape of toric bases is incompatible with this version of the species bound in terms of massless fields, unless the family is cut short at a maximal value of N , as illustrated in figure 3, so the region after the two curves cross is ruled out. In other words, absence of conflict with this strengthened version of the species

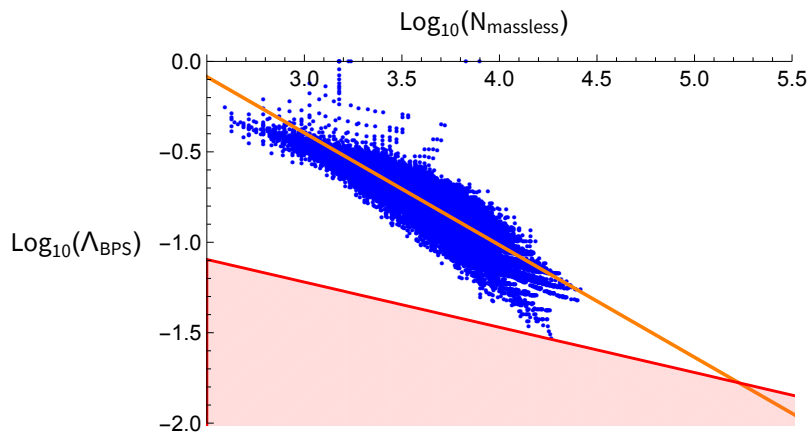


Figure 3. Schematic plot where the toric examples are plotted on top of the restriction required by a qualitative version of the species bound, involving only massless fields. This naive expression would exclude situations where the lightest field is below the blue line, and so is not compatible with the linear trend observed in the data if this is extrapolated indefinitely, suggesting a maximum value for N . There is an uncontrolled $\mathcal{O}(1)$ factor in the species bound line, which manifests in this logarithmic plot as an uncontrolled offset. This is why the plot is only schematic, and we cannot make a sharp prediction for the maximal N .

bound provides yet another hint that the Calabi-Yau landscape is finite, by suggesting an upper bound of the number of massless fields and, consequently, on the Hodge numbers.

It is also interesting to explore the gap set by purely gravitational objects. To compute this gap, we must look at (small) supergravity strings. This can be addressed with an algorithm similar to the one outlined in the previous section, but we can in fact achieve a shortcut based on the structure of the toric mini-Landscape we are studying. As explained in [30], all toric bases we consider are obtained from blow-ups of either \mathbb{CP}^2 or one of the Hirzebruch surfaces \mathbb{F}_m , where $0 \leq m \leq 11$, which hence constitute a set of “minimal models”. For any given toric base B we can always move to a point in the tensor moduli space which is arbitrarily close to the blow-down to one of the minimal models. When doing this no supergravity string becomes tensionless. This means that the global maximum of the minimal supergravity string tension will then be at least as large as the value obtained from one of the 13 minimal models described above.² Having reduced the analysis to these 13 cases, one can compute explicitly that the maximum gap is $\mathcal{O}(1)$, independently of $h^{11}(B)$, similar to what we had seen in the context of theories with higher supersymmetry.

²There is yet another way to define the quantum gravity cutoff namely that defined by the horizon size of the black strings, which we also explored. Since supergravity strings have a nonzero horizon area, we can estimate the cutoff of the EFT from the size of the smallest possible value of the area, as $\Lambda \sim (\text{Horizon area})^{-\frac{1}{3}}$. In turn, the horizon area is related to the central charges of the string [22], which themselves depend on the intersection numbers and the products of the charge vector with the canonical class of the base. We explored this question via a Monte-Carlo algorithm and found that, for moderate $h^{11}(B)$, there is always a curve with horizon area of Planck size, with no noticeable dependence in $h^{11}(B)$. Thus, it seems that the cutoff estimated by this method is also $\mathcal{O}(1)$ in Planck units.

