

The Passport Is Somewhere Safe: Security Is Bought Only With Retrievability, and the Retrieval Aid That Had to Fail

Elsbeth Nunan¹, T. Aurelio Banning¹, Ifeoma Achterberg², and The Standing Panel on Deposited Property³

¹Whitlock Centre for Retrieval Studies, Sanderholm University

²Department of Estate Inventory Science, Cranmoor Institute

³Office of Custody Statistics

Received 3 April 2026; revised 22 July 2026; accepted 1 September 2026

Abstract

The passport is somewhere safe. Everyone in the household agrees on this, nobody can produce it, and both facts follow from the same act of placement. We report the first study of domestic and institutional safekeeping to observe its own denominator, recovering 24,180 placements by exhaustive excavation of premises cleared by a party other than the placer. Recovery does not fall with storage duration; it falls with the care taken at placement. Expecting to repair this, we built and trialled a household Placement Register in which the location is written down as it is chosen. Recovery in the register arm fell (odds ratio 0.71, 95% CI 0.58–0.86), and it fell because the register licensed safer placements rather than better records. We therefore abandon the account we set out to defend. Security and retrievability are one quantity: the $\Lambda \approx 9.7$ bits of placement entropy admitted by a median dwelling, partitioned at the instant of placement and never afterwards enlarged. Every bit withheld from a thief is withheld from the owner, in the same act. A perfectly secure placement is therefore an unrecoverable one, and the same frontier governs hotel safes and national archives alike. “Somewhere safe” is not a location. It is a disposal instruction, and it is always carried out.

Keywords: safekeeping; placement entropy; custody administration; risk compensation; dark logic models

1 Introduction

A household places a passport somewhere safe. Four years later the passport is required, and the household conducts a search which is exhaustive, ill-tempered, and unsuccessful. The episode is universally recognised and almost universally misfiled. It is treated as a lapse of memory, which is to say as a defect in the person rather than a property of the act, and the remedy prescribed is always the same: be more careful, and write it down.

The prescription has never been evaluated. This is partly a measurement problem. The literature on domestic loss samples on the dependent variable — it counts objects that people report having lost, which is to say objects whose absence was noticed [6, 2]. Placements that

were never missed do not appear, and there is no reason to think they are a small share. A study of safekeeping that cannot observe its own denominator can describe the frustration but not the process.

We set out to fix the prescription. Our intended contribution was a retrieval aid: an instrument that would restore to the placer the information that deliberate placement appears to destroy, validated against an excavation-based measure of what is actually recoverable. §3.4 and §4.2 report that instrument and the randomised trial in which it made matters worse. §4.3 onward reports what we found instead, which is not the paper we intended to write.

Three contributions follow. First, an excavation design (§3.1) that observes placements irrespective of whether anyone noticed their loss, yielding 24,180 recovered placements across household removals, deceased-estate clearances, and end-of-tenancy strip-outs. Second, a two-arm randomised trial of the Placement Register (§4.2) whose primary outcome moved in the wrong direction, with pre-specified mediation identifying why. Third, the accounting identity those results force (§4.3), its instantiation as a dark logic model (§5), and its extension to institutional custody (§6), where the regime with the highest measured security is also the one that has most completely disposed of what it holds.

2 Related work

Memory for object location has been studied principally under intentional encoding with short retention intervals, where recall is good and the reported effects are small [9, 7]. The domestic case inverts both conditions: retention intervals run to decades and the encoding, though deliberate, is not rehearsed. Fenn and Whitlock [5] noted the discrepancy and proposed a decay account, which we test and reject in §4.1.

The hiding-place literature is closer in spirit. Work on concealment by householders anticipating burglary consistently finds that placers choose locations by maximising dissimilarity from the object’s habitual position [1], and that the same locations recur across unrelated households at rates far above chance [8]. The second finding is

Table 1: Composition of the Excavation Panel. A stream qualifies only if every admissible location is opened on a known date by a party other than the placer.

Stream	Placem.	Never missed	Median dur. (yr)
Household removals	14,902	46.1%	6.2
Deceased-estate clearance	6,411	63.3%	14.7
End-of-tenancy strip-out	2,867	34.0%	2.9
Total	24,180	49.2%	7.4

usually read as a warning to householders. We read it as evidence that the admissible-placement set is small, enumerable, and shared, which is what makes Λ estimable at all.

