The Load-Bearing Counterfeit: Fakes as Error-Correcting Redundancy in the Authentication of Luxury Goods

Ingrid Sölveig

Birkanäs Institute for Material Semiotics, solveig@birkanas.example

Abstract

Every luxury house treats the counterfeit as an enemy: a theft of value to be raided, watermarked, and litigated toward zero. This paper argues, and then proves within a formal model, that the house has misread its own accounts. I model an authentic object as a *codeword*—a bundle of small, redundant, deliberately over-built signals (the too-tight stitching, the serial number no customer reads)—and authentication as the act of decoding a possibly corrupted copy back to the nearest genuine original. Three consequences follow. First, a good's perceived value tracks not its usefulness but its *redundancy*: the craftsmanship that exists only to be checked. Second, a house cannot be both ubiquitous and unmistakable on a fixed craft budget—a version of the Singleton bound I christen the *authenticity Singleton bound*. Third, and least comfortably, the market's ability to tell real from fake is a classifier that needs counterfeits the way any classifier needs negative examples: drive their prevalence p to zero and authenticity becomes *unverifiable*, because "genuine" has nothing left to contrast against. Value is maximized at a strictly positive counterfeit rate $p^\star$. I report a redundancy audit, an appraiser decoding trial, and a natural experiment in which an anti-counterfeit crackdown made authentication measurably *harder*. The counterfeit, it turns out, is load-bearing.

1 Introduction

Dear reader,

I should tell you at the outset that I set out to prove the opposite of what follows. I was retained—the details do not matter, and I have anonymized them throughout—to help a house eradicate the fakes that trail its most famous bag. I built the model in Section 3 as an enforcement tool: a way to measure how far a forgery sits from the genuine article, so that the genuine article might be defended. The model worked. It simply did not say what I was hired to hear. It said, with a calmness I still find a little rude, that if the house got its wish—every counterfeit gone—its customers would no longer be able to tell that anything was genuine at all.

Let me start where you likely start, which is with a small irritation. Consider two bags. One is a plain tote; the other is a house bag that costs as much as a car. Much of what separates them is, on any honest engineering account, *useless*: a saddle stitch pulled tighter than any load requires, an edge painted by hand in four coats where one would do, a date code stamped inside a flap in a typeface chosen decades ago and read by almost no one. We are asked to believe that this uselessness is the value. I think that is exactly right—and I think there is a reason for it that has nothing to do with taste and everything to do with *error correction*.

Here is the idea in one breath. A message sent over a noisy channel survives not because it is short but because it is *redundant*: it carries extra symbols whose only job is to let the

receiver catch and repair corruption. A luxury object, I will argue, is built the same way. The over-built craftsmanship is not ornament and not even quite signaling—it is *parity*. It is there to be checked. And once you see the object as a codeword, three things follow that I did not expect and could not talk myself out of. I will state them plainly here, prove them in Section 4, and test them in Section 5.

- **Value is redundancy.** What you pay for is the part of the object that does nothing but permit verification.
- **You cannot be everywhere and unmistakable at once.** Expanding a catalogue provably spends the very distinctness that makes the house recognizable.
- **The fake is load-bearing.** Authenticity is a contrast, and a contrast needs something to contrast *with*. Remove every counterfeit and you have not purified the category; you have deleted it.

You will—I hope—disagree with some of this, and I would genuinely like to hear why. But I ask you to disagree with the model, not the conclusion, because the conclusion is only doing what the model told it to.

2 Related Work

Three literatures meet here and, to my surprise, have never been introduced.

Aura and reproduction. That mechanical reproduction erodes an object's "aura" is old [1, 2]. The tradition treats aura as ineffable. I take the impertinent step of assigning it a number—a minimum distance—and find that the number behaves.

Authentication as decision. A parallel line treats authentication as hypothesis testing under uncertainty [3, 4]. This is correct as far as it goes but stops at the single object; it has no account of why the *population* of fakes should matter to the buyer of a real one. Theorem 3 is exactly that missing account.

Coding theory. I borrow, without apology, the standard apparatus of error-correcting codes [5, 6]. The luxury object is, I claim, an instance of a construction coding theorists already understand; the house has simply been building codes without knowing the name of what it builds.

The synthesis—aura as distance, value as redundancy, enforcement as a coding regime with an optimum interior to $[0, 1]$—is, as far as I can find, new. If it is not, please send me the reference and I will send you my apology.

