

Disorder averaging and its UV discontentsJonathan J. Heckman,^{1,*} Andrew P. Turner^{1,†} and Xingyang Yu^{2,‡}¹*Department of Physics and Astronomy, University of Pennsylvania,
Philadelphia, Pennsylvania 19104, USA*²*Center for Cosmology and Particle Physics, New York University, New York, New York 10003, USA*

(Received 11 January 2022; accepted 1 March 2022; published 25 April 2022)

We present a stringy realization of quantum field theory ensembles in $D \leq 4$ spacetime dimensions, thus realizing a disorder averaging over coupling constants. When each member of the ensemble is a conformal field theory with a standard semiclassical holographic dual of the same radius, the resulting bulk can be interpreted as a single asymptotically anti-de Sitter space geometry with a distribution of boundary components joined by wormhole configurations, as dictated by the Hartle–Hawking wave function. This provides a UV completion of a recent proposal by Marolf and Maxfield that there is a high-dimensional Hilbert space for baby universes, but one that is compatible with the proposed swampland constraints of McNamara and Vafa. This is possible because our construction is really an approximation that breaks down both at short distances, but also at low energies for objects with a large number of microstates. The construction thus provides an explicit set of counterexamples to various claims in the literature that holographic and effective field theory considerations can be reliably developed without reference to any UV completion.

DOI: [10.1103/PhysRevD.105.086021](https://doi.org/10.1103/PhysRevD.105.086021)**I. INTRODUCTION**

Disorder averaging is a well-defined procedure in any parametric family of quantum field theories (QFTs). Operationally, one fixes the parameters of the theory, computes correlation functions for this choice, and then at the end, performs a further average over a *classical* probability distribution for these parameters. This is clearly a useful tool for gaining access to “typical” behavior in various systems with a high degree of complexity (see, e.g., [1–3]). It has also appeared in the context of holography (see, e.g., [4–10]) as well as other areas within high energy theory [11–14].

Indeed, recent bottom-up considerations suggest that combining the principles of holography with effective field theory in the context of the Euclidean gravitational path integral naturally results in the appearance of disorder averaging in the conformal field theory (CFT) on the boundary (see, e.g., [7,15–19] as well as [20–23]). One particularly surprising aspect of these considerations is that,

at present, they have resisted an embedding in string theory.¹ It is not hard to see that finding such a completion might be difficult, because the presence of an explicit classical distribution would seem to require treating gravity as an open system.

The hope, then, is that constraints imposed from requiring a UV completion can be sufficiently decoupled from long distance effects in the putative gravity dual. That being said, it is not entirely obvious that this is really the case. For example, in [31], it was argued that the swampland cobordism conjecture (see [32]) precludes the existence of an AdS/CFT correspondence with ensemble averaging for $D > 2$ boundary systems, and in the case of $D \leq 2$, the corresponding systems should be viewed as UV completed in a higher-dimensional system. A closely related point is that for many of these considerations it is actually quite important that the couplings have no position dependence at all. This is required to have a consistent interpretation in terms of the creation of baby universes.²

*jheckman@sas.upenn.edu

†turnerap@sas.upenn.edu

‡xy1038@nyu.edu

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

¹Other examples include the case of “double holography” [24], which involves coupling a large N gauge theory to gravity. For massless gravity, this is in sharp contradiction with all known string constructions and swampland considerations [25,26], but for a massive graviton, it might be possible [27–30].

²This can be seen as a consequence of Gauss's law for translation invariance; an emitted baby universe cannot carry energy or momentum because it has no boundary. We thank H. Maxfield for helpful correspondence on this point.

Our aim in this paper is to engineer ensemble averaged QFTs embedded in string theory. Let us state at the outset that aesthetically the construction we present is a rather contrived UV completion. That being said, it provides us with a general framework for testing the claim that disorder averaging in the context of holography can truly be decoupled from stringy considerations.

The main idea behind our construction is to produce QFTs decoupled from gravity via open strings localized on a small patch of the compactification geometry. To get a statistical ensemble, we simply consider multiple stacks of branes at different locations in the transverse extra dimensions. In particular, by varying the profile of non-normalizable modes in these directions, we can realize different low energy effective field theories with *identical* field content but with different values of the physical parameters. Given a set of K stacks with couplings λ_k for $k = 1, \dots, K$, we can consider a special class of operators O_k for $k = 1, \dots, K$ given by “tracing” over all the different stacks,

$$\mathbb{O} \equiv O_1 + \dots + O_K. \quad (1)$$

The key point is that the correlation functions of these \mathbb{O} s factorize to leading order,

$$\langle \mathbb{O}(x)\mathbb{O}'(y) \rangle \approx \sum_{1 \leq k \leq K} \langle O_k(x)O'_k(y) \rangle, \quad (2)$$

which in turn leads to an averaging over couplings. At short distances, this approximation breaks down because we become sensitive to massive excitations that have been integrated out to reach the effective field theory in the first place.

Depending on the number density of stacks with a given value of λ_k , it is clear that this is building up a “binned” version of an ensemble average. Provided we can engineer a suitable internal profile for the couplings and populate the stacks at distinct values of the couplings, it would appear that we can produce a large class of probability distributions $p(\lambda)$ for ensemble averaging. Note also that taking $K \gg 1$ provides a general way to get a good approximation by the binned distribution of its idealized smooth counterpart.

A particularly important special case is provided by requiring each stack to realize a conventional holographic dual with the same value of the bulk cosmological constant. In this case, we observe that the correlators for the operators $\mathbb{O}(x)$ reconstruct a *single* anti-de Sitter space (AdS) throat region. In this context, the appearance of a disorder average means that asymptotically we do not restrict ourselves to a fixed number of boundary components, but allow this to fluctuate, much as in [19]. Provided we only consider a number of boundary components much smaller than K , we thus make contact with the proposal of Marolf and

Maxfield [19], and in the limit where K is very large, this provides an adequate way to build up an ensemble average and its holographic dual.

However, with an actual construction in hand, we can also identify two general ways in which our UV completion breaks down. First of all, we clearly have a large number of sequestered stacks of branes, so if we proceed to higher energies, we should observe additional contributions beyond the approximate factorization appearing in Eq. (2). This is not altogether surprising, but already points to the fact that the UV completion does place an intrinsic limitation on the sorts of correlation functions we can consider.

Perhaps more surprisingly, there is *another* way in which the approximation can break down, and it is something that occurs even if we restrict to observables deep in the infrared. This is due to the fact that our UV completion implicitly makes reference to a fixed K , and we can actually distinguish between our binned approximation and a smooth distribution after sampling n_* times [see Eq. (12) later on for the precise definition]. In particular, if we consider any bulk object characterized by a density matrix with order n_* or more entries (as, for example, we would need to discuss in constructing the Page curve of a macroscopic black hole [33], see, e.g., [34]), then our putative ensemble average has been pushed beyond its regime of validity.

To make these considerations precise, we present a number of examples illustrating how to generate ensemble averaging for appropriate string-based constructions. Curiously enough, the case where we can maintain the most control is for $D = 4$ superconformal field theories (SCFTs) with an ensemble average over marginal coupling constants. We illustrate this both in terms of compactifications of 6D little string theories (LSTs), brane box configurations, and D3-brane probes of orbifold singularities. In this case, the extra dimensions transverse to the brane stacks provide enough flexibility to produce a nearly arbitrary probability distribution with support on the moduli space of marginal couplings. As another class of examples, we consider various $D = 2$ SCFTs obtained in a similar fashion. In this case, we find that for examples where we can reliably extract an AdS₃ dual description, it is often simplest to consider the fibration of a Calabi–Yau n -fold over a subvariety of its moduli space, resulting in a higher-dimensional (noncompact) Calabi–Yau geometry (see the Appendix). The limitation of this sort of construction is then that our ensemble averages are necessarily restricted to a particular subset of moduli. Similar considerations hold for $D = 1$ “SCFTs” of the sort that appear in the construction of 4D black holes obtained from type II strings on Calabi–Yau threefold backgrounds. Here, we again get an ensemble average, as associated with AdS₂ vacua. It is worth pointing out that at no point do we truly get a 2D gravitational theory; rather, we get an AdS₂ \times S^2 \times X_6 background.

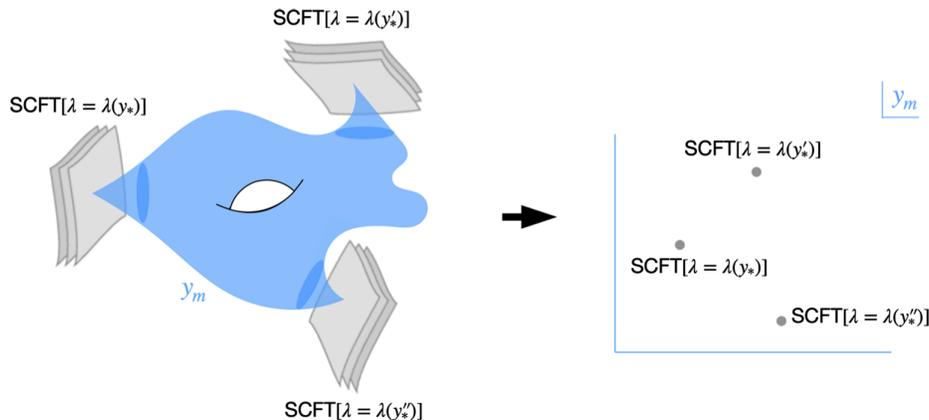


FIG. 1. Copies of the same QFT sector separated from one another in the extra dimensions. QFTs are engineered on the worldvolume of branes (colored gray). The extra-dimensional geometry is colored blue.

Our method meets with less success in the case of AdS_4 , AdS_6 , and AdS_7 vacua, as associated with 3D, 5D, and 6D SCFTs. In the case of 3D SCFTs, this may just be a failure of imagination/stamina. In the case of $D > 4$ SCFTs, we face the fact that there are no marginal deformations available that preserve supersymmetry [35,36].

The rest of this paper is organized as follows. In Sec. II, we present some more details on the general idea of our construction and its regime of validity. We present a holographic interpretation of this construction in Sec. III. Section IV presents explicit string-based examples. We present a brief discussion in Sec. V. Some additional technical details are deferred to the Appendix.

II. ENGINEERING AN ENSEMBLE

In this section, we present the main idea of generating a QFT ensemble average in string theory constructions,³ giving explicit examples later in Sec. IV. We are specifically interested in the case where the coupling constants of the QFT are truly constant in the sense that they have no spacetime dependence. This is the case that has been of primary interest in recent holographic investigations, and also turns out to be the most challenging one to engineer in the context of string constructions. See Ref. [14] for how to get an ensemble average in a gravitational system but with spacetime-dependent couplings.