Risk compensation is the mechanism our trial recovers, and it is well characterised elsewhere: a safety measure that lowers the cost of a hazardous choice reliably increases the rate of that choice, often enough to absorb the measure's benefit entirely [10, 11]. To our knowledge it has not previously been applied to storage. Custody administration has developed its own parallel literature on deposit attrition, largely internal and largely unpublished; the tabulations in [3] are the exception and we rely on them in §6.

Finally, our §5 follows the dark logic model of Dwyer and colleagues [4], which asks an evaluator to specify, before deployment, the pathways by which a well-intentioned intervention might produce the opposite of its intent. We did not build one before deploying the Placement Register. We build one here, and it accounts for the trial result completely. The order of operations was the wrong one.

3 Data and instruments

3.1 The Excavation Panel

A *placement* enters the panel when premises are cleared exhaustively, on a known date, by a party other than the placer, such that every admissible location is opened whether or not anything is thought to be in it. Three recruitment streams satisfy this condition and are summarised in Table 1.

Excavation is what allows the denominator to be observed. A removals crew emptying a dwelling opens the tin inside the drawer inside the wardrobe regardless of whether the household believes anything is in it, and records what was there. Of the 24,180 placements recovered, 11,902 (49.2%) had never been reported missing by anyone, and in 4,106 cases the placer, shown the object, could not state where it had been before being told.

For each placement we record the object class; container nesting depth (drawer $\rightarrow$ tin $\rightarrow$ envelope = 3); metric distance from the object's habitual location; elapsed storage duration; the placer's stated recall, elicited before disclosure; and the reported-missing indicator.

Enumeration of the admissible set follows the protocol of [12]. A location is admissible if two of three independent coders, working from the cleared floor plan and




Figure 1: The excavation design. Λ is fixed by the premises; S and R are elicited from different people by different instruments, so the additivity in Proposition 1 is a testable claim rather than a definition.

a fixed inventory of furniture types, agree that an adult could deliberately put a passport there. Coders never see the recovered placement. Agreement was substantial (Fleiss $\kappa = 0.74$), and the enumerated sets are strikingly consistent between dwellings of similar size, which is the empirical content of the claim in [8] that concealment is convergent. Fixtures that cannot be opened without tools are excluded; the loft is admissible, the cavity wall is not.

3.2 Measuring security: the Blind Searcher Protocol

Figure 1 sets out the design. Let a premises admit N *admissible placements*: locations that a competent adult, asked where one *could* put a passport, would enumerate. Enumeration is performed by the excavation team from the cleared floor plan and is independent of where the object actually was.

Definition 1 (Security). $S = \log_2 m$, where m is the median number of admissible locations opened by a searcher with no knowledge of the placement, matched to the placer on age and dwelling familiarity, proceeding under a fixed protocol until recovery.

Definition 2 (Retrievability). $R = \Lambda - H(p)$, where p is the placer's own posterior over the N admissible placements, elicited as a ranked allocation before any search, and $\Lambda = \log_2 N$.

Both are in bits, both are bounded by Λ , and neither is elicited from the same person. We ran 3,902 blind searches; inter-searcher reliability on repeated placements was $\kappa = 0.81$.

3.3 The Institutional Custody Series

To establish that the phenomenon is not a property of untrained placers, we obtained deposit-level recovery records from six custody regimes whose declared retention periods we could match [3] ($n = 9,044$). These are reported in §6.




Figure 2: Recovery against storage duration. The decay account requires the dashed curve. The panel gives the flat one.

3.4 The Placement Register trial

Two-arm, household-randomised, 18-month follow-up, $n = 1,164$ households (582/582), stratified by dwelling size and occupancy. The treatment is a bound *Placement Register*: a pre-printed household log in which the placer records the object and its location at the moment of placement. Adherence was high — logs were filed for 88% of placements subsequently observed — so what follows is not a null produced by non-compliance.

The primary outcome was the proportion of registered placements produced on demand at 18 months. The pre-registered prediction, derived from the distinctiveness account of §4.2, was that the register would raise it.

4 Three models of placement

4.1 Model A: decay, and why it fails

The folk account, formalised by [5], holds that the object remains where it was put and the memory of it fades, so that retrieval is a race against elapsed time. It predicts that recovery falls with storage duration.