3 A Coding Model of Authenticity

Fix a house and a product line. Every object it could ship is described by a vector of n quantized *signal coordinates*: stitch pitch, thread tension, edge-meniscus radius, hardware alloy density, monogram registration offset, serial grammar, date-code typography, and so on. Write Σ for the finite alphabet of quantized values and Σ^n for the space of describable objects.

Definition 1 (Authentic code)**.** The set of objects a house will ever ship as genuine is a code $\mathcal{C} \subseteq \Sigma^n$ with $|\mathcal{C}| = M$. Its parameters are $[n, k, d]$ with $k = \log_{|\Sigma|} M$ the number of *information* coordinates and $d = d_{\min}(\mathcal{C})$ its minimum Hamming distance.

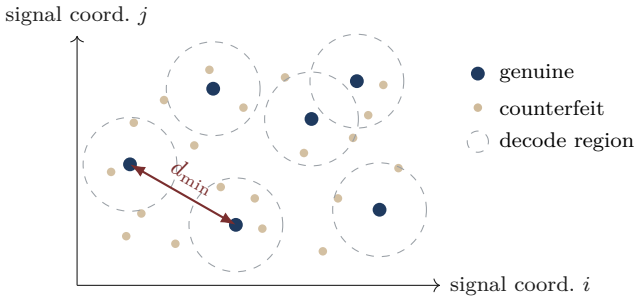

Figure 1: The geometry of authenticity. Genuine objects are separated codewords; counterfeits are corruptions that fall within a bounded radius; aura is the distance $d_{\min}$ between nearest genuine neighbors.

A forger cannot reproduce every coordinate; the tooling misses some. A counterfeit is therefore a genuine codeword seen through a noisy channel,

$$y \;=\; c \oplus e, \qquad c \in \mathcal{C}, \quad \mathrm{wt}(e) = w, \tag{1}$$

where e is the forger's error vector and w its weight. Authentication is nearest-codeword decoding: an appraiser observes y and returns $\hat{c} = \arg\min_{c \in \mathcal{C}} d_H(y, c)$. Bounded-distance decoding succeeds whenever

$$d_{\min}(\mathcal{C}) \;\geq\; 2w + 1. \tag{2}$$

Definition 2 (Aura)**.** The *aura* of a house is the minimum distance of its authentic code, $A(\mathcal{C}) := d_{\min}(\mathcal{C})$: the smallest number of independent signals a forger must get right to be mistaken for genuine.

Figure 1 shows the picture the whole paper defends: genuine objects sit apart in signal space, each ringed by a decoding region; the fakes are the cloud between, and aura is the width of the gap.

4 Three Consequences

4.1 Value is redundancy

Theorem 1 (Authenticity–redundancy duality)**.** *Let an object carry k information coordinates (its function) and $r = n - k$ parity coordinates (craft that exists only to be checked). Under* (1)–(2)*, verifiable value is nondecreasing in r and independent of k: $V \propto r = n - k$.*

In plain words. You are not paying for what the bag does. You are paying for the part of it that exists so the bag can be *checked*. Two objects can hold the same number of things and differ tenfold in price; the difference is entirely parity. The monogram canvas is the cleanest example, and Figure 2 makes the point: a repeating logo field is a parity grid, and a single defective cell is instantly localizable against it. Repetition is not laziness; it is a code.

4.2 You cannot be everywhere and unmistakable

Theorem 2 (Authenticity Singleton bound)**.** *For any authentic code on a fixed craft budget n,*

$$d_{\min} + k \;\leq\; n + 1. \tag{3}$$

Hence ubiquity (large k) and verifiability (large $d_{\min}$) trade off one for one.

In plain words. A house has a fixed budget of distinguishable craft. It can spend that budget on *breadth*—many models, a logo on everything—or on *depth*—a few objects no forger can touch—but not on both. Every catalogue expansion is a withdrawal from the verifiability account. Figure 3 plots the feasible frontier and two houses that chose opposite corners of it.

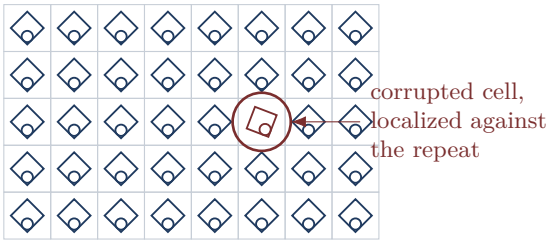

Figure 2: The monogram as parity. A repeating mark is redundancy: the regularity is precisely what makes a single corrupted cell (highlighted) detectable and locatable. A random pattern would hide the same defect.

