

Quantum Darwinism: Entanglement, Branches, and the Emergent Classicality of Redundantly Stored Quantum Information

Robin Blume-Kohout*

*Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545 and
Institute for Quantum Information, California Institute of Technology, Pasadena, CA 91125*

Wojciech H. Zurek†

Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545

(Dated: October 13, 2005)

We lay a comprehensive foundation for the study of redundant information storage in decoherence processes. Redundancy has been proposed as a prerequisite for *objectivity*, the defining property of classical objects. We consider two ensembles of states for a model universe consisting of one system and many environments: the first consisting of arbitrary states, and the second consisting of “singly-branching” states consistent with a simple decoherence model. Typical states from the random ensemble do not store information about the system redundantly, but information stored in branching states has a redundancy proportional to the environment’s size. We compute the specific redundancy for a wide range of model universes, and fit the results to a simple first-principles theory. Our results show that the presence of redundancy divides information about the system into three parts: classical (redundant); purely quantum; and the borderline, undifferentiated or “nonredundant,” information.

PACS numbers: PACS numbers: 03.65.Yz, 03.67.Pp, 03.67.-a, 03.67.Mn

The theory of decoherence [1, 2, 3, 4] has resolved much of the decades-old confusion about the transition from quantum to classical physics (see articles in [5]). It provides a mechanism – weak measurement by the environment – by which a quantum system can be compelled to behave classically. The recent development of quantum information theory has encouraged an information-theoretic view of decoherence, wherein information about a central system “leaks out” into the environment, and thereby becomes classical [6].

In this paper, we pursue a natural extension of the decoherence program, by asking “What happens to the information that leaks out of the system?” That information should be sought in the “rest of the universe” – i.e., the system’s environment. The environment is a witness to the system’s state, and can serve as a resource for measuring or controlling the system. Our particular focus, within this *Environment as a Witness* paradigm, is on how *redundantly* information about the system is recorded in the environment. This is relevant to quantum technology; a detailed picture of how decoherence destroys quantum information may help in designing schemes to correct its effects.

It also illuminates fundamental physics. Massive redundancy can cause certain information to become *objective*, at the expense of other information. The process by which this “fittest” information is propagated through the environment, at the expense of incompatible information, is *Quantum Darwinism*. Two forthcoming papers

([7, 8]) will investigate the *dynamics* of quantum Darwinism.

This paper is focused on the *kinematics* of information storage and the environment-as-a-witness paradigm. It is organized as follows. In Section I, we introduce objectivity and the “environment as a witness” paradigm, show that redundant records indicate objectivity, and propose quantitative and qualitative measures of redundancy. In Section II, we analyze randomly distributed states, show that they do not display redundant information storage, and argue that they do not describe the Universe (see next paragraph) in which we live. In Section III, we propose *singly-branching* states as an alternative description, and use numerics to demonstrate redundant information storage. Section IV presents an analytical model for the numerical results. Finally, we summarize our most important results and discuss future work in Section V.

We use the word “universe” to denote both (a) everything that exists in reality, and (b) a self-contained model of a system and its environment. We distinguish the two by capitalizing usage (a). Thus, while living in the *Universe*, we simulate assorted *universes*.

I. THE ENVIRONMENT AS A WITNESS

Previous studies of decoherence have focused on the system’s reduced density matrix (ρ_S), and on master equations that describe its evolution. To study information flow into the environment, we require a new paradigm.

We begin with a simple observation: **information about a system (\mathcal{S}) is obtained by measuring its environment (\mathcal{E})** (see [1, 9]). Although the standard

*Electronic address: robin@blumekohout.com

†Electronic address: whz@lanl.gov

theories of quantum measurement (see e.g. Von Neumann [10], etc.) presume a direct measurement on the system, real experiments rely on *indirect* measurements. As you read this paper, you measure the albedo of the page – but actually, your eyes are capturing photons from the electromagnetic environment. Information about the page is inferred from assumed correlations between text and photons. A similar argument holds for every physics experiment; the scientist gets information about \mathcal{S} by capturing and measuring a *fragment* of \mathcal{E} .

This motivates us to focus on correlations between \mathcal{S} and individual fragments of \mathcal{E} . In particular, we will seek to determine whether a particular state – or a particular ensemble of states – allows an observer who captures a small fragment of \mathcal{E} to deduce the system’s state. If so, then the system’s state is *objectively* recorded.

A. Objectivity

A property – e.g., the state of a system – is objective when many independent observers agree about it. The observers’ independence is crucial. When many secondary observers are informed by a single primary observer, then only the primary observer’s opinion is objective, *not* necessarily the property which he observed. *Independent* observers, examining a single quantum system, cannot have agreed on a particular measurement basis beforehand. They will generally measure different observables – and therefore will not agree afterward. An isolated quantum system’s state cannot be objective, because measurements of noncommuting observables invalidate each other.

Classical theory, on the other hand, permits observers to measure a system without disturbing it. Properties of classical systems (e.g., classical states) are thus objective. Each observer can record the state in question without altering it, and afterward all the observers will agree on what they discovered. Of course, observers may obtain *different* information – e.g., one observer may make a more effective measurement than another – but not *contradictory* information.

Objectivity provides an excellent criterion for exploring the emergence of classicality through decoherence. A quantum system becomes more classical as its measurable properties become more objective. The use of “measurable” is significant. Nothing can make *every* property of a quantum system objective, because some observables are incompatible with others. Two observers can never simultaneously obtain reliable information about incompatible observables (such as position and momentum) of the same system. Decoherence partially solves this problem by destroying all the observables incompatible with a system’s *pointer observable*. We are thus motivated to explore (a) how the pointer observable becomes objective, and (b) how decoherence and the emergence of objectivity are related.

B. Technical details and assumptions

This “environment as a witness” paradigm [1, 6, 11, 12] is ideally suited to exploring objectivity. In order to make *independent* measurements of \mathcal{S} , multiple observers must partition the environment into fragments. In this paper, we assume that measurements must be made on distinct Hilbert spaces in order to be independent, so we divide the environment into fragments as

$$\mathcal{E} = \mathcal{E}_A \otimes \mathcal{E}_B \otimes \mathcal{E}_C \otimes \dots \quad (1)$$

Several factors limit an observer’s ability to obtain information about \mathcal{S} by measuring a fragment of the environment (\mathcal{E}_A). We can make more or less optimistic assumptions about some of these factors, but **the degree of correlation between \mathcal{S} and \mathcal{E}_A is clearly a limiting factor**. An observer whose particular fragment is not correlated with \mathcal{S} has no way to obtain information about \mathcal{S} . That fragment of \mathcal{E} is irrelevant and, for the purpose of gaining information about \mathcal{S} , might as well not exist. The absolute prerequisite for demonstrating a property’s objectivity is that information about it be recorded in many fragments – that is, *redundantly*.

We quantify redundancy by counting the number of fragments which can provide sufficient information. The redundancy of information about some property is a natural measure of that property’s objectivity [1]. Classical properties are objective because information about them is recorded with [effectively] infinite redundancy. For instance, if we flip a coin, then its final orientation is recorded by trillions of scattered photons. Thousands of cameras, each capturing a tiny fraction of them, could each provide a record. Redundancy is not dependent on *actual* observers. Instead, it is a statement about what observers *could* do, if they existed.

A pertinent question is “Why not allow an observer to measure the system itself?” First, only one observer could be allowed to do so without sacrificing independence. Thus, at most, this would increase redundancy by 1. Furthermore, an observer with access to the central system could measure it in some weird basis, thus destroying its state. Since it’s not then clear what the information obtained by the *other* observers would refer to, we regard the system itself as off limits to observers.

C. The overall program

The work presented here is a natural extension of the decoherence program. However, employing the environment as a communication channel – not just “sink” for information lost to decoherence – is also in a sense “beyond decoherence.” It is the next stage in exploring how classicality emerges from the quantum substrate.

In order to fully understand the role that redundancy and objectivity play in (1) the emergence of classicality, and (2) the destruction of quantum coherence, we’d like to answer the following questions:

1. Given a state $\rho_{S\mathcal{E}}$ for the system and its environment (the “universe”), how do we quantify the redundancy of information (about \mathcal{S}) in \mathcal{E} ?
2. For a particular “universe,” what states are *typical* (that is, likely to exist)? Do they display redundancy? If so, how much?
3. What sorts of (a) initial states, and (b) dynamics lead dynamically to redundancy?
4. Do realistic models of decoherence produce the massive redundancy we expect in the classical regime?
5. For complicated systems, with many independent properties, how do we distinguish *what* property a bit of information is about?
6. When information about an observable is redundantly recorded, is information about incompatible observables inaccessible?

The building blocks of this work – e.g., the reasoning presented in this section – have been laid in recent years by [1, 6, 11, 13, 14]. The first attempt to address items (1) and (3) appeared in [6], and was refined in [11], which also analyzed a particular simple model of decoherence numerically. In the current paper, we answer (1) and (2) in detail, and consider (3) briefly.

D. Computing redundancy

To compute the redundancy (R) of some information (\mathcal{I}), we divide the environment into fragments ($\mathcal{E} = \mathcal{E}_A \otimes \mathcal{E}_B \otimes \dots$), and demand that each fragment supply \mathcal{I} independently. The redundancy of \mathcal{I} is the number of such fragments into which the environment can be divided. A generalized GHZ state is a good example:

$$|\psi\rangle_{S\mathcal{E}} = \alpha |0\rangle_S |0000\dots 0\rangle_{\mathcal{E}} + \beta |1\rangle_S |1111\dots 1\rangle_{\mathcal{E}} \quad (2)$$

We can determine the system’s state by measuring any sub-environment. Each qubit in \mathcal{E} provides all the available information about \mathcal{S} (see, however, note [39]). To extend this analysis to arbitrary states, we need (a) a metric for information, (b) a protocol for dividing the environment into fragments, and (c) an idea of how much of \mathcal{I} is “available”.

1. A metric for information

We use *quantum mutual information* (QMI) as an information metric. QMI is a generalization of the classical mutual information [15]. *Quantum* mutual information is defined in terms of the Von Neumann entropy, $H = -\text{Tr}(\rho \log \rho)$, as:

$$\mathcal{I}_{A:B} = H_A + H_B - H_{AB} \quad (3)$$




FIG. 1: (Color) **Three ways to divide up the universe.** The decoherence paradigm divides the universe into a system (\mathcal{S}) and an environment (\mathcal{E}) as in (a). In the environment-as-a-witness paradigm, we further subdivide \mathcal{E} into *subenvironments*, as in (b). No subenvironment can be further subdivided, and it is easier to measure one \mathcal{E}_i than to make a joint measurement on several. *Fragments* are constructed, so as to provide enough information to infer the state of \mathcal{S} , by combining subenvironments as in (c). Measurements on distinct fragments always commute.

This is simple to calculate, provides a reliable measure of correlation between systems, and has been used previously for this purpose [1, 16, 17]. Unlike classical mutual information, the QMI between system A and system B is not bounded by the entropy of either system. In the presence of entanglement, the QMI can be as large as $H_A + H_B$, which reflects the existence of quantum correlations beyond the classical ones [18].

2. Dividing \mathcal{E} into fragments

A pre-existing concept of locality, usually expressed as a fixed tensor product structure or as a set of allowable structures, is *fundamental* to redundancy analysis. Allowing an *arbitrary* division of \mathcal{E} into fragments would make every state where \mathcal{S} is entangled with \mathcal{E} (see note [40]) equivalent (via re-division of \mathcal{E}) to a GHZ-like state (Eq. 2). Decoherence would be equivalent to redundancy.

The need for a fixed tensor product structure is familiar; both decoherence and entanglement are meaningless without a fixed division between the system and its

environment ([1, 9], see e.g. [19] for a discussion of a possible tensor product structures' origins in measurable observables; an explanation that does not refer to measurements would be needed in the present context). In the environment-as-a-witness paradigm, we divide \mathcal{E} into indivisible *subenvironments*:

$$\mathcal{E} = \mathcal{E}_1 \otimes \mathcal{E}_2 \otimes \mathcal{E}_3 \otimes \dots \otimes \mathcal{E}_{N_{\text{env}}}. \quad (4)$$

These subenvironments can be rearranged into larger *fragments*. A generic fragment consisting of m subenvironments will be written as $\mathcal{E}_{\{m\}}$. The fragment containing the particular subenvironments $\{\mathcal{E}_{i_1}, \mathcal{E}_{i_2}, \dots, \mathcal{E}_{i_m}\}$ is denoted $\mathcal{E}_{\{i_1, i_2, \dots, i_m\}}$.

We assume that each observer captures a *random* fragment of \mathcal{E} . This ensures their strict independence. In essence, we do not allow the observers to caucus over the partition of \mathcal{E} , dividing it up in an advantageous way.

3. How much information is practically available

The maximum information that could be provided about \mathcal{S} is its entropy, $H_{\mathcal{S}}$. In general, no fragment can provide *all* this information [41]. Following the reasoning in [11], we demand that each fragment provide some large fraction, $1 - \delta$ (where $\delta \ll 1$), of the available information about \mathcal{S} . The precise magnitude of the *information deficit* δ should not be important. We denote the redundancy of “all but δ of the available information” by R_{δ} . That is, when we allow a deficit of $\delta = 0.1$, we are computing $R_{0.1}$ or $R_{10\%}$.

To compute R_{δ} , we start by defining N_{δ} as **the number of disjoint fragments \mathcal{E}_i such that $\mathcal{I}_{\mathcal{S}:\mathcal{E}_i} \geq (1 - \delta)\mathcal{I}_{\mathcal{S}:\mathcal{E}}$** . We might just define $R_{\delta} = N_{\delta}$, except for two caveats.

1. A large deficit (δ) in the definition of “sufficient” information could lead to spurious redundancy. Suppose there exist $N = 5$ fragments that provide full information. If $\delta = 0.5$, then we might split each fragment in half to obtain $N_{\delta} = 10$ fragments that each provide “sufficient” information. To compensate for this, we replace N_{δ} with $(1 - \delta)N_{\delta}$.
2. Because of quantum correlations, $\mathcal{I}_{\mathcal{S}:\mathcal{E}_i}$ can be as high as $2H_{\mathcal{S}}$. We allow for this by *assuming* that the information provided by one fragment represents strictly quantum correlations, and throwing this fragment away. This means replacing $(1 - \delta)N_{\delta}$ with $(1 - \delta)N_{\delta} - 1$.