It does not (Figure 2). Across the panel the duration slope is -0.004 per year (95% CI -0.011 to 0.003); placements twenty years old are produced at very nearly the rate of placements two years old.

What recovery does track is *placement deliberation time*: the interval, recorded by the placer or reconstructed from the removals record, between picking the object up and setting it down. Each additional minute of deliberation costs 0.23 bits of R (95% CI 0.18–0.28). The predictor is not the clock. It is the care.

4.2 Model B: distinctiveness, and the trial that refuted it

The second account is better. Placers select locations by maximising dissimilarity from habitual position [1], and a future searcher — including the placer — assigns such locations low prior probability. On this account the bits are simply never written down, and a written record should repair the deficit outright.

This is testable, and it is the instrument we built. The result is Figure 3(a): recovery in the register arm was

lower than control, odds ratio 0.71 (95% CI 0.58–0.86), $p = 0.0006$.

Pre-specified mediation (Figure 3b) identifies the pathway. Register-arm placements were measurably safer: nesting depth $+0.62$ levels (95% CI 0.41–0.83), distance from habitual location $+1.9$ m, and blind-searcher security $+1.4$ bits. The register did not add information. It licensed the placer to spend information they would otherwise have retained — the placer, knowing the location was recorded, chose a location they would not otherwise have dared. This is risk compensation [10], and Model B has no term for it: it cannot represent a treatment that improves encoding and reduces recovery, because it has no channel by which the placer’s behaviour responds to the aid.

4.3 Model C: the custody budget

Both failures point the same way. Consider what a placement actually does. The location of an object is a random variable over the N admissible placements, carrying $\Lambda = \log_2 N$ bits. The placer’s advantage over an adversary is exactly the mutual information between the placement and the placer’s memory of it, and the placement act allocates that information once, between two parties, in a single decision.

Proposition 1 (Additivity of custody). *For any placement,*


$$S + R = \Lambda = \log_2 N,$$

with Λ fixed at the instant of placement and never subsequently enlarged.


Sketch. Let X be the placement, uniform on the admissible set, so $H(X) = \Lambda$, and let M be the placer’s memory trace. The chain rule gives $H(X) = I(X; M) + H(X | M)$ exactly, and the placer’s retained advantage is $I(X; M)$, which is R . The substantive step is the identification of S with $H(X | M)$: a searcher who is not the placer holds the enumerated set and nothing further, so the residual uncertainty a placement imposes on every party other than the placer is precisely the part of $H(X)$ that M does not carry. A placement that gains a bit of security must move X into a region the searcher weights less, and the placer distinguishes that region only through M ; the increment is therefore drawn from $I(X; M)$ and from nowhere else. Substituting, $S + R = \Lambda$. $\square$

The identity is testable, because S and R are elicited from different people by different instruments (Figure 4). Regressing $S + R$ on $\hat{\Lambda}$ across the panel gives a slope of 0.98 (95% CI 0.94–1.02) and an intercept indistinguishable from zero, with $R^2 = 0.91$. We estimate $\hat{\Lambda} = 9.7$ bits (95% CI 9.4–10.0) for a median dwelling: about 830 places a passport could be.

Model C accounts for all three observations. Nothing decays, because the bits were never held (§4.1). Deliberation is expensive because deliberation *is* the act of spending (§4.1). And the register made things worse because it raised the ceiling on affordable security without raising Λ , which is fixed by the premises (§4.2).



(a) Primary outcome. OR 0.71 (95% CI 0.58–0.86), $p = 0.0006$.



(b) Pre-specified mediation: blind-searcher security, distance from the habitual location, and container nesting depth. Every safety term moved up; recovery moved down.

Figure 3: The Placement Register trial. The register was issued to raise recovery and lowered it; (b) reports the mediation.




Figure 4: Additivity of custody. Each point is one recovered placement, with S measured on a blind searcher and R on the placer. Grey lines are the iso- Λ family; the bold line is the median dwelling. Regressing $S + R$ on $\hat{\Lambda}$ gives slope 0.98 (95% CI 0.94–1.02), $R^2 = 0.91$. Triangles are domestic keys in median dwellings, which overshoot the identity by 1.6 bits. We do not explain them.

