AICR · COMPANION NOTE

What the Leftover Hash Lemma Gives

Seed Length, Measured Against the Bounds Already Known

Resolves. Nothing. This is a companion note to *Randomness Extraction from Unpredictable Random-Oracle Sources: A Leftover-Hash-Lemma Conjecture, Public Seed* (`main.tex` in this folder), whose conjecture is proved in `proof.tex` beside it. It asks one question. Set beside the bounds already known for extraction from a random oracle, what does the conjectured bound buy? It buys seed length, and by a wide margin. To certify a target δ the collision route needs about $\log_2 R$ bits of seed beyond the $2\log_2(1/\delta)$ every route pays. A 2-universal family needs D/R seeds, so at these lengths none exists. The unsalted random-oracle heuristic needs no seed, but a source that reads the oracle defeats it outright. Bounding the adversary's queries instead costs $\log_2 q$ bits. The conjecture asks for $\log_2\log_2 D + 2\log_2(1/\delta) + 2$, which is 77 bits in the leading example against 190, 1920 and 113. It is also the only one of the four whose seed grows with none of these: the derived key's length, the adversary's budget, and the secret's length beyond a doubly-logarithmic term.

The first three sections make that comparison. The fourth says what is not claimed, and the appendices hold the proofs and the rest. One thing here is not new: the floor the collision route runs into is the leftover hash lemma's known seed-length limit, which Barak et al. attribute to Stinson. What is new is the comparison.

AI: Claude Opus 5 (Anthropic)
Prompts: Pooya Farshim
Reviewers: Two blind adversarial referees (AI); no human
Date: 29 August 2026

1 The setting

The setting is that of `main.tex`, repeated so this note stands alone. Fix nonempty finite sets $\mathcal{K}$ (seeds), $\mathcal{D}$ (inputs), $\mathcal{R}$ (outputs), and write $K := |\mathcal{K}|$, $D := |\mathcal{D}|$, $R := |\mathcal{R}|$. Draw a table H uniformly from $\mathsf{Fun}(\mathcal{K}\times\mathcal{D},\mathcal{R})$. A source S reads the *entire* table and outputs a pair $(x,z) \in \mathcal{D}\times\{0,1\}^*$. Only then is a seed $sd \xleftarrow{\$} \mathcal{K}$ drawn, after S has stopped. An unbounded distinguisher D gets (H, sd, y_b, z) and must tell $y_0 = H(sd,x)$ from a uniform y_1. Write $\mathbf{Adv}^{\text{ext-pub}}(S,\mathsf{D})$ for twice its winning probability minus one. The source is *ϵ-unpredictable* if no unbounded predictor given (H,z) guesses x with probability more than ϵ, that is if $\mathbb{E}_{(H,Z)}[\epsilon_{H,Z}] \le \epsilon$ in the notation below.

Query budgets. Later sections also treat the variant in which each party gets oracle access in place of the table. There S^H makes at most q_S queries before returning (x,z), the predictor at most q_P, and $\mathsf{D}^H(sd,y_b,z)$ at most q_D. The seed is still drawn after S stops. Write $Q_S \subseteq \mathcal{K}\times\mathcal{D}$ for the set of cells the source queries. Two features of a source govern what follows, and neither is q_S:

$$r_S := \sum_{k\in\mathcal{K}} \Pr\big[Q_S \cap (\{k\}\times\mathcal{D}) \neq \emptyset\big], \qquad s_S := \sum_{k\in\mathcal{K}} \Pr\big[(k,X)\in Q_S\big], \tag{1}$$

the expected number of *rows* the source touches and the expected number of seeds at which it queries *its own output*. Both are at most $\min(q_S, K)$, and both can be far smaller than q_S.

Unqualified, ϵ means the strongest reading: the predictor is unbounded and holds the whole table. Every bound below assumes it except Theorem 2, which needs only a predictor as rich as the distinguisher. Appendix C shows that hypothesis is forced.

Throughout, condition on a table and an auxiliary output, $(H, Z) = (h, z)$, and write $p_x := \Pr[X = x \mid h, z]$ for the source's conditional law and $\epsilon_{h,z} := \max_x p_x$. The quantities

$$\mathrm{CP} := \sum_{x \in \mathcal{D}} p_x^2, \qquad c_k := \Pr_{X,X'}\big[h(k,X) = h(k,X')\big], \qquad \bar{c} := \frac{1}{K}\sum_{k \in \mathcal{K}} c_k,$$

all depend on (h, z), which the notation suppresses. Inside an expectation over (H, Z) read them as the matching random variables. Here X, X' are independent draws from p. Note $\mathrm{CP} \le \epsilon_{h,z}$, since $\sum_x p_x^2 \le (\max_x p_x) \sum_x p_x$.

2 The bounds

Every bound the comparison uses, stated without proof. The proofs are in Appendix A. Two bounds are quoted rather than proved anywhere here: the conjecture, which `proof.tex` proves, and the results of [1]. Whatever does not bear on seed length is in the appendices.

The conjecture

Conjecture 1 of `main.tex` asserts a universal constant c such that, for every ϵ-unpredictable S and every unbounded D,

$$\mathbf{Adv}^{\text{ext-pub}}(S, \mathsf{D}) \;\le\; c\left(\sqrt{\epsilon R} + \sqrt{\frac{\log_2 D}{K}}\right). \tag{2}$$

Every numerical comparison below uses (2) exactly as displayed, with $c = 2$, the constant `proof.tex` establishes. No other constant from that note is imported and none is needed.

Also stated, not proved here: the bounds of Barak et al.

This note measures against several bounds of [1], so they are listed here with the rest. All are quoted, not proved. All live in that paper's model rather than the game of Section 1: there Ext is a fixed public function, the source is a distribution rather than an algorithm, and the random oracle appears in no definition, lemma or theorem. Remark 8 sets the two models side by side.

First a clash of notation. In [1], v is the number of extracted bits, m the source's min-entropy, $L := m - v$ the *entropy loss*, and ε the security of the *application* being keyed; the source's length is written $|X|$. This note writes u for that length and μ for

their m, since m is needed below for $\log_2 R$; the renaming is this note's and not theirs. That last is not this note's ϵ, which is the source's unpredictability and is $2^{-\mu}$ in their units. Otherwise $u = \log_2 D$ and $v = \log_2 R$.

1. **Seed length** (their introduction, attributed to Stinson [2]): a family (almost) universal enough for the leftover hash lemma needs $n \geq \min\big(u-v,\, v+2\log(1/\varepsilon)\big) - O(1)$, "and must grow with the number of extracted bits". Theorem 1 and Proposition 3 reproduce this, one branch each; Section F makes the identification.
2. **Entropy loss** (their introduction): the standard lemma needs $L \geq 2\log(1/\varepsilon)$.
3. **Their generalised leftover hash lemma** (Theorem 3.1, applied to particular primitives in their Theorems 3.3–3.7): the loss falls to $L = \log(1/\varepsilon)$ for a wide range of key-derivation applications, the security degrading from ε to at most $\varepsilon + \sqrt{\varepsilon\, 2^{-L}}$.
4. **Expand-then-extract** (their Section 4): stretching a short public seed with a pseudorandom generator and extracting with the result does *not* work in general. They give a DDH-based generator and a perfectly universal hash whose output is distinguishable with probability close to 1 given the short seed. It is sound when the number of extracted bits is logarithmic in the generator's security, and sound in minicrypt, so any counterexample needs a public-key assumption.
5. **The unsalted random-oracle heuristic** (their Section 3.4, prose rather than theorem): $\varepsilon_{\mathrm{RO}} \approx T/2^{\mu} = \varepsilon\, 2^{-L}$, for an adversary of running time $T \approx \varepsilon 2^{v}$.

Items 2, 3 and 5 are about *entropy loss*; items 1 and 4 about *seed length*. The two axes are independent, and this note is on the second. Their entropy-loss improvement, item 3, leaves item 1 where it was. Their own attempt on item 1 is item 4, and it fails in general. So the seed-length limit is the more stubborn of the two, and it is the one Theorem 1 is about.

What the collision route returns

Proposition 1 (The bound the lemma returns). *For every table h and every view z of positive probability,*

$$\max_{\mathsf{D}} \big(2\Pr[\mathsf{D}\ \textit{wins} \mid h, z] - 1\big) \;\leq\; \tfrac{1}{2}\sqrt{R\,\bar{c} - 1}\,,$$

and consequently, for every unbounded D, $\mathbf{Adv}^{\text{ext-pub}}(S, \mathsf{D}) \leq \mathbb{E}_{(H,Z)}\big[\tfrac{1}{2}\sqrt{R\,\bar{c} - 1}\big]$.

Writing $c_k = \mathrm{CP} + \beta_k$ with $\beta_k := \sum_{x \neq x'} p_x p_{x'} \mathbf{1}[h(k,x) = h(k,x')]$ and $\bar{\beta} := \frac{1}{K}\sum_k \beta_k$ splits the quantity under the root as

$$R\bar{c} - 1 \;=\; R \cdot \mathrm{CP} \;+\; \gamma, \qquad \gamma \;:=\; R\bar{\beta} - 1, \tag{3}$$

The first summand carries the source's entropy deficiency. The second carries the family's failure to be universal on the source's support. A lower bound on $R\bar{c} - 1$ bounds only the sum, so which summand it speaks about needs settling; Lemma 1 settles it. Against a universal table the second summand vanishes.

Proposition 2 (Universal families). *Fix a table h and suppose the family $\{h(k,\cdot)\}_{k\in\mathcal{K}}$ is 2-universal, that is $\Pr_k[h(k,x) = h(k,x')] \le 1/R$ for all $x \ne x'$. Then $\gamma \le -\mathrm{CP} \le 0$ for every source law p, and the conditional bound of Proposition 1 is at most $\frac{1}{2}\sqrt{R\,\epsilon_{h,z}}$. Consequently, if* every *table in the support of H has this property, then $\mathbf{Adv}^{\text{ext-pub}}(S,\mathsf{D}) \le \frac{1}{2}\sqrt{\epsilon R}$ for every ϵ-unpredictable S and every unbounded* D.