By assuming the worst case, we have obtained a *lower bound* for the true redundancy:

$$R_{\delta} \geq (1 - \delta)N_{\delta} - 1. \quad (5)$$

For small δ , this is fairly tight, as N_{δ} is clearly an upper bound. Since our current toolset, subject to the caveats




FIG. 2: (Color) **Three profiles for partial information plots (\mathcal{I} vs. m)**. (a): the behavior of *independent* environments. (b): information is stored *redundantly*. (c): information is *encoded* in multiple environments.

mentioned above, does not permit a more precise determination of R_{δ} , we report the lower bound throughout. Thus, when we report “ $R_{10\%} = 9$,” we really mean “ $R_{10\%}$ is at least 9, and not much more.”

E. Identifying qualitative redundancy

The actual *amount* of redundancy is often less important than the qualitative observation that information is stored very redundantly (e.g., $R \gg 1$). Whether $R = 100$ or $R = 1000$, the information in question is certainly objective – but if $R \sim 1$, then its objectivity is in doubt. We also wish to consider more general questions: e.g., how much does R_{δ} depend on δ ? or *why* does a state display virtually no redundancy?

For these purposes, we plot the amount of information about \mathcal{S} supplied by a fragment of size m ($\mathcal{I}_{\mathcal{S}:\mathcal{E}_{\{m\}}}$), against m . Since there are very many fragments of a given size, we average $\mathcal{I}_{\mathcal{S}:\mathcal{E}_{\{m\}}}$ over a representative sample of fragments to obtain $\overline{\mathcal{I}}(m)$. The plot of $\overline{\mathcal{I}}(m)$, which shows the partial information yielded by a partial environment, is a *partial information plot* (PIP). When the universe is in a pure state (see [20], and Appendix A), the PIP must be anti-symmetric around its center (see Fig. 2). Together with the observation that $\overline{\mathcal{I}}(m)$ must be strictly non-decreasing (capturing *more* of the environment cannot *decrease* the amount of information obtained), this permits the three basic profiles shown in Figure 2.

Redundancy (see Fig. 2b) is characterized by a rapid rise of $\overline{\mathcal{I}}$ at relatively small m , followed by a long “classical plateau”. In this region, all the easily available information has been obtained. Additional environments confirm what is already known, but provide nothing new. Only by capturing *all* the environments can an observer manipulate quantum correlations. The power to do so is indicated by the sharp rise in $\overline{\mathcal{I}}$ at $m \sim N_{\text{env}}$.

II. INFORMATION STORAGE IN RANDOM STATES

Redundant information storage is ubiquitous in the classical world. We might naïvely expect that randomly chosen states of a model universe – e.g., a D_S -dimensional system in contact with a bath of N_{env} D_E -dimensional systems – would display massive redundancy. To test this hypothesis, we compute partial information plots for random states, and average them over the uniform ensemble. This was first done in [20], for qubits. In this work, we extend the analysis to systems and environments with arbitrary sizes.

A. The uniform ensemble

For any [finite] D -dimensional Hilbert space, there exists a unitarily invariant uniform distribution over states, usually referred to as *Haar measure*. We examine the behavior of typical random states by averaging PIPs over this uniform ensemble. This average can be obtained analytically, using a formula for the average entropy of a subspace that was conjectured by Page [21], then proved by Sen [22] and others [23, 24].

Page’s formula [21, 22, 23, 24] for the mean entropy $\overline{H}(m, n)$ of an m -dimensional subsystem of an mn -dimensional system (where $m \leq n$) is

$$\overline{H}(m, n) = \sum_{k=n+1}^{mn} \frac{1}{k} - \frac{m-1}{2n} \quad (6)$$

$$= \Psi(mn) - \Psi(n+1) - \frac{m-1}{2n}, \quad (7)$$

where the latter expression is given in terms of the *digamma* Ψ function. For a D_S -dimensional system in contact with N_{env} environments of size D_E , the average mutual information between the system and m subenvironments is

$$\begin{aligned} \overline{\mathcal{I}_{\mathcal{S}:\mathcal{E}_{\{m\}}}} &= \overline{H}(D_S, D_E^{N_{\text{env}}}) \\ &+ \overline{H}(D_E^m, D_S D_E^{N_{\text{env}}-m}) \\ &- \overline{H}(D_S D_E^m, D_E^{N_{\text{env}}-m}). \end{aligned} \quad (8)$$

B. Partial information plots (PIPs)

Our results (Figs. 3-5) demonstrate that **typical states from the uniform ensemble do not display redundancy**. Figure 3a illustrates typical behavior. As an observer captures successively more subenvironments (increasing m), he gains virtually no information about \mathcal{S} . $\mathcal{I}_{\mathcal{S}:\mathcal{E}_{\{m\}}}$ remains close to zero. When approximately 50% of the subenvironments have been captured, the observer begins to gain information. $\overline{\mathcal{I}}$ rises rapidly, through H_s and onward nearly to $2H_s$.

Information about \mathcal{S} is *encoded* in the environment (as in Fig. 2c), much as a classical bit can be encoded in the parity of an ancillary bitstring. In the classical example, however, *every* bit of the ancilla must be captured to deduce the encoded bit.

This encoding, or “anti-redundancy”, is related to quantum error correction [25, 26, 27, 28]. In an encoding state, any majority subset of the \mathcal{E}_i has nearly-complete information. The recorded information is unaffected by the loss of any minority subset. States with this behavior can be used as a quantum code to protect against bit loss. Our results show that *generic* states – i.e., states selected randomly from the whole \mathcal{SE} Hilbert space – form a nearly-optimal error-correction code for bit-loss errors. Shannon noted similar behavior for classical codewords [29].

Figures 3b and 4 extend this result to larger systems. The results are consistent; information is still encoded, and only the total amount of encoded information changes.

C. Conclusions

Our first main result is that **typical states selected randomly from the uniform ensemble display no redundant information storage**. Instead, they display encoding or anti-redundancy. This is not to say that *all* states are “antiredundant”, merely that redundant information storage is rare. As m declines from $\frac{N_{\text{env}}}{2}$, $\overline{\mathcal{I}}(m)$ declines exponentially. For large N_{env} , states where information is *not* encoded this way are vanishingly rare. If even a small fixed fraction ϵ of states displayed the opposite “redundant” behavior, then $\overline{\mathcal{I}}(m)$ would have to be $O(\epsilon)$ at small m . The fact that $\overline{\mathcal{I}}(m)$ is exponentially close to zero implies that the fraction of non-“encoding” states must decline exponentially with N_{env} .

The obvious conclusion is that the Universe does not evolve into random states. Our observations of ubiquitous redundancy in the real Universe are inconsistent with the random-state model. This is interesting, but not terribly surprising. There is no good reason to expect that the Universe’s state *would* be random – we are not, for instance, in thermodynamic equilibrium. The interactions of systems with their environments must select states that *are* characterized by greater redundancy. In the next section, we suggest and analyze such an ensemble.

III. DECOHERENCE AND BRANCHING STATES

Decoherence – the loss of information to the environment – is a prerequisite for redundancy. The simplest models of decoherence [30] are essentially identical to those for quantum measurements. A set of pointer states for the system, $\{|n\rangle\}$, are singled out, and the environ-

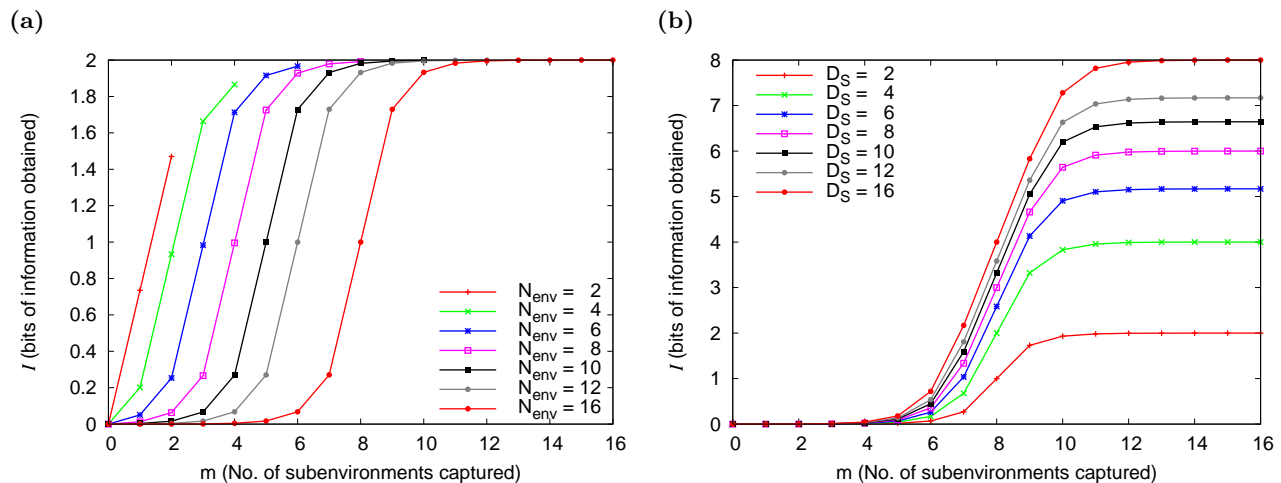


FIG. 3: (Color) **Partial information plots (PIPs) for the uniform ensemble.** We plot the average information ($\bar{\mathcal{I}}$) obtainable from a fragment ($\mathcal{E}_{\{m\}}$), against the fragment's size (m). $\mathcal{I}(m)$ is averaged over **all states in the uniform ensemble**. (a): A qubit system coupled to environments consisting of $N_{\text{env}} = 2 \dots 16$ qubits. (b): Systems with sizes $D_S = 2 \dots 16$ coupled to a 16-qubit environment. **Discussion:** No significant information is obtained until almost half the subenvironments have been captured. Once $m > \frac{N_{\text{env}}}{2}$, virtually all possible information (both quantum and classical) is available. Because more than half the environment is required to obtain useful information, there is no redundant information storage in typical uniformly-distributed states. Instead, the information is *encoded* throughout the environment.

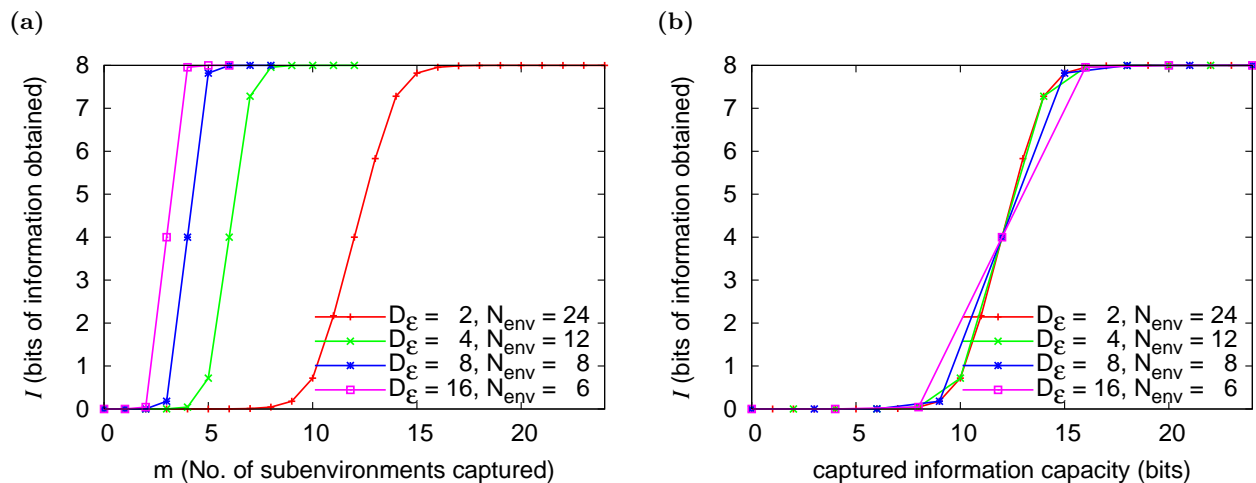


FIG. 4: (Color) **Equivalent environments** When the state of the universe is chosen randomly, the environment's Hilbert space dimension determines its information-recording properties. (a): PIPs for a 16-d system coupled to several equivalent environments with $D_{\text{total}} = 2^{24}$. The subenvironments are $\{2, 4, 8, 16\}$ -dimensional, and N_{env} is scaled appropriately. The plots are essentially identical – only the scaling of the m -axis changes. (b): The same data, but with the *captured fraction* of the environment plotted on the independent axis.

ment “measures” which $|n\rangle$ the system is in, by evolving from some initial state ($|\varepsilon_0\rangle$) into a *conditional* state, $|\varepsilon_n\rangle$. If ρ_S is written out in the pointer basis, its diagonal elements (ρ_{nn}) remain unchanged. Coherences between different pointer states (e.g., ρ_{nm}) are reduced by a *de-coherence factor*:

$$\gamma_{nm} \equiv \langle \varepsilon_n | \varepsilon_m \rangle. \quad (9)$$

We presume that (a) the subenvironments are initially unentangled, (b) each subenvironment “measures” the same basis of the system, and (c) the state of the universe

is pure. In this simple model, the universe is initially in a product state:

$$|\Psi_0\rangle = |s_0\rangle \otimes |\varepsilon_0^{(1)}\rangle \otimes |\varepsilon_0^{(2)}\rangle \otimes \dots \otimes |\varepsilon_0^{(N_{\text{env}})}\rangle. \quad (10)$$

The subenvironments do not interact with each other, and the system does not evolve on its own. Letting the system's initial state be $|s_0\rangle = \sum_n s_n |n\rangle$, the universe evolves over time into:

$$|\Psi_t\rangle = \sum_n s_n |n\rangle_S \otimes |\varepsilon_n^{(1)}\rangle \otimes |\varepsilon_n^{(2)}\rangle \otimes \dots \otimes |\varepsilon_n^{(N_{\text{env}})}\rangle, \quad (11)$$