5.4 Relation to the Tadpole Conjecture

Our results can also be of interest for other questions currently under scrutiny. For instance, there is a strong connection to the Tadpole Conjecture [34, 35]. This conjecture claims that, in vacua with a large number of moduli, the number of flux quanta that would be required to stabilize all of them is too large to satisfy the tadpole conditions. This yields an obstruction to stabilize all moduli and puts into question the existence of a large landscape of flux vacua. In the context of Type IIB Calabi-Yau flux compactifications, the conjecture claims that the flux contribution to the D3-brane tadpole condition scales as

$$N_{\text{flux}} = \int F_3 \wedge H_3 > 2\alpha(h_{2,1} + 1) \quad \text{for large } h_{2,1} \quad (5.23)$$

where $h_{2,1}$ is the number of complex structure moduli and α is conjectured to be $1/3$. Reference [36] checked this conjecture in the large complex structure regime by using Kahler moduli data of the mirror dual Calabi-Yau threefold in the Kreuzer-Skarke database [37] presented in [31]. The large complex structure regime maps to the stretched Kahler cone of the mirror defined by setting all curve volumes ≥ 1 to keep computational control. In this regime, minimization conditions of the self-dual flux potential imply that N_{flux} scales as the volume of the mirror plus some additional terms. Hence, checking the conjecture reduces to finding the minimal value of the volume within the stretched Kahler cone. This is exactly the same optimization problem that we have attacked in this paper! In [36] this minimal volume was estimated, although not computed explicitly, using the results of [31] to grow as $N_{\text{flux}} \gtrsim \min(\text{vol}(CY_3)) \sim h_{2,1}^{6,6}$. We cannot apply directly our results to this setup, since we have only computed explicitly the volume of the base of the elliptically fibered threefolds, and not of the total Calabi-Yau. However, assuming they scale similarly, we obtain a power much smaller than the one estimated in [36] (namely $\sim h_{2,1}^{1,4}$) which is consistent but slightly stronger than the Tadpole Conjecture.

In the case of $d = 5$ $\mathcal{N} = 1$ EFTs the optimization problem to obtain the maximal BPS gap resembles even more the one required to check the Tadpole Conjecture at large complex structure in [36], since it is the overall volume of the CY that enters in the BPS mass formula. Interestingly, the naive estimation of the BPS mass gap performed in section 4 using genericity arguments (yielding $\mathcal{V}_{\text{min}} \sim N$) would precisely reproduce the scaling of the Tadpole Conjecture. However, the lesson from studying actual examples in 6d is that these genericity arguments can be too simplistic and Λ_{BPS} can grow even faster with the number of massless fields. We will promote this observation to the “BPS desert conjecture” in the Conclusions, which would imply that the Tadpole Conjecture is satisfied with room to spare. It should be emphasized that arguments of this kind refer to generic features of the Calabi-Yau landscape, and do not exclude the possibility of special cases that evade them.

We suspect that the connection with the Tadpole Conjecture runs deeper than what this analysis in the large complex structure regime (or the stretched Kahler cone of the mirror) suggests. Recall that the maximum BPS mass gap occurs at a special point deep in the interior of the moduli space far away from all boundaries. Taking into account that N_{flux} for self-dual fluxes increases when approaching the boundaries [36], it is natural to expect that the minimal value of N_{flux} will happen near (or perhaps even at the same)

point that maximizes the BPS mass gap. This provides further motivation to study these special points in more detail.

6 Conclusions and finiteness of the CY landscape

In this paper we have explored the desert scenario in the string landscape; namely searching for points on moduli of supergravity theories where there are no massive states until mass scales close to the Planck scale. We have found that for a number of examples the points on moduli where we obtain a maximal gap are well in the interior of the moduli space, and in a sense almost by definition among the most interesting and difficult points to analyze in the moduli space of supergravity theories. As such, regardless of the motivation of the present paper, *it would be interesting to have a more systematic study of such special points in the string landscape*. In addition we have focused on the mass scales set by BPS data which allow us to compute the maximal cutoff set by BPS states Λ_{BPS} exactly. Here we have found two surprises: the first surprise is that the naive supersymmetric arguments for theories with 8 supercharges lead to the Λ_{BPS} scaling with the number of massless modes N exactly as predicted by the species scale. The second surprise is that this naive argument is not consistent with a large class of actual examples which suggest that Λ_{BPS} becomes too small too quickly with increasing N . To avoid contradiction with the species scale we propose that there are massive states lower than Λ_{BPS} not associated with BPS objects leading to $\Lambda < \Lambda_{BPS}$ which changes the slope of this dependence. However, this can only work in a finite range of N due to the fact that the BPS cutoff sets an upper bound for the cutoff of the full theory. This therefore leads to an upper bound on the number of massless modes on Calabi-Yau manifolds. Turning this around, if we could find a theoretical reason behind the higher slope for the dependence of Λ_{BPS} on the number of massless modes, it would lead to a theoretical explanation of the upper bound on the number of massless modes in these theories.