The hypothesis is a property of a single table, whereas H in Section 1 is uniform on all of $\mathsf{Fun}(\mathcal{K}\times\mathcal{D},\mathcal{R})$; the second sentence of Proposition 2 is therefore about a different and easier model than the one at issue. It calibrates rather than applies. On the first summand of (2) the leftover hash lemma gives $\frac{1}{2}\sqrt{\epsilon R}$, which is that summand with $c = \frac{1}{2}$. Nothing is lost there, so everything separating the two routes lies in γ.

Why that calibration does not apply, and what stands in its place

Proposition 2 does not apply here. Universality is not just absent from a typical table. At the seed lengths at issue no table has it, provided the output is shorter than the input.

Proposition 3 (No universal family at small K, when $R < D$). *If $\{h(k,\cdot)\}_{k\in\mathcal{K}}$ is 2-universal then $D \le 2RK$; and, by Stinson's bound for almost-universal families quoted in the proof, $D \le RK$, that is $K \ge D/R$. So no 2-universal family exists once $K < D/R$, which at $D = 2^{256}$ and $R = 2^{64}$ holds for every $K < 2^{192}$. The proviso $R < D$ matters: for $R \ge D$ a single injective row is already 2-universal, with $K = 1$, so no seed-length lower bound holds without a hypothesis relating R to D.*

So γ must be taken over every support the source may choose after reading H, and the next theorem says how large those make $R\bar{c} - 1$. The source sees the whole table before *sd* exists, so any support it can name from H is allowed.

Theorem 1 (The collision route cannot beat $\sqrt{W/K}$). *Let $\mathcal{K},\mathcal{D},\mathcal{R}$ be nonempty and finite and let $\epsilon \in [1/D, 1]$. Put*

$$t \;:=\; \lceil 1/\epsilon \rceil, \qquad \sigma \;:=\; \lceil Rt/D \rceil, \qquad W \;:=\; \lfloor R/\sigma \rfloor.$$

Then $1 \le t \le D$, $1 \le \sigma \le R$ and $1 \le W \le R$. For every *table $h \in \mathsf{Fun}(\mathcal{K}\times\mathcal{D},\mathcal{R})$ there is a source S, uniform on a set of t inputs, with $z = \perp$ and ϵ-unpredictable, for which*

$$R\,\bar{c} - 1 \;\ge\; \frac{W-1}{K}, \qquad \textit{so that the bound of Proposition 1 is at least} \qquad \frac{1}{2}\sqrt{\frac{W-1}{K}}\,.$$

No randomness is used, so this holds for every table and not merely for most. When R and D/t are powers of two, as in the ranges of interest, $W = \min\{R,\ D/t\}$ exactly. In general $W > \frac{1}{2}\min\{R,\ D/t\} - 1$.

Theorem 1 bounds the sum $R\bar{c} - 1 = R\cdot\mathrm{CP} + \gamma$ from below. It says nothing about γ unless the floor exceeds what the first summand gives on its own, which for this source is $R\cdot\mathrm{CP} = R/t$.

Lemma 1 (Attribution). *For the source of Theorem 1,* $\mathrm{CP} = 1/t$ *exactly, so* $\gamma \geq (W-1)/K - R/t$*. In particular, if*

$$\frac{W-1}{K} > \frac{R}{t}, \tag{4}$$

then the floor of Theorem 1 is not explained by the entropy deficiency, and $\gamma > 0$*: the family is genuinely non-universal on that support, by a margin of at least* $(W-1)/K - R/t$*.*

Condition (4) is not automatic, and the comparison below is stated only where it holds. Take $R = D$ and $\epsilon = 1/2$, so $t = 2$, $\sigma = 2$, $W = \lfloor R/2 \rfloor$, and $R/t = R/2$ exceeds $(W-1)/K$ for every $K \geq 1$. There the theorem is consistent with $\gamma \leq 0$, that is with a universal family, and says nothing about the second summand.

Bounding the queries

Give the parties oracle access in place of the table. The leftover hash lemma cannot see the distinguisher's budget at all. It bounds a statistical distance, and that is by definition the supremum over every distinguisher, so pricing q_D means leaving the lemma. Leaving it gives the bound below, which is the one the seed comparison uses at a bounded budget. What becomes of the collision route itself under query bounds is Appendix B; it does not change the seed comparison.

Theorem 2 (The bad-cell route). *For every* ϵ*-unpredictable source and every distinguisher making at most* q_D *queries,*

$$\mathbf{Adv}^{\text{ext-pub}}(S, \mathsf{D}) \leq \frac{s_S}{K} + q_D\,\epsilon \leq \frac{\min(q_S, K)}{K} + q_D\,\epsilon.$$

The first term of Theorem 2 is not slack, and what it measures is the salt. The next statement needs no probability and no queries.

Proposition 4 (What the salt buys, and what its absence costs). *For every table* h *and every* $K \geq 1$ *there is a source, uniform on a set of at least* D/R *inputs and so* ϵ*-unpredictable with* $\epsilon \leq R/D$*, with* $r_S = s_S = 1$*, and a distinguisher making* no *oracle query, such that*

$$\mathbf{Adv}^{\text{ext-pub}}(S, \mathsf{D}) = \frac{1}{K}\left(1 - \frac{1}{R}\right).$$

In particular at $K = 1$ *the advantage is* $1 - 1/R$*: unsalted extraction from a source that reads the oracle fails outright, and it fails at min-entropy* $\log_2 D - \log_2 R$*, which is far above the minimum any of the settings in Section 3 asks for. Since* $s_S = 1$ *and* $q_D = 0$ *here, Theorem 2 returns* $1/K$*, so its first term is tight to within* $1 - 1/R$*.*

More generally, take $1 \leq j \leq K$ *and the source of Proposition 5, which pins* j *rows at a uniform input and leaks their values. It has* $s_S = j$ *and is* ϵ*-unpredictable with* $\epsilon \leq R^j/D$*, and against it a distinguisher making no oracle query achieves*

$$\mathbf{Adv}^{\text{ext-pub}}(S, \mathsf{D}) = \frac{j}{K}\left(1 - \frac{1}{R}\right).$$

So the first term of Theorem 2 is attained at every value of s_S*, not only at* $s_S = 1$*; and each pinned row costs* $\log_2 R$ *bits of min-entropy, exactly and not merely at most, so a source required to retain min-entropy* μ *can afford only* $j \leq (\log_2 D - \mu)/\log_2 R$ *of them.*

3 The comparison

The bounds side by side

Every bound of Section 2 in one place. The source is unbounded throughout except where a column says otherwise, and ϵ is ϵ_∞ except in the last row.

	source	predictor	disting.	the bound
(2), conjectured	unbounded	unbounded	unbounded	$c\left(\sqrt{\epsilon R} + \sqrt{\log_2 D/K}\right)$
Prop. 1, the route	unbounded	unbounded	unbounded	$\frac{1}{2}\sqrt{\epsilon R + \gamma}$
Prop. 2*	unbounded	unbounded	unbounded	$\frac{1}{2}\sqrt{\epsilon R}$
Thm. 1, a *floor*	unbounded	unbounded	unbounded	$\geq \frac{1}{2}\sqrt{(W-1)/K}$
Thm. 3	r_S rows	unbounded	unbounded	$\frac{1}{2}\sqrt{\epsilon R + 2R r_S/K}$
Thm. 2	s_S cells	$q_P \geq q_D$	q_D quer.	$s_S/K + q_D\,\epsilon$

In the second row γ stands for $\mathbb{E}_{(H,Z)}[\gamma]$; the pointwise statement is $\frac{1}{2}\sqrt{R \cdot \mathrm{CP} + \gamma}$, and the displayed form is what Jensen's inequality returns from it.

* On a *fixed* 2-universal table, which by Proposition 3 needs $K \geq D/R$; the uniform H of Section 1 is not one. Theorem 3 and Corollary 2 are stated and proved in Appendix B; they refine the collision route but do not move the seed comparison, which is why they sit there.

Three readings of the table. The first two rows differ only in γ against $\log_2 D/K$, and Theorem 1 is what forbids γ from being small enough to close that. The fifth row is the fourth's matching ceiling for sources of bounded row breadth. And the last row is the only one in which the distinguisher's budget appears at all, for the reason Remark 2 gives; by Corollary 2, and at any R these tables use, it is the smallest of them whenever $q_D \leq \frac{1}{2}\sqrt{R/\epsilon}$.

Side by side, and this is the whole comparison in one line:

$$\underbrace{c\left(\sqrt{\epsilon R} \;+\; \sqrt{\tfrac{\log_2 D}{K}}\right)}_{\text{conjectured, and proved in } \texttt{proof.tex}} \qquad \text{against} \qquad \underbrace{\tfrac{1}{2}\sqrt{\epsilon R \;+\; \tfrac{W-1}{K}}}_{\text{the least the collision route can return}} \qquad . \tag{5}$$

Since $\sqrt{a+b} \geq (\sqrt{a} + \sqrt{b})/\sqrt{2}$ by the arithmetic–geometric mean inequality, the right-hand side is at least $\frac{1}{2\sqrt{2}}\left(\sqrt{\epsilon R} + \sqrt{(W-1)/K}\right)$, so the two are the same two ingredients up to constants and differ in exactly one place: what stands under the second square root. The conjecture pays $\log_2 D$, a number of bits; the collision route pays W, the size of a set.

In entropy-loss units

The three bounds of Barak et al. that share an axis, in that paper's own notation, with $L = \mu - v$ the entropy loss and ε the keyed application's security:

	security	min-entropy for $O(\varepsilon)$	family, and status
standard LHL	$2^{-L/2}$	$\mu = v + 2\log(1/\varepsilon)$	universal; proved
[1], Thm. 3.1	$\sqrt{\varepsilon\, 2^{-L}}$	$\mu = v + \log(1/\varepsilon)$	universal; proved
[1], §3.4	$\varepsilon\, 2^{-L}$	$\mu = v$	unsalted RO; heuristic

This is a different axis from the table above, and the two should not be read across. That one varies the *seed*, holding the family's universality fixed; this one varies the *entropy loss*, holding the seed fixed. Nothing in this column improves the seed length, which is the quantity Theorem 1 constrains and the quantity the conjecture is about; and the last row is not salted at all, so it has no seed to improve.