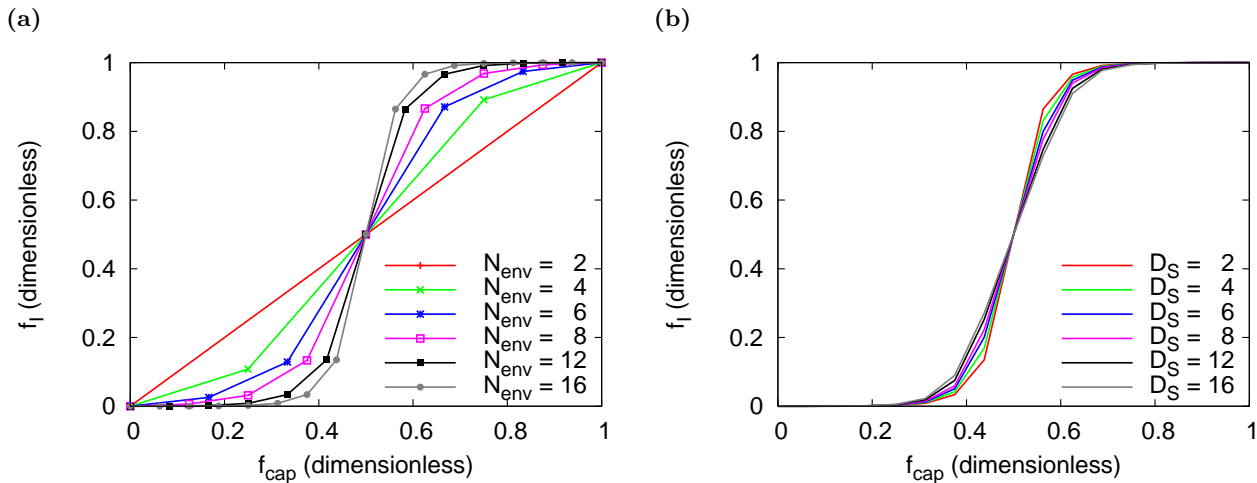


FIG. 5: (Color) **Scaled versions (SPIPs) of the plots in Fig. 3.** SPIPs are useful for comparing environments with different numbers of subenvironments, and for computing R_δ , the redundancy for a given fraction $1 - \delta$ of the total information. To estimate redundancy, simply draw a horizontal line at $f_I = \frac{1-\delta}{2}$, and note the value of f_{cap} where it intersects the PIP. This provides a good estimate of $1/R_\delta$. It is not a perfect estimate for several reasons; most importantly, the PIP and SPIP plot *the average* \mathcal{I} obtained from a given-sized fragment of the environment. This is not the same as *the average fragment size* (\bar{m}) required to obtain \mathcal{I} , since we average the same data over different variables. In these plots, of course, no redundancy is evident – we are looking ahead to the next section.

where $|\varepsilon_n^{(j)}\rangle$ is the *conditional* state into which the j th subenvironment evolves *if* the system is in state $|n\rangle$. Different conditional states of a given subenvironment will *not* generally be orthogonal to one another, except in highly simplified (e.g. C-NOT) models.

A. The branching-state ensemble

We refer to the states defined by Eq. 11 as *singly-branching states*, or simply as *branching states*. In Everett’s many-worlds interpretation [31], a branching state’s wavefunction has D_S branches. Each branch is perfectly correlated with a particular pointer state of the system. The subenvironments are not entangled with each other, only correlated (classically) via the system. In contrast, a typical random state from the uniform ensemble has D_{universe} branches, with a new branching at every subsystem.

In dynamical models of decoherence, the universe at a given time will be described by a *particular* branching state that depends on the environment’s initial state, and on its dynamics. In this paper, we sidestep the difficulties of specifying these parameters, by considering the ensemble of *all* branching states. We select the conditional $|\varepsilon_n^{(j)}\rangle$ at random from each subenvironment’s uniform ensemble. Each pointer state of the system is correlated with a randomly chosen product state of all the environments.

The amount of available information is set by the system’s initial state (i.e., the s_n coefficients). The eigenvalues of ρ_S after complete decoherence, which determine its maximum entropy, are $\lambda_n = |s_n|^2$. Since we cannot

examine *all* possible states, we focus on maximally “measurable” generalized Hadamard states:

$$s_n = \frac{1}{\sqrt{D_S}} \quad \forall n. \quad (12)$$

To verify that our results are generally valid, we also treat (briefly) another class of initial states.

By examining the branching-state ensemble, we are *not* conjecturing that the Universe is found exclusively in branching states. Branching states form an interesting and physically well-motivated ensemble to explore. We shall see that, unlike the uniform ensemble, the branching-state ensemble displays redundancy consistent with observations of the physical Universe. Our Universe might well tend to evolve into similar states, but we are not ready to establish such a conjecture. Characterizing the states in which the physical Universe (or a fragment thereof) *is* found is a substantially more ambitious project.

B. Numerical analysis of branching states

We begin our exploration of branching states by examining typical PIPs, for various systems and environments. We average these PIPs over the branching-state ensemble, so there are only three adjustable parameters: D_S , D_E , and N_{env} . Our results confirm that information is stored redundantly. Next, we examine a quantitative measure of redundancy (R_δ), and its dependence on D_S , D_E , and N_{env} . Finally, we derive some analytical approximations, compare them with numerical data, and discuss the implications of our results.

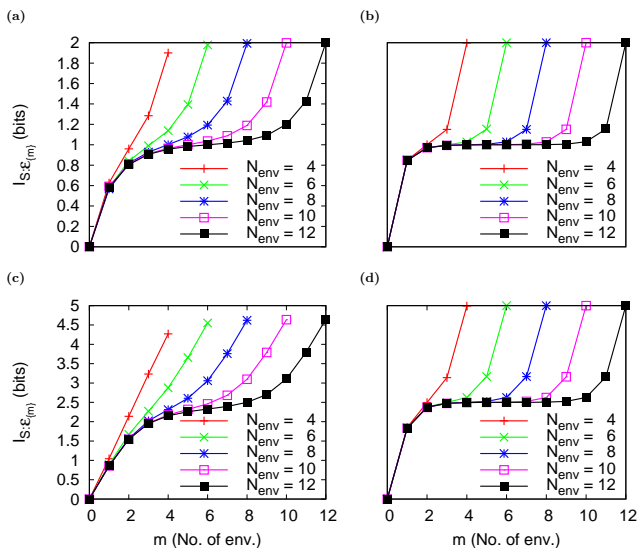


FIG. 6: (Color) **PIPs for ensembles of singly-branching states.** The system is initialized in a Hadamard state, and decohered by N_{env} subenvironments. We plot the average information ($\bar{\mathcal{I}}$) available from a collection of m subenvironments. (a): A qubit is decohered by qubits. (b): A qubit is decohered by 5-dimensional subenvironments. (c): A 5-d system is decohered by qubits. (d): A 5-d system is decohered by 5-d subenvironments. **Discussion:** As N_{env} is increased from 4 to 12, a “classical plateau” appears. This indicates redundant information storage. In the regime $m \ll N_{\text{env}}$, the PIP converges to an asymptotic form. When \mathcal{S} is larger than \mathcal{E} (see (c)), the environment is barely sufficient to decohere the system, and there is no redundancy (see also Fig. 9).

1. Partial information plots

Information is redundant when small fragments yield nearly-complete information – that is, when the PIP looks like Fig. 2b. PIPs for branching states (Fig. 6) show exactly this profile. $\bar{\mathcal{I}}(m)$ rises rapidly from $\bar{\mathcal{I}}(0) = 0$, then approaches $H_{\mathcal{S}}$ asymptotically to produce a “classical plateau” centered at $m = \frac{N_{\text{env}}}{2}$.

As N_{env} grows, the interesting regimes at $m \sim 0$ and $m \sim N_{\text{env}}$ do not change; the classical plateau simply extends to connect them. The initial bits of information that an observer gains about a system are extremely useful, but eventually a point of diminishing returns is reached, where further information is redundant. The degree of redundancy should therefore scale with N_{env} .

2. Non-Hadamard states for \mathcal{S}

Non-Hadamard states provide a different spectrum of information for \mathcal{E} to capture. We consider states defined by

$$s_n \propto \frac{1}{\sqrt{2^n}}, \quad (13)$$

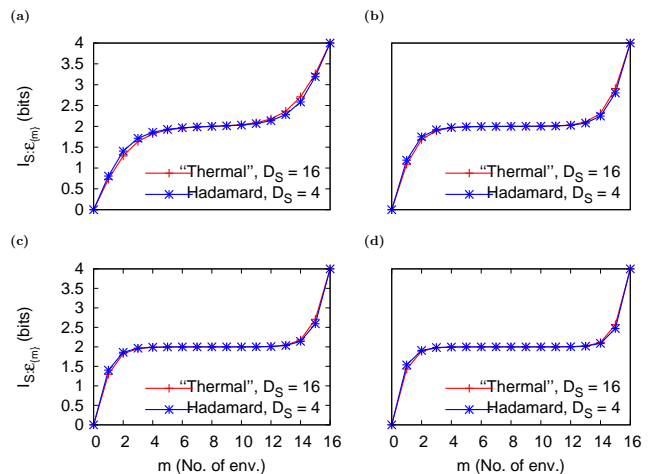


FIG. 7: (Color) **PIPs for non-Hadamard states:** $D_{\mathcal{E}} = 2, 3, 4, 5$ in plots (a), (b), (c), (d), respectively. The system is 16-dimensional, and initialized in a “thermal” state, where $s_n \propto \frac{1}{\sqrt{2^n}}$. The entropy of this density matrix is ~ 2 bits (as opposed to 4 bits for a $D_{\mathcal{S}} = 16$ Hadamard state). We compare the PIPs for “thermal” states with $D_{\mathcal{S}} = 16$ to PIPs for Hadamard $D_{\mathcal{S}} = 4$ states, which also develop 2 bits of entropy, varying the subenvironments’ size. These PIPs confirm that our observations apply to non-Hadamard states, and that $H_{\mathcal{S}}$ characterizes how information about the system is stored.

The post-decoherence spectrum of $\rho_{\mathcal{S}}$ is non-degenerate – in fact, it is exactly that of a thermal spin – i.e., a particle with a Hamiltonian $\mathbf{H} = \mathbf{J}_z$, in equilibrium with a bath at finite temperature. We refer to these states as “thermal” branching states (and retain quotation marks to emphasize that our justification of this nomenclature is unphysical).

Our general approach is to assume that the system’s maximum entropy determines its informational properties. The entropy of a decohered “thermal” state does not increase logarithmically with $D_{\mathcal{S}}$, but asymptotes to $H_{\mathcal{S}} = 2$ bits. This is exactly the entropy of a $D_{\mathcal{S}} = 4$ Hadamard state, so in the limit $D_{\mathcal{S}} \rightarrow \infty$, “thermal” states should behave much the same as a $D_{\mathcal{S}} = 4$ Hadamard state.

This conjecture is confirmed in Fig. 7, which compares PIPs for “thermal” states with $D_{\mathcal{S}} = 16$ to PIPs for Hadamard states with $D_{\mathcal{S}} = 4$. The plots’ similarity indicates that $H_{\mathcal{S}}$ is the major factor in how information about \mathcal{S} is recorded. Further numerical results use Hadamard states for specificity’s sake.

3. How PIPs scale with the composition of \mathcal{E}

As the number of subenvironments in \mathcal{E} grows, comparing PIPs for different environments becomes difficult. Re-parameterizing the axes, and plotting the *fraction* of \mathcal{I} available from a *fraction* of \mathcal{E} , allows direct comparison of different universes. Scaled PIPs (SPIPs) for environ-

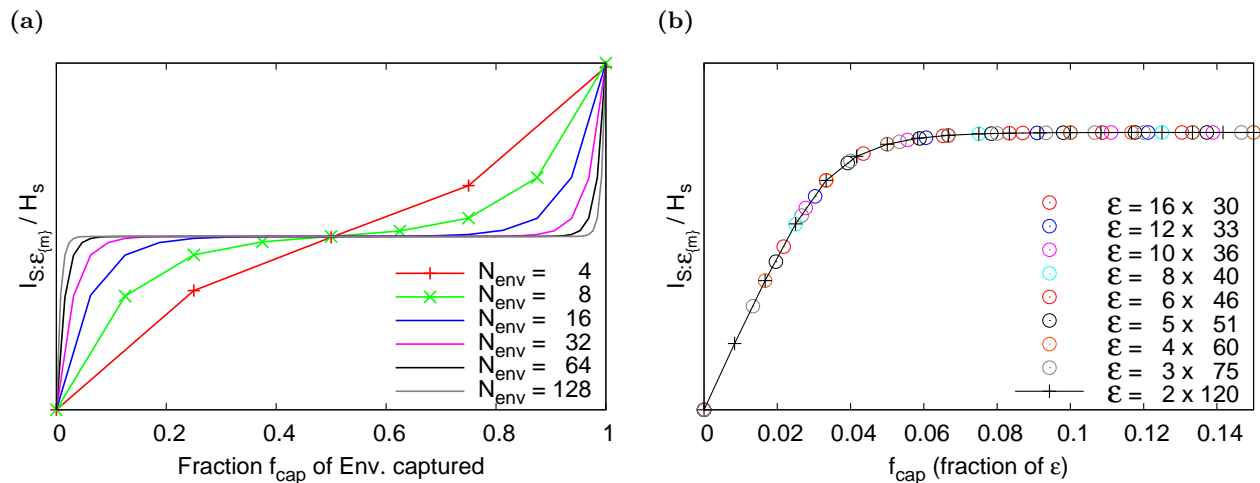


FIG. 8: (Color) **Scaled partial information plots (SPIPs)** compare information storage in different environments. **(a)**: A qutrit system coupled to $N_{\text{env}} = 4 \dots 128$ qutrit environments. **(b)**: A qutrit system coupled to nine different environments with the same *information capacity*. **Discussion**: As N_{env} increases, redundancy (indicated by sharp curvature) grows (plot **(a)**). If N_{env} and $D_{\mathcal{E}}$ are scaled so that total Hilbert space dimension ($D_{\mathcal{E}}^{N_{\text{env}}}$) remains constant, then the SPIP remains unchanged (plot **(b)**). Plot **(b)** also illustrates the difference between the regime of *linear* information gain (here, $f_{\text{cap}} < 0.04$) and the exponential convergence to the “classical plateau” thereafter.

ments with $N_{\text{env}} = 4 \dots 128$ (Fig. 8a) show that the information about \mathcal{S} becomes more redundant as N_{env} grows.