Our goal is to mimic the effects of disorder averaging in a QFT via a string construction. More precisely, we suppose that our QFT depends on a set of couplings $\{\lambda\} \equiv \lambda$, and that we have a smooth classical distribution $p_{\text{smooth}}(\lambda)$. Given operators $O^{(1)}, \dots, O^{(m)}$ of the QFT, the disorder averaged correlator is obtained by evaluating the correlation function with respect to a fixed value of λ , and then performing a further averaging with respect to $p_{\text{smooth}}(\lambda)$,

³For the reader interested in learning more about string theory, the authors recommend [37,38]. It is a fascinating subject.

$$\overline{\langle O^{(1)} \dots O^{(m)} \rangle} \equiv \int d\lambda p_{\text{smooth}}(\lambda) \langle O^{(1)} \dots O^{(m)} \rangle. \quad (3)$$

Our plan is to engineer an ensemble of QFTs that are decoupled from gravity. Loosely speaking, this is arranged by taking a limit in which we try to retain some localized degrees of freedom while decoupling the gravitational degrees of freedom.⁴ This in turn means that some degrees of freedom are non-normalizable, and these descend to coupling constants $\{\lambda\}$ of the QFT sector. An important comment here is that even though these degrees of freedom are constants in $D < 10$ spacetime dimensions, they are still dynamical in the full $D = 10$ (resp., $D = 11$) spacetime associated with string theory (resp., M-theory).⁵

Now, precisely because we are dealing with a non-compact extra-dimensional geometry, we are free to consider multiple copies of the same QFT sector, just separated from one another in the extra dimensions. This stringy setup is roughly illustrated in Fig. 1. The parameters on each local sector are then specified by the internal profile of the non-normalizable modes in the vicinity of each local model. Introducing K in such local sectors, we can label the corresponding couplings as λ_k for $k = 1, \dots, K$, where $\lambda_k \equiv \{\lambda_k\}$ denotes the full set of couplings in each local model. Since the field content is assumed to be identical, we also have the same set of operators $\{O_k\}$ for each local model. The different sectors are decoupled only in an approximate sense because we can also consider degrees of freedom that stretch from one sector to the other (e.g., open strings in

⁴For example, open strings attached to a stack of D-branes.

⁵Indeed, in Ref. [14], it was noted that in a string compactification with multiple QFT sectors coupled to gravity, performing a partial trace over all but one such QFT sector results in a system characterized by a mixed state for position-dependent coupling constants. A subtlety with this approach is that the process of decoupling gravity also tends to force the previously obtained mixed state back into a pure state, and we will ultimately need to take a different tack to generate a position-independent ensemble average.

D-brane models), and this introduces a UV cutoff Λ_{UV} for this approximate factorization of the Hilbert space of states,

$$\mathbb{H} \approx \mathcal{H}_1 \otimes \cdots \otimes \mathcal{H}_K. \quad (4)$$

In the effective action for the full system, the action breaks up into a set of distinct contributions that then mix via higher-dimension operators with suppression scale Λ_{UV} ,

$$\mathbb{S} = S_1 + \cdots + S_K + S_{\text{mix}}. \quad (5)$$

Given a specific operator of a single sector, there is a natural subset of operators obtained by forming a sum over all such sectors,

$$\mathbb{O} \equiv O_1 + \cdots + O_K. \quad (6)$$

Normalized correlation functions for n such operators are then evaluated as

$$\langle \mathbb{O}^{(1)} \cdots \mathbb{O}^{(n)} \rangle_{\text{normalized}} \equiv \frac{1}{K} \langle \mathbb{O}^{(1)} \cdots \mathbb{O}^{(n)} \rangle_{\mathbb{H}}, \quad (7)$$

where on the right-hand side the correlation function is evaluated with respect to the ground state of the full system. The prefactor of $1/K$ can be understood as the statement that we just require a well-behaved large K limit. Another way to understand the same requirement is that we are just measuring all correlation functions in units of the traced identity operator,

$$\mathbb{I} = \mathbf{id}_1 + \cdots + \mathbf{id}_K. \quad (8)$$

Correlation functions for the \mathbb{O} operators close to leading order. To see why, we note that

$$\langle \mathbb{O}^{(1)} \cdots \mathbb{O}^{(n)} \rangle_{\text{normalized}} \approx \frac{1}{K} \sum_{1 \leq k \leq K} \langle O_k^{(1)} \cdots O_k^{(n)} \rangle_{\mathcal{H}_k}, \quad (9)$$

where to leading order the different factors \mathcal{H}_i are decoupled from one another. There can also be cross terms between the stacks, but these are all subleading contributions. Consider first connected correlators. In this case, sequestering the stacks suppresses such contributions. For disconnected correlators (for example, $\langle O_k O_k \rangle \langle O_l O_l \rangle$ with $k \neq l$), we further note that such contributions are kinematically suppressed relative to their connected counterparts (as is clear by passing to momentum space).⁶ In this sense, it is a consistent truncation.

Our main claim is that this can be used to build up a discretized approximation to a disorder averaging by a continuous distribution $p_{\text{smooth}}(\lambda)$. To see why, we observe that each expectation value over \mathcal{H}_k makes reference to the

couplings $\{\lambda_k\}$ of that sector. If we have a total of $K(\lambda)$ sectors with a particular set of couplings, then the probability associated with this choice is

$$p_{\text{disc}}(\lambda) = \frac{K(\lambda)}{K}. \quad (10)$$

As it stands, this is a discrete probability distribution, but it is important to note that in any actual string construction, there is an intrinsic ‘‘spread’’, as associated with the overall tension of a brane/region of localization in the internal geometry.⁷ For this reason, it is more appropriate to view our construction as building up a continuous distribution, but one that has been suitably ‘‘binned’’. More precisely, introducing an indicator function $I_{\lambda', \varepsilon_{\lambda'}}(\lambda)$, which has unit area and has support on a small region of size $\varepsilon_{\lambda'}$ centered around λ' , the probability of drawing λ builds up a histogram composed of small bins of size $\varepsilon_{\lambda'}$,

$$p_{\text{bin}}(\lambda) = \sum_{\lambda'} I_{\lambda', \varepsilon_{\lambda'}}(\lambda) \frac{K(\lambda')}{K}. \quad (11)$$

The values of the $\varepsilon_{\lambda'}$ depend on the UV cutoff Λ_{UV} , as well as specific details of the model and construction.

The approximation just developed enables us to closely mimic the disorder average we would get for a smooth probability distribution, but there are clearly some limitations. One of these is already apparent from the general setup: in the correlation functions for the \mathbb{O} operators, we observe that factorization will begin to break down when any correlators probe a short distance scale of size Λ_{UV} . Indeed, this is just the statement that we have a UV cutoff. Provided we work at long distances close to an infrared fixed point, we can hope to neglect such contributions.

There is also an entropic breakdown as associated with sampling the distribution $p_{\text{bin}}(\lambda)$ a large number of times. To see the issue, we consider two sorts of observers, a ‘‘daemonic observer’’ who has access to the full set of operators O_i , and another ‘‘ignorant observer’’ who is confined to making do with just the \mathbb{O} operators.⁸ For the daemonic observer, they can, after performing n measurements associated with correlations between the O_i and the \mathbb{O} operators, extract n independent and identically distributed draws from p_{bin} . The discrepancy between $p_{\text{smooth}}(\lambda)$ and $p_{\text{bin}}(\lambda)$ is captured by the relative entropy/Kullback–Leibler divergence, which reflects the amount of information we would gain upon switching from the binned distribution $p_{\text{bin}}(\lambda)$ to the ‘‘true’’ (although, from our perspective, unobtainable) distribution p_{smooth} ,

⁷For example, there is a minimal length scale $\ell_{\text{min}} \sim (\frac{1}{T_p})^{\frac{1}{p+1}}$ that a D p -brane of tension $T_p \sim (g_s \ell_{\text{st}}^{p+1})^{-1}$ can probe (see, e.g., [39]).

⁸This is related to the broader question of how well a low energy observer can ever hope to reconstruct a given UV completion, see, e.g., [14,40–44] as well as [45–52].

⁶We thank H. Maxfield for correspondence on this point.

$$D_{\text{KL}}(p_{\text{smooth}} \| p_{\text{bin}}) = \int d\lambda p_{\text{smooth}}(\lambda) \log \frac{p_{\text{smooth}}(\lambda)}{p_{\text{bin}}(\lambda)} \approx c\epsilon^2 \equiv \frac{1}{n_*}, \quad (12)$$

where c is an order one constant dependent on the details of the distribution, and ϵ is a representative value of the size of our histogram bins. We can view ϵ^{-1} as telling us the total number of distinct histogram bins. An observer who samples the distribution order $n \sim n_*$ times will be able to detect the difference between the string construction and the smooth “idealization”. Let us note that typically $\epsilon^{-1} \ll K$, since we necessarily need to group multiple values of the couplings in a single histogram bin to get an adequate approximation of $p_{\text{smooth}}(\lambda)$ in the first place.

Similar considerations hold for the “ignorant observer”, but their strategy for inferring the existence of a distribution over couplings is somewhat different. In this case, the point is to sample a number of different correlation functions over just the \mathbb{O} s, and in so doing infer the higher moments of the probability distribution $p_{\text{bin}}(\lambda)$. Here, distinguishability is governed by the number of moments of the distribution they are able to extract.

While the details of a particular model will dictate the specific operators to consider, we can always consider the partition function in Euclidean signature, as obtained by placing our QFT on some D -dimensional background. For a fixed value of the couplings λ , we have, in each of our local sector path integrals,

$$Z_\lambda = \int \mathcal{D}\phi e^{-S_\lambda[\phi]}, \quad (13)$$

in the standard notation. Observe that in our stringy construction, the quantity \bar{Z} is obtained from just performing a sum over each individual local sector,

$$\bar{Z} = \frac{1}{K} \sum \int \mathcal{D}\phi_k e^{-S_{\lambda_k}[\phi_k]}. \quad (14)$$

Once we engineer $p_{\text{bin}}(\lambda)$, we also implicitly get a probability distribution over just Z . To see why, we can similarly consider the quantities Z_k, Z_k^2, \dots , and compute the corresponding moments for the partition function.

In terms of our \mathbb{O} operator formalism, we construct the corresponding operators as follows. First of all, we build a copied unnormalized thermal density matrix,

$$\mathbb{P}^{(1)} \equiv \rho_1 + \dots + \rho_K, \quad (15)$$

with $\rho_k \equiv \exp(-\beta H_k) \otimes \rho_{\text{GND},k_\perp}$. Here, H_k is the Hamiltonian on the k th stack, and we explicitly tensor by $\rho_{\text{GND},k_\perp}$, the ground state associated with all the other Hilbert space factors, since as far as the evaluation of the

partition function is concerned, an observer there has no access to any other states.⁹ Upon evaluating $K^{-1}\text{Tr}\mathbb{P}^{(1)}$, we then get just \bar{Z} , as in Eq. (14). To get the higher order terms such as Z_k^m , we need a corresponding operator acting on a Hilbert space of states, so we consider m replicas of \mathcal{H}_k , namely, the m -fold tensor product $\mathcal{H}_k^{\otimes m}$. With respect to this, we introduce the replica density matrix for the k th stack,

$$\rho_k^{(m,\text{rep})} \equiv \rho_k^{(1)} \rho_k^{(2)} \dots \rho_k^{(m-1)} \rho_k^{(m)}, \quad (16)$$

and then, the corresponding copied thermal density matrix including all the replicas is

$$\mathbb{P}^{(m)} \equiv \rho_1^{(m,\text{rep})} + \dots + \rho_K^{(m,\text{rep})}. \quad (17)$$

Upon evaluating $K^{-1}\text{Tr}\mathbb{P}^{(m)}$, we then get just \bar{Z}^m . Continuing in this fashion, we can clearly treat the partition function itself as a random variable and evaluate its moments.