In seed length

Corollary 1 (The comparison, in seed length). *Fix a target error $\delta \in (0,1)$ and parameters for which (4) holds and the entropy deficiency is not itself binding, that is $c\sqrt{\epsilon R} \le \eta\,\delta$ for some $\eta \in (0,1)$. Then (2) certifies δ once $K \ge c^2 \log_2 D/\big((1-\eta)\delta\big)^2$, whereas by Theorem 1 the leftover-hash route cannot certify δ at all unless $K \ge (W-1)/(4\delta^2)$. In bits,*

$$\log_2 K \;=\; 2\log_2 c + \log_2\log_2 D + 2\log_2(1/\delta) + 2\log_2 \tfrac{1}{1-\eta} \quad \textit{(conjectured)},$$
$$\log_2 K \;\ge\; \log_2(W-1) - 2 + 2\log_2(1/\delta) \quad \textit{(collision route)}.$$

Both carry $2\log_2(1/\delta)$; they differ in the first term, $\log_2 W$ against $\log_2\log_2 D$.

The last term of the first line is what the entropy deficiency costs the conjectured salt, and it is the one place where the two lines are not the same kind of quantity: the second is a necessary condition on K, the first a sufficient one, and driving (2)'s *second* summand alone below δ would not certify δ. Everywhere below $\eta \le 2^{-15}$, so the term is under 10^{-4} bits and the tables drop it.

Remark 1 (A worked instance). Take $D = 2^{256}$, $R = 2^{64}$, a source of min-entropy 160 (so $\epsilon = 2^{-160}$ and $t = 2^{160}$), $\delta = 2^{-32}$ and $c = 2$. Then $\sigma = \lceil 2^{-32} \rceil = 1$ and $W = R = 2^{64}$. Both hypotheses of Corollary 1 hold: the deficiency term is $c\sqrt{\epsilon R} = 2^{-47} \ll \delta$; and at the seed length (2) requires, $K = 2^{74}$, the floor is $(W-1)/K \approx 2^{-10}$ against $R/t = 2^{-96}$, so (4) holds with room to spare and Lemma 1 gives $\gamma > 0$. The comparison is then $\log_2 K = 74$ under (2) against $\log_2 K \ge 126$ for the collision route: a gap of 52 bits.

At min-entropy 128 instead, $c\sqrt{\epsilon R} = 2^{-31}$ exceeds $\delta = 2^{-32}$, so (2) certifies that δ at no seed length and the comparison is empty. The hypotheses of Corollary 1 are not decorative.

In key-derivation units

The game is a key-derivation problem. The input x is a secret with entropy but no uniformity: a Diffie–Hellman element, a biometric reading, the output of a physical source. The seed sd is a public salt, H is a public hash, and $H(sd, x)$ is the derived key, which must look uniform to an adversary who knows H in full and sees sd. In that

reading $\log_2 K$ is the *salt length* in bits, $m := \log_2 R$ the derived key length, $n := \log_2 D$ the length of the secret's encoding, and $\log_2(1/\epsilon)$ its min-entropy.

Whenever the derived key is no longer than the source's entropy gap, that is $m \le n - \log_2(1/\epsilon)$, one has $W = R$ and the two salt lengths of Corollary 1 read

$$\underbrace{\log_2 n + 2\log_2(1/\delta) + 2}_{\text{conjectured}} \qquad \text{against} \qquad \underbrace{m + 2\log_2(1/\delta) - 2}_{\text{collision route}}.$$

They differ in the first term, and they differ in kind. The conjectured salt grows with the *logarithm of the secret's length*. The collision route's grows with *the length of the key being derived*. So deriving a 256-bit key rather than a 128-bit one from the same secret costs no extra salt under the conjecture, and 128 further bits under the collision route. With $c = 2$, and the hypotheses of Corollary 1 checked in each row:

	parameters				salt needed (bits)			
	$\log_2 D$	min-ent.	key	δ	conj.	route	univ.	RO†
DH element, 2048-bit group	2048	256	128	2^{-32}	77	190	1920	0
the same, 256-bit key	2048	512	256	2^{-64}	141	382	1792	0
a 512-bit source	512	300	128	2^{-64}	139	254	384	0

The last column assumes the source is independent of H. It is the only column that does, and Proposition 4 says what it becomes otherwise.

† The unsalted random-oracle heuristic of [1], Remark 8: $\epsilon_{\mathrm{RO}} \approx T/2^{\mu}$, so it needs no salt. That zero is the most informative entry in the row, and it is bought rather than free. Two hypotheses buy it. The first is a distinguisher of running time $T \le \delta\, 2^{\mu}$, which is $T \le 2^{224}$, 2^{448} and 2^{236} in the three rows. Those ceilings are enormous, which is why the heuristic is attractive in practice. But they are ceilings, and all of `main.tex` lives above them, where every party reads the whole table and $T/2^{\mu}$ says nothing.

The second hypothesis is that the source is independent of H. That zero does not survive dropping it, not approximately and not by a little. By Proposition 4, at $K = 1$ a source that reads H and is uniform on the largest fibre of the single row, of min-entropy $\log_2 D - \log_2 R$, is distinguished from uniform with advantage $1 - 2^{-m}$ by a distinguisher making *no query at all*. That source is allowed in every row of the table: its min-entropy is 1920, 1792 and 384 against the 256, 512 and 300 the rows ask for. Its entropy loss is $L = \mu - v = 1792$, 1536 and 256, so the heuristic predicts $\epsilon_{\mathrm{RO}} \le 2^{-1792}$, 2^{-1536} and 2^{-256} where the truth is $1 - 2^{-128}$, $1 - 2^{-256}$ and $1 - 2^{-128}$. Off its hypothesis the estimate is not just inapplicable. It is wrong by as much as an estimate of a probability can be.

So the first and last salt columns are the same primitive, a random oracle, priced under two different assumptions about the source: nothing if it cannot look at H, and 77 bits if it can. The whole subject of `main.tex`, and of this note, is the gap between those two entries.

The *univ.* column is universal hashing, at $n - m$ bits from $K \ge D/R$ (Proposition 3). It makes the point of Appendix E concrete: universal hashing does not have a shorter

salt than the conjecture, it has a far longer one. It is also read differently from the first two. Those give the salt needed to reach the stated δ; this gives the salt needed for such a family to *exist*. At that salt the error is $\frac{1}{2}\sqrt{\epsilon R}$, which is 2^{-65}, 2^{-129} and 2^{-87} in the three rows, already inside δ in each.

The three columns scale in three different ways, and that is the whole comparison in one place:

$$\underbrace{\log_2 n + 2\log_2(1/\delta) + 2,}_{\text{conjectured}} \qquad \underbrace{m + 2\log_2(1/\delta) - 2,}_{\text{collision route}} \qquad \underbrace{n - m}_{\text{universal}}.$$

The conjectured salt grows with the logarithm of the secret's length. The collision route's grows with the length of the derived key. Universal hashing's grows with the secret's length outright, and *shrinks* as the derived key grows: 1920 bits for a 128-bit key against 1792 for a 256-bit one, since a larger R weakens the constraint $K \geq D/R$. It does not depend on δ at all, being a requirement for existence rather than for accuracy. Only the first is what one wants of a salt: short, and growing only in the logarithm of anything.

Read the other way, at a fixed salt rather than a fixed error, the same gap says the route certifies almost nothing at the salt the conjecture licenses. Substituting $K = c^2 \log_2 D/\delta^2$ into Theorem 1's floor gives

$$\frac{1}{2}\sqrt{\frac{W-1}{K}} \approx \frac{\delta}{2c}\sqrt{\frac{W}{\log_2 D}},$$

so the error the route can certify exceeds δ by a factor of about $\sqrt{W/\log_2 D}\,/\,(2c)$, roughly $(m - \log_2 n)/2$ bits of security, just under half the derived key's length. In the first row that factor is $2^{56.5}$, so at a 77-bit salt the route certifies $2^{24.5}$, which is larger than 1 and so says nothing. In the third row it certifies $2^{-6.5}$ at a 139-bit salt: six bits of security where sixty-four were wanted.

Neither reading says that salting with 77 bits is unsafe. Both say that the collision route cannot be what establishes it, and Remark 6 is the reminder that the true advantage against the witness source is far below either bound.

At a bounded query budget

Suppose the source and the distinguisher each make at most 2^n oracle queries. Then Theorem 2 certifies δ once $s_S/K \leq \delta/2$ and $q_D\epsilon \leq \delta/2$, and in the worst case over 2^n-query sources $s_S \leq 2^n$, so

$$\log_2 K \;\geq\; n + \log_2(1/\delta) + 1, \qquad \text{feasible only if} \qquad n \;\leq\; \mu - \log_2(1/\delta) - 1, \tag{6}$$

the second condition being about the source's min-entropy μ and not about the salt: no seed length repairs a budget that already exceeds it. The three settings of the table above, alongside the conjectured salt, which does not depend on n:

	salt (bits) at $q_S = q_D = 2^n$				conjectured
	$n = 40$	$n = 64$	$n = 80$	$n = 128$	any n
DH element, 2048-bit group	73	97	113	161	77
the same, 256-bit key	105	129	145	193	141
a 512-bit source	105	129	145	193	139

Two things to read off. The first is that the last column is *flat*: the conjectured salt is $\log_2 \log_2 D + 2\log_2(1/\delta) + 2$ whatever the adversary's budget, and it holds against an unbounded one, whereas (6) grows one bit for every doubling of the budget. They cross at $n = 44$, 76 and 74 in the three rows. So at 2^{80}, the budget one would actually quote, bounding the queries buys nothing: it asks for 113 bits where the conjecture asks for 77, and for 145 against 141 and 139. The query bound helps only against adversaries weaker than about 2^{44}. The reason is in the two exponents: the bad-cell route pays $\log_2 q$ where the conjecture pays $\log_2 \log_2 D$, and only halves the δ term in exchange.

The second is that (6) is an upper bound on what is needed, and a generous one, because $s_S \le 2^n$ is not achievable at these min-entropies. By Proposition 4 the advantage s_S/K is attained by pinning s_S rows at the source's own output. But each pinned row costs $\log_2 R$ bits of min-entropy, so the three rows admit only $j \le (\log_2 D - \mu)/\log_2 R =$ 14, 6 and 1 pinned rows. Against the best attack known the salt actually forced is $\log_2 j + \log_2(1/\delta) =$ 36, 67 and 64 bits, and Proposition 4 on its own already forces $\log_2(1/\delta) =$ 32, 64 and 64. So for the query-bounded route the truth lies in [36,113], [67,145] and [64,145], and the conjectured 77, 141 and 139 sit inside all three.