Motivated by this example we make the more general conjecture, the *BPS desert conjecture*, that $\Lambda_{BPS} \sim N^{-a}$ where N is the number of massless modes with $a > \frac{1}{d-2}$ in all cases and in all dimensions for supergravity theories. This would then give a theoretical reason for the existence of an upper bound on the number of massless modes in a supergravity theory, which is a requirement for finiteness of the string landscape. Given that Λ_{BPS} is in principle computable with known techniques it would be interesting to check the general BPS desert conjecture. This conjecture also has other important implications. For example considering the case of $d = 5$ and applying it to the case of M-theory on arbitrary (not necessarily elliptic) CY 3-folds we obtain that if we fix the volume of holomorphic 2-cycles to be larger than 1, then the volume of CY will scale as $h_{1,1}^b$ with $b > 1$ (since the volume of 2-cycles would scale as N^{-a} with $a > \frac{1}{3}$ if we fixed the overall volume to be 1) and this leads to a stronger form of the tadpole conjecture at large complex structure³ using the setup of [36].

It is exciting to try to apply the general lessons we have learned from this supersymmetric study to the Standard Model, even though we have to be cautious due to the lack

³It is natural to speculate that the 4d version of the BPS desert conjecture will also lead to a similar bound for an arbitrary point in complex structure in flux compactifications.

of supersymmetry in the real world. One of the typical scenarios considered plausible for the SM is precisely the desert scenario, where there are no extra massive states other than SM fields until we reach the GUT scale. At the GUT scale we get new massive states which include massive fields needed to complete the Standard Model to a GUT (perhaps in a higher dimensional sense as is often the case in string theory [38]). Viewing the Standard Model states as the analog of massless states, in a sense this is very similar to what we have found at special interesting points in the moduli of the supersymmetric theories, where the BPS gap or the species scale are set by massive field theory states. We can push the massive states Λ very close to the Planck scale, differing from it by N^{-a} where N is the number of massless modes and $a > \frac{1}{d-2}$ which for $d = 4$ leads to $a > 1/2$. Taking into account that the number of massless modes for the Standard Model is about $N \sim 10^2$, this suggests that $\Lambda_{GUT} < 10^{-1}M_P$. In particular if we had $a \sim 1$ we would rather be close to the expected value of $\Lambda_{GUT} \sim 10^{-2}M_P$. An alternative explanation of the mismatch could be the existence of a dark sector with a large number of degrees of freedom, with $N \sim 10^4$, that would also lower the scale.

Acknowledgments

We would like to thank Severin Luest for useful discussions. We are grateful to the Simons Center for Theoretical Physics, Stony Brook, for hospitality during the early stages of this work in the context of its Summer Workshop. The work of MM, CV is supported by a grant from the Simons Foundation (602883, CV) and by the NSF grant PHY-2013858, which also supported IV in the early stages of this work. The work of IV is now supported by the grants RYC2019-028512-I from the MCI (Spain), PGC2018-095976-B-C21 from MCIU/AEI/FEDER, UE and the grant IFT Centro de Excelencia Severo Ochoa SEV-2016-0597.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited. SCOAP³ supports the goals of the International Year of Basic Sciences for Sustainable Development.

