Identity Localization: Numerical Sketches

exploratory figures for atomic photon absorption

Jonathan R. Landers

Notation guide. The word "alphabet" means the list of possible photon alternatives in the experiment. If the photon could be one of $N = 16$ distinguishable time bins, polarizations, or modes, then the label set has size $N = 16$. A balanced set of N perfectly distinguishable alternatives carries $\log_2 N$ bits of identity. The symbol A denotes the part of the atom we inspect, and D denotes every other degree of freedom we ignore: recoil, emitted fields, unobserved motional modes, or any external record. The dimensions d_A and d_D count how many orthogonal quantum states those two spaces can hold. The entropy $S(\rho)$ is the von Neumann entropy, measured in bits, and $h_2(p)$ is the binary entropy.

Model. Think of the incoming photon as carrying a label x: which polarization, time bin, spectral mode, or other distinguishable alternative arrived. There are N equally likely labels, so x ranges over N possible values. The absorption process is modeled as a perfectly reversible map into the full output $A'D$, where A' is the final atom after absorption and D is everything we do not read out:

$$|x\rangle \longmapsto |\Psi_x\rangle_{A'D}, \qquad \langle\Psi_x|\Psi_y\rangle = \delta_{xy}.$$

Here $|x\rangle$ is the incoming photon label state, $|\Psi_x\rangle_{A'D}$ is the corresponding final joint state, and δ_{xy} is 1 when the labels are the same and 0 otherwise. The inner product $\langle\Psi_x|\Psi_y\rangle$ measures whether two final global states are distinguishable. The orthogonality condition says the full universe still knows which label arrived. Nothing is globally erased. The question is local: how much of that label can be recovered from the atom alone? To answer that, we trace out D:

$$\rho_A^x = \mathrm{Tr}_D\, |\Psi_x\rangle\langle\Psi_x|.$$

The state ρ_A^x is what the atom alone looks like after label x was absorbed. The trace operation Tr_D means "ignore D and keep only the atomic state." The retained internal identity is the Holevo information

$$\chi_A = S\left(\frac{1}{N}\sum_x \rho_A^x\right) - \frac{1}{N}\sum_x S(\rho_A^x), \qquad \eta_A = \frac{\chi_A}{\log_2 N}.$$

Here χ_A is the number of label bits readable from the atom, and η_A is the fraction of the incoming identity budget that localized in the atom. If $\eta_A = 1$, the atom is a perfect identity memory. If $\eta_A = 0$, the atom's final state carries no information about which photon alternative was absorbed.

Since A has dimension d_A,

$$\chi_A \le \log_2 d_A, \qquad \eta_A \le \min\left(1, \frac{\log_2 d_A}{\log_2 N}\right).$$

The simulations below sample Haar-random isometries into $A'D$. That means: choose a generic reversible absorption map, uniformly at random, with no special memory design built in. The figures compare what such generic maps usually do with what the dimension bound says is possible in principle.

Interpretive warning. These are not laboratory data and not a model of a particular atom. They are numerical thought experiments. Their job is to show the geometry of the identity accounting: conserved globally, partly or mostly lost from the atom when D is ignored.

Figure 1: one bit, one overlap. This is the clean analytic case from the paper: two possible incoming labels and two final atomic states.

What is being simulated. We vary only the overlap $c_A = |\langle e_0|e_1\rangle|$ between the two final atomic states $|e_0\rangle$ and $|e_1\rangle$. These are the two possible states of the atom after absorbing label 0 or label 1. No random channels are involved here.

How to read it. The horizontal axis measures how similar the two atomic outputs are. At $c_A = 0$ they are orthogonal and easy to tell apart. At $c_A = 1$ they are identical and impossible to tell apart from the atom alone. The vertical axis is the retained identity fraction.

Point of the plot. Identity localization is just distinguishability. If the atom ends in two orthogonal states, it remembers which label arrived. If both labels drive the same atomic state, the atom absorbed the energy but forgot the label.