III. HOLOGRAPHIC INTERPRETATION

Let us now specialize to the case where each local model is a conformal field theory with a semiclassical AdS gravity dual with the same bulk value of the cosmological constant. In this case, we clearly obtain a large number of AdS throats in which bulk fields Φ_k^{bulk} have a boundary condition set by the particular values of the couplings on the boundary,

$$\Phi_k^{\text{bulk}} \rightarrow \lambda_k. \quad (18)$$

Our aim is to understand the sense in which the construction just presented can be interpreted in terms of a *single* AdS throat region. In the process, we make contact with the baby universe interpretation of ensemble averaging proposed in [19], but one that respects the swampland constraints of [32].

Our approach to this question is to focus on the operator subsector defined by the \mathbb{O} s of our QFT (now a CFT) sector. Along these lines, consider a local operator $\mathbb{O}(x)$,

$$\mathbb{O}(x) = O_1(x) + \dots + O_K(x), \quad (19)$$

where each summand has the same field content on its respective stack. By the standard rules of [53], we know that for each $O_k(x)$, we can (in principle) write down a

⁹It is of course tempting to also consider a different class of operators defined as $\mathcal{O} = \sum_k O_k \otimes \rho_{\text{GND},k_\perp}$, which share many of the same properties as the \mathbb{O} operators. One issue is that it is physically rather awkward to create an excitation in one stack, and simultaneously enforce a projection onto the ground state of the other stacks. As already mentioned, however, this is quite appropriate in constructing a corresponding partition function.

corresponding bulk field profile in an $\text{AdS}_{(k)}$ geometry. Said differently, each bulk quantity “casts a shadow” corresponding to a specific dual in the CFT. On the other hand, precisely because the O_k are built from the same fields, we see that even though we are dealing with a large number of AdS throats, the collective motion described by $\mathbb{O}(x)$ only accesses a single AdS. Said differently, because the connected correlators for the \mathbb{O} operators close, according to the standard AdS/CFT dictionary [53], they reconstruct a single AdS throat.

Another way to arrive at the same conclusion is to consider the geometric entanglement entropy for the ground state between a ball B and its complement B^c . Again, if we were initially dealing with a single local sector of our construction, we would simply introduce the pure state $\rho_k = |0\rangle_{kk}\langle 0|$. As is well known, this has a gravity dual description in terms of a “minimal area surface” homologous to B , and its “area” tracks with the entanglement entropy [54,55]. In the present setting, the ground state is given by the tensor product,

$$\rho_{\text{GND}} = \rho_1 \otimes \cdots \otimes \rho_K. \quad (20)$$

In the limit where the additional throats are sequestered from each other, the partial trace collapses to a single “diagonal” contribution. Again, the interpretation is that for this set of states, we are building up a single bulk “minimal area surface”.

Summarizing the discussion so far, we have seen that even though our stringy geometry is building up a large number of AdS throats, the closed subsector defined by the \mathbb{O} operators only reconstructs a single throat, and this is the one that produces an ensemble-averaged CFT. In other words, only a single bulk AdS geometry is needed to match to the subsector associated with the \mathbb{O} s. See Fig. 2 for an illustrative depiction.

In the UV complete realization in terms of multiple AdS throats, one might of course ask whether there could be wormhole configurations that join these individual throats, perhaps via some generalization of the construction presented in [56]. The general point is that our construction only mimics an ensemble average provided we have sufficient statistics. For this reason, it is natural to expect that the dominant contribution from saddle points of the Euclidean path integral instead comes from wormhole configurations that have fractionated; i.e., they join many boundaries.

Of course, the notion of ensemble averaging in AdS/CFT has recently been a topic of much interest, and so it is natural to ask how the present description fits with this. To this end, we next turn to a brief summary of the proposal of Marolf and Maxfield (MM) [19] in terms of baby universes, and then explain why it can be a valid approximation compatible with the considerations of McNamara and Vafa (McV) [32].

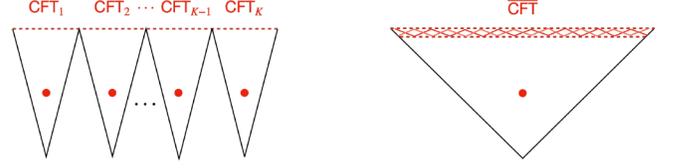


FIG. 2. On the left, the string theory construction builds a large number of D -dimensional CFTs, each of which has its own dual AdS throat. The closed subset of operators defined by the \mathbb{O} s, however, only reconstructs a single AdS throat, which is dual to an ensemble-averaged CFT, shown on the right.

A. Baby universe disintegration

To frame the discussion to follow, we first provide a brief summary of the MM proposal for how to capture the effects of ensemble averaging in AdS/CFT from the perspective of the gravity dual. Following [19], consider an AdS gravity theory with a set of fields denoted as Φ (including the metric), with boundary conditions labeled by J : $\Phi \sim J$. Note that the boundary can have more than one component generically. The Euclidean gravitational path integral defined by an asymptotic boundary with n connected components is then

$$\langle Z[J_1] \cdots Z[J_n] \rangle \equiv \int_{\Phi \sim J} \mathcal{D}\Phi e^{-S_{\text{grav}}[\Phi]}, \quad (21)$$

where J_1, \dots, J_n correspond to different components of the asymptotic boundary. As a point of notation, let us emphasize that here and throughout this subsection, J_a really are just the boundary couplings, but to emphasize that they are *not* associated with a particular set of stacks in our UV completion, we write J_a , with a having no relation to the indexing of all the λ_k .

The path integral defined in Eq. (21) cannot generically be factorized into those of disconnected boundaries [57,58], e.g., with $n = 2$,

$$\langle Z[J_1]Z[J_2] \rangle \neq \langle Z[J_1] \rangle \langle Z[J_2] \rangle. \quad (22)$$

This is because of the presence of Euclidean wormholes corresponding to the bulk manifold whose connected component includes the two boundaries (see Fig. 3). From the bulk perspective, this nonfactorization comes from the dynamical interactions of two boundaries connected by Euclidean wormholes. From the boundary-CFT point of view, nevertheless, this nonfactorization should be rather interpreted as the ensemble average over a classical probability distribution as follows.

By cutting open the above integral carefully so that the intermediate slice intersects no asymptotically AdS boundaries, one defines the baby universe Hilbert space \mathcal{H}_{BU} for the complete set of intermediate states separating “past” and “future”. The set of boundary conditions $\{J_1, \dots, J_n\}$ is then associated to a state,

$$\langle Z[J_1]Z[J_2] \rangle = \begin{array}{c} \langle Z[J_1] \rangle \\ \text{---} \\ \langle Z[J_2] \rangle \end{array} + \begin{array}{c} \text{---} \\ \text{---} \end{array}$$

FIG. 3. The presence of Euclidean wormholes results in non-factorization of the gravitational path integral.

$$|Z[J_1] \cdots Z[J_n]\rangle \in \mathcal{H}_{\text{BU}}. \quad (23)$$

One special state is the Hartle–Hawking state with no boundary. Its norm gives rise to the “cosmological partition function”,

$$\langle \text{HH} | \text{HH} \rangle = \int_{\text{no boundary}} \mathcal{D}\Phi e^{-S_{\text{grav}}[\Phi]}. \quad (24)$$

Next, introduce operators $\widehat{Z}[J]$ on \mathcal{H}_{BU} for any boundary condition J so that

$$\widehat{Z}[J]|Z[J_1] \cdots Z[J_n]\rangle = |Z[J]Z[J_1] \cdots Z[J_n]\rangle. \quad (25)$$

The eigenstates of $\widehat{Z}[J]$ then form a basis of \mathcal{H}_{BU} and are defined by

$$\widehat{Z}[J]|\alpha\rangle = Z_\alpha[J]|\alpha\rangle \quad \forall J, \quad (26)$$

with $\langle \alpha' | \alpha \rangle = \delta_{\alpha'\alpha}$. Note that any boundary condition can be derived from the corresponding operators acting on Hartle–Hawking state as

$$|Z[J_1] \cdots Z[J_n]\rangle = \widehat{Z}[J_1] \cdots \widehat{Z}[J_n] | \text{HH} \rangle. \quad (27)$$

The gravitational path integral in Eq. (21) can then be expressed as

$$\begin{aligned} \langle Z[J_1] \cdots Z[J_n] \rangle &= \langle \text{HH} | \widehat{Z}[J_1] \cdots \widehat{Z}[J_n] | \text{HH} \rangle \\ &= \sum_{\alpha_0, \dots, \alpha_n} \langle \text{HH} | \alpha_0 \rangle \langle \alpha_0 | Z[J_1] | \alpha_1 \rangle \\ &\quad \cdots \langle \alpha_{n-1} | Z[J_n] | \alpha_n \rangle \langle \alpha_n | \text{HH} \rangle \\ &= \langle \text{HH} | \text{HH} \rangle \sum_{\alpha} p_\alpha Z_\alpha[J_1] \cdots Z_\alpha[J_n], \end{aligned} \quad (28)$$

where p_α is the probability for each α state, computed by $p_\alpha = \frac{|\langle \text{HH} | \alpha \rangle|^2}{\langle \text{HH} | \text{HH} \rangle}$.

The appearance of the $|\alpha\rangle$ states is rather disturbing, especially in the context of string theory where we have no

evidence at all for such tunable parameters. Indeed, in string constructions coupled to gravity, all known examples of coupling constants descend from dynamical moduli. This point was significantly sharpened in [32] where they showed that in even more general terms, the cobordism hypothesis of the swampland program is enough to require the baby universe Hilbert space \mathcal{H}_{BU} to be one dimensional; namely, we can indeed speak of fixing the boundary values of the coupling constants, just as we would in “standard” AdS/CFT. From this perspective, ensemble averaging really has no general meaning in a UV complete framework such as string theory.

We now argue that in spite of appearances, our embedding in string theory provides a way to make contact with both proposals.

First of all, the very fact that we have an ensemble average means that any putative gravity dual will likely have to match on to the characterization provided by the MM picture. Indeed, assuming we can work out the “standard” AdS dual for a single throat in our construction, we already know the bulk field content, and thus, in principle, can discuss the computation of the Hartle–Hawking wave function for this theory.