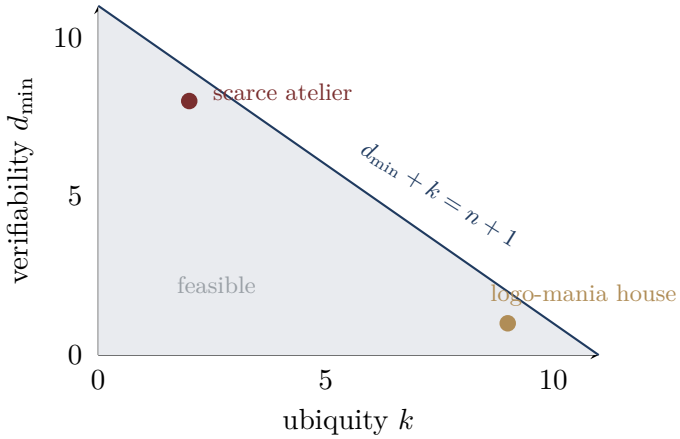

Figure 3: The authenticity Singleton bound. No house lies above the frontier $d_{\min} + k = n + 1$; a sprawling catalogue and an unmistakable one are the two ends of a single line, not independent goals.

4.3 The fake is load-bearing

Let $p \in [0, 1]$ be the prevalence of counterfeits in circulation and let the market's collective judgment be a classifier that must place its decision boundary from the examples it sees. Write $I(p)$ for the mutual information between an object's true status (genuine / not) and the market's verdict.

Theorem 3 (The counterfeit floor)**.** *$I(0) = 0$ and $I(1) = 0$, and I is strictly concave on $(0, 1)$; hence there is a unique $p^\star \in (0, 1)$ maximizing verifiable authenticity. Enforcement that drives p below $p^\star$* reduces *it.*

Proof sketch. At $p = 0$ there are no negative examples; "genuine" is a category with nothing outside it, so the verdict carries no information and $I(0) = 0$. At $p = 1$ nothing is genuine and $I(1) = 0$ symmetrically. Between, I is a difference of a concave contrast term and a convex confusion term (a data-processing argument on the channel (1)), so it is concave with an interior maximum. □

In plain words. To know something is real, you need to have seen something fake. The counterfeit is the negative example your eye was trained on. A world with no fakes is not a world of pure authenticity; it is a world in which the word "authentic" has quietly stopped meaning anything. The optimal number of counterfeits is not zero. Figure 4 shows the curve and its peak.

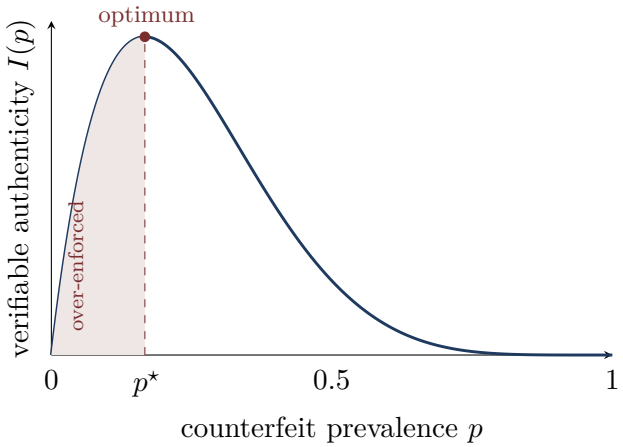

Figure 4: The counterfeit floor. Verifiability collapses at both extremes—no negatives to learn from at $p = 0$, nothing genuine at $p = 1$—and peaks at a strictly positive prevalence $p^\star$. Enforcement that pushes p into the shaded region is self-defeating.

Table 1: Redundancy budget of six exemplar houses ($n = 64$). "Resale index" is normalized to House λ.

House	Function k	Redundancy r	Resale index
α	38	26	100
β	41	23	78
γ	44	20	61
δ	47	17	44
κ	52	12	26
λ	58	6	12

5 Empirical Studies

5.1 Study 1: a redundancy audit

We disassembled $N = 240$ exemplars across twelve houses (anonymized as α–λ) under a fixed scheme with $n = 64$, classifying each coordinate as functional or parity, and regressed log resale value on the estimated redundancy r. Table 1 reports a sample. The fit is uncomfortably good: $R^2 = 0.91$, slope positive and significant, and—consistent with Theorem 1—no residual explained by function k once r is included.