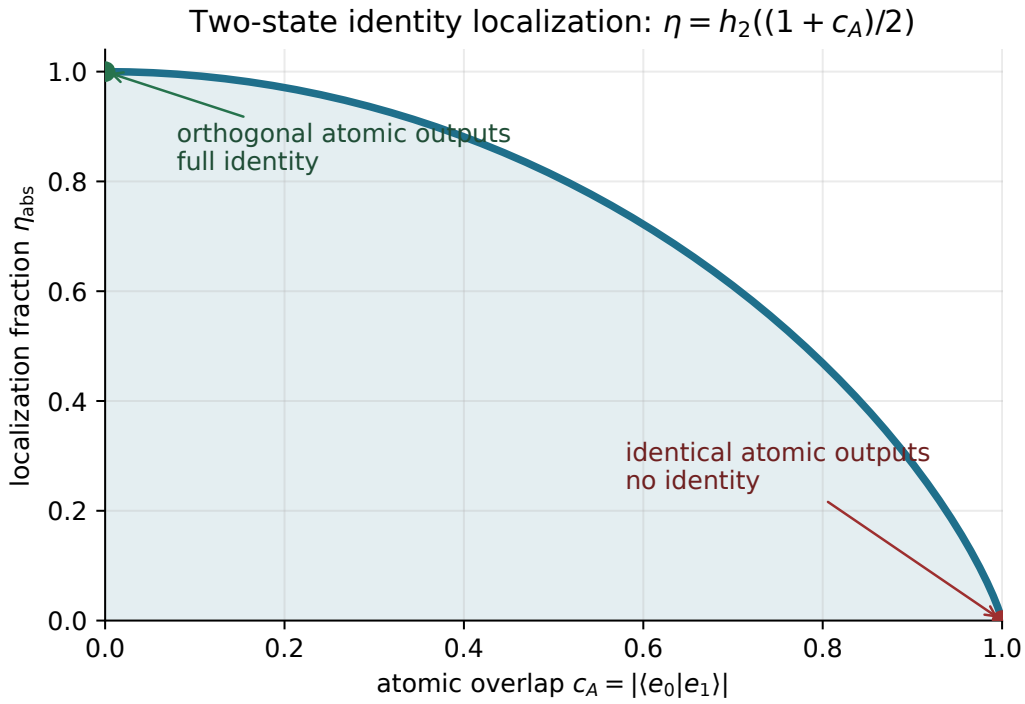

Figure 1: **Analytic two-state curve.** For a binary label, atomic overlap alone fixes the retained identity: $\eta = h_2((1 + c_A)/2)$, where η is the same retained fraction as η_A in this two-state case and h_2 is the binary entropy. Orthogonal outputs keep the bit; identical outputs erase it from the atom.

Figure 2: the best possible finite memory. This plot is a capacity calculation, not a random simulation. It asks how much identity an ideal atom could retain if its internal Hilbert space had dimension d_A.

What is being simulated. For each number N of possible photon labels, the incoming identity budget is $\log_2 N$ bits. The atom can carry at most $\log_2 d_A$ bits internally.

How to read it. Each curve is one internal dimension d_A. When the curve equals 1, the atom has enough capacity to store the whole label. Once the curve falls below 1, there are too many possible photon labels for the atom's internal manifold.

Point of the plot. This is the no-go in visual form. A two-level atom can hold one internal bit. It can perfectly localize a binary label, but not a 16-way label. The missing identity must live outside the internal atom.