Different environments, whose total Hilbert space dimensions are the same, act equivalently (see also Sec. II B). We have simulated a 16-dimensional system coupled to nine different, but equivalent, environments (Fig. 8b). Although the *number* and *size* of the subenvironments are varied, the redundancy of the available information depends only on \mathcal{E} ’s total information capacity: $c \equiv \log[\dim(\mathcal{H})]$. Each \mathcal{E} in Fig. 8b has $c \simeq 120$ bits, so their SPIPs are essentially identical.

4. Redundancy: numerical values

Branching states are natural generalizations of GHZ states, so we expect redundant information storage. Figure 9 confirms this over a wide range of parameters. The amount of redundancy is proportional to the size of the environment, which agrees with the classical intuition that very large environments should store many copies of information about the system. Larger subenvironments (measured by $D_{\mathcal{E}}$) increase redundancy by storing more information in each subenvironment. Conversely, larger systems have more properties to measure, which in turn require more space for information storage. The total amount of redundancy is reduced for large $D_{\mathcal{S}}$.

The other important feature of the plots in Fig. 9 is the relatively weak dependence of R_{δ} on the information deficit (δ). As we vary δ from 2% to 25% (a full order of magnitude), R_{δ} changes by less than a factor of 2. The distinction between classical (massively redundant) and quantum (nonredundant) information is largely indepen-

dent of δ .

IV. THEORETICAL ANALYSIS OF BRANCHING STATES

The numerical analysis in the previous section offers compelling evidence that

1. Information is stored redundantly in branching states,
2. The amount of redundancy scales with N_{env} , and
3. R_{δ} is relatively insensitive to δ .

In this section, we construct theoretical models for PIPs and redundancy, which confirm these hypotheses.

A. Structural properties of branching states

We begin by using the structure inherent to branching states to compute a quantity of fundamental interest,

$$\mathcal{I}_{S:\mathcal{E}_{\{m\}}} = H_{\mathcal{S}} + H_{\mathcal{E}_{\{m\}}} - H_{S\mathcal{E}_{\{m\}}}, \quad (14)$$

the mutual information between the system and a partial environment $\mathcal{E}_{\{m\}}$.

We require the entropies of $\rho_{\mathcal{S}}$, $\rho_{\mathcal{E}_{\{m\}}}$, and $\rho_{S\mathcal{E}_{\{m\}}}$. Tracing over the rest of the universe is simplified by the structure that Eq. 11 implies. Each relevant density matrix (regardless of its actual dimension) has only $D_{\mathcal{S}}$ nonzero eigenvalues. That is, the reduced states for S , $\mathcal{E}_{\{m\}}$, and $S\mathcal{E}_{\{m\}}$ are all “virtual qudits” with $D = D_{\mathcal{S}}$.




FIG. 9: (Color) **Redundancy for an assortment of branching-state ensembles.** (a): $R_{10\%}$ for a D -dimensional system decohered by D -dimensional subenvironments. (b): $R_{10\%}$ for a 5-dimensional system decohered by $D_{\mathcal{E}} = 2 \dots 5$ -dimensional subenvironments. (c): $R_{10\%}$ for a $D_{\mathcal{S}} = 2 \dots 16$ -dimensional system decohered by 4-dimensional subenvironments. (d): R_{δ} for $\delta = 0.001 \dots 0.25$ and $D_{\mathcal{S}} = D_{\mathcal{E}} = 5$. **Discussion:** Each plot shows the ensemble-average of R_{δ} , as a function of N_{env} . R_{δ} increases linearly with the number of environments. R_{δ} increases with $D_{\mathcal{E}}$, but decreases with $D_{\mathcal{S}}$. Larger environments – but larger systems have more information to be stored. Information is stored with slightly greater efficiency for large $D_{\mathcal{S}}$ and $D_{\mathcal{E}}$ (plot (a)). Note that if \mathcal{S} is larger than \mathcal{E} (e.g., $D_{\mathcal{S}} = 16$ in plot (c)), there may be no redundancy. Finally, δ affects redundancy (plot (d)) – but varying δ by a full order of magnitude (from 2% to 25%) changes R_{δ} by less than 50%.

Each ρ , when reduced to its $D_{\mathcal{S}}$ -dimensional support, is *spectrally* equivalent to a partially decohered variant of the system’s initial state:

$$|s_0\rangle\langle s_0| = \sum_{nm} s_n s_m^* |n\rangle\langle m|. \quad (15)$$

In other words, we can obtain $\rho_{\mathcal{S}}$, $\rho_{\mathcal{E}_{\{m\}}}$, or $\rho_{\mathcal{S}\mathcal{E}_{\{m\}}}$ by taking $|s_0\rangle\langle s_0|$ and suppressing the off-diagonal elements according to a specific rule.

To determine this rule, we define (for each subenvironment) a *multiplicative decoherence factor*, γ :

$$\gamma_{ij}^{(k)} = \langle \mathcal{E}_j^{(k)} | \mathcal{E}_i^{(k)} \rangle, \quad (16)$$

and an associated *additive decoherence factor*, d :

$$d_{ij}^{(k)} \equiv -\log \gamma_{ij}^{(k)}. \quad (17)$$

Now, $\gamma_{ij}^{(k)}$ quantifies how much \mathcal{E}_k contributes to decohering $|i\rangle$ from $|j\rangle$. The γ -factors from different \mathcal{E}_k combine multiplicatively; the d -factors provide a convenient additive representation. Each relevant density matrix ρ_X (for

$X \in \{\rho_{\mathcal{S}}, \mathcal{E}_{\{m\}}, \mathcal{S}\mathcal{E}_{\{m\}}\}$) is given by:

$$\langle i | \rho_{\mathcal{S}} | j \rangle = (s_i s_j^*) e^{-d_{ij}^{(X)}}. \quad (18)$$

The d -factor for each subsystem is a sum over d -factors for the component \mathcal{E}_k :

$$d_{ij}^{(\mathcal{E}_{\{m\}})} = \sum_{k \in \mathcal{E}_{\{m\}}} d_{ij}^{(k)} \quad (19)$$

$$d_{ij}^{(\mathcal{S})} = \sum_{k \in \mathcal{E}} d_{ij}^{(k)} \quad (20)$$

$$d_{ij}^{(\mathcal{S}\mathcal{E}_{\{m\}})} = \sum_{k \notin \mathcal{E}_{\{m\}}} d_{ij}^{(k)}. \quad (21)$$

Thus, each ρ appears to have been decohered by a different subset of \mathcal{E} :

- $\rho_{\mathcal{S}}$ has been decohered by **every** subenvironment,
- $\rho_{\mathcal{S}\mathcal{E}_{\{m\}}}$ has been decohered by all the subenvironments *not* in $\mathcal{E}_{\{m\}}$,
- $\rho_{\mathcal{E}_{\{m\}}}$ has been decohered by all the subenvironments in $\mathcal{E}_{\{m\}}$.

Note: If the last point seems counter-intuitive, recall that for any bipartite decomposition of $|\psi\rangle_{AB}$, the reduced ρ_A and ρ_B are spectrally equivalent. Thus $\rho_{\mathcal{E}_{\{m\}}}$ is equal to $\rho_{\overline{\mathcal{E}_{\{m\}}}}$, where $\overline{\mathcal{E}_{\{m\}}}$ contains all the environments *not* in $\mathcal{E}_{\{m\}}$.

Computing $\mathcal{I}_{\mathcal{S}:\mathcal{E}_m}$ (in terms of the entropy of these three states) can be done *exactly* via numerical diagonalization. For qubit systems, it can also be done analytically (see [20] for extensive details). For our model, we now derive an approximation for $H(\rho)$.

B. Theoretical PIPs: averaging $\mathcal{I}(m)$

As a particular ρ is decohered by more and more subenvironments, its off-diagonal elements decline rapidly toward zero. We will treat the off-diagonal elements of a *partially* decohered state, $\rho = \sum_{ij} s_i s_j^* \gamma_{ij} |i\rangle\langle j|$, as a perturbation around the *fully* decohered state ρ_0 , which has eigenvalues $\lambda_i = |s_i|^2$ and entropy H_0 .

1. Average entropy of partially decohered states

Let $\rho = \rho_0 + \Delta$, where Δ is a small off-diagonal perturbation to ρ_0 , and expand its entropy as $H(\rho) \approx H(\rho_0) + O(\Delta)$. An intuitively appealing starting point is the MacLaurin expansion of $H(x) = -x \ln(x)$, which yields

$$H(\rho_0 + \Delta) \approx H(\rho_0) - \text{Tr}[\Delta(\mathbb{1} - \ln(\rho_0))] - \frac{1}{2} \frac{\Delta^2}{\rho_0} + \frac{1}{6} \frac{\Delta^3}{\rho_0^2} \dots \quad (22)$$

The first order term in Eq. 22 vanishes, because Δ is purely off-diagonal and $\mathbb{1} - \ln(\rho)$ is purely diagonal. The leading term is thus $\frac{\Delta^2}{2\rho_0}$ – but the matrix quotient $\frac{\Delta^{k+1}}{\rho_0^k}$ is ill-defined when Δ and ρ_0 do not commute.

A more involved expansion of $H(\rho)$ around $\rho = \mathbb{1}$ (see Appendix C) yields a series for $H(\rho_0 + \Delta)$. It is equivalent to Eq. 22 for scalars, but for matrices it involves (1) expanding ρ_0^{-k} in a power series, and (2) taking a totally symmetric product between Δ^{k+1} and the resulting power series.

To leading order in Δ ,

$$H(\rho) \approx H(\rho_0) - \frac{|\overline{\gamma}|^2}{2} (h(\rho_0) - 1), \quad (23)$$

where $|\overline{\gamma}|^2$ is the average of $|\gamma_{ij}|^2$ over all $i \neq j$, and $h(\rho_0)$ is a nontrivial function,

$$h(\rho_0) = \sum_{n,p=0}^{\infty} \frac{\text{Tr}[\rho_0(\mathbb{1} - \rho_0)^p] \text{Tr}[\rho_0(\mathbb{1} - \rho_0)^n]}{n + p + 1}. \quad (24)$$

2. Effective Hilbert space dimension

In general, $h(\rho_0)$ cannot be simplified further. However, it is well approximated by the *effective Hilbert space dimension* of ρ_0 . To see this, we consider the special case where ρ_0 has D identical eigenvalues, $\lambda_i = \frac{1}{D}$. When reduced to its support, $\rho_0 = \frac{\mathbb{1}}{D}$. The summation can be done explicitly:

$$\begin{aligned} h(\rho_0) &= \sum_{n,p=0}^{\infty} \frac{\text{Tr}\left[\frac{\mathbb{1}}{D}(\mathbb{1} - \frac{\mathbb{1}}{D})^p\right] \text{Tr}\left[\frac{\mathbb{1}}{D}(\mathbb{1} - \frac{\mathbb{1}}{D})^n\right]}{n + p + 1} \\ &= \sum_{n,p=0}^{\infty} \frac{((1 - D^{-1})^p) ((1 - D^{-1})^n)}{n + p + 1} \\ &= \sum_{n,p=0}^{\infty} \frac{(1 - D^{-1})^{n+p}}{n + p + 1} \\ &= \sum_{n+p=0}^{\infty} (1 - D^{-1})^{n+p} \\ &= D \end{aligned} \quad (25)$$

Note that D appeared *only* based on the eigenvalue spectrum of ρ_0 . In the example above, the $H_0 = H(\rho_0) = \log(D)$. Since the total range of $\overline{\mathcal{I}}(m)$ is proportional to H_0 , a logical generalization is

$$h(\rho_0) \approx e^{H_0}, \quad (26)$$

$$H(\rho) \approx H(\rho_0) - \frac{|\overline{\gamma}|^2}{2} (e^{H_0} - 1). \quad (27)$$

Numerical experimentation, and an analytic calculation in $D_{\mathcal{S}} = 2$, confirm that Eq. 26 is a good approximation everywhere, in addition to being exact for (1) maximally mixed states, and (2) pure states.

3. Average decoherence factors

The γ_{ij} depend on the details of $\psi_{\mathcal{S}\mathcal{E}}$. However, when they are small enough to count as a perturbation on ρ , the environment's Hilbert space is very large. The $|\gamma_{ij}|^2$ can then be treated as independent random variables, so $|\overline{\gamma}|^2$ is equal to an average over the entire branching state ensemble:

$$\begin{aligned} \overline{|\gamma|^2} &= \overline{\langle \psi | \psi' \rangle \langle \psi' | \psi \rangle} \\ &= \text{Tr} \left(\overline{|\psi\rangle\langle\psi|} \overline{|\psi'\rangle\langle\psi'|} \right) \\ &= \frac{\text{Tr}(\mathbb{1}\mathbb{1})}{D_{\mathcal{E}}^2} \\ &= D_{\mathcal{E}}^{-1} \end{aligned} \quad (28)$$

This is the mean value of $|\overline{\gamma}|^2$ for a *single* subenvironment. For a collection of m subenvironments, m such γ factors are multiplied together, so the mean value of $|\overline{\gamma}|^2$ becomes $D_{\mathcal{E}}^{-m}$.