Note that the probability distribution $p_{\text{bin}}(\lambda)$ we engineer in string theory does not directly correspond to the probability p_α for α states, although the two notions are clearly implicitly related in some way. From the perspective of the bulk gravitational system, it is tempting to say that the specific details of the path integral dictate a particular “preferred choice” for $p_{\text{smooth}}(\lambda)$ (see, e.g., [17–19]), although even this relies on having enough data in the form of a specific set of α states, and a specific choice of bulk gravitational action with which to construct the Hartle–Hawking state in the first place. Indeed, turning the discussion around, there seems to be little constraint on what sort of $p_{\text{bin}}(\lambda)$ s we can end up generating, and so we leave it as an interesting question to determine precisely how to fill in this entry of the AdS/CFT correspondence. From the perspective of the present construction, we take this to mean that there are ambiguities in specifying the Euclidean path integral for quantum gravity, and resolving these ambiguities in different ways can result in different choices for the ensemble average in the boundary theory.

Once we accept the existence of the $|Z[J_1] \cdots Z[J_n]\rangle$ states, the appearance of the baby universe states $|\alpha\rangle$ would appear to follow. At this point, however, we recall that our ensemble average picture can break down, both at high energies, but also entropically whenever n gets sufficiently large. Probing either regime of validity shatters the illusion, and we can no longer work in terms of a single AdS throat region. For example, when the number of boundary components n becomes sufficiently large [i.e., of order n_* of Eq. (10)], then we have already seen there can be an entropic breakdown. Observe that when n is really of order K , we can even resolve the individual AdS throats, and so

in this limit, we just recover the standard AdS/CFT dictionary, with a single one-dimensional baby universe Hilbert space for each throat, much as in [31].

Of course, one of the main reasons to seek out an AdS/CFT interpretation of our ensemble average system is its potential use in studying aspects of quantum gravity in anti-de Sitter space. Along these lines, it is also natural to ask about whether the approximation is reliable enough to provide access to the microscopic details of a black hole. To arrange this, we follow the procedure in [59,60] and consider the CFT on the background $S^1 \times S^{D-1}$. The size of the thermal circle sets a corresponding temperature (and thus size) for an AdS–Schwarzschild geometry, but one can of course consider more elaborate configurations with various chemical potentials switched on. Now, suppose we are interested in probing the n_{micro} microstates of this black hole, perhaps as captured by Hawking radiation quanta. To construct a quantity such as the Page curve, a boundary observer will need to sample order n_{micro}^2 times [33]. However, if n_{micro} exceeds $\sqrt{n_*} \sim \epsilon^{-1}$ of Eq. (12), then the ensemble interpretation becomes problematic. In light of this, it is unclear (at least to us) how we can use this setup to learn about the Page curve of a macroscopic black hole.

It is what it is.

IV. EXAMPLES

Having demonstrated the main ideas behind ensemble averaging in the context of string theory, we now turn to explicit examples. Our aim here is not to be exhaustive, but rather to showcase a few different methods, including their advantages and disadvantages.

For starters, we focus on supersymmetric quantum field theories since these are the ones over which we have maximal control. We also primarily focus on superconformal field theories (SCFTs), since these have a chance (in a suitable large- N limit) of having a holographic dual. Within this setting, we identify two general methods for building an ensemble, one that we refer to as various brane box models (and their dual incarnations), and another based on building up new compactification geometries by fibering an existing geometry over a noncompact subvariety of its moduli space.

One way to build a large class of QFTs is with D-branes that end on NS5-branes, so-called “brane box” models. Arranging these NS5-branes in various repeating patterns such as d -dimensional generalizations of a cube, we can produce a rich set of possible SCFTs. For our purposes, the important point is that for $D \leq 4$ systems the relative positions between the NS5-branes descend to non-normalizable parameters of the resulting quantum field theory. We get an ensemble average by repeating this construction in some of the directions transverse to the original d dimensions used to make a single instance of the SCFT.

For example, if we attempt to engineer a D -dimensional QFT using spacetime-filling branes in the geometry $\mathbb{R}^{D-1,1} \times \mathbb{R}^d \times Y_{10-D-d}$, then the entire configuration sits at a single point of the transverse Y geometry, which in many cases of interest is just \mathbb{R}^{10-D-d} . Moving to a different point of Y , we can then arrange for a different choice of couplings. By a chain of T dualities, these constructions can also be related to the worldvolume theory of branes probing singularities, and these in turn can often be generated by appropriate compactifications of 6D SCFTs and little string theories (LSTs). An important feature of this method of construction is that precisely because we can tune the moduli in the transverse Y directions, we can use this to engineer an essentially arbitrary probability distribution over the couplings of the model. The disadvantage of this approach is that in some cases it is difficult to guarantee that we generate a theory with a candidate AdS dual.

Another way to generate examples consists of taking a system of branes wrapped on subspaces of a Calabi–Yau d -fold X_{2d} (a $2d$ -real-dimensional space). More precisely, we assume that the geometry takes the form $\mathbb{R}^{D-1,1} \times Y_{2m} \times X_{2d}$, where Y is an m -complex-dimensional geometry transverse to the branes. In particular, our branes sit at a particular point of Y , and the parameters of the QFT descend from the geometric moduli of X . To get an ensemble average, we consider \mathcal{M}_X , the moduli space of X . Observe that we can consider the total space \mathcal{X} as defined by $X \rightarrow \mathcal{X} \rightarrow B$, where $B \subset \mathcal{M}_X$. Cutting out an m -complex-dimensional subspace B of \mathcal{M}_X that has no singular fibers then generates a noncompact Calabi–Yau of real dimension $2m + 2d$.¹⁰ We can also entertain more general fibrations, possibly with singular fibers, and we present some explicit examples of precisely this sort in the Appendix. An important advantage of this approach is that such brane constructions in Calabi–Yau compactifications often come with readily defined AdS duals. A drawback of this approach is that it does not, in general, allow us to engineer an arbitrary ensemble average. This is simply because the best we can do is to sweep out a probability distribution with support on an m -complex-dimensional subspace of the full moduli space.

In the remainder of this Section, we turn to some particular examples, illustrating the pitfalls, the possibilities, the perils, and the promise of generating ensemble averages.

A. $D=4$

We now engineer some ensembles of $D=4$ SCFTs. We begin by constructing an ensemble for $\mathcal{N}=4$ Super Yang–Mills theory, and then turn to examples with lower supersymmetry.

¹⁰We thank T. Pantev for discussions on this point.

1. Warmup: Approximating $\mathcal{N}=4$ SYM

Perhaps the simplest case to consider is that of type IIB string theory on the background $\mathbb{R}^{3,1} \times \mathbb{C}^3$ with a stack of N_c spacetime-filling D3-branes sitting at a point of \mathbb{C}^3 . At low energies, the open string degrees of freedom realize $\mathcal{N}=4$ Super Yang–Mills (SYM) theory with gauge group $U(N_c)$. The value of the complexified gauge coupling τ is controlled by the background value of the type IIB axiodilaton,

$$\tau = C_0 + i \exp(-\phi). \quad (29)$$

In type IIB supergravity, this is characterized by the extremal brane solution with a constant axiodilaton profile, localized source for the self-dual five-form flux, and metric (see, e.g., [61]),

$$\begin{aligned} ds^2 &= H^{-1/2} ds_{\mathbb{R}^{3,1}}^2 + H^{1/2} ds_{\mathbb{C}^3}^2, \quad \text{with} \\ H &= 1 + \frac{4\pi g_s N_c \alpha'^2}{r^4}, \end{aligned} \quad (30)$$

where r is the distance from the D3-brane stack. As is well known, in the near-horizon limit, this produces an $\text{AdS}_5 \times S^5$ geometry with N_c units of self-dual five-form flux threading the two factors [59]. The AdS radius and S^5 radius L are correlated and related to the open string parameters as $L^4 = 2g_{\text{YM}}^2 N_c \alpha'^4$. We are, of course, free to consider moving these D3-brane stacks to separate points in \mathbb{C}^3 . In all these local sectors, the value of the axiodilaton is always the same, and we generate a rather trivial probability distribution of values for the axiodilaton. The supergravity approximation is the same, the only change being the harmonic function H , which is now given by

$$H = 1 + \sum_k \frac{4\pi g_s N_c \alpha'^2}{|\vec{r} - \vec{r}_k|^4}, \quad (31)$$

with \vec{r}_k the position of the k th stack.

To get a more general class of distributions, we now introduce an additional source as specified by a stack of D7-branes that sits at a point of the middle factor in $\mathbb{R}^{3,1} \times \mathbb{C}_\perp \times \mathbb{C}^2$ and fills the remaining eight directions. The brane configuration is as follows:

	$\mathbb{R}^{3,1}$				\mathbb{C}_\perp		\mathbb{C}^2			
	0	1	2	3	4	5	6	7	8	9
D3	×	×	×	×						
D7	×	×	×	×			×	×	×	×

Doing so produces a position-dependent profile for the axiodilaton, but also breaks half the supersymmetry in the

system. Indeed, if we now have a stack of D3-branes located at distinct points of the geometry, then each can experience a different value of the axiodilaton τ_k , and the low-energy effective action on each stack is of the form

$$S_k = S_{\mathcal{N}=4}(\tau_k) + S_{\mathcal{N}=2}, \quad (32)$$

where the contribution from explicit $\mathcal{N}=2$ breaking terms is captured by a collection of higher-dimension operators. The precise form of these contributions can be worked out by noting that this D3/D7 system is just engineering a 4D $\mathcal{N}=2$ gauge theory. There are two distance scales that control the strength of these higher-dimension operators. One is the relative separation between the D3-branes,

$$\Lambda_{k_1, k_2} \equiv \frac{\text{dist}(D3_{k_1}, D3_{k_2})}{\alpha'}, \quad (33)$$

as measured in the full \mathbb{C}^3 factor transverse to all the D3-branes, and the other is the relative separation between the D3-branes and the D7-branes,

$$\Lambda_{k, D7} \equiv \frac{\text{dist}(D3_k, D7)}{\alpha'}, \quad (34)$$

as measured in the \mathbb{C}_\perp factor transverse to the D7-branes. Provided we only ask questions at low energies compared with these cutoffs, we get an adequate approximation to ensemble averaging in $\mathcal{N}=4$ SYM, but one that has an $\mathcal{N}=2$ UV completion.