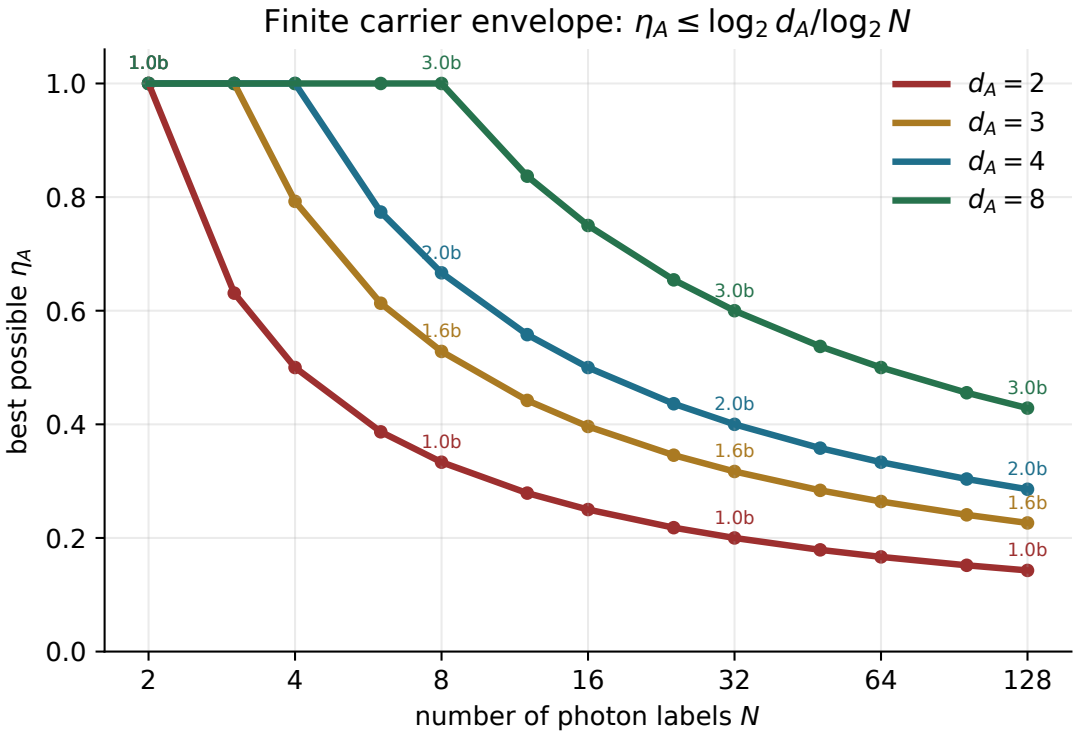

Figure 2: **Best possible finite-carrier envelope.** Once $N > d_A$, perfect internal localization is impossible. Even an ideal internal memory can keep at most $\log_2 d_A$ bits, so the fraction falls like $\log_2 d_A / \log_2 N$.

Figure 3: generic absorption is much worse than ideal memory. The previous figure showed what an ideal internal carrier could do. This one asks what a generic reversible absorption map usually does.

What is being simulated. For each N and d_A, we sample many random isometries $|x\rangle \mapsto |\Psi_x\rangle_{A'D}$ with discarded dimension fixed at $d_D = 16$. In plain terms: we try many generic reversible maps from photon labels into atom plus unobserved degrees of freedom. For every sample we trace out D and compute η_A.

How to read it. Solid curves are median retained fractions from random samples. Shaded bands show typical spread. Dashed curves are the dimension-envelope limits from Figure 2.

Point of the plot. The bound is a ceiling, not the typical behavior. A random coupling usually spreads the label into D in a way that leaves the atom with very little distinguishing information. A good quantum memory is therefore highly structured, not generic.