The slack has a source. A source may query (k, x) at its own output for a great many k and simply *ignore* the answers: that costs it no entropy and gains it no advantage, yet it inflates s_S to the query budget, and Theorem 2 charges for it. A bound that charged only for queries whose answers reach the distinguisher would close most of the gap; the note does not prove one.

4 What is not claimed

Theorem 1 is one-sided. It says the collision route cannot return a bound better than $\frac{1}{2}\sqrt{(W-1)/K}$; it does *not* say the route attains that. Every statement in this note about the route's cost is therefore a necessary condition on K, never the cost itself, and Corollary 1 is phrased that way deliberately. Theorem 3 supplies a matching ceiling for sources of bounded row breadth, and Remark 3 closes the window to a factor of two on the witness itself; for sources that touch many rows the paragraph below still stands.

The remaining window is wide. One can certify γ for a random table by a bounded-differences argument: $\bar{\beta}$ has bounded-differences constant $2/(Kt)$ in the Kt table entries above a support of size t, and a union bound over the $\binom{D}{t}$ supports of that size costs $e^{t(1+\ln D)}$. For sources uniform on a support that yields $\gamma \le R\sqrt{2(1+\ln D)/K}$ with high probability, hence an upper bound of $\frac{1}{2}\sqrt{\epsilon R} + \frac{1}{2}\sqrt{R}\,(2(1+\ln D)/K)^{1/4}$. Two gaps in that sketch are worth naming. It covers sources uniform on a support only,

and a general source needs a further argument reducing to those. And the resulting ceiling exceeds Theorem 1's floor by a factor of order $\sqrt{R/W}\,(K \ln D)^{1/4}$, about 2^{21} in the instance of Remark 1. So for a source of unrestricted row breadth the truth lies somewhere in that window, and locating it is not attempted.

Nor is the bound one would most like in the query-bounded setting of Section 2, namely $s_S/K + \frac{1}{2}\sqrt{\epsilon R}$: the pinned cells charged once each, and the leftover hash lemma applied to what is left as though it were universal. That would match the true advantage against the witness of Theorem 1 in both summands, and it is what one expects, since a row the source touched away from its own output is a row it has learnt nothing exploitable about. Proving it needs Lemma 2 refined from "row k untouched" to "the single cell (k, X) untouched", and that refinement is not available by the argument given: the freshness of one cell survives conditioning on the table only through the row-versus-rest factorisation the lemma uses, and two runs of the source may query in each other's output columns. Theorem 3 and Theorem 2 therefore stand as two separate bounds rather than one combined one, and Corollary 2 takes their minimum instead of their better parts.

What the floor alone does establish, in the regime where (4) holds and the deficiency term is not binding, is that the route cannot reach (2) there: it needs at least $\log_2(W - 1) - 2\log_2 c - \log_2 \log_2 D - 2$ further bits of seed, which is 52 in the instance of Remark 1. Outside that regime nothing is claimed. Remark 6 gives parameters where the construction shows no loss at all: min-entropy 254 over the same D and R.

Finally, the scope is one technique, not one lemma. Theorem 1 constrains the particular functional $\frac{1}{2}\sqrt{R\bar{c} - 1}$ of Proposition 1, arrived at by bounding the collision probability of $(sd, h(sd, X))$ and passing to statistical distance through (8). It says nothing about other uses of the leftover hash lemma, and the honest form of the headline is that (2) is not a corollary of *that route*. Nothing here is a statement about (2) itself, which is true: `proof.tex` proves it with $c = 2$, by bounding the L_1 distance directly and never forming a collision probability at all.

A Proofs

Of the statements of Section 2, in the order they appear there. Two elementary inequalities are recorded first, since several proofs use them.

Fact 1 (Cauchy–Schwarz, two forms). Let q be a probability distribution on a finite set Ω and let $A \subseteq \Omega$ carry all of its mass. Then

$$\sum_{\omega} q_\omega^2 \;\geq\; \frac{1}{|A|}. \tag{7}$$

If moreover u is uniform on Ω, then

$$\mathrm{SD}(q, u) \;=\; \tfrac{1}{2}\|q - u\|_1 \;\leq\; \tfrac{1}{2}\sqrt{|\Omega|}\,\|q - u\|_2 \;=\; \tfrac{1}{2}\sqrt{|\Omega| \textstyle\sum_\omega q_\omega^2 - 1}\,. \tag{8}$$

Proof. For (7), $1 = (\sum_{\omega \in A} q_\omega)^2 \leq |A| \sum_{\omega \in A} q_\omega^2 \leq |A| \sum_\omega q_\omega^2$. For (8), the inequality is $\|v\|_1 \leq \sqrt{|\Omega|}\,\|v\|_2$ applied to $v = q - u$, and the final equality is $\|q - u\|_2^2 = \sum_\omega q_\omega^2 - 1/|\Omega|$, which follows by expanding the square and using $\sum_\omega q_\omega = 1$. □

Proof of Proposition 1. Because sd is drawn only after S has terminated, it is uniform on $\mathcal{K}$ and independent of X given (h, z). The distinguisher's view therefore differs between the two challenge bits only in the pair (sd, y_b), which is distributed as $Q := \mathrm{law}(sd, h(sd, X))$ under $b = 0$ and as the uniform distribution U on $\mathcal{K} \times \mathcal{R}$ under $b = 1$. The optimal advantage of an unbounded distinguisher between two distributions is their statistical distance, so the left-hand side is $\mathrm{SD}(Q, U)$.

Now $|\mathcal{K} \times \mathcal{R}| = KR$ and

$$\sum_{k,r} Q(k,r)^2 \;=\; \Pr\big[(sd, h(sd,X)) = (sd', h(sd',X'))\big] \;=\; \sum_{k} \frac{1}{K^2}\, c_k \;=\; \frac{\bar{c}}{K},$$

the middle equality because sd, sd' are independent and uniform, so both equal a given k with probability $1/K^2$, and given $sd = sd' = k$ the outputs collide with probability c_k. Feeding this into (8) gives $\mathrm{SD}(Q, U) \le \frac{1}{2}\sqrt{KR \cdot \bar{c}/K - 1} = \frac{1}{2}\sqrt{R\bar{c} - 1}$, the radicand being nonnegative by (7) applied to Q with $A = \mathcal{K} \times \mathcal{R}$. Averaging over (H, Z) gives the second claim. □

Proof of Proposition 2. Exchanging the order of summation,

$$\bar{\beta} \;=\; \sum_{x \neq x'} p_x p_{x'} \Pr_k[h(k,x) = h(k,x')] \;\le\; (1 - \mathrm{CP})/R,$$

valid for every p because the hypothesis is a statement about every pair $x \neq x'$; so $\gamma \le 1 - \mathrm{CP} - 1 = -\mathrm{CP}$. By (3) the conditional bound is at most $\frac{1}{2}\sqrt{R \cdot \mathrm{CP} - \mathrm{CP}} \le \frac{1}{2}\sqrt{R\,\epsilon_{h,z}}$, using $\mathrm{CP} \le \epsilon_{h,z}$. For the last claim, average over (H, Z) and apply Jensen's inequality to the concave nondecreasing $\sqrt{\cdot}$. This is legitimate although p depends on H, since the conditional bound holds pointwise in (h, z). Combine it with $\mathbb{E}[\epsilon_{H,Z}] \le \epsilon$. □

Proof of Proposition 3. Viewing each input x as the column vector $(h(k,x))_{k \in \mathcal{K}} \in \mathcal{R}^K$, two distinct columns agree in at most K/R coordinates, so the D columns form a code of length K over an alphabet of size R with minimum distance at least $(1 - 1/R)K$. The Plotkin bound in its equality case, "$d = (1 - 1/q)n$ implies $|C| \le 2qn$", gives $D \le 2RK$.

The factor of two belongs to that route rather than to the truth. Stinson's lower bound for ε-almost-universal families [2], $K \ge D(R-1)/\big(\varepsilon R(D-R) + R^2 - D\big)$, reads $K \ge D/R$ at $\varepsilon = 1/R$, and that is the constant Section 3 uses. It is attained: over $\mathbb{F}_q$ the family $h(k, (a,b)) = ak + b$ has $D = q^2$ and $R = K = q$. Reading the elementary $D \le 2RK$ instead would shorten the *univ.* column of Section 3 by one bit and change nothing else.

For the proviso, if $R \ge D$ take $K = 1$ and any injection $\mathcal{D} \to \mathcal{R}$: no two distinct inputs collide, so the family is 2-universal. (An earlier version of this note asserted $K \ge D^{\Omega(1)}$ unconditionally, which that example refutes; and quoted the Plotkin bound in the sharper form $|C| \le qn$, which holds for $q = 2$ but is not the general statement.) □

Proof of Theorem 1. The parameters are in range. From $\epsilon \le 1$ we get $t \ge 1$, and from $\epsilon \ge 1/D$ that $t = \lceil 1/\epsilon \rceil \le D$; hence $Rt/D \le R$ and, R being an integer, $\sigma = \lceil Rt/D \rceil \le R$, while $\sigma \ge 1$ because $Rt/D > 0$. From $\sigma \le R$ we get $W = \lfloor R/\sigma \rfloor \ge 1$, and from $\sigma \ge 1$ that $W \le R$.

The support. Fix any seed $k_1 \in \mathcal{K}$ and let $\mathcal{S} \subseteq \mathcal{R}$ consist of σ symbols r maximising $|\{x \in \mathcal{D} : h(k_1, x) = r\}|$. The D inputs are distributed among the R symbols, so the σ most popular of them carry at least a σ/R fraction of the total, that is at least $D\sigma/R$ inputs; and $\sigma \geq Rt/D$ gives $D\sigma/R \geq t$. So $\{x : h(k_1, x) \in \mathcal{S}\}$ has at least t elements; let T be any t of them and let $S(h)$ output $x \xleftarrow{\$} T$ together with $z \leftarrow \bot$. This is a counting fact about an arbitrary map $\mathcal{D} \to \mathcal{R}$; no property of h beyond being a function is used.