Let us now turn to the class of ensembles we can actually engineer in this setting. First of all, the whole point of introducing a stack of D7-branes is that we can thus generate a position-dependent axiodilaton. More broadly, this and more general choices of nonperturbative bound states of 7-branes can be understood in F theory [62–64] by considering a noncompact elliptically fibered K3 surface with the minimal Weierstrass model,

$$y^2 = x^3 + f(z)x + g(z), \quad (35)$$

where in the present setting $f(z)$ and $g(z)$ are treated as polynomials in the holomorphic coordinate z of the \mathbb{C}_\perp factor. The possible values of the axiodilaton are implicitly encoded in the $SL(2, \mathbb{Z})$ -invariant j function,

$$j = 1728 \frac{4f^3}{4f^3 + 27g^2}, \quad (36)$$

which has the weak-coupling expansion in $q = \exp(2\pi i \tau)$ given by

$$j = q^{-1} + 744 + 196884q + \dots \quad (37)$$

For example, $j = \infty$ corresponds to weak coupling at $\tau = i\infty$, and $j = 1728$ corresponds to $\tau = i$, while $j = 0$

corresponds to $\tau = \exp(2\pi i/6)$. In the case of a single stack of M D7-branes sitting at $z = 0$, we just have

$$j_{D7} = z^{-M}. \quad (38)$$

To get a particular value of the axiodilaton (or, more precisely, its j invariant), we simply consider D3-branes at the desired value of z . Note that this intrinsically comes with some limitations, because to populate the distribution near the $\tau = i\infty$ region of moduli space, we necessarily must move close to the stack of D7-branes, which in turn lowers the UV scale $\Lambda_{k,D7}$ in our effective field theory. The other issue we face is how to sequester the D3-brane stacks from one another. This is less problematic, because even if they sit at the same point of \mathbb{C}_\perp , we are free to move them away from each other in the \mathbb{C}^2 factor. We thus conclude that the UV cutoff is set by $\Lambda_{k,D7}$, and this in turn depends on what sort of distribution we wish to engineer.

Consider next the supergravity background generated by our D3/D7 system. Since we are dealing with D3-brane probes of an F-theory geometry, the main change is that the metric on the \mathbb{C}_\perp factor is controlled by that of the noncompact elliptically fibered K3 space $K3 \rightarrow \mathbb{C}_\perp$. In particular, we observe that in the near-horizon limit for each local stack, we indeed get a collection of individual AdS throats, but of different sizes as dictated by the local profile of the τ_k . Note also that the subleading $\mathcal{N} = 2$ breaking terms amount to a deviation away from a pure $\text{AdS}_5 \times S^5$ geometry.

Precisely because the value of the bulk cosmological constant is different for each local sector, there is no sense in which we can give a holographic interpretation in terms of a single AdS geometry. This is in accord with the fact that consistency of the MM proposal requires the baby universe Hilbert space interpretation to be trivial (i.e., one dimensional) in this special case [19,31].

2. Quiver gauge theory ensemble

We now proceed to engineer an ensemble average of quiver gauge theories that enjoys a holographic dual. The model we consider consists of N_c D3-branes probing the orbifold singularity $\mathbb{C}^2/\mathbb{Z}_M$ with group action $(u, v) \mapsto (\xi u, \xi^{-1} v)$ on the holomorphic coordinates, where ξ is a primitive M th root of unity. As is well known from [65,66], the worldvolume theory for this model results in a 4D $\mathcal{N} = 2$ SCFT described by a quiver gauge theory with gauge groups arranged in a circular ring, joined by hypermultiplets in bifundamental representations (see Fig. 4). The model also comes with a collection of marginal parameters, captured by the holomorphic couplings $(\tau^{(1)}, \dots, \tau^{(M)}) \equiv \vec{\tau}$. These are encoded in the choice of closed string moduli.

To better understand the sense in which these closed string moduli are tunable, it is helpful to consider some dual realizations of the same low-energy effective field theory. One way to proceed is to observe that there is a

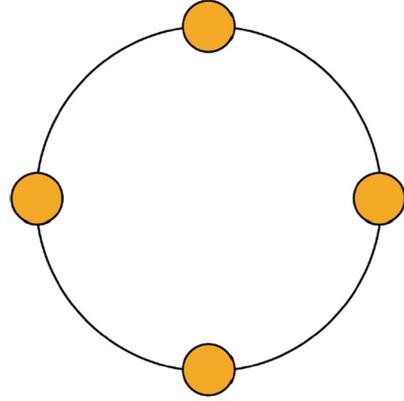


FIG. 4. Quiver diagram of the 4D $\mathcal{N} = 2$ SCFT obtained from N_c D3-branes probing $\mathbb{C}^2/\mathbb{Z}_M$, for $M = 4$. Each node represents an $SU(N_c)$ gauge group, and links between them denote bifundamental hypermultiplets.

6D little string theory obtained via F theory from a configuration of collapsing -2 curves arranged in a circular ring (see, e.g., [67]). Wrapping N_c D7-branes on each -2 curve (a Kodaira I_{N_c} fiber) then results in a 6D quiver gauge theory. Compactifying on a further T^2 then produces the desired 4D $\mathcal{N} = 2$ SCFT. In this realization, the gauge couplings descend from the complexified Kähler volume of $T^2 \times \Sigma_m$, for $m = 1, \dots, M$ labeling the different -2 curves. In particular, each of these is a tunable complexified gauge coupling in the 4D field theory.

The tensor branch of the 6D little string theory can also be realized from a configuration of M NS5-branes arranged in a circular ring with N_c D6-branes suspended in between each neighboring pair. In this picture, the relative distance between each NS5-brane sets the value of 6D gauge coupling, and further compactification on a T^2 again results in the same 4D gauge theory.

To be more explicit, let us consider T dualizing the T^2 wrapped by D6-branes; the 4D gauge theory is then realized alternatively by the following D4/NS5 system:

	0	1	2	3	4	5	6	7	8	9
N_c D4s	×	×	×	×				×		
M NS5s	×	×	×	×	×	×				

(see Fig. 5 for an illustration of this brane configuration in the 4, 5, and 6 directions). Observe also that a T duality on direction 6 directly connects this construction to that of N_c D3-branes probing the same $\mathbb{C}^2/\mathbb{Z}_M$ singularity. In all these cases, then, we have a geometric characterization of the resulting moduli.

Let us now turn to the construction of an ensemble average. Returning to the picture of D3-branes probing an orbifold, observe that we can generate a large ensemble of local singularities via the hypersurface in \mathbb{C}^3 given by

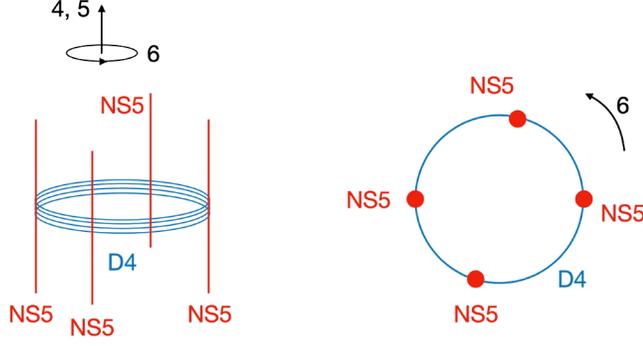


FIG. 5. D4/NS5 system leading to the 4D quiver gauge theory. Each D4-brane segment corresponds to a gauge group, whose coupling is set by the relative distance in the 6 direction between the two boundary NS5-branes as $\frac{1}{g_i} = \frac{x_{i+1}^6 - x_i^6}{g_s \sqrt{\alpha'}}$. The effective 4D SCFT on the D4-brane worldvolume is the same quiver gauge theory shown in Fig. 4.

$$y^2 = x^2 + \prod_k (z - z_k)^{M_k}. \quad (39)$$

Near $x = y = (z - z_k) = 0$, this generates a $\mathbb{C}^2/\mathbb{Z}_{M_k}$ singularity, so if we take $M_k = M$ for all k , and place N_c D3-branes near each singularity, we produce an ensemble of such systems. Observe that because the closed string moduli of each singularity are decoupled, we can arrange a nearly arbitrary ensemble. The same is clear from working with any of the other local pictures. Note also that there is a natural UV completion available in terms of a stack of KN_c D3-branes probing a $\mathbb{C}^2/\mathbb{Z}_{KM}$ singularity.

In the large- N_c limit, we can also arrange for a holographic dual description. To get the same size AdS, we just require that the value of the string theory dilaton (in the D3-brane probe picture) is the same for each stack. In the quiver gauge theory, this is controlled by the particular combination of parameters,

$$\tau_{\text{IB}} = \tau^{(1)} + \dots + \tau^{(M)}, \quad (40)$$

at each local model, which we take to be equal for each value of $k = 1, \dots, K$. In this description, the remaining $M - 1$ degrees of freedom are tunable marginal couplings over which our ensemble runs.

Clearly, one can consider far more elaborate examples, but the essential point is that at least for $D = 4$ SCFTs, we can engineer a wide variety of ensemble averages.

B. $D = 3$

Let us briefly comment on the construction of $D = 3$ SCFTs and issues with engineering ensemble averages in this setting. First of all, we can just take a $D = 4$ SCFT and compactify it on a circle. In many cases, this can also result in a 3D SCFT. From the perspective of the examples just considered, we could alternatively start from a 6D LST or

6D SCFT and compactify it on a three-manifold with negative sectional curvature. The primary challenge from this perspective is how to build a class of compactification geometries where we can vary the marginal couplings.¹¹

To this general point, one might attempt to realize examples using M theory on an eight-manifold with a large number of singularities. The theory of N_c M2-branes probing any one singularity will produce an $\text{AdS}_4 \times X_7$ background, and in principle, the geometric moduli of X_7 can be traced to corresponding deformations in the 3D SCFT. One can consider, for instance, an ensemble average over the Chern–Simons levels of the ABJ(M) model [71,72] by introducing K copies of N_c M2-branes probing $\mathbb{C}^4/\mathbb{Z}_{M_k}$ singularities. The difficulty we now face is that the radius of the dual AdS_4 scales as $L^2 \propto N_c/M_k$ with other parameters fixed. On the other hand, if we attempt to average over the Chern–Simons couplings, we must then simultaneously adjust the value of N_c . It would be interesting to build an explicit ensemble average in this case, but this would seem to require first obtaining a better understanding of how the geometry of a given background descends to marginal parameters of the field theory.¹²

C. $D = 2$

Let us now turn to the construction of $D = 2$ SCFTs, and, when appropriate, their AdS_3 duals. We begin with an example of a brane box model where we have a great deal of freedom in generating an ensemble average. The downside, however, is that the existence of an AdS_3 dual is not always apparent. The other general method we consider yields an $\text{AdS}_3 \times S^3 \times X_4$ dual, but at the expense of a less-flexible class of possible probability distributions.

1. Brane box examples

Similarly to 4D SCFTs, in addition to branes probing geometries, one can also engineer 2D theories using intersecting branes. Let us now consider 2D $\mathcal{N} = (0, 2)$ theories engineered by the brane box model [77,78].

¹¹A general class of 3D SCFTs can be obtained by compactification of 6D $\mathcal{N} = (2, 0)$ SCFTs on a three-manifold, which are known as the $T[M_3]$ theories (see, e.g., [68,69]). When $M_3 = \Sigma \times S^1$, where Σ is a Riemann surface, $T[M_3]$ admits marginal deformations. However, the holographic dual of $T[M_3]$ in this case is not AdS_4 gravity, but rather a “gravitational domain wall” separating two AdS regions (see, e.g., [70]). Even if one focuses on the CFT side, the top-down approach to building an ensemble average of these 3D SCFTs is not clear to us. We thank S. Gukov for discussions on this point.

¹²See references [73–76] for some examples of how these marginal couplings are realized in various AdS_4 gravity duals. However, the geometric or brane realization counterparts of these deformations are not yet known, so it is unclear to us how to build an ensemble with random couplings from a top-down approach. We thank N. Bobev and A. Tomasiello for helpful correspondence on this point.