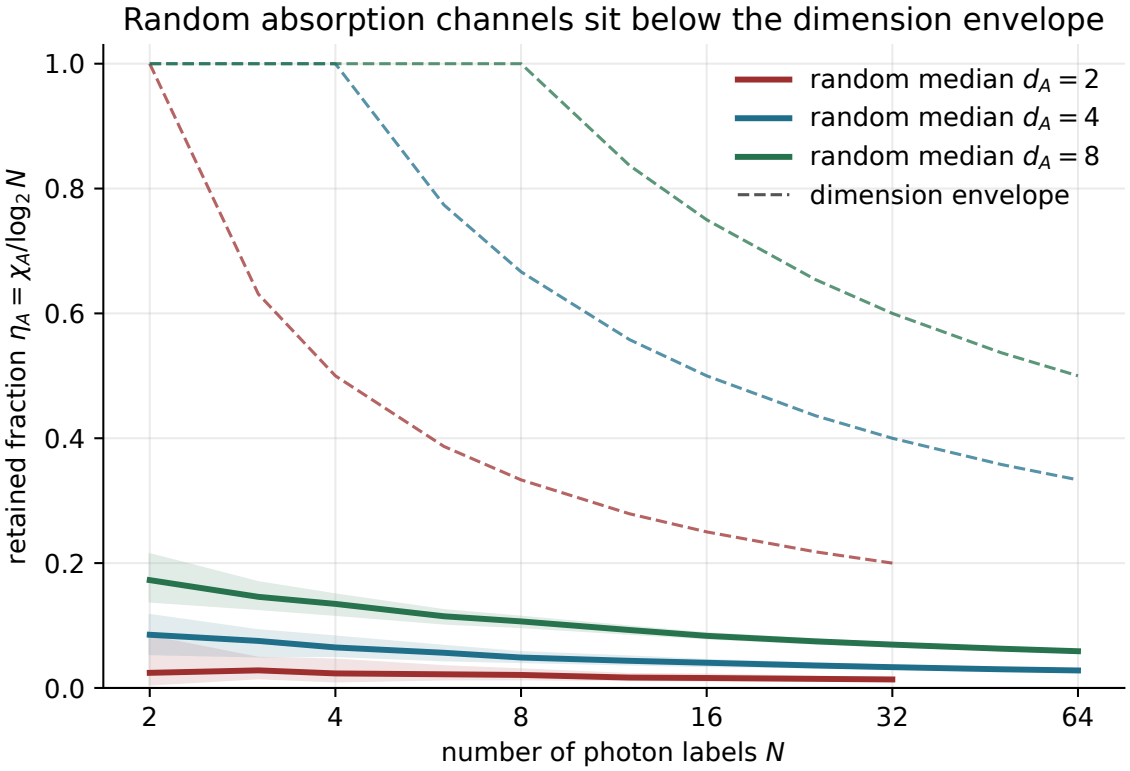

Figure 3: **Random absorption channels.** Haar-random isometries usually sit well below the ideal envelope. The dashed curves are dimension bounds; shaded bands show the central 80% of sampled retained fractions.

Figure 4: dimension pressure across many cases. This heatmap compresses the random-channel experiment into one table.

What is being simulated. Rows change the atom's internal dimension d_A. Columns change the number N of possible photon labels. Each cell averages η_A over random isometries with $d_D = 16$.

How to read it. Brighter cells mean the atom retained more of the label. Darker cells mean most identity moved into D. Blank cells are cases where the total output dimension $d_A d_D$ was too small to embed all N orthogonal labels.

Point of the plot. Bigger atoms retain more; larger sets of possible photon labels are harder to retain. But even when the dimension bound would allow more, generic maps still retain only a small fraction.