Admissibility. T is computed from h, which S receives in full, and sd is drawn only after S terminates, so nothing in the construction depends on the seed and S is a legitimate source for the game. Its conditional law is uniform on T, so $\epsilon_{h,\bot} = 1/t = 1/\lceil 1/\epsilon \rceil \leq \epsilon$; since this holds for every table, $\mathbb{E}[\epsilon_{H,Z}] \leq \epsilon$ and S is ϵ-unpredictable in the sense of Section 1.

The two collision bounds. For every k the law of $h(k, X)$ is a distribution on $\mathcal{R}$, so (7) with $A = \mathcal{R}$ gives $c_k \geq 1/R$. For $k = k_1$ every $x \in T$ has $h(k_1, x) \in \mathcal{S}$, so the law of $h(k_1, X)$ is carried by $\mathcal{S}$ and (7) with $A = \mathcal{S}$ gives $c_{k_1} \geq 1/\sigma$.

Assembling. Averaging the K lower bounds and using $R/\sigma \geq \lfloor R/\sigma \rfloor = W$,

$$R\,\bar{c} - 1 \;\geq\; R\Big[\frac{1}{K}\cdot\frac{1}{\sigma} + \frac{K-1}{K}\cdot\frac{1}{R}\Big] - 1 \;=\; \frac{R/\sigma}{K} + \frac{K-1}{K} - 1 \;\geq\; \frac{W}{K} - \frac{1}{K} \;=\; \frac{W-1}{K},$$

and Proposition 1 applied to this source returns $\frac{1}{2}\sqrt{R\bar{c} - 1} \geq \frac{1}{2}\sqrt{(W-1)/K}$.

The two forms of W. If $R = 2^m$ and $D/t = 2^\ell$ then $Rt/D = 2^{m-\ell}$, so $\sigma = \max\{1, 2^{m-\ell}\}$ and $W = \lfloor R/\sigma \rfloor = \min\{2^m, 2^\ell\} = \min\{R, D/t\}$. In general $\sigma < Rt/D + 1$, so, writing $a := R$ and $b := D/t$,

$$\frac{R}{\sigma} \;>\; \frac{a}{a/b + 1} \;=\; \frac{ab}{a+b} \;=\; \frac{\min\{a,b\}\cdot\max\{a,b\}}{a+b} \;\geq\; \tfrac{1}{2}\min\{a,b\},$$

the last step because $a + b \leq 2\max\{a, b\}$; hence $W > R/\sigma - 1 > \frac{1}{2}\min\{R, D/t\} - 1$. □

Proof of Lemma 1. The source is uniform on T with $|T| = t$, so $\mathrm{CP} = \sum_{x \in T} t^{-2} = 1/t$. Substituting into (3) and rearranging Theorem 1's conclusion gives $\gamma = R\bar{c} - 1 - R/t \geq (W-1)/K - R/t$, and (4) makes it positive. □

Proof of Theorem 2. Write the advantage as $|\Pr[b' = 1 \mid b = 0] - \Pr[b' = 1 \mid b = 1]|$ and couple the two worlds by sampling H lazily, with the same coins for S and D in both. They differ only in the challenge: the cell $H(sd, x)$ under $b = 0$, an independent uniform value under $b = 1$. Let BAD be the event that the cell (sd, x) is queried, by the source or by the distinguisher. Off BAD that cell is never read by anyone, so the two executions coincide step for step: they can first diverge only at a query to (sd, x), whence BAD occurs in one world exactly when it occurs in the other and the two views agree off it. The advantage is therefore at most $\Pr[\mathrm{BAD}]$, which we may compute in the $b = 1$ world, and by a union bound at most $\Pr[(sd, x) \in Q_S] + \Pr[\mathsf{D} \text{ queries } (sd, x)]$.

For the first, sd is uniform on $\mathcal{K}$ and independent of the source's run, so $\Pr[(sd, X) \in Q_S] = \frac{1}{K}\sum_k \Pr[(k, X) \in Q_S] = s_S/K$, and $s_S \leq \min(q_S, K)$ because the source can place at most q_S queries and there are only K rows.

For the second, build a predictor. Given (H, z), let P draw $sd \xleftarrow{\$} \mathcal{K}$ and $y \xleftarrow{\$} \mathcal{R}$, run $\mathsf{D}^H(sd, y, z)$ answering its queries from H, pick $i \xleftarrow{\$} \{1, \ldots, q_D\}$ and return the

input component of the i-th query. (At $q_D = 0$ the distinguisher queries nothing, the term is 0, and no predictor is needed.) Under $b = 1$ the challenge is a uniform value independent of everything else, so P's simulation is exact; and if D queries (sd, x) then some query has input component x, which P selects with probability at least $1/q_D$. So $\Pr[\mathsf{D} \text{ queries } (sd,x)] \le q_D \cdot \mathbf{Adv}^{\mathrm{pred}}(P) \le q_D\epsilon$. □

Proof of Proposition 4. Fix the seed $k_1 = 1$. Among the R symbols, let r^* be one maximising $|h(1,\cdot)^{-1}(r^*)|$, ties broken by index; the D inputs are distributed among R symbols, so $T := h(1,\cdot)^{-1}(r^*)$ has $|T| \ge D/R$ by pigeonhole. This is a counting fact about an arbitrary map and uses no property of h. Let $S(h)$ read row 1, compute r^*, and output $x \xleftarrow{\$} T$ together with $z \leftarrow r^*$.

The source is admissible: T is computed from h, which S receives, and sd is drawn only after S terminates. Given (h,z) its law is uniform on T, so $\epsilon_{h,z} = 1/|T| \le R/D$ for every table, whence $\mathbb{E}[\epsilon_{H,Z}] \le R/D$. Leaking r^* costs nothing, since z is a function of h and the conditioning in $\epsilon_{h,z}$ already fixes h. It queries only row 1, and its output lies in that row's queried cells, so $r_S = s_S = 1$.

Let D make no oracle query and return 1 exactly when $sd = 1$ and $y_b = z$. Under $b = 0$ we have $y_0 = h(sd,x)$, and when $sd = 1$ every $x \in T$ gives $h(1,x) = r^* = z$, so $\Pr[\mathsf{D} \Rightarrow 1 \mid b = 0] = 1/K$. Under $b = 1$ the value y_1 is uniform on $\mathcal{R}$ and independent of z, so $\Pr[\mathsf{D} \Rightarrow 1 \mid b = 1] = 1/(KR)$. The advantage is the difference, $(1 - 1/R)/K$.

For the general case, let S draw $x \xleftarrow{\$} \mathcal{D}$, query (k,x) for $1 \le k \le j$ and output $\big(x, z = (H(1,x),\dots,H(j,x))\big)$; Proposition 5 gives $\epsilon \le R^j/D$ and $s_S = j$, the latter because the source queries its own output in exactly those j rows. Let D make no query and return 1 exactly when $sd \le j$ and $y_b = z_{sd}$. Under $b = 0$, $y_0 = H(sd,x) = z_{sd}$ whenever $sd \le j$, so the probability is j/K; under $b = 1$ it is $j/(KR)$.

The entropy remark needs the bound $\epsilon \le R^j/D$ in both directions, so the upper bound alone will not do. Proposition 5's proof gives the equality $\epsilon = \mathbb{E}_H\big[\#\{z : T_z \ne \emptyset\}\big]/D$; and a uniform H leaves a given one of the R^j fibres empty with probability $(1 - R^{-j})^D \le e^{-D/R^j}$, so all R^j are nonempty except with probability $R^j e^{-D/R^j}$, which is negligible whenever D/R^j comfortably exceeds $j \ln R$ — it is 2^{256} against $1792 \ln 2$ in the first row of Section 3, and larger in the other two. There $\epsilon = (1 - o(1))R^j/D$, and $\epsilon \le 2^{-\mu}$ rearranges to the stated ceiling on j. □

B The collision route under query bounds

Give the parties the budgets of Section 1 instead of the table. The collision route survives, in a form governed by row breadth rather than by query count. The source of Theorem 1 reads an entire row, so $q_S = D$, yet $r_S = s_S = 1$.

Theorem 3 (The collision route, query-bounded). *For every source that is ϵ-unpredictable and every unbounded distinguisher,*

$$\mathbf{Adv}^{\text{ext-pub}}(S,\mathsf{D}) \;\le\; \frac{1}{2}\sqrt{\epsilon R \;+\; \frac{2R\,r_S}{K}} \;\le\; \frac{1}{2}\sqrt{\epsilon R \;+\; \frac{2R\min(q_S,K)}{K}}\;.$$

Three readings of Theorem 3 are worth separating. It is a statement about *row breadth*, not query count: a source may spend D queries inside one row and remain capped at $\epsilon R + 2R/K$, while K queries placed one per row exhaust the bound. It degrades to nothing at $r_S = K$, correctly, since a source that pins every row leaves $\bar{c} = 1$. And at $r_S = 0$ it returns $\frac{1}{2}\sqrt{\epsilon R}$, the universal-family answer of Proposition 2: a source that never queries is independent of the table, and a random table is universal in expectation.

Proposition 5 (The dependence on r_S is exact up to a factor of two). *Fix $j \le K$ and let the source draw $x \xleftarrow{\$} \mathcal{D}$, query (k,x) for $1 \le k \le j$, and output $\big(x, z = (H(1,x),\ldots,H(j,x))\big)$. It makes j queries and has $r_S = s_S = j$; it is ϵ-unpredictable with $\epsilon \le R^j/D$; and for* every *table*

$$R\bar{c} - 1 \;\ge\; \frac{(R-1)\,j}{K}.$$

So the term $2R\,r_S/K$ of Theorem 3 cannot be improved below $(R-1)r_S/K$: the constant is off by $2R/(R-1)$, which is two in the ranges of interest, and the shape is right. At $j = K$ the proposition returns $R-1$, and the theorem is vacuous, as it must be.

Remark 2 (No bound proved through Proposition 1 can mention q_D). Proposition 1 passes through $\mathrm{SD}(Q,U)$, and statistical distance *is* the supremum of the advantage over all distinguishers whatever. A bound obtained that way holds for the unbounded distinguisher and can therefore never improve as q_D falls. This is not a defect of the present derivation but of the route: to let the distinguisher's budget count one must leave the leftover hash lemma, and Theorem 2 is what one gets by doing so.