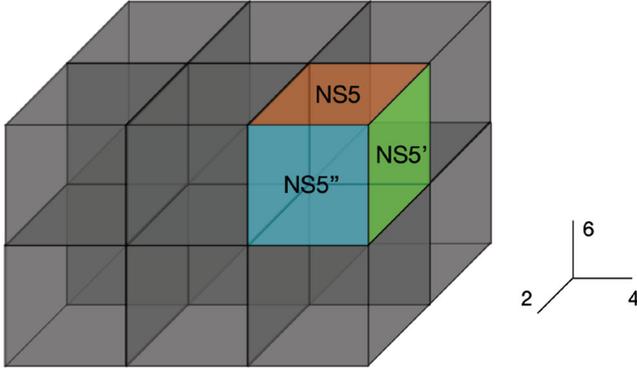


FIG. 6. Brane box model on T^3 with $m = 2, m' = 3$ and $m'' = 2$. Every box bounded by NS5-, NS5'-, and NS5''-branes is filled by M D4-branes, corresponding to one $U(M)$ gauge group in the 2D effective field theory.

The brane configuration consists of M D4-branes, m NS5-branes, m' NS5'-branes, and m'' NS5''-branes, as follows:

	0	1	2	3	4	5	6	7	8	9
D4	×	×	×		×		×			
NS5	×	×	×	×	×	×				
NS5'	×	×	×	×				×	×	
NS5''	×	×			×	×	×	×		

where all branes sit on the same position in the 8,9-plane. The D4-branes are finite in the 2, 4, and 6 directions as a T^3 and bounded by three types of NS5-branes. For each pair of parallel NS5-branes, the brane configuration in directions 2, 4, and 6 is a 3D box filled by D4-branes. With m NS5-branes, m' NS5'-branes, and m'' NS5''-branes, we have a 3D grid containing $m \cdot m' \cdot m''$ boxes, as illustrated in Fig. 6.

The effective field theory on D4-branes in the non-compact 0 and 1 directions is a 2D $\mathcal{N} = (0, 2)$ gauge theory with $U(M)^{m \cdot m' \cdot m''}$ gauge symmetry. Each $U(M)$ gauge group corresponds to one brane box on the T^3 comprising the 2, 4, and 6 directions. Performing T duality on this T^3 , the brane configuration becomes D1-branes probing a $\mathbb{C}^4 / (\mathbb{Z}_m \times \mathbb{Z}_{m'} \times \mathbb{Z}_{m''})$ singularity, with orbifold group action,

$$\begin{aligned}
 (z_1, z_2, z_3, z_4) &\rightarrow (z_1, z_2, e^{\frac{2\pi i}{m}} z_3, e^{\frac{2\pi i}{m}} z_4), \\
 (z_1, z_2, z_3, z_4) &\rightarrow (z_1, e^{\frac{2\pi i}{m'}} z_2, z_3, e^{\frac{2\pi i}{m'}} z_4), \\
 (z_1, z_2, z_3, z_4) &\rightarrow (e^{\frac{2\pi i}{m''}} z_1, z_2, z_3, e^{\frac{2\pi i}{m''}} z_4),
 \end{aligned} \tag{41}$$

which indeed preserves $\mathcal{N} = (0, 2)$ supersymmetry on the 2D worldvolume of the D1-branes.

Due to the nonvanishing elliptic genera of these 2D $\mathcal{N} = (0, 2)$ theories computed in [79], it is natural to conjecture that these theories flow to SCFTs in the IR.¹³ Here, we

assume the existence of the SCFT in the IR and build the ensemble average.¹⁴

One of the marginal couplings in 2D is the Fayet-Iliopoulos (FI) term for the $U(1)$ factor of each $U(N)$ gauge group. The corresponding FI parameters are encoded in the positions of the NS5-branes in several directions. Namely, for a given $U(M)$ gauge group associated with the (m, m', m'') -th brane box, its FI parameter is given in terms of the separations of the NS5-, NS5'-, and NS5''-branes in directions 7, 5, and 3, respectively (see, e.g., [77]),

$$r_{m,m',m''} = \frac{(x_{m+1}^7 - x_m^7) + (x_{m'+1}^5 - x_{m'}^5) + (x_{m''+1}^3 - x_{m''}^3)}{\sqrt{\alpha}}. \tag{42}$$

We can build an ensemble of brane box models by repeating this construction at different points in the 8,9-plane. Indeed, tuning the relative positions of the NS5-, NS5'-, and NS5''-branes in each instance of the construction, we can engineer an ensemble of 2D $\mathcal{N} = (0, 2)$ SCFTs in which the FI parameters are drawn from a probability distribution.

A drawback of this approach is that although this provides us with a way to generate ensemble averaging in 2D SCFTs, the existence of a putative AdS dual is somewhat unclear. We now turn to examples for which we understand the holographic dual more clearly.

2. Holographic example

To generate some examples with a holographic dual, we consider type IIB string theory on the background $\mathbb{R}_{\text{time}} \times \mathbb{C}^2 \times S^1 \times K3$, with N_5 D5-branes wrapped on $S^1 \times K3$ and N_1 D1-branes wrapped on the S^1 factor. We keep the D1/D5 system coincident at the same point of \mathbb{C}^2 . This engineers a 2D SCFT with $\mathcal{N} = (4, 4)$ supersymmetry on the spacetime $\mathbb{R}_{\text{time}} \times S^1$. The brane configuration is as follows:

	\mathbb{R}_{time}	S^1	\mathbb{C}^2					$K3$			
	0	1	2	3	4	5	6	7	8	9	
D1	×	×									
D5	×	×					×	×	×	×	

The corresponding supergravity solution describes an extremal black string, and in the near-horizon limit, it is given by the geometry $\text{AdS}_3 \times S^3 \times K3$, which figures

¹³To our knowledge, the explicit SCFTs at the IR fixed point and the AdS gravity duals for these $\mathcal{N} = (0, 2)$ theories have not yet been constructed, which by itself is a problem deserving further investigation.

¹⁴We thank N. Benjamin, S. Franco, and S. Gukov for helpful discussions on this point.

prominently in the $\text{AdS}_3/\text{CFT}_2$ correspondence [59] (see also [80]). The size of the AdS radius in Planck units is proportional to $N_1 N_5$, which is in turn proportional to the central charge of the 2D SCFT $c = 6N_1 N_5$. For our purposes, the important point is that this 2D SCFT comes with a set of marginal couplings controlled by the metric moduli of the K3 surface, and geometrically, these can be understood (via duality) as Narain moduli of the sort that appear in [8–10].

To generate an ensemble, we begin by observing that $\mathbb{C}^2 = \mathbb{C}_{\text{seq}} \times \mathbb{C}_{\text{base}}$ can be written as a product of two complex lines. We use the \mathbb{C}_{seq} factor to sequester the different QFTs from one another and use the \mathbb{C}_{base} factor to build a K3-fibered Calabi–Yau threefold $X_3 \rightarrow \mathbb{C}_{\text{base}}$ with fiber given by our K3 surface. Populating a different number of local sectors at each value of the moduli, we can again construct a probability distribution in a one-parameter subspace of the K3 moduli space. Let us also note that, much as in our discussion of the D3/D7 system, approaching a singular point in the moduli space of K3s limits the regime of validity of our approximation.

D. $D = 1$

Finally, let us turn to the case of one-dimensional quantum mechanical systems and their approximate AdS_2 duals. Along these lines, we observe that type IIA strings on $\mathbb{R}_{\text{time}} \times \mathbb{R}^3 \times X$, with X a Calabi–Yau threefold, produces a 4D $\mathcal{N} = 2$ supergravity background. Wrapping D0-, D2-, D4-, and D6-branes on holomorphic cycles of X can, for suitable charges, result in a 4D black hole solution with near-horizon limit $\text{AdS}_2 \times S^2$, with the radius set by the background value of the charges (see, e.g., [81]). The brane configuration is shown below,

	\mathbb{R}_{time}	\mathbb{R}^3			X					
	0	1	2	3	4	5	6	7	8	9
D0	×									
D2	×				[2 of 6 directions]		
D4	×				[4 of 6 directions]		
D6	×				×	×	×	×	×	×

Here, the bracketed directions denote the subset of directions containing the cycles wrapped by D2- and D4-branes, but do not specify the exact directions those cycles fill out. In the 4D theory, the D0- and D2-brane charges give rise to an electric charge vector \vec{Q} , while the D4- and D6-brane charges give rise to a magnetic charge vector \vec{P} . In general, it is a challenging problem to construct multicenter black hole solutions, but in the special case where all the centers have the *same* charge $(\vec{Q}_k, \vec{P}_k) = (\vec{Q}, \vec{P})$, the configurations are mutually BPS, and in particular can be separated from one another in

arbitrary directions of the \mathbb{R}^3 factor. For our purposes, we find it convenient to write $\mathbb{R}^3 = \mathbb{R}_{\text{seq}} \times \mathbb{C}_{\text{base}}$.

The worldvolume theory of this configuration is also challenging to describe, but at least at weak string coupling (away from the supergravity limit), there is a corresponding quiver quantum mechanics description available (see, e.g., [82]). In the type IIA description, the superpotential couplings of this model are controlled by the complex structure moduli of X , and in the mirror type IIB description, they are captured by the Kähler moduli. An important open problem is to explain the sense in which this quiver quantum mechanics actually “flows” to a 1D SCFT.¹⁵

To generate an ensemble average in the 1D system, we return to the decomposition $\mathbb{R}^3 = \mathbb{R}_{\text{seq}} \times \mathbb{C}_{\text{base}}$. Since we can retain a supersymmetric multicenter configuration no matter how we separate the branes in the \mathbb{R}^3 direction, we use the \mathbb{R}_{seq} factor to sequester the black holes, and treat \mathbb{C}_{base} as the base of a noncompact Calabi–Yau fourfold $X_4 \rightarrow \mathbb{C}_{\text{base}}$, with fiber given by X at a specified value of the moduli. Much as in our 2D holographic example, this provides us with an ensemble average over a one-parameter subspace of the moduli space.¹⁶

An important subtlety with this construction is that as we vary the Calabi–Yau moduli, we may be forced to deal with a further fragmentation of a single center solution into a multicenter solution (the supergravity analog of wall crossing), as studied for example in [84,85]. For our present purposes, this introduces some additional complications into the CFT_1 description and any putative AdS_2 dual. One can view this as either imposing additional restrictions on what sorts of distributions we can engineer, or alternatively, as an opportunity to come to grips with multiple large AdS_2 throats right from the start.¹⁷ As an additional comment, let us note that if we move away from the truncated subsector of operators, we get a large number of $\text{AdS}_2 \times S^2$ throats, much as in the stringy baby universe construction of Ref. [86].