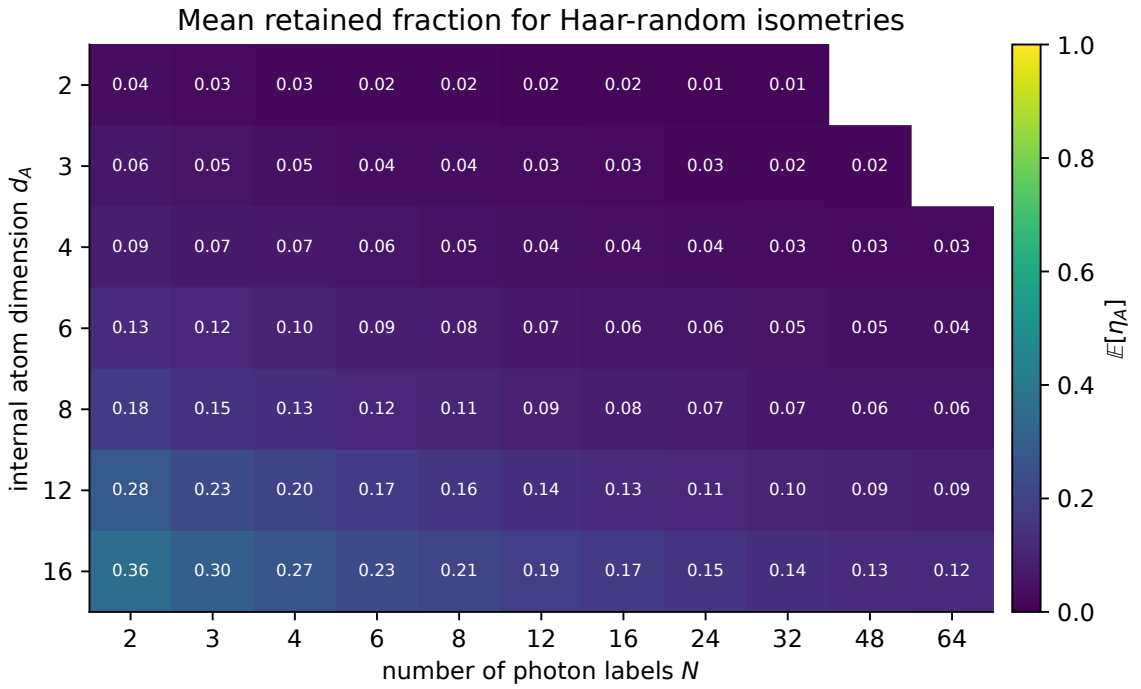

Figure 4: **Dimension pressure.** Mean retained fraction increases with internal dimension d_A and decreases as the number N of possible photon labels grows. The plot is not a new bound; it shows what a random absorption channel typically does inside the allowed region.

Figure 5: the bookkeeping line. This plot shows that the simulations respect the same accounting identity as the paper.

What is being simulated. We fix $N = 16$, so the incoming label budget is $\log_2 16 = 4$ bits. For each random channel, we compute how many bits the atom retains, χ_A, and define the rest as leakage, $L_A = 4 - \chi_A$.

How to read it. Every point is one random absorption map. The diagonal line is the identity budget: retained bits plus leaked bits equals four. Vertical dashed lines show the internal capacity $\log_2 d_A$.

Point of the plot. The atom and discarded modes divide a fixed conserved total. Moving right means the atom remembers more; moving up-left means more identity is outside the atom. Internal dimension caps how far right the atom can go.

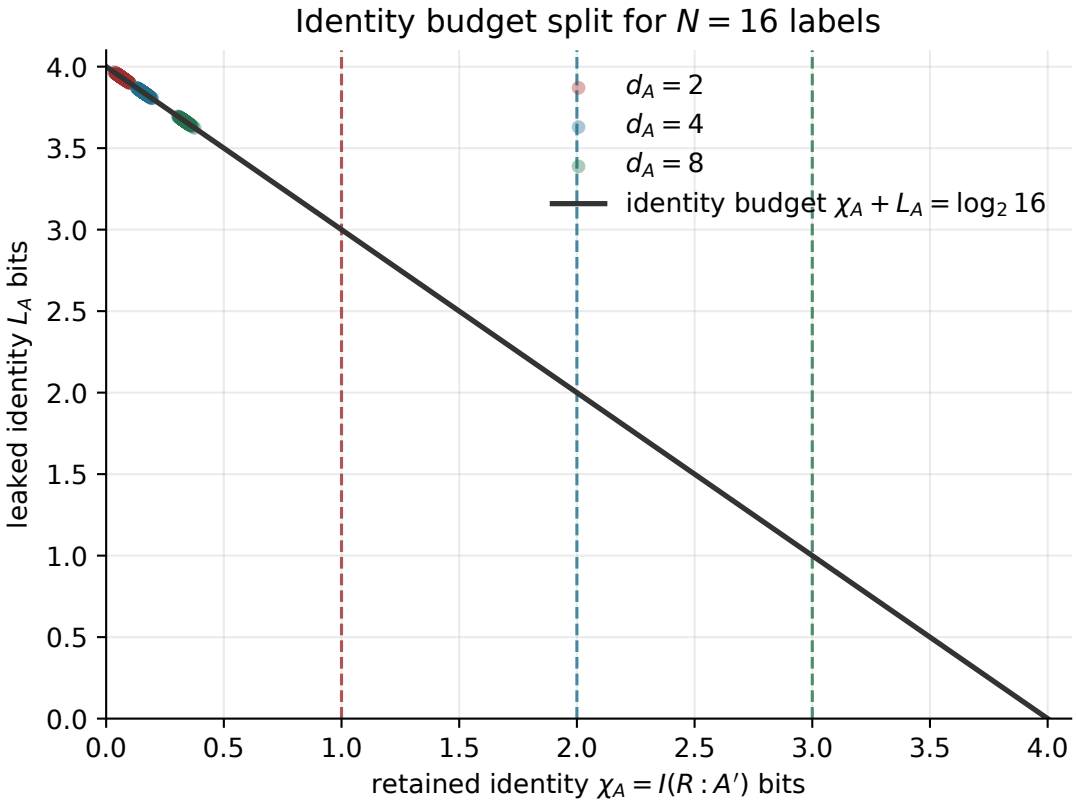

Figure 5: **Identity budget split.** For orthogonal labels, the full absorbed identity budget is $\mathcal{J}_{\text{abs}} = \log_2 N$ bits. The atom retains χ_A and the complement is leakage: $L_A = \log_2 N - \chi_A$. Random channels move along this accounting line while the internal dimension caps how far right they can go.