4. The result

Putting this all together, the average entropy of a $D_{\mathcal{S}}$ -dimensional system decohered by m $D_{\mathcal{E}}$ -dimensional environments is

$$\overline{H} \simeq H_0 - \frac{e^{H_0} - 1}{2} D_{\mathcal{E}}^{-m}, \quad (29)$$

and the average mutual information between the system and m subenvironments is

$$\begin{aligned} \overline{\mathcal{I}}(m) &\approx H_0 - \frac{e^{H_0} - 1}{2} \left(D_{\mathcal{E}}^{-m} - D_{\mathcal{E}}^{-(N_{\text{env}} - m)} \right) \\ &= H_0 + (e^{H_0} - 1) \sinh \left[\left(m - \frac{N_{\text{env}}}{2} \right) \ln(D_{\mathcal{E}}) \right]. \end{aligned} \quad (30)$$

Equation 30 is only a good approximation only near the classical plateau, where $\overline{\mathcal{I}} \simeq H_0$. Around $m = 0$ and $m = N_{\text{env}}$, $\overline{\mathcal{I}}$ rises linearly, not exponentially. Each subenvironment can provide only $\log_2 D_{\mathcal{E}}$ bits of information, so until the information starts to become redundant, we're in a different regime (see Fig. 8b).

Once the information capacity of the captured environments ($m \log D_{\mathcal{E}}$) becomes greater than the amount of information in the system (H_0), Eq. 30 becomes valid. It describes the slow approach to “perfect” information about the system, as m increases. Figure 10 compares exact (numerical) results for $\overline{\mathcal{I}}(m)$ to the approximation in Eq. 30.

C. Theoretical redundancy: averaging $m(\mathcal{I})$

Branching states develop when each subenvironment interacts independently with \mathcal{S} . The data in Section

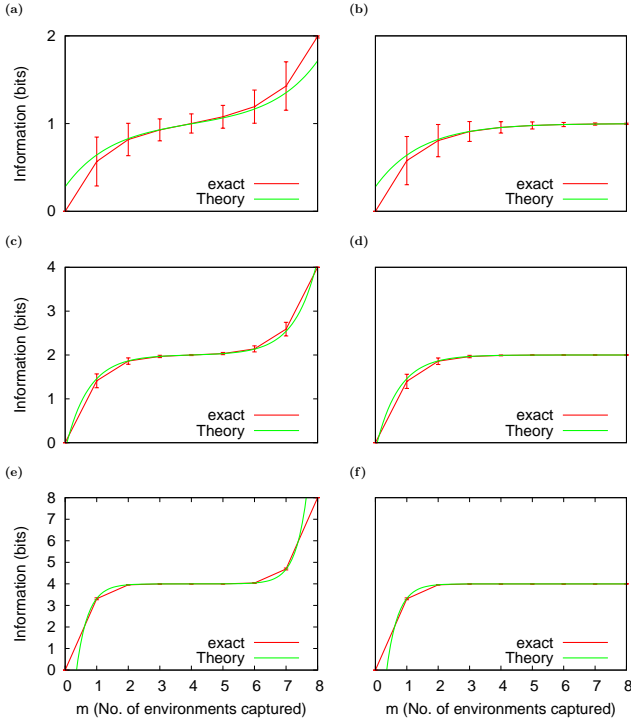


FIG. 10: (Color) **Numerical PIPs vs. Theory:** We compare the approximation derived in Sec. IV B with numerics. Error bars on numerics represent typical fluctuations over the branching-state ensemble. **(a):** $D_S = D_E = 2$, $N_{\text{env}} = 8$. **(b):** $D_S = D_E = 2$, $N_{\text{env}} = 32$. **(c):** $D_S = D_E = 4$, $N_{\text{env}} = 8$. **(d):** $D_S = D_E = 4$, $N_{\text{env}} = 32$. **(e):** $D_S = D_E = 16$, $N_{\text{env}} = 8$. **(f):** $D_S = D_E = 16$, $N_{\text{env}} = 32$. **Discussion:** The approximation is virtually perfect near the classical plateau. For small m , the rate of information gain is more nearly linear, and the approximation fails. Although it works well at $m = 0$ for $D_S = 4$ (plots **(b),(e)**), it fails spectacularly near $m = 0$ for large D_S (plots **(c),(f)**)

III B 4 (esp. Fig. 9) confirm that redundancy in branching states is proportional to N_{env} . A certain number of subenvironments (m_δ) is enough to provide sufficient information.

To capture this scaling, we define *specific redundancy* as

$$r_\delta = \lim_{N_{\text{env}} \rightarrow \infty} \frac{R_\delta}{N_{\text{env}}} = \frac{1 - \delta}{m_\delta} \quad (31)$$

In this section, we use specific redundancy to examine precisely how D_S , D_E , and δ affect information storage in branching states. We derive an approximate formula for r_δ , and compare its predictions to numerical data.

In the previous section, we computed the average information yielded by m environments. Now, we compute the average m required to achieve a given \mathcal{I} .

When N_{env} is large, $H_S \mathcal{E}_{\{m\}} \simeq H_S \simeq H_0$, so $\mathcal{I}_{S:\mathcal{E}_{\{m\}}} \simeq H \mathcal{E}_{\{m\}}$. We take Eq. 27,

$$\mathcal{I}_{S:\mathcal{E}_{\{m\}}} \approx H_S - \frac{1}{2} \overline{|\gamma|^2} (e^{H_S} - 1), \quad (32)$$

as a starting point. For the fragment to provide “sufficient” information, $\mathcal{I} - H_S$ must be less than δH_S , which requires

$$\frac{\sum_{i \neq j} |\gamma_{ij}|^2}{D_S(D_S - 1)} (e^{H_S} - 1) \leq 2\delta H_S. \quad (33)$$

Assuming ρ_0 is maximally mixed (i.e., $e^{H_0} = D_S$), and replacing the γ_{ij} with independent random variables γ_n , we obtain the following condition on a “sufficiently large” fragment:

$$\left[\frac{D_S(D_S - 1)}{2} \sum_{n=1}^{D_S} |\gamma_n|^2 \right] \leq \delta D_S H_S \quad (34)$$

The interaction of $\frac{1}{2} D_S(D_S - 1)$ independent γ -factors makes it difficult to solve Eq. 34 rigorously. We begin instead by considering a qubit system, which has only one off-diagonal γ .

1. Specific redundancy for qubit systems

For a single qubit, there is only one decoherence factor: d_{01} , which we’ll refer to simply as d . Eq. 34 simplifies to:

$$d \geq d_\delta \equiv -\frac{1}{2} \log(2\delta H_S) \quad (35)$$

The increase in d with m can be approximated as a biased random walk, where each step has a mean length (\bar{d}) and a variance (Δd). After m environments are added to the fragment, d obeys a normal distribution ($p_m(d)$), whose mean and variance are $m\bar{d}$ and $\sqrt{m}\Delta d$, respectively. We postpone the calculation of \bar{d} and Δd for the moment.

Let $p_{\text{suff}}(m)$ be the probability that a fragment consisting of m subenvironments provides sufficient information (i.e., satisfies equation 35). Then

$$p_{\text{suff}}(m) = \int_{d_\delta}^{\infty} p_m(d) dd, \quad (36)$$

and the probability that m environments are *required* is

$$p_{\text{req}}(m) = p_{\text{suff}}(m) - p_{\text{suff}}(m - 1) \quad (37)$$

$$= \int_{m-1}^m \frac{\partial}{\partial n} p_{\text{suff}}(n) dn, \quad (38)$$

and the expected fragment size (\bar{m}) is

$$\begin{aligned}
\bar{m} &= \sum_{m=0}^{\infty} m p_{\text{req}}(m) \\
&= \sum_{m=0}^{\infty} m \int_{m-1}^m \frac{\partial}{\partial n} p_{\text{suff}}(n) dn \\
&\simeq \int_0^{\infty} \left(m + \frac{1}{2}\right) \frac{\partial}{\partial m} p_{\text{suff}}(m) dm \\
&= \frac{1}{2} + \int_0^{\infty} m \frac{\partial}{\partial m} p_{\text{suff}}(m) dm \\
&= \frac{1}{2} + \int_0^{\infty} (1 - p_{\text{suff}}(m)) dm \\
&= \frac{1}{2} + \int_0^{\infty} dm \int_{-\infty}^{d_\delta} p_m(d) dd. \tag{39}
\end{aligned}$$

We interchange the order of integration, substitute the appropriate normal distribution for $p_m(d)$, and end up with

$$\bar{m} = \frac{d_\delta}{\bar{d}} + \frac{\Delta^2}{2\bar{d}^2} + \frac{1}{2}. \tag{40}$$

2. Specific redundancy for general D_S

Whereas Eq. 35 (for qubits) has one $|\gamma|^2$ term, Eq. 34 involves a sum of $\frac{1}{2}D_S(D_S + 1)$ such terms. Deriving an analyzing a probability distribution for this sum is very difficult, so we take a simpler route. We replace the sum over terms with a single term, $\frac{1}{2}D_S(D_S + 1) \cdot \gamma^2$, where γ^2 represents all the off-diagonal terms. The new condition for sufficient information is:

$$\begin{aligned}
\frac{D_S(D_S - 1)}{2} \gamma^2 &\leq \delta D_S H_S \\
\gamma^2 &\leq \frac{2\delta H_S}{D_S - 1} \\
d &\geq d_\delta \equiv -\frac{1}{2} \log \left(\frac{2\delta H_S}{D_S - 1} \right). \tag{41}
\end{aligned}$$

D_S has been incorporated into a redefinition of d_δ . Equation 40 is still valid for qubits, but it generalizes to

$$\bar{m} = \frac{\log(D_S - 1) - \log(2\delta H_S)}{2\bar{d}} + \frac{\Delta d^2}{2\bar{d}^2} + \frac{1}{2}. \tag{42}$$

We combine this expression with Eq. 31 to obtain a general estimate for specific redundancy:

$$r_\delta = \frac{2\bar{d}^2(1 - \delta)}{\Delta^2 + \bar{d}^2 + \bar{d}(\log(D_S - 1) - \log(2\delta H_S))} \tag{43}$$

3. Dependence of mean decoherence factor (d) on D_E

The computation of \bar{d} and Δd in terms of D_E is somewhat tedious. Details can be found in Appendix D, where

we calculate:

$$\bar{d} = \frac{1}{2}(\Psi(D_E) + \gamma_{\text{EM}}), \tag{44}$$

$$\Delta d^2 = \frac{\pi^2}{24} - \frac{\Psi_1(D_E)}{4}, \tag{45}$$

in terms of the digamma ($\Psi(n)$) and trigamma ($\Psi_1(n)$) functions [32, 33], and the Euler-Mascheroni constant $\gamma_{\text{EM}} = 0.577\dots$. These functions may not be familiar to all readers, so we present the first few values in Table I.

D_E	2	3	4	5	6	8
\bar{d}	$\frac{1}{2}$	$\frac{3}{4}$	$\frac{11}{12}$	$\frac{25}{24}$	$\frac{137}{120}$	$\frac{363}{280}$
Δd	$\frac{1}{2}$	$\frac{\sqrt{5}}{4}$	$\frac{7}{12}$	$\frac{\sqrt{205}}{24}$	$\frac{\sqrt{5269}}{120}$	$\frac{\sqrt{266681}}{840}$

TABLE I: The table shows the first few values of \bar{d} and Δd , for environments of size $D_E \in [2, 3, 4, 5, 6, 8]$. See Appendix D for details on the calculation.

For larger D_E , we can safely approximate Eqs. 44-45 as:

$$\bar{d} \simeq \frac{1}{2}(\log(D_E) + \gamma_{\text{EM}}) \tag{46}$$

$$\Delta d \simeq \frac{\pi}{\sqrt{24}}. \tag{47}$$

4. How good is the estimate?

In Figure 11, we compare numerical results to the approximation of Eq. 43. The analytical estimate is very good for qubit systems, but loses some fidelity for larger D_S . A more sophisticated treatment of the multiple γ_{ij} terms – each representing an independent observable which the environment must record – would eliminate this error.

To get an intuitive feel for the dependence of r_δ on its parameters, we consider the regime of large systems, large environments, and small deficit – i.e., $H_0 \gg 1$, $\bar{d} \sim \frac{1}{2} \log(D_E)$, $\Delta d \sim \frac{\pi}{\sqrt{24}}$, and $\delta \ll 1$. In this regime, we can ruthlessly simplify Eq. 43 to obtain a simple prediction:

$$r_\delta \approx \frac{\log(D_E)}{\log(D_S) - \log(\delta)}. \tag{48}$$

The plots in Fig. 12 show the ratio between numerical r_δ data and the simple predictions of Eq.48. They confirm that Eq.48 is a good rule of thumb.

Eq. 48 can be interpreted as a capsule summary of how redundancy scales in the “random-state” model of decoherence.