Three robustness checks were pre-specified. Restricting to placements whose admissible set was enumerated by all three coders unanimously moves the slope to 0.97. Dropping the deceased-estate stream, the least representative of the three, moves it to 1.01. Substituting a rank-based estimator of R for the allocation-based one moves it to 0.96. The identity is not an artefact of any one instrument, and it survives substitution of weaker ones.

Corollary 1 (Disposal). $S = \Lambda \implies R = 0$. A perfectly secure placement is unrecoverable.

Corollary 1 is the model's content rather than a pathology of it. Safekeeping does not trade off against retrieval at the margin; it is the same operation described from the other side. An object placed with maximum care has been disposed of, and the disposal is complete at the moment of placement. What remains is only the delay before anyone establishes it.

Corollary 2 (Premises). $\Lambda = \log_2 N$ is increasing in N ,

so larger premises carry larger budgets and can afford both terms.

This holds: recovery rises with $\log_2 N$ at $\beta = 0.31$ per bit. It does not help. Placement rate rises with N at very nearly the compensating rate, and total items lost per household is flat across dwelling size. Households in larger dwellings do not lose fewer things. They lose the same things more safely.

5 A dark logic model of safekeeping

Following [4], we specify in Figure 5 the pathways by which safekeeping produces the outcome it is adopted to prevent. All four canonical pathways are instantiated, which is unusual.

Perverse incentive. The placement is rewarded immediately, in felt security, and punished later, diffusely, at retrieval. The placer optimises the term available to them at the moment of choice, and that term is S .

Displacement. The cost does not fall on the placer. In the deceased-estate stream, 71% of excavated placements were recovered by a party who did not know the object existed. Safekeeping exports its loss in time, and frequently past the placer's own death.

Dependency trap. Households practising deliberate safekeeping lose the habitual-location competence that made retrieval work without it. When register arm logs were withdrawn at month 18, those households recovered at rates below their own pre-trial baseline.

Backlash. This one is terminal. A completed Placement Register is a document that states, for every entry, exactly where the household's valuables are. It collapses S to near zero across its whole contents. The register works — that is precisely the difficulty — and it works by destroying the property it was issued to protect. It cannot be distributed, and a household that has kept one faithfully cannot safely retain it.

We note, without proposing a remedy, that the register must itself be placed somewhere safe.




Figure 5: A dark logic model of safekeeping, after [4]. All four canonical pathways are instantiated. The fourth is terminal: the only instrument that can restore retrievability is a written record of every placement, and such a record is precisely what the placements were made to withstand.

Table 2: The Institutional Custody Series, ordered by measured security. Regimes are matched on declared retention period. S is the blind-searcher measure of §3.2, applied to staff not on the depositing shift.

Regime	Deposits	S (bits)	Returned
Hotel in-room safe	2,410	4.1	94.7%
Left luggage, staffed	1,988	5.0	91.1%
School office	1,502	5.6	86.8%
Hospital property store	1,447	7.2	74.2%
Bank safe-deposit box	1,102	9.4	59.6%
National archive, closed stack	595	14.8	10.4%
Total / range	9,044	4.1–14.8	$r = -0.99$

6 Institutional custody

If the identity is real it should not care whether the placer is a householder or an institution whose entire function is custody. Table 2 and Figure 6 report six regimes.

Measured security and recovery rate are collinear across the series at $r = -0.99$. The ordering is the one the identity requires and the one professional intuition inverts: the hotel in-room safe, the least secure regime in the set, returns property most reliably, and the national closed-stack archive, which achieves the highest S we have measured anywhere, approaches the disposal limit of Corollary 1 more closely than anything else in the series. Its holdings are secure to a standard that no other institution in the series approaches. On the account developed here they have been disposed of, and the finding aids are a record of the disposal rather than a means of reversing it.

The archive is the case in which the identity is least intuitive and most cleanly satisfied. Closed-stack holdings are protected by exactly the measures the identity charges for: restricted access, indirection through a catalogue, physical separation of the finding aid from the item, and a staff rotation that ensures no individual retains a durable memory of where anything is. Each is a delib-




Figure 6: Six custody regimes on the same frontier ($r = -0.99$). The dashed line is the least-squares fit; the archive is administered best.

erate transfer from R to S , and the transfers are audited, minuted, and regarded as achievements. The institution has done everything correctly. It has, in consequence, the lowest return rate in the series by a factor of nearly six.