E. $D > 4$

For SCFTs in $D = 5$ or 6 ,¹⁸ it was demonstrated in [35,36] that there is no supersymmetry-preserving marginal deformation. It would be interesting to address this issue in the context of the MM picture, since it implicitly imposes

¹⁵Of course, in one dimension, we face the issue that the CFT condition of a traceless stress tensor would appear to trivialize the theory altogether. So, we must already be prepared to work in terms of an approximate notion of conformal invariance where we introduce an explicit IR cutoff from the start, much as in [7].

¹⁶The disorder averaged quiver quantum mechanics has been studied in [83], where the probability distribution is chosen to be Gaussian. Our construction in this section can be regarded as why/how the ensemble arises in the quiver quantum mechanics.

¹⁷We thank N. Bobev for comments.

¹⁸Note that $D = 6$ is the highest spacetime dimension where SCFTs exist [87].

restrictions on the sort of boundary conditions for the bulk gravitational path integral.

V. DISCUSSION

In this paper, we have presented a string-based method for engineering ensemble averaged QFTs. The main idea in our construction is to first build multiple copies of the same QFT, but in which the non-normalizable modes of the background geometry vary in the transverse directions. Truncating to the subsector of operators that are distributed over all the local sectors, we have shown that this can be used to build up an approximation to ensemble averaging with a UV completion. The construction can also be used to realize models with an approximate AdS dual. There is an intrinsic regime of validity that can be detected just by sampling macroscopic objects a large number of times. This is important because it cuts to the core issue of whether holography and effective field theory can actually be combined without reference to any putative UV completion. Our (admittedly contrived) construction serves as a counterexample to the claim that a UV completion is “not necessary”.

In light of these considerations, it would be interesting to revisit the computation of a Page curve for a macroscopic black hole (see, e.g., [24,88,89]). For some recent critiques of these calculations, see, e.g., [27,29,30].

It is also natural to ask whether we could directly couple our ensemble of QFTs back to gravity. At least in the context of string theory, this is deeply problematic because in all known constructions, there are sharp upper bounds on the total number of extra sectors we can introduce. Moreover, once we recouple to gravity, all the moduli will once again become dynamical and must be stabilized in some way.

Aside from these intrinsic limitations of our method of construction, it is also interesting to ask about the extent to which we can generate more general classes of probability distributions. For example, in our examples based on fibering a Calabi–Yau manifold over its moduli space, we are automatically restricted in the sorts of distributions we can build. It would be helpful to pinpoint whether there is a deeper reason for such constraints.

ACKNOWLEDGMENTS

We thank V. Balasubramanian, N. Benjamin, N. Bobev, S. Franco, S. Gukov, H. Maxfield, J. McNamara, T. Pantev, D. Pei, M. Porrati, A. Tomasiello, and C. Vafa for helpful discussions and correspondence. We thank N. Bobev, J. Magan, and E. Shaghoulian for helpful comments on an earlier draft. J. J. H. and X. Y. thank the 2021 Simons summer workshop at the Simons Center for Geometry and Physics for kind hospitality during part of this work. X. Y. thanks the UPenn theory group for hospitality during this work. X. Y. also thanks the Caltech Particle Theory Group for hospitality during part of this work. The work of J. J. H. and A. P. T. is supported by the DOE (HEP) Award No. DE-SC0013528.

APPENDIX: CALABI–YAU MODULI SPACE FIBRATION

In this Appendix, we discuss how to produce a Calabi–Yau $(d + m)$ -fold as a fibration of a Calabi–Yau d -fold X over some m -dimensional subvariety $B \subset \mathcal{M}_X$ of its moduli space \mathcal{M}_X .

As discussed in the main text, adjunction tells us that we can produce a Calabi–Yau total space by fibering a Calabi–Yau over a Calabi–Yau subvariety of its own moduli space.¹⁹ However, this story is complicated by singularities in the fiber, and so we also consider here some more specific examples.

Explicitly, let us consider the case that the d -fold is defined as the vanishing locus of a homogeneous degree- $(d + 2)$ polynomial,

$$p_{d+2} = \sum_{i_0+\dots+i_{d+1}=d+2} a_{i_0,\dots,i_{d+1}} x_0^{i_0} \cdots x_{d+1}^{i_{d+1}}, \quad (\text{A1})$$

in \mathbb{P}^{d+1} . To fiber this over an m -complex-dimensional base B , we promote $a_{i_0,\dots,i_{d+1}}$ and x_j to sections of line bundles,

$$\begin{aligned} x_j &\in \Gamma(Y, \pi^*(\mathcal{L}_j) \otimes \mathcal{O}_Y(1)), \\ a_{i_0,\dots,i_{d+1}} &\in \Gamma(B, \mathcal{L} \otimes \bigotimes_j \mathcal{L}_j^{-i_j}), \end{aligned} \quad (\text{A2})$$

where $Y = \mathbb{P}^{d+1}(\mathcal{L}_0 \oplus \cdots \oplus \mathcal{L}_{d+1})$ is the ambient space, π is the projection map of the fibration $\pi: Y \rightarrow B$, and $\mathcal{L}, \mathcal{L}_j$ are line bundles over B . The total space \mathcal{X} of the fibration is then a hypersurface $p_{d+2} = 0$ in the bundle Y over B . The total Chern class of this bundle is

$$c(Y) = c(B) \prod_j (1 + c_1(\mathcal{L}_j) + c_1(\mathcal{O}_Y(1))). \quad (\text{A3})$$

By adjunction, we can then compute the total Chern class of \mathcal{X} as

$$c(\mathcal{X}) = \frac{c(Y)}{1 + c_1(\mathcal{L}) + (d + 2)c_1(\mathcal{O}_Y(1))}, \quad (\text{A4})$$

giving the first Chern class,

$$c_1(\mathcal{X}) = c_1(B) + \sum_j c_1(\mathcal{L}_j) - c_1(\mathcal{L}). \quad (\text{A5})$$

Thus, the Calabi–Yau condition is simply

$$c_1(\mathcal{L}) = c_1(B) + \sum_j c_1(\mathcal{L}_j). \quad (\text{A6})$$

As an explicit example, consider the case of the quintic threefold in \mathbb{P}^4 given by the hypersurface equation

¹⁹For our purposes, we actually consider the pullback of such a fibration by a bijection from a subspace of spacetime onto this subvariety of the moduli space.

$$\sum_{i=0}^4 \mu x_i^5 - 5\psi x_0 x_1 x_2 x_3 x_4 = 0, \quad (\text{A7})$$

with $[x_0 : x_1 : x_2 : x_3 : x_4]$ homogeneous coordinates on \mathbb{P}^4 . This is a generalization of the familiar Dwork family of quintics, which corresponds to $\mu = 1$. In this case, by homogeneity we have

$$\begin{aligned} x_j &\in \Gamma(Y, \pi^*(\mathcal{L}_x) \otimes \mathcal{O}_Y(1)), \\ \mu, \psi &\in \Gamma(B, \mathcal{L} \otimes \mathcal{L}_x^{-5}). \end{aligned} \quad (\text{A8})$$

Thus, the Calabi–Yau condition becomes

$$c_1(\mathcal{L}) = c_1(B) + 5c_1(\mathcal{L}_x). \quad (\text{A9})$$

-
- [1] P. W. Anderson, Absence of diffusion in certain random lattices, *Phys. Rev.* **109**, 1492 (1958).
- [2] S. Sachdev and J. Ye, Gapless Spin Fluid Ground State in a Random, Quantum Heisenberg Magnet, *Phys. Rev. Lett.* **70**, 3339 (1993).
- [3] G. Parisi, *Field Theory, Disorder and Simulations* (World Scientific, Singapore, 1992).
- [4] S. A. Hartnoll and J. E. Santos, Disordered Horizons: Holography of Randomly Disordered Fixed Points, *Phys. Rev. Lett.* **112**, 231601 (2014).
- [5] O. Aharony, Z. Komargodski, and S. Yankielowicz, Disorder in large-N theories, *J. High Energy Phys.* **04** (2016) 013.
- [6] A. Kitaev, A simple model of quantum holography, in *Proceedings of the KITP Program: Entanglement in Strongly-Correlated Quantum Matter* (2015), <http://online.kitp.ucsb.edu/online/entangled15/kitaev/>, <http://online.kitp.ucsb.edu/online/entangled15/kitaev2/>.
- [7] J. Maldacena and D. Stanford, Remarks on the Sachdev–Ye–Kitaev model, *Phys. Rev. D* **94**, 106002 (2016).
- [8] N. Afkhami-Jeddi, H. Cohn, T. Hartman, and A. Tajdini, Free partition functions and an averaged holographic duality, *J. High Energy Phys.* **01** (2021) 130.
- [9] A. Maloney and E. Witten, Averaging over Narain moduli space, *J. High Energy Phys.* **10** (2020) 187.
- [10] N. Benjamin, C. A. Keller, H. Ooguri, and I. G. Zadeh, Narain to Narnia, *Commun. Math. Phys.* **390**, 425 (2022).
- [11] I. Z. Rothstein, Gravitational Anderson Localization, *Phys. Rev. Lett.* **110**, 011601 (2013).
- [12] D. Green, Disorder in the early universe, *J. Cosmol. Astropart. Phys.* **03** (2015) 020.
- [13] N. Craig and D. Sutherland, Exponential Hierarchies from Anderson Localization in Theory Space, *Phys. Rev. Lett.* **120**, 221802 (2018).
- [14] V. Balasubramanian, J. J. Heckman, E. Lipeles, and A. P. Turner, Statistical coupling constants from hidden sector entanglement, *Phys. Rev. D* **103**, 066024 (2021).
- [15] J. Maldacena, D. Stanford, and Z. Yang, Diving into traversable wormholes, *Fortschr. Phys.* **65**, 1700034 (2017).
- [16] G. Penington, S. H. Shenker, D. Stanford, and Z. Yang, Replica wormholes and the black hole interior, [arXiv:1911.11977](https://arxiv.org/abs/1911.11977).
- [17] P. Saad, S. H. Shenker, and D. Stanford, JT gravity as a matrix integral, [arXiv:1903.11115](https://arxiv.org/abs/1903.11115).
- [18] D. Stanford and E. Witten, JT gravity and the ensembles of random matrix theory, *Adv. Theor. Math. Phys.* **24**, 1475 (2020).
- [19] D. Marolf and H. Maxfield, Transcending the ensemble: Baby universes, spacetime wormholes, and the order and disorder of black hole information, *J. High Energy Phys.* **08** (2020) 044.
- [20] S. B. Giddings and A. Strominger, String wormholes, *Phys. Lett. B* **230**, 46 (1989).
- [21] S. B. Giddings and G. J. Turiaci, Wormhole calculus, replicas, and entropies, *J. High Energy Phys.* **09** (2020) 194.
- [22] W. Cottrell, B. Freivogel, D. M. Hofman, and S. F. Lokhande, How to build the thermofield double state, *J. High Energy Phys.* **02** (2019) 058.
- [23] J. Sonner and M. Vielma, Eigenstate thermalization in the Sachdev–Ye–Kitaev model, *J. High Energy Phys.* **11** (2017) 149.
- [24] A. Almheiri, R. Mahajan, J. Maldacena, and Y. Zhao, The page curve of hawking radiation from semiclassical geometry, *J. High Energy Phys.* **03** (2020) 149.
- [25] H. Ooguri and C. Vafa, On the geometry of the string landscape and the swampland, *Nucl. Phys.* **B766**, 21 (2007).
- [26] J. J. Heckman and C. Vafa, Fine tuning, sequestering, and the swampland, *Phys. Lett. B* **798**, 135004 (2019).
- [27] H. Geng and A. Karch, Massive islands, *J. High Energy Phys.* **09** (2020) 121.
- [28] C. F. Uhlemann, Islands and page curves in 4d from type IIB, *J. High Energy Phys.* **08** (2021) 104.
- [29] H. Geng, A. Karch, C. Perez-Pardavila, S. Raju, L. Randall, M. Riojas, and S. Shashi, Inconsistency of islands in theories with long-range gravity, *J. High Energy Phys.* **01** (2022) 182.
- [30] S. Raju, Failure of the split property in gravity and the information paradox, [arXiv:2110.05470](https://arxiv.org/abs/2110.05470).
- [31] J. McNamara and C. Vafa, Baby universes, holography, and the swampland, [arXiv:2004.06738](https://arxiv.org/abs/2004.06738).
- [32] J. McNamara and C. Vafa, Cobordism classes and the swampland, [arXiv:1909.10355](https://arxiv.org/abs/1909.10355).
- [33] D. N. Page, Information in Black Hole Radiation, *Phys. Rev. Lett.* **71**, 3743 (1993).
- [34] D. N. Page, Time dependence of hawking radiation entropy, *J. Cosmol. Astropart. Phys.* **09** (2013) 028.
- [35] J. Louis and S. Lüst, Supersymmetric AdS₇ backgrounds in half-maximal supergravity and marginal operators of (1, 0) SCFTs, *J. High Energy Phys.* **10** (2015) 120.
- [36] C. Cordova, T. T. Dumitrescu, and K. Intriligator, Deformations of superconformal theories, *J. High Energy Phys.* **11** (2016) 135.
- [37] B. Zwiebach, *A First Course in String Theory* (Cambridge University Press, Cambridge, England, 2006).