Figure 6: bigger unobserved space makes the atom less informative. This is the clearest numerical stress test of the role of D.

What is being simulated. We hold the label size fixed at $N = 16$ and vary the discarded dimension d_D. The atom dimension d_A is fixed along each curve.

How to read it. The horizontal axis is how much unobserved space the absorption map can use. The vertical axis is the mean retained identity fraction in the atom.

Point of the plot. For generic maps, enlarging D gives the identity more room to delocalize outside the atom. The full output still contains the label perfectly, but the atom alone becomes less diagnostic of which label arrived.

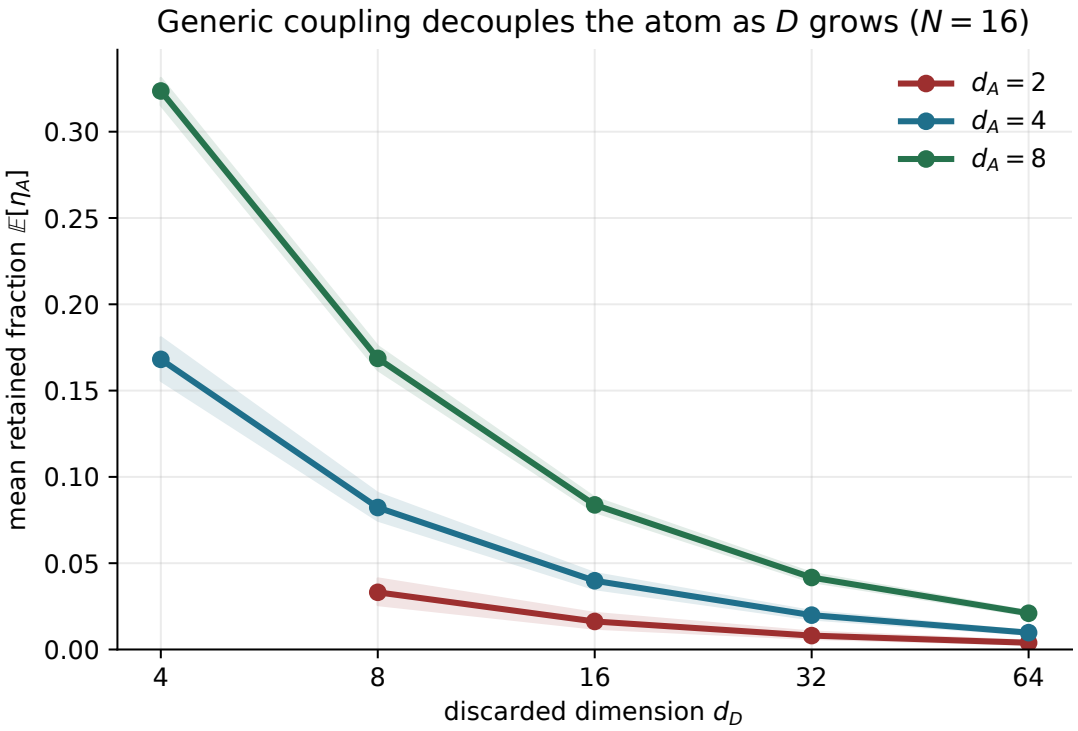

Figure 6: **Generic coupling is not a memory.** With $N = 16$ fixed, increasing the discarded dimension d_D makes the reduced atomic states more alike. The atom retains less label information even though the full $A'D$ output still contains the complete identity.

Core takeaway. The simulations do not alter the paper's theorem. They make its geometry visible: global identity is conserved by the isometry, but finite internal dimension and generic coupling leave most high-N identity outside the atom. The atom can absorb the energy while the reference label spreads into D.