References

- [1] C. Vafa, *The string landscape and the swampland*, [hep-th/0509212](https://arxiv.org/abs/hep-th/0509212) [HUTP-05-A043] [[INSPIRE](https://inspirehep.net/literature/46070)].
- [2] B.S. Acharya and M.R. Douglas, *A finite landscape?*, [hep-th/0606212](https://arxiv.org/abs/hep-th/0606212) [IC-2006-42] [[INSPIRE](https://inspirehep.net/literature/57333)].
- [3] M.R. Douglas, *Spaces of quantum field theories*, *J. Phys. Conf. Ser.* **462** (2013) 012011 [[arXiv:1005.2779](https://arxiv.org/abs/1005.2779)] [[INSPIRE](https://inspirehep.net/literature/87001)].
- [4] Y. Hamada, M. Montero, C. Vafa and I. Valenzuela, *Finiteness and the swampland*, *J. Phys. A* **55** (2022) 224005 [[arXiv:2111.00015](https://arxiv.org/abs/2111.00015)] [[INSPIRE](https://inspirehep.net/literature/2000000)].
- [5] G. Dvali and M. Redi, *Black hole bound on the number of species and quantum gravity at LHC*, *Phys. Rev. D* **77** (2008) 045027 [[arXiv:0710.4344](https://arxiv.org/abs/0710.4344)] [[INSPIRE](https://inspirehep.net/literature/15111)].

- [6] G. Dvali, *Black holes and large N species solution to the hierarchy problem*, *Fortsch. Phys.* **58** (2010) 528 [[arXiv:0706.2050](#)] [[INSPIRE](#)].
- [7] G. Dvali and C. Gómez, *Species and strings*, [arXiv:1004.3744](#) [CERN-PH-TH-2010-069] [[INSPIRE](#)].
- [8] H. Ooguri and C. Vafa, *On the geometry of the string landscape and the swampland*, *Nucl. Phys. B* **766** (2007) 21 [[hep-th/0605264](#)] [[INSPIRE](#)].
- [9] N. Arkani-Hamed, S. Dimopoulos and S. Kachru, *Predictive landscapes and new physics at a TeV*, [hep-th/0501082](#) [SLAC-PUB-10928] [[INSPIRE](#)].
- [10] T.W. Grimm, E. Palti and I. Valenzuela, *Infinite distances in field space and massless towers of states*, *JHEP* **08** (2018) 143 [[arXiv:1802.08264](#)] [[INSPIRE](#)].
- [11] B. Heidenreich, M. Reece and T. Rudelius, *The weak gravity conjecture and emergence from an ultraviolet cutoff*, *Eur. Phys. J. C* **78** (2018) 337 [[arXiv:1712.01868](#)] [[INSPIRE](#)].
- [12] J. Polchinski, *String theory. Volume 2: superstring theory and beyond*, Cambridge University Press, Cambridge, U.K. (2007) [[INSPIRE](#)].
- [13] H.-C. Chang and L.-C. Wang, *A simple proof of Thue's theorem on circle packing*, [arXiv:1009.4322](#).
- [14] O. Aharony, Z. Komargodski and A. Patir, *The moduli space and M (atrix) theory of $9d$ $N = 1$ backgrounds of M /string theory*, *JHEP* **05** (2007) 073 [[hep-th/0702195](#)] [[INSPIRE](#)].
- [15] T. Ortin, *Gravity and strings*, Cambridge University Press, Cambridge, U.K. (2015) [[INSPIRE](#)].
- [16] H.-C. Kim, G. Shiu and C. Vafa, *Branes and the swampland*, *Phys. Rev. D* **100** (2019) 066006 [[arXiv:1905.08261](#)] [[INSPIRE](#)].
- [17] M. Montero and C. Vafa, *Cobordism conjecture, anomalies, and the string lamppost principle*, *JHEP* **01** (2021) 063 [[arXiv:2008.11729](#)] [[INSPIRE](#)].
- [18] M. Dierigl and J.J. Heckman, *Swampland cobordism conjecture and non-Abelian duality groups*, *Phys. Rev. D* **103** (2021) 066006 [[arXiv:2012.00013](#)] [[INSPIRE](#)].
- [19] M. Cvetič, M. Dierigl, L. Lin and H.Y. Zhang, *String universality and non-simply-connected gauge groups in $8d$* , *Phys. Rev. Lett.* **125** (2020) 211602 [[arXiv:2008.10605](#)] [[INSPIRE](#)].
- [20] Y. Hamada and C. Vafa, *$8d$ supergravity, reconstruction of internal geometry and the swampland*, *JHEP* **06** (2021) 178 [[arXiv:2104.05724](#)] [[INSPIRE](#)].
- [21] A. Bedroya, Y. Hamada, M. Montero and C. Vafa, *Compactness of brane moduli and the string lamppost principle in $d > 6$* , *JHEP* **02** (2022) 082 [[arXiv:2110.10157](#)] [[INSPIRE](#)].
- [22] H.-C. Kim, H.-C. Tarazi and C. Vafa, *Four-dimensional $N = 4$ SYM theory and the swampland*, *Phys. Rev. D* **102** (2020) 026003 [[arXiv:1912.06144](#)] [[INSPIRE](#)].
- [23] A. Font et al., *Exploring the landscape of heterotic strings on T^d* , *JHEP* **10** (2020) 194 [[arXiv:2007.10358](#)] [[INSPIRE](#)].
- [24] A. Dabholkar and J.A. Harvey, *Nonrenormalization of the superstring tension*, *Phys. Rev. Lett.* **63** (1989) 478 [[INSPIRE](#)].
- [25] O. Bergman, M.R. Gaberdiel and G. Lifschytz, *String creation and heterotic type I' duality*, *Nucl. Phys. B* **524** (1998) 524 [[hep-th/9711098](#)] [[INSPIRE](#)].
- [26] C. Vafa, *Evidence for F theory*, *Nucl. Phys. B* **469** (1996) 403 [[hep-th/9602022](#)] [[INSPIRE](#)].