- [38] L. E. Ibanez and A. M. Uranga, *String Theory and Particle Physics: An Introduction to String Phenomenology* (Cambridge University Press, Cambridge, England, 2012).
- [39] S. H. Shenker, Another length scale in string theory?, [arXiv:hep-th/9509132](#).
- [40] J. J. Heckman, Statistical inference and string theory, *Int. J. Mod. Phys. A* **30**, 1550160 (2015).
- [41] V. Balasubramanian, J. J. Heckman, and A. Maloney, Relative entropy and proximity of quantum field theories, *J. High Energy Phys.* **05** (2015) 104.
- [42] J. J. Heckman, J. G. Bernstein, and B. Vigoda, MCMC with strings and branes: The suburban algorithm (extended version), *Int. J. Mod. Phys. A* **32**, 1750133 (2017).
- [43] J. J. Heckman, J. G. Bernstein, and B. Vigoda, MCMC with strings and branes: The suburban algorithm, [arXiv:1605.06122](#).
- [44] R. E. Fowler III, Information theoretic interpretations of renormalization group flow, Ph.D. thesis, North Carolina University, 2020.
- [45] V. Balasubramanian, Statistical inference, occam's razor and statistical mechanics on the space of probability distributions, [arXiv:cond-mat/9601030](#).
- [46] P. Mehta and D. J. Schwab, An exact mapping between the variational renormalization group and deep learning, [arXiv:1410.3831](#).
- [47] J. Erdmenger, K. T. Grosvenor, and R. Jefferson, Information geometry in quantum field theory: Lessons from simple examples, *SciPost Phys.* **8**, 073 (2020).
- [48] J. Halverson, A. Maiti, and K. Stoner, Neural networks and quantum field theory, *Mach. Learn. Sci. Tech.* **2**, 035002 (2021).
- [49] J. Stout, Infinite distance limits and information theory, [arXiv:2106.11313](#).
- [50] J. Erdmenger, K. T. Grosvenor, and R. Jefferson, Towards quantifying information flows: Relative entropy in deep neural networks and the renormalization group, [arXiv:2107.06898](#).
- [51] H. Erbin, V. Lahoche, and D. O. Samary, Nonperturbative renormalization for the neural network-QFT correspondence, *Mach. Learn. Sci. Tech.* **3**, 015027 (2022).
- [52] R. Fowler and J. J. Heckman, Misanthropic entropy and renormalization as a communication channel, [arXiv:2108.02772](#).
- [53] S. S. Gubser, I. R. Klebanov, and A. M. Polyakov, Gauge theory correlators from noncritical string theory, *Phys. Lett. B* **428**, 105 (1998).
- [54] S. Ryu and T. Takayanagi, Holographic Derivation of Entanglement Entropy from AdS/CFT, *Phys. Rev. Lett.* **96**, 181602 (2006).
- [55] V. E. Hubeny, M. Rangamani, and T. Takayanagi, A covariant holographic entanglement entropy proposal, *J. High Energy Phys.* **07** (2007) 062.
- [56] V. Balasubramanian, M. Decross, and G. Sárosi, Knitting wormholes by entanglement in supergravity, *J. High Energy Phys.* **11** (2020) 167.
- [57] E. Witten and S.-T. Yau, Connectedness of the boundary in the AdS/CFT correspondence, *Adv. Theor. Math. Phys.* **3**, 1635 (1999).
- [58] J. M. Maldacena and L. Maoz, Wormholes in AdS, *J. High Energy Phys.* **02** (2004) 053.
- [59] J. M. Maldacena, The large N limit of superconformal field theories and supergravity, *Adv. Theor. Math. Phys.* **2**, 231 (1998).
- [60] E. Witten, Anti-de Sitter space, thermal phase transition, and confinement in gauge theories, *Adv. Theor. Math. Phys.* **2**, 505 (1998).
- [61] G. T. Horowitz and A. Strominger, Black strings and P-branes, *Nucl. Phys.* **B360**, 197 (1991).
- [62] C. Vafa, Evidence for F-theory, *Nucl. Phys.* **B469**, 403 (1996).
- [63] D. R. Morrison and C. Vafa, Compactifications of F-theory on Calabi-Yau threefolds—I, *Nucl. Phys.* **B473**, 74 (1996).
- [64] D. R. Morrison and C. Vafa, Compactifications of F-theory on Calabi-Yau threefolds—II, *Nucl. Phys.* **B476**, 437 (1996).
- [65] M. R. Douglas and G. W. Moore, D-branes, quivers, and ALE instantons, [arXiv:hep-th/9603167](#).
- [66] A. E. Lawrence, N. Nekrasov, and C. Vafa, On conformal field theories in four dimensions, *Nucl. Phys.* **B533**, 199 (1998).
- [67] L. Bhardwaj, M. Del Zotto, J. J. Heckman, D. R. Morrison, T. Rudelius, and C. Vafa, F-theory and the classification of little strings, *Phys. Rev. D* **93**, 086002 (2016); Erratum, *Phys. Rev. D* **100**, 029901 (2019).
- [68] T. Dimofte, D. Gaiotto, and S. Gukov, Gauge theories labelled by three-manifolds, *Commun. Math. Phys.* **325**, 367 (2014).
- [69] A. Gadde, S. Gukov, and P. Putrov, Fivebranes and 4-manifolds, *Progr. Math.* **319**, 155 (2016).
- [70] S. Gukov, Counting RG flows, *J. High Energy Phys.* **01** (2016) 020.
- [71] O. Aharony, O. Bergman, D. L. Jafferis, and J. Maldacena, $\mathcal{N} = 6$ superconformal Chern-Simons-Matter theories, M2-branes and their gravity duals, *J. High Energy Phys.* **10** (2008) 091.
- [72] O. Aharony, O. Bergman, and D. L. Jafferis, Fractional M2-branes, *J. High Energy Phys.* **11** (2008) 043.
- [73] J. P. Gauntlett, S. Lee, T. Mateos, and D. Waldram, Marginal deformations of field theories with AdS₄ duals, *J. High Energy Phys.* **08** (2005) 030.
- [74] E. Imeroni, On deformed gauge theories and their string/M-theory duals, *J. High Energy Phys.* **10** (2008) 026.
- [75] N. Bobev, P. Bomans, F. F. Gautason, and V. S. Min, Marginal deformations from type IIA supergravity, *SciPost Phys.* **10**, 140 (2021).
- [76] N. Bobev, F. F. Gautason, and J. van Muiden, The holographic conformal manifold of 3d $\mathcal{N} = 2$ S-fold SCFTs, *J. High Energy Phys.* **07** (2021) 221.
- [77] H. Garcia-Compean and A. M. Uranga, Brane box realization of chiral gauge theories in two-dimensions, *Nucl. Phys.* **B539**, 329 (1999).
- [78] S. Franco, S. Lee, and R.-K. Seong, Brane brick models, toric Calabi-Yau 4-folds and 2d (0,2) quivers, *J. High Energy Phys.* **02** (2016) 047.
- [79] S. Franco, D. Ghim, S. Lee, and R.-K. Seong, Elliptic genera of 2d (0,2) gauge theories from brane brick models, *J. High Energy Phys.* **06** (2017) 068.
- [80] J. D. Brown and M. Henneaux, Central charges in the canonical realization of asymptotic symmetries: An example from three-dimensional gravity, *Commun. Math. Phys.* **104**, 207 (1986).

- [81] A. Simons, A. Strominger, D. M. Thompson, and X. Yin, Supersymmetric branes in $\text{AdS}_2 \times S^2 \times CY_3$, *Phys. Rev. D* **71**, 066008 (2005).
- [82] F. Denef, Quantum quivers and hall/hole halos, *J. High Energy Phys.* **10** (2002) 023.
- [83] D. Anninos, T. Anous, and F. Denef, Disordered quivers and cold horizons, *J. High Energy Phys.* **12** (2016) 071.
- [84] F. Denef, B. R. Greene, and M. Raugas, Split attractor flows and the spectrum of BPS D-branes on the quintic, *J. High Energy Phys.* **05** (2001) 012.
- [85] F. Denef and G. W. Moore, Split states, entropy enigmas, holes and halos, *J. High Energy Phys.* **11** (2011) 129.
- [86] R. Dijkgraaf, R. Gopakumar, H. Ooguri, and C. Vafa, Baby universes in string theory, *Phys. Rev. D* **73**, 066002 (2006).
- [87] W. Nahm, Supersymmetries and their representations, *Nucl. Phys.* **B135**, 149 (1978).
- [88] G. Penington, Entanglement wedge reconstruction and the information paradox, *J. High Energy Phys.* **09** (2020) 002.
- [89] A. Almheiri, N. Engelhardt, D. Marolf, and H. Maxfield, The entropy of bulk quantum fields and the entanglement wedge of an evaporating black hole, *J. High Energy Phys.* **12** (2019) 063.