We stress that no regime in the series is administered badly. Each is administered well, and the identity does not distinguish between competent and incompetent custody. It says only that custody is expenditure, and that an institution which spends everything on security has nothing left to spend on return.

7 Discussion

We did not set out to argue this. The Placement Register was the intended contribution and it failed, in a direction we had not thought to pre-specify, and the identity in Proposition 1 is what remained once the two accounts we preferred had been eliminated.

An anomaly we cannot explain. Domestic keys violate additivity. Across 2,214 recovered key placements, $S + R$ exceeds $\hat{\Lambda}$ by a mean of 1.6 bits, which the identity does not permit. The excess survives every specification

we have tried, including re-enumeration of the admissible set by an independent team, and it is not present in any other object class. We have no account of it and we report it because the alternative is to conceal it.

Limitations. Four are material. First, Λ is estimated from an admissible-placement set enumerated by the excavation team rather than by the placer; if placers entertain locations the team does not enumerate, Λ is understated and additivity is being fitted against a moving denominator. We cannot exclude this. Second, the register trial cannot be blinded — a household knows whether it has been given a bound log — so expectancy effects are confounded with the mechanism we claim. Third, the estate stream samples placements made by people who subsequently died, which is a different population from placements at large, and we have not modelled how the two differ. Fourth, the keys anomaly is unexplained, and we make no claim that the identity is exact until it is resolved.

What follows administratively. If safekeeping is a disposal method then it is currently the only one practised at scale without record, consent, or notification of the eventual finder. Deposit regimes already regulate destruction; they do not regulate custody as destruction. We suggest the distinction is not sustainable, and that the reclassification is a matter of bookkeeping rather than of principle.

8 Conclusion

Security and retrievability are drawn from one account, and the account is settled at the moment the object is set down. Nothing is later lost, because nothing was later held; the bits denied to a thief were denied to the owner in the same motion, and the household's four-year search was for information that was spent before the drawer closed. The instrument we built to recover it made matters worse for the reason it had to, and every custody regime we measured, from the hotel safe to the national archive, sits on the same frontier in the same order.

"Somewhere safe" is not a location. It is a disposal instruction, and it is always carried out.

Acknowledgements

We thank the removals and estate-clearance firms who permitted observation of their working practice, and the households who allowed a stranger to open every drawer they own. The Placement Registers issued during the trial have been recalled and destroyed.

References

- [1] Arbuthnot, R. and Nkemdirim, S. (2018). Dissimilarity maximisation in householder concealment: evidence from 604 pre-burglary interviews. *Journal of Residential Security*, 31(2), 145–168.
- [2] Bekele, T. (2021). What the missing-item literature is missing. *Review of Household Studies*, 12(4), 401–419.
- [3] Committee on Deposited Property (2024). *Deposit attrition across six custody regimes: consolidated tabulations 1998–2023*. Office of Custody Statistics, Report CS-2024-11.
- [4] Dwyer, C., Halloran, P. and Weiss, M. (2019). Dark logic models: specifying the pathways to harm before deployment. *Evaluation and Programme Theory*, 44(3), 275–299.
- [5] Fenn, A. and Whitlock, J. (2015). Decay accounts of long-interval object localisation. *Memory and Cognition Quarterly*, 58(1), 33–57.
- [6] Harrower, D. (2019). Sampling on the dependent variable in studies of domestic loss. *Methods in Applied Enumeration*, 7(2), 88–104.
- [7] Laskaris, E. and Vondra, M. (2020). Intentional encoding of object position under short retention. *Cognitive Reports*, 22(6), 611–630.
- [8] Marisol, B., Achterberg, I. and Prynne, K. (2022). Convergent concealment: the forty-one recurring locations. *Journal of Residential Security*, 35(1), 1–29.
- [9] Oduya, F. (2017). Location memory in naturalistic settings: a review. *Annual Review of Applied Memory*, 9, 213–248.
- [10] Peltzman, S. (1975). The effects of automobile safety regulation. *Journal of Political Economy*, 83(4), 677–726.
- [11] Serrano, L. and Duffield, H. (2019). Compensating behaviour after protective adoption: a meta-analysis of 88 trials. *Risk Analysis Letters*, 14(2), 159–181.
- [12] Nunan, E. and Banning, T. A. (2025). Excavation as a sampling frame for household property. *Methods in Applied Enumeration*, 13(3), 240–266.