1. Redundancy is proportional to N_{env} , the number of independent subenvironments. **More environments produce more redundancy.**

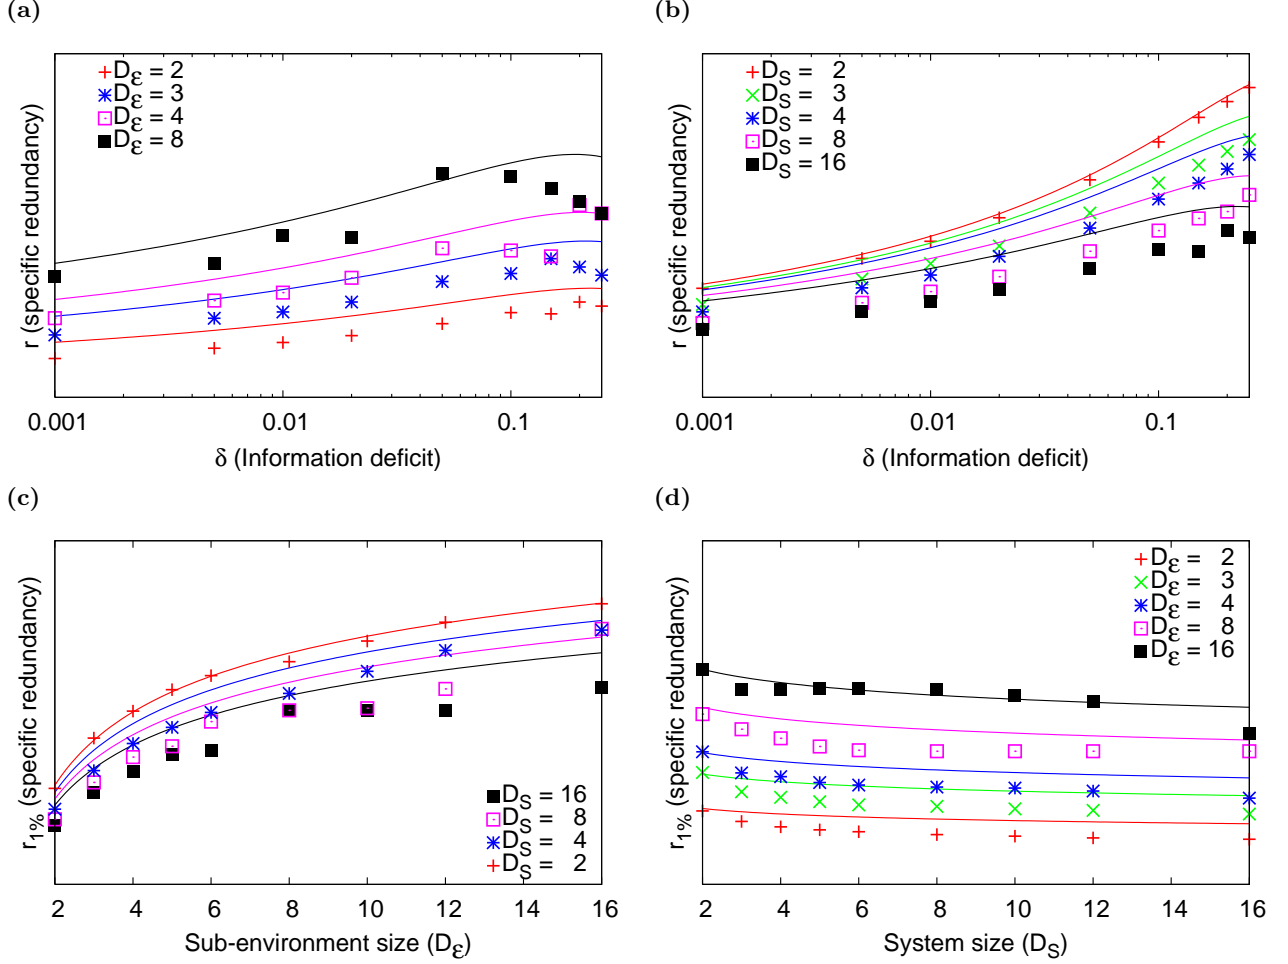


FIG. 11: (Color) **Specific redundancy** ($r_\delta \equiv R_\delta/N_{\text{env}}$): numerical data (symbols) compared with theory (Eq. 43, solid lines). (a): r_δ vs. δ , for a 16-d system coupled to 2, 3, 4, 8-dimensional subenvironments. (b): r_δ vs. δ , for 2, 3, 4, 8, 16-d systems coupled to qubit subenvironments. (c): $r_{1\%}$ vs. $D_\mathcal{E}$. (d): $r_{1\%}$ vs. $D_\mathcal{S}$. **Discussion:** Theory predicts the overall behavior of redundancy well. It is nearly perfect for $D_\mathcal{S} = 2$, but overestimates r for larger systems. As δ increases, r_δ saturates and even *declines* because of the $(1 - \delta)$ prefactor in Eq. 5. When δ is large, the theory breaks down (see (a)), because a single subenvironment can provide sufficient information.

2. Redundancy is proportional to \bar{d} , the *mean decoherence factor* of a single subenvironment, which grows as $\log D_\mathcal{E}$. **Larger environments produce more redundancy, in proportion to their information capacity.**
3. Redundancy is (roughly) inversely proportional to $H_\mathcal{S}$, the total information available about the system. **Larger systems require more space in the environment.**
4. The deficit (δ) appears as a logarithmic addition to $H_\mathcal{S}$. Reducing the amount of “ignorable” information is equivalent to making the system bigger. **Redundancy depends only weakly (logarithmically) on the deficit, δ .**

V. CONCLUSIONS AND DISCUSSION

‘There is no information without representation’: information has to be stored somewhere. To retrieve it, we must measure the systems where it is stored. To understand the properties of information, we look at the properties of this retrieval process. We have focused on the question: **How easily can information about a system be retrieved from its environment?**

The answer is strongly dependent on how the system became correlated with its environment. Random interactions between \mathcal{S} and all of \mathcal{E} leave no useful correlations – to learn about \mathcal{S} we must measure most of \mathcal{E} . However, when localized parts of \mathcal{E} interact independently with \mathcal{S} , an observer can learn about \mathcal{S} by measuring a small fragment of \mathcal{E} . Furthermore, the information that he learns is objective – another independent observer will arrive at the same conclusions.

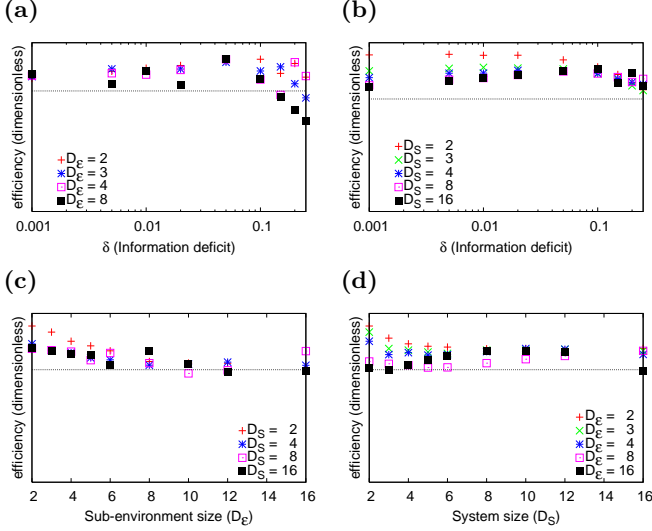


FIG. 12: (Color) “Efficiency”: **specific redundancy rescaled by information capacity**. Equation 48 provides a simple approximation for redundancy, based on the relative information capacity of the system (with a correction for δ) and its environment. We reproduce the data of Fig. 11, but use Eq. 43 to rescale specific redundancy. **Discussion:** Efficiency is consistently near 1: when the universe is in a random branching state, information about \mathcal{S} is efficiently recorded in \mathcal{E} . Equation 43 is accurate for large D_S and D_E (and small δ). When the system or the subenvironments are small, Eq. 43 underestimates information storage efficiency.

This redundant imprinting of selected observables on the environment is *quantum Darwinism*. It leads to objective reality in a quantum Universe. Typical PIPs for branching states (see Fig. 13) illustrate how different sorts of information are selected or deprecated. The information in \mathcal{E} about \mathcal{S} divides naturally into three parts.

$$\mathcal{I}_{\mathcal{S}:\mathcal{E}} = \mathcal{I}_R + \mathcal{I}_{NR} + \mathcal{I}_Q. \quad (49)$$

The *redundant information* (\mathcal{I}_R) is classical – it can be obtained easily, by many independent observers. Its selective proliferation is the essence of quantum Darwinism. Ollivier et. al. showed, in [11], that \mathcal{I}_R is not only easy to obtain, but difficult to ignore. An observer who succeeds in extracting \mathcal{I}_R , and continues to probe, finds a “classical plateau”. Measurements on additional subenvironments increase his knowledge of \mathcal{S} only slightly – mostly, they only confirm what he already knows. Only a *perfect* and *global* measurement of *everything* can reveal more than the redundant information.

Purely *quantum information* (\mathcal{I}_Q) represents observables that are incompatible with the pointer observable. This is the information that quantum Darwinism selects *against*. It is (a) encoded amongst the environments, much as a classical bit can be encoded in the parity of many ancilla bits; (b) accessible only through a global measurement on *all* of \mathcal{E} ; and (c) easily destroyed when \mathcal{E} decoheres.




FIG. 13: (Color) *Quantum Darwinism* selects certain observable properties of the system and propagates information about them throughout the environment. The preferred observable[s] become redundant **at the expense of incompatible observables**. As shown here, PIPs illustrate the results of Quantum Darwinism. Information about \mathcal{S} becomes divided into three parts: redundant information (\mathcal{I}_R), quantum information (\mathcal{I}_Q), and non-redundant information (\mathcal{I}_{NR}). Redundant information is objective, and therefore classical. It can be obtained with relative ease. Quantum information represents the non-preferred observables, marginalized by Quantum Darwinism, which can only be measured by capturing all of \mathcal{E} . Non-redundant information (determined by the slope of $\overline{\mathcal{I}}(m)$ at $m = \frac{N_{env}}{2}$) represents the ambiguous borderline, undifferentiated as yet into classical and quantum fractions. When \mathcal{I}_{NR} is small, the central region of the PIP becomes flat. This “classical plateau” indicates that an observer can obtain full information without capturing the entire environment.

Finally, *non-redundant information* (\mathcal{I}_{NR}) represents a grey area – the border between the classical and quantum domains. It exists only when the classical plateau in $\overline{\mathcal{I}}(m)$ has a nonzero slope. This is why we allow for a deficit (δ) when computing redundancy.

Information storage in randomly selected *arbitrary* states of the model universe is dramatically different from information storage in randomly selected *singly-branching* states. The contrast between these two cases emphasizes the importance of the environment’s structure. Overly simple thermodynamic arguments (e.g., maximum entropy in absence of gravity) indicate that the physical Universe should evolve into states that are uniformly distributed. Our results, however, show that objects which display the redundancy characteristic of *our* Universe must have *structured* correlations with their environments.

Decoherence theory emphasizes the role of the environment in the quantum-to-classical transition, but only as a reservoir where unwanted quantum superpositions and correlations can be hidden, out of sight. Even this view – which now seems somewhat narrow – has produced important advances in our understanding over the past quarter century. Examples include einselection, the special role of pointer states, and the view of classical-

ity as an emergent phenomenon. Nevertheless, it is clear from our discussion above and from related recent work [11, 12], that “tracing out \mathcal{E} ” obscures crucial aspects of the environment’s role.

The environment is a witness – a communication channel through which observers acquire the vast majority (if not all) of their information about the Universe. Surprisingly, this realization has taken more than 75 years since the formulation of quantum mechanics in its present form. It goes against a strong classical tradition of looking for solutions of fundamental problems in *isolated* settings. This tradition is incompatible with the role of *states* in quantum theory.

Quantum states, unlike classical states, do not define what “exists objectively”. They are too malleable – too easily perturbed and redefined by measurements. Moreover, in quantum mechanics, what is *known* about a system’s state is inextricably intertwined with what it *is*. Classical states, in contrast, have existence independently of the knowledge of them. To put it tersely (and in the spirit of complementarity), quantum states play both ontic (describing what is) and epistemic (describing what is known to be) roles [42]. Thus, for many purposes, it makes no sense to talk about a state of a completely isolated quantum system.

Our Universe is ‘quantum to the core’ (see e.g. Ref. [35] for an up-to-date review of the experimental evidence), so the only place to look for objective classicality is within the quantum theory itself. Decoherence has certainly supplied part of the answer: Only *some* of the states in an open system’s Hilbert space are stable. Those that are not stable, cannot “exist objectively”. Even these einselected pointer states, however, are vulnerable to perturbation by an observer who measures directly. Yet, objectivity implies that many different (and initially ignorant) observers can independently find out the state.

The environment-as-a-witness point of view solves this problem by recognizing that we gain essentially all of our information indirectly, from the environmental degrees of freedom (with the possible exception of specific laboratory experiments). As the environment is the “channel”, and as only a part of it can be intercepted, the obvious question is: **How is information is deposited in \mathcal{E} ? and what kind of information?**

Quantum Darwinism, which we have begun to analyse here and elsewhere [1, 6, 11, 12], aims to supply the answer. Our basic conclusion is that the redundancy evident in our Universe is not a generic property of randomly selected states in large multipartite (system plus multi-component environment) Hilbert spaces. However, when states in that Hilbert space are created by the interactions usually invoked in discussions of environment-induced superselection, redundancy appears. Thus, objectivity can arise through the dynamics of decoherence. In that sense, decoherence is the mechanism that delivers quantum Darwinism – a more complete view of classicality’s emergence.

While we have already witnessed the birth of this new

point of view, it is still far from mature. In particular, our conclusion about redundancy and the typical structure of entanglement was reached without analyzing dynamics *per se*. We have laid the foundation for a full-fledged study of quantum Darwinism by analysing kinematic properties of states, and postponed the study of evolution in specific models to forthcoming publications [7, 8]. Moreover, by employing von Neumann entropy, we have focused on the amount of information (rather than on what this information is about). Differences between various definitions of mutual information exist (see “discord”, Ref. [18]), and are symptomatic of the “quantumness” of the underlying correlations. Less “quantum” definitions of mutual information, involving conditional information, *de facto* presume a measurement. They have also been used [1, 11, 12]), along with other tools ([36, 37]), to show that the familiar pointer observables are the “fittest” in the (quantum) Darwinian sense. Studying the dynamics of quantum Darwinism, and the connections with various definitions of information, are the obvious next steps.

Acknowledgments

We thank Harold Ollivier and David Poulin for vigorous discussions. This research was supported in part by NSA and ARDA.

APPENDIX A: PROPERTIES OF QMI: THE SYMMETRY THEOREM

The *symmetry theorem* for QMI is important for understanding the shape of PIPs (partial information plots). It says, in essence, that the amount of information that can be gained from the *first* few environments to be captured, is mirrored by the amount of information that can be gained from the *last* few environments. Thus, when capturing a small fraction of \mathcal{E} yields much information, an equivalent amount of information *cannot* be gained without capturing the last outstanding bits of \mathcal{E} .