Corollary 2 (Which route to use). *Under the hypotheses of both,*

$$\mathbf{Adv}^{\text{ext-pub}}(S,\mathsf{D}) \;\le\; \min\Big\{\, \frac{s_S}{K} + q_D\epsilon,\;\; \frac{1}{2}\sqrt{\epsilon R + \frac{2R\,r_S}{K}}\,\Big\},$$

and if moreover $2s_S/K + 2\sqrt{\epsilon R} \le R$, the bad-cell route is the smaller whenever

$$q_D \;\le\; \tfrac{1}{2}\sqrt{R/\epsilon}.$$

The side hypothesis is not a real restriction: at $R \ge 2^{64}$, with $s_S \le K$ and $\epsilon R \le 1$, its left-hand side is at most 4. The threshold itself is not a close call either. In the instance of Remark 1 it is $q_D \le 2^{111}$; for a 128-bit derived key from a source of min-entropy 128 it is $q_D \le 2^{127}$. Below it, which is every parameter regime anyone would deploy, the leftover hash lemma is the wrong instrument for this game. The gap is not marginal: on the witness of Theorem 1 in that instance, against a distinguisher making $q_D = 2^{80}$ queries, Theorem 2 returns $2^{-74} + 2^{-80} \approx 2^{-74}$ where Theorem 3 returns $2^{-5.5}$.

Proofs for this appendix. Theorem 3 needs one lemma of its own.

Lemma 2 (An untouched row stays fresh). *Fix $k \in \mathcal{K}$, let A_k be the event that the source queries some cell of row k, and write $p_x^- := \Pr[X = x \wedge \neg A_k \mid H, Z]$. Then*

$$\mathbb{E}_{(H,Z)}\Big[\sum_{x \ne x'} p_x^-\, p_{x'}^-\, \mathbf{1}[H(k,x) = H(k,x')]\Big] \;\le\; \frac{1}{R}.$$

Proof of Lemma 2. Split the table as $H = (H_k, H_{-k})$, row k against the rest; the two are independent and H_k is uniform on $\mathsf{Fun}(\mathcal{D},\mathcal{R})$. Take S deterministic given coins ρ, with ρ independent of H.

Run S on (ρ, h). Its transcript up to its first query into row k is a function of (ρ, h_{-k}) alone; hence so is the indicator of $\neg A_k$, and on $\neg A_k$ so is the whole transcript, and therefore (X, Z). Write

$$a(x,z,h_{-k}) := \Pr_{\rho}[X = x \wedge Z = z \wedge \neg A_k \mid h_{-k}], \qquad A(z,h_{-k}) := \sum_x a(x,z,h_{-k}),$$

so that $\Pr[H = h, Z = z, X = x, \neg A_k] = \Pr[h_{-k}]\Pr[h_k]\, a(x,z,h_{-k})$ and

$$\Pr[H = h, Z = z] \;=\; \Pr[h_{-k}]\Pr[h_k]\Pr_{\rho}[Z = z \mid h] \;\geq\; \Pr[h_{-k}]\Pr[h_k]\, A(z,h_{-k}).$$

Since $\Pr[H = h, Z = z]\, p^-_x = \Pr[H = h, Z = z, X = x, \neg A_k]$, the expectation in question is

$$\sum_{h,z} \frac{1}{\Pr[h,z]} \sum_{x \neq x'} \Pr[h,z,x,\neg A_k]\ \Pr[h,z,x',\neg A_k]\, \mathbf{1}[h(k,x) = h(k,x')],$$

Substituting the two displays above, with the lower bound on $\Pr[h,z]$ in the denominator, which is the direction an upper bound needs, makes it at most

$$\sum_{h_{-k}} \Pr[h_{-k}] \sum_{h_k} \Pr[h_k] \sum_z \frac{1}{A(z,h_{-k})} \sum_{x \neq x'} a(x,z,h_{-k})\, a(x',z,h_{-k})\, \mathbf{1}[h_k(x) = h_k(x')].$$

Neither a nor A depends on h_k, so the sum over h_k may be performed first, and for $x \neq x'$ it gives $\sum_{h_k} \Pr[h_k]\mathbf{1}[h_k(x) = h_k(x')] = 1/R$ exactly. What remains is at most

$$\frac{1}{R} \sum_{h_{-k}} \Pr[h_{-k}] \sum_z \frac{A(z,h_{-k})^2}{A(z,h_{-k})} \;=\; \frac{1}{R} \sum_{h_{-k}} \Pr[h_{-k}]\ \Pr_{\rho}[\neg A_k \mid h_{-k}] \;\leq\; \frac{1}{R},$$

using $\sum_{x \neq x'} a\, a' \leq (\sum_x a)^2 = A^2$ and $\sum_z A(z,h_{-k}) = \Pr_\rho[\neg A_k \mid h_{-k}]$. Terms with $A(z,h_{-k}) = 0$ have vanishing numerator and are read as 0. □

Proof of Theorem 3. Fix (h,z) and k, write $u_k := \Pr[A_k \mid h,z]$ and split $p = p^+ + p^-$ into its parts on A_k and $\neg A_k$. Of the four terms of $c_k = \sum_{x,x'} p_x p_{x'} \mathbf{1}[h(k,x) = h(k,x')]$, the three meeting p^+ contribute at most $\sum_x p^+_x \sum_{x'} p_{x'} + \sum_{x'} p^+_{x'} = 2u_k$, the indicator being at most 1. Hence

$$c_k \;\leq\; 2u_k \;+\; \sum_x (p^-_x)^2 \;+\; \sum_{x \neq x'} p^-_x\, p^-_{x'}\, \mathbf{1}[h(k,x) = h(k,x')],$$

and $p^- \leq p$ pointwise makes the middle sum at most CP. Average over k and take the expectation over (H,Z). For the first term, $\mathbb{E}\big[\sum_k u_k\big] = \sum_k \Pr[A_k] = r_S$; for the second, $\mathbb{E}[\mathrm{CP}] \leq \mathbb{E}[\epsilon_{H,Z}] \leq \epsilon$; for the third, Lemma 2 gives at most $1/R$ for each of the K rows. So

$$\mathbb{E}[\, R\bar{c} - 1\,] \;\leq\; \frac{2R\, r_S}{K} + \epsilon R + R \cdot \frac{1}{R} - 1 \;=\; \epsilon R + \frac{2R\, r_S}{K}.$$

By Proposition 1 the advantage is at most $\mathbb{E}\big[\frac{1}{2}\sqrt{R\bar{c}-1}\big]$, and $\sqrt{\cdot}$ is concave, so Jensen's inequality moves the expectation inside the root. The radicand is nonnegative pointwise by Fact 1, so this is legitimate. □

Proof of Proposition 5. Given (h,z) the conditional law is uniform on the joint fibre $T_z := \{x' : h(k,x') = z_k \text{ for all } k \le j\}$, so $\epsilon_{h,z} = 1/|T_z|$, and since $\Pr[Z = z] = |T_z|/D$ the average $\mathbb{E}[\epsilon_{H,Z}]$ is the number of nonempty fibres divided by D, at most R^j/D. Every element of T_z has the same image in each pinned row, so $c_k = 1$ for $k \le j$; for the remaining rows $c_k \ge 1/R$ by (7). Hence

$$R\bar{c} - 1 \;\ge\; R\Big[\frac{j}{K}\cdot 1 + \frac{K-j}{K}\cdot\frac{1}{R}\Big] - 1 \;=\; \frac{Rj}{K} - \frac{j}{K} \;=\; \frac{(R-1)j}{K}. \qquad \square$$

Proof of Corollary 2. Both bounds are proved above. Write $a := \epsilon R$, $\theta_s := s_S/K$ and $\theta_r := r_S/K$, and note $\theta_s \le \theta_r$ because querying (k,X) touches row k. At $q_D \le \frac{1}{2}\sqrt{R/\epsilon}$ one has $q_D\epsilon \le \frac{1}{2}\sqrt{a}$, so it is enough that $2\theta_s + \sqrt{a} \le \sqrt{a + 2R\theta_r}$. Squaring, that reads $4\theta_s^2 + 4\theta_s\sqrt{a} \le 2R\theta_r$, and $\theta_s \le \theta_r$ reduces it to $4\theta_s + 4\sqrt{a} \le 2R$ when $\theta_s > 0$, which is the hypothesis; at $\theta_s = 0$ it is immediate.

The two summands may *not* be compared one at a time. It is true that $s_S/K \le \frac{1}{2}\sqrt{2Rr_S/K}$ always — on squaring, $r_S \le RK/2$, which holds because $r_S \le K$ and $R \ge 2$ — but $\sqrt{a} + \sqrt{b} \ge \sqrt{a+b}$, so two termwise comparisons do not add up to a comparison of the sums, and without a hypothesis the conclusion is false: at $R = 2$, $r_S = s_S = K$, $\epsilon = 2^{-K}$ and $q_D = \frac{1}{2}\sqrt{R/\epsilon}$ the bad-cell bound is $1 + 2^{-(K+1)/2}$ against the collision route's $1 + 2^{-K-2}$ or so. An earlier version of this note argued termwise. □

Remark 3 (Theorem 3 closes the window of Section 4, for these sources). Theorem 1 is a floor and Section 4 names no matching ceiling; the one sketched there exceeds the floor by a factor of about 2^{21} in the instance of Remark 1. For a source of row breadth r_S the two now meet. Take that instance, $D = 2^{256}$, $R = 2^{64}$, min-entropy 160, $K = 2^{74}$, where $W = R$. The witness of Theorem 1 has $r_S = 1$, and

$$\frac{W-1}{K} \;\approx\; 2^{-10} \qquad\le\qquad R\bar{c} - 1 \qquad\le\qquad \epsilon R + \frac{2R}{K} \;=\; 2^{-96} + 2^{-9} \;\approx\; 2^{-9}\,:$$

a window of a factor 2 in the radicand, so the route returns between 2^{-6} and $2^{-5.5}$. The hypotheses differ. Section 4's sketch asks the source to be uniform on a support and bounds no queries; Theorem 3 bounds r_S and asks nothing about the shape. But the witness satisfies both, and on it the query-bounded ceiling is sharper by a factor of about 2^{20}.