- [27] V. Kumar, D.R. Morrison and W. Taylor, *Global aspects of the space of 6D $N = 1$ supergravities*, *JHEP* **11** (2010) 118 [[arXiv:1008.1062](#)] [[INSPIRE](#)].
- [28] S. Katz, H.-C. Kim, H.-C. Tarazi and C. Vafa, *Swampland constraints on 5d $N = 1$ supergravity*, *JHEP* **07** (2020) 080 [[arXiv:2004.14401](#)] [[INSPIRE](#)].
- [29] N. Seiberg and W. Taylor, *Charge lattices and consistency of 6D supergravity*, *JHEP* **06** (2011) 001 [[arXiv:1103.0019](#)] [[INSPIRE](#)].
- [30] D.R. Morrison and W. Taylor, *Toric bases for 6D F-theory models*, *Fortsch. Phys.* **60** (2012) 1187 [[arXiv:1204.0283](#)] [[INSPIRE](#)].
- [31] M. Demirtas, C. Long, L. McAllister and M. Stillman, *The Kreuzer-Skarke axiverse*, *JHEP* **04** (2020) 138 [[arXiv:1808.01282](#)] [[INSPIRE](#)].
- [32] M. Gross, *A finiteness theorem for elliptic Calabi-Yau threefolds*, *Duke Math. J.* **74** (1994) 271 [[alg-geom/9305002](#)] [[PRINT-93-0392](#)] [[INSPIRE](#)].
- [33] S. Lanza, F. Marchesano, L. Martucci and I. Valenzuela, *The EFT stringy viewpoint on large distances*, *JHEP* **09** (2021) 197 [[arXiv:2104.05726](#)] [[INSPIRE](#)].
- [34] I. Bena, J. Blåbäck, M. Graña and S. Lüst, *The tadpole problem*, *JHEP* **11** (2021) 223 [[arXiv:2010.10519](#)] [[INSPIRE](#)].
- [35] I. Bena, J. Blåbäck, M. Graña and S. Lüst, *Algorithmically solving the tadpole problem*, *Adv. Appl. Clifford Algebras* **32** (2022) 7 [[arXiv:2103.03250](#)] [[INSPIRE](#)].
- [36] E. Plauschinn, *The tadpole conjecture at large complex-structure*, *JHEP* **02** (2022) 206 [[arXiv:2109.00029](#)] [[INSPIRE](#)].
- [37] M. Kreuzer and H. Skarke, *Complete classification of reflexive polyhedra in four-dimensions*, *Adv. Theor. Math. Phys.* **4** (2000) 1209 [[hep-th/0002240](#)] [[INSPIRE](#)].
- [38] C. Beasley, J.J. Heckman and C. Vafa, *GUTs and exceptional branes in F-theory — I*, *JHEP* **01** (2009) 058 [[arXiv:0802.3391](#)] [[INSPIRE](#)].