Theorem 1 (Mutual Information Symmetry Theorem). *Let the universe be in a pure state $|\psi\rangle_{\mathcal{SE}}$, and let the environment \mathcal{E} be partitioned into two chunks \mathcal{E}_A and \mathcal{E}_B . Then the total mutual information between the system and its environment is equal to the sum of the mutual informations between \mathcal{S} and \mathcal{E}_A and between \mathcal{S} and \mathcal{E}_B : that is, $\mathcal{I}_{\mathcal{S}:\mathcal{E}} = \mathcal{I}_{\mathcal{S}:\mathcal{E}_A} + \mathcal{I}_{\mathcal{S}:\mathcal{E}_B}$.*

Proof. We simply expand each mutual information as $\mathcal{I}_{x:y} = H_x + H_y - H_{xy}$, and use the fact that if a bipartite system $x \otimes y$ has a pure state $|\psi\rangle_{xy}$, then the entropies

of the parts are equal; $H_x = H_y$.

$$\begin{aligned}
\mathcal{I}_{S:\mathcal{E}_A} + \mathcal{I}_{S:\mathcal{E}_B} &= H_S + H_A - H_{SA} + H_S + H_B - H_{SB} \\
&= H_S + H_A - H_B + H_S + H_B - H_A \\
&= H_S + H_S \\
&= H_S + H_{AB} - 0 \\
&= \mathcal{I}_{S:\mathcal{E}}
\end{aligned}$$

Corollary 1. *Under no circumstances can two sub-environments both have $\mathcal{I} > H_S$ information about the system.*

If the universe is in a pure state, then the Symmetry Theorem states that any bipartite division of the environment will yield two chunks, at least one of which has $\mathcal{I} \leq H_S$. Additionally, we note that a chunk has at least as much \mathcal{I} about the system as any of its sub-chunks (that is, decreasing the size of a chunk cannot increase its \mathcal{I}). If we could find two chunks A and B with $\mathcal{I} > H_S$, then by subsuming the remainder of \mathcal{E} into A we would have a bipartite division into A' and B , each of which has $\mathcal{I} > H_S$ – but this contradicts the Symmetry Theorem.

The proof for a mixed state of the universe follows from the ‘‘Church of the Larger Hilbert Space’’ argument. We purify $\rho_{S\mathcal{E}}$ by enlarging the environment from \mathcal{E} to \mathcal{E}' , and follow the same steps to show that \mathcal{E}' cannot have two subenvironments with $\mathcal{I} > H_S$. Since \mathcal{E} is a subset of \mathcal{E}' , it too cannot have two such subenvironments.

Corollary 2. *For a pure state $|\psi\rangle_{S\mathcal{E}}$ of the universe, the partial information plot (PIP) must be antisymmetric around the point $(m = \frac{N}{2}, \mathcal{I} = H_S)$.*

This follows straightforwardly from the Symmetry Theorem. For each chunk $\mathcal{E}_{\{m\}}$ of the environment that contains m individual environments, there exists a complementary chunk $\mathcal{E}_{\{N-m\}}$, containing the complement of $\mathcal{E}_{\{m\}}$, with $N - m$ individual environments. The Symmetry Theorem implies that $\mathcal{I}_{S:\mathcal{E}_{\{m\}}} + \mathcal{I}_{S:\mathcal{E}_{\{N-m\}}} = \mathcal{I}_{S:\mathcal{E}} = 2H_S$. By averaging this equation over all possible chunks $\mathcal{E}_{\{m\}}$, we obtain an equation for the PIP: $\overline{\mathcal{I}}(m) + \overline{\mathcal{I}}(N - m) = 2H_S$. This equation is equivalent to the stated Corollary.

APPENDIX B: PERFECT STATES

The primary intuition that we obtain from the $\overline{\mathcal{I}}(m)$ plots is that most states are ‘‘encoding’’ states, but an important sub-ensemble of states are ‘‘redundant’’ states. We are naturally led to ask whether ‘‘perfect’’ examples of each type of state exist – that is, a state that encodes information more redundantly than any other state, or a state that hides the encoded information better than any other state.

The answer is somewhat surprising: whereas perfectly redundant states exist for any N and any $D_S, D_{\mathcal{E}}$, perfect coding states apparently exist only for certain N (at least

for $D_S = D_{\mathcal{E}} = 2$). The perfectly redundant states are easy to understand; they are the generalized GHZ (and GHZ-like) states of the form:

$$|\Psi_{S\mathcal{E}}\rangle = \alpha |0\rangle_S \bigotimes_i |0\rangle_{\mathcal{E}_i} + \beta |1\rangle_S \bigotimes_i |1\rangle_{\mathcal{E}_i}, \quad (\text{B1})$$

with the obvious generalizations to higher $D_S, D_{\mathcal{E}}$. Of course, it’s necessary that $D_{\mathcal{E}} \geq D_S$.

A true GHZ state is invariant under interchange of any two subsystems; however, since mutual information is invariant under local unitaries, we only require that the states $|0\rangle_{\mathcal{E}_i}$ and $|1\rangle_{\mathcal{E}_i}$ be orthogonal. Clearly, such states exist for all N . Any sub-environment with $0 < m < N$ has exactly $H(S)$ information, but only by capturing the entire environment ($m = N$) can we obtain the full $\mathcal{I} = 2H(S)$. Thus, the information is stored with N -fold redundancy.

A perfect coding state, on the other hand, would be one where $\overline{\mathcal{I}}(m) = 0$ for any $m < N/2$, and $\overline{\mathcal{I}}(m) = \mathcal{I}_{S\mathcal{E}}$ for $m > N/2$. An equivalent condition, for qubit universes, is the existence of two orthogonal states of N qubits, each of which is maximally entangled under all possible bipartite divisions. If such pairs of states exist, then the system states $|0\rangle$ and $|1\rangle$ can be correlated with them to produce the perfect coding state. It is known (as detailed in [28]) that such states only exist for $N = 2, 3, 5, 6$, and possibly for $N = 7$ (for $N = 6$, only a single state exists[38]). Thus, while for large N almost every state is an excellent coding state, perfect examples seem not to exist except for $N = 2, 3, 5, (7?)!$ We are not aware of any results for non-qubit systems.

APPENDIX C: ENTROPY OF A NEAR-DIAGONAL DENSITY MATRIX

Suppose that the pure state $\hat{\pi} = |\psi\rangle\langle\psi|$, whose components in the pointer basis are

$$\langle i|\psi\rangle = s_i, \quad (\text{C1})$$

is subjected to decoherence. The off-diagonal elements are reduced according to

$$\pi_{i,j} \longrightarrow \sigma_{i,j} = \gamma_{i,j} \pi_{i,j}, \quad (\text{C2})$$

where $\gamma_{i,i} = 1$ for all i . The limiting point of the process, where $\gamma_{i,j} = 0$ for all $i \neq j$, is ρ :

$$\rho_{i,j} = \delta_{ij} |s_i|^2. \quad (\text{C3})$$

As the $\gamma_{i,j}$ approach zero, σ converges to ρ . The partially decohered σ can be written as

$$\sigma = \rho + \Delta, \quad (\text{C4})$$

where Δ is strictly off-diagonal. Δ is defined by

$$\Delta_{i,j} = (1 - \delta_{ij}) \gamma_{i,j} s_i s_j^*. \quad (\text{C5})$$

As σ approaches ρ , its entropy approaches the entropy of ρ . Our goal here is to write $H(\sigma)$ as a power series (in Δ) around $H(\rho)$.

The entropy of σ is

$$H(\sigma) = -\text{Tr}(\sigma \ln \sigma) = \text{Tr}(\tilde{H}(\sigma)) \quad (\text{C6})$$

where

$$\tilde{H}(\sigma) \equiv -\sigma \ln \sigma. \quad (\text{C7})$$

The difference between $H(\sigma)$ and $H(\rho)$ is

$$\delta H = \text{Tr}(\delta \tilde{H}) = \text{Tr}(\tilde{H}(\rho + \Delta) - \tilde{H}(\rho)). \quad (\text{C8})$$

We will seek a power series for $\delta \tilde{H}$. Keeping in mind that its *trace* is the relevant quantity, we will discard traceless terms.

1. A naïve approach to expanding $H(\rho + \Delta)$

It's tempting to begin by expanding Eq. C7 around $\sigma = \rho$. Using the MacLaurin series for $-\sigma \ln \sigma$ gives

$$\begin{aligned} \tilde{H} &= -\Delta(\mathbb{1} + \ln \rho) - \sum_{n=0}^{\infty} \frac{(-1)^n}{(n+1)(n+2)} \frac{\Delta^{n+2}}{\rho^{n+1}} \quad (\text{C9}) \\ &\approx -\frac{\Delta^2}{2\rho} + \frac{\Delta^3}{6\rho^2} \dots \end{aligned} \quad (\text{C10})$$

We discarded the first term because it is traceless. Unfortunately, matrix quotients are not well-defined. $\frac{\Delta}{\rho}$ could mean either $\Delta\rho^{-1}$ or $\rho^{-1}\Delta$ – and, in fact, both are non-symmetric and therefore incorrect. Other symmetric orderings, such as $\rho^{-\frac{1}{2}}\Delta\rho^{-\frac{1}{2}}$, also give incorrect results. The expansion in Eq. C10 is an inappropriate generalization of a *scalar* expansion, and is ill-defined. We will take a different approach which (a) gives the correct result, and (b) defines the correct representation of matrix quotients.

2. The correct approach

Instead of expanding $\tilde{H}(\sigma)$ around $\sigma = \rho$, we expand *both* $\tilde{H}(\sigma)$ and $\tilde{H}(\rho)$ around the identity.

$$\begin{aligned} \delta \tilde{H} &= \tilde{H}(\rho + \Delta) - \tilde{H}(\rho) \\ &= \tilde{H}(\mathbb{1} - (\mathbb{1} - \rho - \Delta)) - \tilde{H}(\mathbb{1} - (\mathbb{1} - \rho)). \end{aligned}$$

The expansion around $\mathbb{1}$ is always well-defined, because $\mathbb{1}$ and its inverse commute with everything:

$$\tilde{H}(\mathbb{1} - x) = x - \sum_{n=0}^{\infty} \frac{x^{n+2}}{(n+1)(n+2)}. \quad (\text{C11})$$

Using this expansion in $\delta \tilde{H}$ yields

$$\delta \tilde{H} = -\Delta + \sum_{n=0}^{\infty} \frac{[(\mathbb{1} - \rho)^{n+2} - (\mathbb{1} - \rho - \Delta)^{n+2}]}{(n+1)(n+2)}. \quad (\text{C12})$$

We once again discard Δ because it is traceless, leaving only the sum. The two matrix powers within the sum can be rewritten using the identity

$$(\mathbb{1} + x)^n = \sum_{j=0}^n \binom{n}{j} x^j, \quad (\text{C13})$$

which yields

$$\delta \tilde{H} = - \sum_{n=0}^{\infty} \sum_{j=0}^{n+2} (-1)^j \binom{n+2}{j} ((\rho + \Delta)^j - \rho^j). \quad (\text{C14})$$

In order to simplify this, we must introduce a new notation. Consider $(x + y)^p$, where x and y may be either scalars or matrices. For scalar x and y ,

$$(x + y)^p = \sum_{k=0}^p \binom{p}{k} a^k b^{p-k}, \quad (\text{C15})$$

whereas for matrices, $\binom{p}{k} x^k y^{p-k}$ is replaced by a sum over $\binom{p}{k}$ orderings of k x 's and $p - k$ y 's. We define the notation $x^k \circ y^{p-k}$ to describe this sum: e.g.,

$$x^2 \circ y^2 = \frac{x^2 y^2 + xyxy + xy^2 x + yx^2 y + yxyx + y^2 x}{6}, \quad (\text{C16})$$

but when x and y are scalars

$$x^2 \circ y^2 = x^2 y^2. \quad (\text{C17})$$

Using this definition of a *totally symmetric product*,

$$(\rho + \Delta)^j = \sum_{k=0}^j \binom{j}{k} \Delta^k \circ \rho^{j-k}, \quad (\text{C18})$$

and the entropy difference operator $\delta \tilde{H}$ is

$$\delta\tilde{H} = -\sum_{n=0}^{\infty} \sum_{j=0}^{n+2} \sum_{k=0}^{j-1} \frac{(-1)^j}{(n+1)(n+2)} \binom{n+2}{j} \binom{j}{k+1} \Delta^{k+1} \circ \rho^{j-k-1} \quad (\text{C19})$$

$$= -\sum_{k=1}^{\infty} \sum_{n=0}^{\infty} \sum_{j=0}^n \frac{(-1)^{j+k+1}}{(n+k)(n+k+1)} \binom{n+k+1}{j+k+1} \binom{j+k+1}{k+1} \Delta^{k+1} \circ \rho^j \quad (\text{C20})$$

$$-\sum_{n=0}^{\infty} \sum_{j=0}^{n+1} \frac{(-1)^j}{(n+1)(n+2)} (j+1) \binom{n+2}{j+1} \Delta \circ \rho^j \quad (\text{C21})$$

The $k = 0$ term can be discarded because $\text{Tr}(\Delta \circ \rho^j) = \text{Tr}(\Delta \rho^j) = 0$. We then perform the sum over j to obtain

$$\delta\tilde{H} = -\sum_{k=1}^{\infty} \sum_{n=0}^{\infty} \frac{(-1)^k}{(n+k)(n+k+1)} \binom{n+k+1}{k+1} \Delta^{k+1} \circ (\mathbb{1} - \rho)^n. \quad (\text{C22})$$

Expanding the binomial coefficients and simplifying leads to the following result:

$$\delta\tilde{H} = \sum_{k=1}^{\infty} \frac{(-1)^k}{k(k+1)} \Delta^{k+1} \circ \sum_{n=0}^{\infty} \binom{k+n-1}{n} (\mathbb{1} - \rho)^n. \quad (\text{C23})$$

We have come full circle. The sum over n in Eq. C23 is just the MacLaurin expansion for ρ^{-k} around $\rho = \mathbb{1}$. Equation C23 can thus be written *symbolically* as

$$\delta\tilde{H} = \sum_{k=1}^{\infty} \frac{(-1)^k}{k(k+1)} \Delta^{k+1} \circ [\rho^{-k}], \quad (\text{C24})$$

if the symmetric product $\Delta^{k+1} \circ \rho^{-k}$ is interpreted as “take the symmetric product of Δ^{k+1} with the power series representing ρ^{-k} .”

Essentially, what we have derived is the “correct” interpretation of the matrix quotient $\frac{\Delta^{k+1}}{\rho^k}$. This result is interesting in its own right, but for now we are interested only in the leading order (i.e., Δ^2) term. Truncating the series at $k = 1$, we obtain the following simple result:

$$\delta H \approx -hf \sum_{n=0}^{\infty} \text{Tr} [\Delta^2 \circ (\mathbb{1} - \rho)^n] + O(\Delta^3). \quad (\text{C25})$$

This is the simplest possible *general* form for δH . In order to perform the traces, we need to take advantage of the form of the symmetric product.