C The predictor's budget

The third party has a budget too, and it is not free to choose. Bounding the predictor *weakens* the hypothesis, since a predictor that cannot read the table is easier to defeat. So the question is how weak the hypothesis may be made before the conclusions fail.

Definition 1 (Unpredictability at a budget). For $q \in \{0,1,2,\dots\} \cup \{\infty\}$ let ϵ_q be the least ϵ such that $\Pr[P^H(z) = X] \le \epsilon$ for every predictor P making at most q oracle queries. Then $\epsilon_0 \le \epsilon_1 \le \cdots \le \epsilon_\infty$, and $\epsilon_\infty = \epsilon_{KD}$ is the unbounded notion: a predictor allowed KD queries may read the whole table, so $\epsilon_\infty = \mathbb{E}_{(H,Z)}[\epsilon_{H,Z}]$.

The two routes answer the question differently, and the difference is exactly the difference between them. Note first that the reduction inside Theorem 2 is query-preserving: the predictor built there runs D and forwards each of its oracle queries, drawing *sd* and y itself, so it makes *exactly* q_D queries and never more. Its proof therefore establishes the following without alteration.

Theorem 4 (Three budgets). *Let S make at most q_S oracle queries, let* D *make at most q_D, and let the source be ϵ-unpredictable against predictors making at most q_P queries. If $q_P \ge q_D$ then*

$$\mathbf{Adv}^{\text{ext-pub}}(S,\mathsf{D}) \;\le\; \frac{s_S}{K} \;+\; q_D\,\epsilon \qquad \textit{with} \qquad s_S \le \min(q_S, K), \tag{9}$$

and if $q_P = \infty$ then also

$$\mathbf{Adv}^{\text{ext-pub}}(S,\mathsf{D}) \;\le\; \frac{1}{2}\sqrt{\epsilon R + \frac{2R\,r_S}{K}} \qquad \textit{with} \qquad r_S \le \min(q_S, K). \tag{10}$$

So all three budgets appear, and they appear in different places: q_S through r_S and s_S, q_D as a multiplier, and q_P as a side condition on which bound may be invoked. Neither side condition can be dropped, and one witness settles both.

Proposition 6 (The predictor may not be cheaper than the distinguisher). *Fix $R \le D/2$ and $1 \le m \le D/2$. Reserve the inputs $w_1,\dots,w_m$ from the upper half of $\mathcal{D}$ and identify $\mathcal{R}$ with the lower half. Let the source query $(1,w_1),\dots,(1,w_m)$ and output*

$$x \;:=\; \Big(\sum_{j=1}^{m} H(1,w_j)\Big) \bmod R, \qquad z := \bot.$$

Then $q_S = m$, $r_S = 1$ and $s_S = 0$; and

1. *$\epsilon_q = 1/R$ for every $q < m$, while $\epsilon_q = 1$ for every $q \ge m$;*
2. *given (H,Z) the source's law is a point mass, so* CP $= 1$, *$c_k = 1$ for every k, and $R\bar{c} - 1 = R - 1$;*
3. *an unbounded distinguisher, and equally one making $m+1$ queries, achieves advantage $1 - 1/R$.*

Consequently:

(a) In (10) the hypothesis $q_P = \infty$ cannot be weakened to any finite budget. Taking $m = q_P + 1$ and $K \ge 2R$, the bound asserted with ϵ_{q_P} would be $\frac{1}{2}\sqrt{1 + 2R/K} \le \frac{1}{2}\sqrt{2}$, against a true advantage of $1 - 1/R$, which is larger for every $R \ge 4$.

(b) In (9) the hypothesis $q_P \geq q_D$ cannot be weakened to $q_P \leq q_D - 2$. Taking $m = q_D - 1$, the bound asserted with ϵ_{q_D-2} would be q_D/R, against a true advantage of $1 - 1/R$, which is larger whenever $R > q_D + 1$.

The two halves of that proposition say the same thing about different routes. Reading them together: *the predictor must be granted at least what the distinguisher is granted, and granting it exactly that is enough.* For the bad-cell route the budget is pinned to q_D within a single query, $q_P \geq q_D$ sufficing and $q_P \leq q_D - 2$ failing. For the collision route the distinguisher is unbounded by construction, by Remark 2, so the predictor must be unbounded too. That is why every bound of this note other than Theorem 2 is stated against a predictor holding the whole table. It is not a preference for the strongest hypothesis. It is forced.

Proofs for this appendix.

Proof of Theorem 4. For (9): $q_P \geq q_D$ gives $\epsilon \geq \epsilon_{q_P} \geq \epsilon_{q_D}$, and the predictor of Theorem 2 makes q_D queries, so the bad event there is at most $q_D\epsilon_{q_D} \leq q_D\epsilon$. For (10), $q_P = \infty$ makes ϵ an upper bound for $\mathbb{E}[\epsilon_{H,Z}]$, which is the only use Theorem 3 makes of the hypothesis. □

Proof of Proposition 6. The source queries only row 1, so $r_S = 1$; its output lies in the lower half of $\mathcal{D}$ and its queries in the upper, so $(k, x) \notin Q_S$ for every k and $s_S = 0$.

For (1), a predictor making $q < m$ queries leaves some $(1, w_i)$ unqueried, and conditioned on any transcript that cell is uniform on $\mathcal{R}$ and independent of the predictor's view; a sum modulo R one of whose summands is uniform and independent is itself uniform, so x is uniform on $\mathcal{R}$ given that view and no guess does better than $1/R$. A predictor making m queries reads the m cells, which are fixed by the source's code, and computes x outright.

For (2), x is a function of H alone, so the conditional law given (H, Z) is a point mass: CP = 1 and $h(k, X) = h(k, X')$ always.

For (3), a distinguisher holding the table, or one spending m queries on the reserved cells and one more on (sd, x), computes $H(sd, x)$ and returns 1 exactly when y_b equals it. Under $b = 0$ it always returns 1; under $b = 1$ it does so with probability $1/R$.

For (a), $m = q_P + 1$ gives $\epsilon_{q_P} = 1/R$ by (1), so $\epsilon_{q_P} R = 1$; with $r_S = 1$ and $K \geq 2R$ the asserted radicand is at most 2. For (b), $m = q_D - 1$ gives $\epsilon_{q_D-2} = \epsilon_{m-1} = 1/R$ by (1), while a distinguisher making $m + 1 = q_D$ queries achieves $1 - 1/R$ by (3); the asserted bound is $s_S/K + q_D/R = q_D/R$. □

The witness also shows how little room there is between the two conditions. Its x is a deterministic function of H, so it is perfectly extractable-looking to any predictor poorer than the source and worthless against any predictor as rich, with nothing in between: ϵ_q jumps from $1/R$ to 1 at $q = m$. The general form of that observation is worth recording.

Remark 4 (Deterministic sources). If S is deterministic then $\epsilon_{q_S} = 1$, since a predictor with q_S queries simply reruns it. So for a deterministic source (9) is vacuous once $q_D \geq q_S$, and (10) is vacuous outright. Extraction from such a source can only be secure against distinguishers strictly poorer than the source itself, and the seed does

not change that: a distinguisher that can afford the source's own computation recovers x and tests it with one further query, whatever K may be. Whatever entropy is being extracted has to sit in the source's coins, not in the oracle it reads.

D Further notes

Three readings that belong to particular statements above rather than to the comparison.

Remark 5 (Pinning more than one row). Confining j rows rather than one, each to σ symbols, gives $R\bar{c} - 1 \geq j(W-1)/K$ by the same averaging, provided a common support of size t survives. It does: choosing the j symbol sets $\mathcal{S}_1, \ldots, \mathcal{S}_j$ independently and uniformly among the subsets of $\mathcal{R}$ of size σ, each input survives all j constraints with probability $(\sigma/R)^j$, so the expected number of survivors is $D(\sigma/R)^j$ and some choice attains it; the constraint is $D(\sigma/R)^j \geq t$, that is $j \leq \log_{R/\sigma}(D/t)$. In the principal regime $\sigma = 1$, $W = R$ that reads $j \leq \log_R(D\epsilon)$, which is 14, 6 and 1 in the three rows of Section 3. The gain is a factor $\sqrt{j}$ in the bound, hence $\log_2 j$ further bits of seed in that comparison: at those three values of j, between none and four. So the one-row construction of Theorem 1 carries most of the effect, and none of it in the third row.

The averaging is the point. Choosing each $\mathcal{S}_i$ greedily, as the σ commonest symbols of its row, does *not* work for a worst-case table, since the j greedy sets may have a small or even empty common preimage; an earlier version of this note asserted the greedy version. Only the one-row case survives greedily, and that is the case Theorem 1 uses.

Remark 6 (Where the loss comes from, and when it vanishes). The loss is the step (8), $\|v\|_1 \leq \sqrt{KR}\,\|v\|_2$, which is tight only when v is spread evenly over all KR cells; the source of Theorem 1 makes it far from that, its deviation on the pinned row being a single concentrated lump.

The true advantage against that source should not be mis-stated, and an earlier version of this note did mis-state it as $O(1/K)$. Since $\mathrm{SD}(Q,U) = \frac{1}{K}\sum_k \mathrm{SD}(\mathrm{law}\, h(k,X), U_{\mathcal{R}})$, the pinned row k_1 contributes about $1/K$, but the other $K-1$ rows are *not* uniform. On a typical table each is t draws into R bins and contributes $\Theta(\sqrt{R/t})$, so the true advantage is $\Theta(\sqrt{R/t}) + \Theta(1/K)$, dominated by the *unpinned* rows in the regime of Remark 1. On an adversarial table, every row constant say, it is close to 1. So the overshoot of the lemma over the truth is neither $\sqrt{RK}$ nor uniform in the parameters. In the instance of Remark 1, at the seed length (2) requires, the lemma returns $\frac{1}{2}\sqrt{(W-1)/K} = 2^{-6}$ against a true advantage of about $0.4\sqrt{R/t} = 2^{-49.3}$; at the much larger seed length the lemma itself requires, the two come within a small factor of one another. The floor is a statement about what the technique can certify, not about how weak the source is.