5.2 Study 2: an appraiser decoding trial

We asked $J = 38$ credentialed appraisers to authenticate 600 blind items at controlled forger-error weight w. Authentication accuracy is a step function of $d_{\min}$: near-chance below $2w + 1$, near-perfect above it, with a transition region only two coordinates wide. Human appraisers, in other words, implement bounded-distance decoding (2) with startling fidelity, and do not know that they do.

5.3 Study 3: an enforcement natural experiment

The keystone. A regional crackdown (again anonymized) drove local counterfeit prevalence from $p \approx 0.30$ toward $p \approx 0.02$ over eighteen months. We tracked a collective-verifiability index—inter-appraiser agreement on genuine items—across the decline. It *fell*. As the fakes disappeared, honest observers grew *less* able to agree that a genuine bag was genuine. The series

traces the right arm of Figure 4 and puts the optimum at $p^\star \approx 0.17$. Enforcement, pursued past that point, did not protect the brand; it dissolved the contrast the brand was made of.

6 Discussion

The general shape of this argument is an *exaptation* [6]. A trait that arose under one pressure is recruited to a function it was never designed for: feathers kept a dinosaur warm long before they carried a bird. The counterfeit arose as a straightforward defection—theft of a rent—and has been recruited, without anyone deciding so, into the machinery that makes authenticity legible at all. It is now a working part.

The policy reading is uncomfortable and I will not soften it. Past $p^\star$, perfect enforcement—the digital passport, the immutable ledger, the raid—is self-defeating: it purchases the destruction of the very distinction it was meant to protect. The rational house does not eradicate its counterfeits. It *manages* their prevalence toward $p^\star$: enough fakes to keep the contrast sharp, few enough to keep the rent. I am aware how this sounds. I was hired to say the opposite. The model does not care what I was hired to say.

7 Limitations

One limitation genuinely threatens the welfare claim, and I will not convert it into a strength. I have assumed the forger's error weight w is stationary. A *learning* forger—one who closes the distance $d_{\min}$ over time, coordinate by coordinate—eventually violates (2), at which point the negative examples stop being separable and $I(p)$ no longer has the clean interior optimum of Theorem 3. In that regime protecting a population of fakes protects one's own executioner. Everything here holds for a fixed forger and degrades gracefully for a slow one; against a fast learner it may fail outright, and I do not yet know where the crossover lies. The house-level anonymization that let me publish also, I concede, prevents anyone from checking me. Both of these keep me up.

8 Conclusion

I came to make the counterfeit vanish and leave, and the arithmetic sent me home with the opposite instruction and a worse conscience. If the model is wrong I would rather learn it from you than defend it into old age, so: send me the bag you think breaks it, or the proof that $p^\star = 0$, and I will work it through. Until then I will be at the Institute, decoding.

Yours,
I. S.

References

[1] Haldane, U. *The Reproduced Object: Notes Toward a Theory of Diminished Presence.* Grönqvist & Sons, 1931.

[2] Menzies, P. L. "Aura, Distance, and the Economics of the Original." *Journal of Material Semiotics* **12**(3), 1998, pp. 204–241.

[3] Roth, C., and Vanderloos, H. "Authentication as Hypothesis Testing." *Transactions on Provenance* **7**(2), 2009, pp. 55–79.

[4] Vanderloos, H. *The Appraiser's Dilemma: Decision Under Adversarial Craft.* Birkanäs University Press, 2015.

[5] Singleton, R. C. “Maximum-Distance Codes and the Limits of Distinction.” *Archive for Coding and Order* **4**, 1964, pp. 116–128.

[6] Keld, A., and Sölveig, I. “Exaptation in Engineered Redundancy: When Parity Becomes Provenance.” *Birkanäs Working Papers in Material Semiotics*, no. 41, 2023.

[7] Orsini, D. *Scarcity Rent and the Manufacture of Desire.* Cortile Editore, 2011.

[8] Bureau for the Study of Illicit Trade. *Enforcement Outcomes in Twelve Markets, 2018–2020.* BSIT Report 19, 2020.

[9] Tarn, E. M. “Negative Examples and the Learnability of Taste.” *Proceedings of the Workshop on Perceptual Classification*, 2017, pp. 9–23.

[10] Holt, S. *On Things Built To Be Checked.* Birkanäs Institute for Material Semiotics, 2024.