From the definition of the symmetric product, we can write out explicit expressions for $\Delta^k \circ M^n$, for particular small values of k .

$$\Delta \circ M^n = \frac{1}{n+1} \sum_{p=0}^n M^p \Delta M^{n-p} \quad (\text{C26})$$

$$\Delta^2 \circ M^n = \frac{2}{(n+1)(n+2)} \sum_{p=0}^n \sum_{q=0}^{n-p} M^q \Delta M^p \Delta M^{n-p-q} \quad (\text{C27})$$

The second case (for Δ^2) is the useful one. We need the trace of the symmetric product, which can be simplified using the cyclic property of trace,

$$\text{Tr} [\Delta^2 \circ M^n] = \frac{1}{n+1} \sum_{p=0}^n \text{Tr} [\Delta M^p \Delta M^{n-p}]. \quad (\text{C28})$$

Together with Eq. C25, this formula yields an *explicit* expression for δH :

$$\delta H \approx -\frac{1}{2} \sum_{n=0}^{\infty} \frac{1}{n+1} \sum_{p=0}^n \text{Tr} [\Delta (\mathbb{1} - \rho)^p \Delta (\mathbb{1} - \rho)^{n-p}] \quad (\text{C29})$$

We now insert specific forms for ρ and Δ , from Eqs. C3 and C5:

$$\text{Tr} [\Delta M^p \Delta M^{n-p}] = \sum_{i,j,k,l=0}^{D_S-1} \Delta_{ij} (\mathbb{1} - \rho)_{jk}^p \Delta_{kl} (\mathbb{1} - \rho)_{li}^{n-p} \quad (\text{C30})$$

$$= \sum_{i,j,k,l=0}^{D_S-1} s_i s_j^* s_k s_l^* \gamma_{ij} \gamma_{kl} \delta_{jk} \delta_{il} (1 - |s_j|^2)^p (1 - |s_i|^2)^{n-p} \quad (\text{C31})$$

$$= \sum_{i,j \neq i} |s_i|^2 (1 - |s_i|^2)^{n-p} |s_j|^2 (1 - |s_j|^2)^p |\gamma_{ij}|^2. \quad (\text{C32})$$

Since the goal is to average over an ensemble of states, we replace $|\gamma_{ij}|^2$ with an average, $\overline{|\gamma|^2}$,

$$\begin{aligned} \text{Tr} [\Delta M^p \Delta M^{n-p}] &= \overline{|\gamma|^2} \left[\sum_i (|s_i|^2 (1 - |s_i|^2)^p) \sum_j (|s_j|^2 (1 - |s_j|^2)^{n-p}) \right. \\ &\quad \left. - \sum_k (|s_k|^4 (1 - |s_k|^2)^n) \right] \\ &= \overline{|\gamma|^2} [\text{Tr} [\rho(\mathbb{1} - \rho)^p] \text{Tr} [\rho(\mathbb{1} - \rho)^{n-p}] - \text{Tr} [\rho^2(\mathbb{1} - \rho)^n]] \end{aligned} \quad (\text{C33})$$

Inserting this expression into Eq. C29 yields

$$\delta H \approx -\frac{\overline{|\gamma|^2}}{2} \sum_{n=0}^{\infty} \frac{1}{n+1} \sum_{p=0}^n [\text{Tr} [\rho(\mathbb{1} - \rho)^p] \text{Tr} [\rho(\mathbb{1} - \rho)^{n-p}] - \text{Tr} [\rho^2(\mathbb{1} - \rho)^n]]. \quad (\text{C34})$$

Finally, we can simplify this expression slightly by (1) taking advantage of the identity $\sum_{n=0}^{\infty} (\mathbb{1} - \rho)^n = \rho^{-1}$, and (2) rearranging the summation variables.

$$\delta H \approx -\frac{\overline{|\gamma|^2}}{2} \left[\sum_{n=0}^{\infty} \sum_{p=0}^n \frac{\text{Tr} [\rho(\mathbb{1} - \rho)^p] \text{Tr} [\rho(\mathbb{1} - \rho)^{n-p}]}{n+1} - \sum_{n=0}^{\infty} \text{Tr} [\rho^2(\mathbb{1} - \rho)^n] \right] \quad (\text{C35})$$

$$= -\frac{\overline{|\gamma|^2}}{2} \left[\sum_{n=0}^{\infty} \sum_{p=0}^{\infty} \frac{\text{Tr} [\rho(\mathbb{1} - \rho)^p] \text{Tr} [\rho(\mathbb{1} - \rho)^n]}{n+p+1} - \text{Tr} \left[\rho^2 \sum_{n=0}^{\infty} (\mathbb{1} - \rho)^n \right] \right] \quad (\text{C36})$$

$$= -\frac{\overline{|\gamma|^2}}{2} \left[\sum_{n=0}^{\infty} \sum_{p=0}^{\infty} \frac{\text{Tr} [\rho(\mathbb{1} - \rho)^p] \text{Tr} [\rho(\mathbb{1} - \rho)^n]}{n+p+1} - 1 \right] \quad (\text{C37})$$

Equation C37 is the simplest form we have been able to achieve, except in very special cases, for $H(\rho + \Delta) - H(\rho)$.

APPENDIX D: PROBABILITY DISTRIBUTIONS FOR ADDITIVE DECOHERENCE FACTORS

If $|\psi\rangle$ and $|\psi'\rangle$ are selected at random from the uniform ensemble of $D_{\mathcal{E}}$ -dimensional quantum states, then the probability that $|\langle \psi | \psi' \rangle| = \gamma$ (for $\gamma \in [0 \dots 1]$) is

$$p(\gamma) = 2(D_{\mathcal{E}} - 1)\gamma(1 - \gamma^2)^{D_{\mathcal{E}}-2} \quad (\text{D1})$$

The additive decoherence factor d is given by $d = -\log(\gamma)$, so that $\gamma = e^{-d}$ and $d \in [0 \dots \infty]$. The probability distribution transforms as

$$\begin{aligned} p(d) \mathbf{d}d &= p(\gamma) \mathbf{d}\gamma \\ p(d) &= p(\gamma) \left| \frac{\mathbf{d}\gamma}{\mathbf{d}d} \right| \\ &= e^{-d} p(\gamma) \\ &= 2(D_{\mathcal{E}} - 1) e^{-2d} (1 - e^{-2d})^{D_{\mathcal{E}}-2} \end{aligned} \quad (\text{D2})$$

The decoherence factor for a collection of subenvironments is simply the sum of $d^{(i)}$ over the contributing subenvironments. Ideally, we could obtain exact distributions $p_m(d)$ for a sum of m such d -factors. For an environment composed of qubits ($D_{\mathcal{E}} = 2$), $p(d)$ is a 1st-order Poisson distribution, so $p_m(d)$ is just the m th order Poisson distribution (for details, see [20]).

For larger subenvironments ($D_{\mathcal{E}} > 2$), no such simple description exists. However, the distribution functions $p(d)$ are well-approximated by Gaussian distributions. We can treat the summing problem as a biased random walk, where the addition of another subenvironment represents a step forward with an approximately Gaussian-distributed stepsize.

To compute the mean and variance of an m -step random walk, we first compute the mean value \bar{d} and variance $\Delta d = \sqrt{(\overline{d^2} - \bar{d}^2)}$ for a single subenvironment. Ex-

trapolating to a collection of m systems requires setting $\bar{d}_m = m\bar{d}$ and $\Delta d_m = \sqrt{m}\Delta d$.

For a single subenvironment, the mean \bar{d} is given by $\bar{d} = \int_0^\infty dp(d)dl$. This integral is somewhat nontrivial, involving an expansion in binomial coefficients:

$$\begin{aligned} \bar{d} &= 2(D_\mathcal{E} - 1) \int_0^\infty de^{-2d} (1 - e^{-2d})^{D_\mathcal{E}-2} dl \\ &= 2(1 - D_\mathcal{E}) \int_0^\infty de^{-2d} \sum_{k=0}^{D_\mathcal{E}-2} \binom{D_\mathcal{E}-2}{k} (e^{-2kd}) dl \\ &= 2(1 - D_\mathcal{E}) \sum_{k=0}^{D_\mathcal{E}-2} (-)^k \binom{D_\mathcal{E}-2}{k} \int_0^\infty d (e^{-2(k+1)d}) dl \\ &= \frac{D_\mathcal{E} - 1}{2} \sum_{k=0}^{D_\mathcal{E}-2} \frac{(-)^k (D_\mathcal{E} - 2)!}{(k+1)^2 k! (D_\mathcal{E} - 2 - k)!} \\ &= \frac{1}{2} (\Psi(D_\mathcal{E}) + \gamma_{\text{EM}}) \end{aligned} \quad (\text{D3})$$

where $\Psi(D_\mathcal{E})$ is the digamma function, and $\gamma_{\text{EM}} = 0.5772\dots$ is the Euler-Mascheroni constant. A virtually identical calculation for \bar{d}^2 yields

$$\Delta d^2 = \frac{\pi^2}{24} - \frac{\Psi_1(D_\mathcal{E})}{4} \quad (\text{D4})$$

in terms of the trigamma function $\Psi_1(D_\mathcal{E})$.

-
- [1] W. H. Zurek, *Reviews of Modern Physics* **75**, 715 (2003).
[2] J. P. Paz and W. H. Zurek, *Les Houches Summer School Session* **72**, 535 (2001).
[3] E. Joos, H. D. Zeh, C. Kiefer, D. Giulini, J. Kupsch, and I.-O. Stamatescu, *Decoherence and the appearance of a classical world in quantum theory* (New York: Springer, 2003), 2nd ed.
[4] M. Schlosshauer, *Reviews of Modern Physics* **76**, 1267 (2004).
[5] J. A. Wheeler and W. H. Zurek, *Quantum Theory and Measurement* (Princeton University, Princeton, NJ, 1983).
[6] W. H. Zurek, *Annalen der Physik* **9**, 855 (2000).
[7] R. Blume-Kohout and W. H. Zurek (2005), in prep.
[8] R. Blume-Kohout and W. H. Zurek (2005), in prep.
[9] W. H. Zurek, *Progress of Theoretical Physics* **89**, 281 (1993).
[10] J. von Neumann, *Mathematical Foundations of Quantum Mechanics* (Princeton University, Princeton, NJ, 1955).
[11] H. Ollivier, D. Poulin, and W. H. Zurek, *Physical Review Letters* **93**, 220401 (2004).
[12] H. Ollivier, D. Poulin, and W. H. Zurek, arxiv.org/quant-ph **04**, 0408125 (2004).
[13] W. H. Zurek, *Philosophical Transactions of the Royal Society, Series A* **356**, 1793 (1998).
[14] W. H. Zurek, arxiv.org/quant-ph **03**, 0308163 (2003).
[15] T. H. Cover and J. A. Thomas, *Elements of Information Theory* (Wiley-Interscience, 1991).
[16] W. H. Zurek, *Physical Review D* **26**, 1862 (1982).
[17] W. H. Zurek, in *Quantum Optics, Experimental Gravitation, and Measurement Theory*, edited by P. Meystre and M. O. Scully (Plenum, New York, 1983), p. 87.
[18] H. Ollivier and W. H. Zurek, *Physical Review Letters* **88**, 017901 (2002).
[19] P. Zanardi, D. A. Lidar, and S. Lloyd, *Physical Review Letters* **92**, 060402 (2004).
[20] R. Blume-Kohout and W. H. Zurek, arxiv.org/quant-ph **04**, 0408147 (2004).
[21] D. N. Page, *Physical Review Letters* **71**, 1291 (1993).
[22] S. Sen, *Physical Review Letters* **77**, 1 (1996).
[23] S. K. Foong and S. Kanno, *Physical Review Letters* **72**, 1148 (1994).
[24] J. Sanchez-Ruiz, *Physical Review E* **52**, 5653 (1995).
[25] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, 2000).
[26] E. Knill and R. Laflamme, *Physical Review A* **55**, 900 (1997).
[27] D. Gottesman, *Physical Review A* **57**, 127 (1998).
[28] A. J. Scott, *Physical Review A* **69**, 052330 (pages 10) (2004).
[29] C. E. Shannon and W. Weaver, *The Mathematical Theory of Communication* (University of Illinois Press, Urbana, 1949).
[30] W. H. Zurek, *Physical Review D* **24**, 1516 (1981).
[31] H. Everett, *Reviews of Modern Physics* **29**, 454 (1957).
[32] E. W. Weisstein, in *Mathworld - A Wolfram Web Resource* (Wolfram Research, Inc., 1999-2005), <http://mathworld.wolfram.com/DigammaFunction.html>.
[33] E. W. Weisstein, in *Mathworld - A Wolfram Web Resource* (Wolfram Research, Inc., 1999-2005), <http://mathworld.wolfram.com/TrigammaFunction.html>.
[34] W. H. Zurek, arxiv.org/quant-ph p. 0405161 (2004).
[35] M. Schlosshauer, arxiv.org/quant-ph p. 0506199 (2005).
[36] D. A. R. Dalvit, J. Dziarmaga, and W. H. Zurek, *Physical Review Letters* **86**, 373 (2001).
[37] D. A. R. Dalvit, J. Dziarmaga, and W. H. Zurek, arxiv.org/quant-ph p. 0509174 (2005).
[38] A. R. Calderbank, E. M. Rains, P. W. Shor, and N. J. A. Sloane, *IEEE Transactions on Information Theory* **44**, 1369 (1998).
[39] We must make the *right* measurement – in this case, one which distinguishes $|0\rangle$ from $|1\rangle$ – in order to get the information. In this work, the amount of information that one subenvironment has is always *maximized* over all possible measurements.
[40] In this work, we assume that the universe is in a pure state. Any correlation between \mathcal{S} and \mathcal{E} is due to en-

tanglement. Similar conclusions seem to apply when the environment is initially mixed, but we have not investigated these cases exhaustively.

[41] The GHZ state in Eq. 2 is the exception that proves the rule. Such states are measure-zero in Hilbert space. Per-

fect C-NOT interactions are required to make them.

[42] See [1] and especially [34] for further discussion of quantum states' *epiontic* nature.