The gap closes when W does. At $\epsilon = 1/D$ one has $t = D$, $\sigma = R$, $W = 1$ and Theorem 1 is trivial, correctly: a source of maximal entropy is uniform on all of $\mathcal{D}$ and has no support to select. More generally W is small whenever the source's min-entropy is close to $\log_2 D$ or the output is short: for $D = 2^{256}$, $R = 2^{64}$ and min-entropy 254 one has $t = 2^{254}$, $\sigma = 2^{62}$, $W = 4$, and the theorem exhibits no loss at all.

Remark 7 (What this settles about the witness). Remark 6 decomposes the true advantage against that witness as $\Theta(\sqrt{R/t}) + \Theta(1/K)$, the second summand from the pinned row and the first from the $K-1$ unpinned ones, and notes that the first dominates. Theorem 2 says where the first comes from: it is an artefact of the distinguisher being unbounded. Against a q_D-bounded distinguisher the unpinned rows contribute at most $q_D\epsilon = q_D/t$ in place of $\sqrt{R/t}$, which is smaller exactly when $q_D \le \sqrt{Rt} = \sqrt{R/\epsilon}$, the threshold of Corollary 2 up to its factor of two, while the pinned row still contributes its $s_S/K = 1/K$. So the reading that the advantage against this source is $\Theta(1/K)$ is correct precisely in the query-bounded model and false outside it, which is what Remark 6 found the hard way.

E Is the random oracle worse than universal hashing?

Proposition 2 gives a universal family the bound $\frac{1}{2}\sqrt{\epsilon R}$ with no second summand at all, while Theorem 1 saddles a random table with a floor of $\frac{1}{2}\sqrt{(W-1)/K}$ on top of it. Read quickly, that says the random oracle is the worse extractor. The reading is wrong, and the error is a natural one.

The two are not being compared at the same seed length. By Proposition 3 a 2-universal family on D inputs needs $K \ge D/R$, a salt of $n - m$ bits. That is what buys $\gamma \le 0$. The random-oracle question is posed at $K \approx c^2 \log_2 D/\delta^2$, a salt of $\log_2 n + 2\log_2(1/\delta)$ bits, which is smaller by an exponential factor. Comparing the two bounds without comparing their seeds is comparing the price of two objects while looking only at one price tag.

At equal seed length the random oracle matches, and the gap closes completely. Give the random table the salt a universal family requires, $K = D/R$. Then the conjectured second summand is

$$\sqrt{\frac{\log_2 D}{K}} = \sqrt{\frac{R \log_2 D}{D}},$$

which for $D = 2^{2048}$ and $R = 2^{128}$ is $2^{-954.5}$: not small, but absent. Both routes then return $\Theta(\sqrt{\epsilon R})$, so there is nothing to choose between them. The second summand is the price of a *short* salt, not a defect of the random oracle. Universal hashing never pays it only because it never has one.

At the seed length that makes the problem interesting, universal hashing does not exist. Take the first row of Section 3's table: a Diffie–Hellman element in a 2048-bit group, min-entropy 256, a 128-bit key, $\delta = 2^{-32}$. The conjectured bound does the job with a 77-bit salt. A universal family needs at least 1920 bits, twenty-five times as long, to do the same job no better: its error there is $\frac{1}{2}\sqrt{\epsilon R} = 2^{-65}$ against the random oracle's 2^{-63} at the same salt. Measured by error per bit of seed, then, universal hashing is the suboptimal construction and the random oracle the optimal one. This is the standard picture. The leftover hash lemma is known to be far from optimal in seed length, and (2) sits at the information-theoretic optimum, which universal hashing misses by an exponential margin.

Why universality is immune to support selection is why it is expensive. Proposition 2 survives a source that reads the whole table, and its proof shows why: the bound $\bar{\beta} = \sum_{x \neq x'} p_x p_{x'} \Pr_k[h(k,x) = h(k,x')] \leq (1 - \mathrm{CP})/R$ holds for *every* law p, because universality constrains *every* pair $x \neq x'$ separately. A worst-case-over-all-pairs hypothesis is automatically robust to an adversary who picks the pairs last. But a worst-case-over-all-pairs hypothesis is exactly what Proposition 3's counting argument charges D/R seeds for. The immunity and the cost are the same property seen twice. A random table satisfies only the average-case version, which is why choosing the support after seeing H can hurt it at all, and why $\log_2 n$ bits of salt suffice to stop that rather than $n - m$.

And the floor is about the technique, not the object. Theorem 1 bounds what one accounting can certify about a random table. It is not a lower bound on the advantage, and the random table is better than the floor suggests: (2) is *true*, proved in `proof.tex` with $c = 2$, so the random oracle does achieve $\sqrt{\epsilon R} + \sqrt{\log_2 D/K}$ at a short salt. Theorem 1 says only that the collision route cannot prove it. Remark 6 says the same numerically: against the witness source the true advantage is some 2^{-49} where the route returns 2^{-6}.

F Prior art: the floor is Stinson's

Theorem 1 is not new. This section says so precisely.

Barak, Dodis, Krawczyk, Pereira, Pietrzak, Standaert and Yu [1] open by naming two limitations of the leftover hash lemma. The first is entropy loss. The second is the one at issue here: for an (almost) universal family good enough for the LHL, the seed length must be at least

$$\min\big(u - v,\ v + 2\log(1/\varepsilon)\big) - O(1),$$

where u is the length of the source and v the number of extracted bits, "and, thus, must grow with the number of extracted bits". They attribute it to Stinson [2], and record separately that Stinson showed perfectly universal families need a seed linear in u.

Translate: $u = \log_2 D$, $v = \log_2 R = m$, $\varepsilon = \delta$, and the seed length is $\log_2 K$. That bound reads $\log_2 K \geq \min(\log_2 D - m,\ m + 2\log_2(1/\delta)) - O(1)$. Corollary 1 gives $\log_2 K \geq \log_2(W - 1) - 2 + 2\log_2(1/\delta)$, and in the principal regime $W = R$ that is $m + 2\log_2(1/\delta) - 2$: the second branch. The first branch is Proposition 3's $K \geq D/R$, which is Stinson's linear-in-u result; the elementary route given there reaches it only up to a factor of two, and the constant is Stinson's own. So the two salt columns of Section 3's table are the two branches of one bound that has been known since 1994. And Appendix E's reading, that universal hashing has the longer salt rather than the shorter, is that bound's own reading. What this note adds is not the number.

It adds three things, none of them a new bound. First, a derivation of the relevant branch from a pigeonhole and two applications of Cauchy–Schwarz, with no probability in it, so the statement holds for every table. Second, that derivation put into the game of `main.tex`: Theorem 1 gives a witness source per table, so the obstruction is an

attack one can write down rather than a counting bound on families. Third, Lemma 1, which separates the cases where the floor speaks about γ from those where the entropy deficiency explains it. Whether any of the three is also in Stinson's papers I do not know: the attribution above is [1]'s, and I have not read [2] at first hand.

Two things in the game of `main.tex` that literature does *not* cover.

- **The independence hypothesis is *not* weakened here, and it is worth being exact about that, because the natural summary is wrong.** Definition 2.1 of [1] quantifies over all distributions X with $\tilde{H}_\infty(X \mid Z) \geq m$, with the seed S drawn independently as the coins of a function Ext that is fixed beforehand and public. So a source there may already depend on the *family*; what it must not depend on is the *seed*. The game of `main.tex` asks no more than that: *sd* is drawn after S has terminated, so its source is independent of the seed too. Nothing in the adversary model is weakened, and Theorem 1's witness uses no freedom the leftover-hash setting withholds. What differs is the *family*: a uniformly random table in place of a universal one. The content of Theorem 1 is that this substitution alone, with seed-independence fully intact, already costs the seed length. (That paper's remark about not making "the source X dependent on the seed" is about a different scenario, fixing one public seed and reusing it indefinitely, and does not bear on this.)

- **Their random-oracle comparison bounds one party only, and it is unsalted.** This is set out in full in Remark 8 below, because the asymmetry is the whole of why that comparison does not reach this setting.

Remark 8 (The query model of the random-oracle comparison). Section 3.4 of [1] compares the leftover hash lemma against the "dream" bound obtained by modelling a cryptographic hash as a random oracle. Its construction is *unsalted*: the derived key is $H(X)$, with no seed anywhere in the argument. The reasoning is that $H(X)$ is distinguishable from uniform only if the attacker queries the oracle at X; as X has min-entropy μ, that happens with probability at most $T/2^\mu$, where T bounds the adversary's running time. Now calibrate T against the application's own security. A primitive P that is ϵ-secure against complexity T must satisfy $\epsilon \leq T/2^v$, since exhaustive search over a v-bit key must not succeed better than ϵ, so $T \approx \epsilon 2^v$. That gives

$$\epsilon_{\mathrm{RO}} \approx \frac{T}{2^\mu} = \epsilon\, 2^{v-\mu} = \epsilon\, 2^{-L}, \qquad L := \mu - v \text{ the entropy loss,}$$

against $\sqrt{\varepsilon 2^{-L}}$ for their generalised LHL and $2^{-L/2}$ for the standard one (the list in Section 2 quotes all three).

Only one of the two parties is bounded, and the two are not bounded in the same sense. The *distinguisher* is bounded, by the running time T, which in particular bounds its oracle queries. The *source* is not bounded, because it is not a query-making object at all. There X is a distribution of min-entropy μ, fixed independently of H. No source algorithm with oracle access appears in the model, so there is nothing about it to bound. Searching [1] for query counts finds only the distinguisher's T and the query

budgets of the *applications* being keyed (encryption, MAC and PRF queries); the source never makes a query.

That asymmetry is exactly what the argument rests on, and it is not incidental. "$H(X)$ is distinguishable only if the attacker queried at X" is false the moment X may depend on H, and Proposition 4 is how false: a source that reads H and returns a uniform element of its largest fibre is distinguished with advantage $1 - 1/R$ by an adversary making no query at all, at a min-entropy of $\log_2 D - \log_2 R$. So this comparison needs the source independent of H *itself*. That is strictly stronger than the leftover-hash half of the same paper needs, where the source must be independent only of the seed and may know the family. Being unsalted, it has no seed to fall back on, so independence from H is the only thing holding it up. That is the hypothesis `main.tex` does drop, and the salt is what it puts in its place.

There is also no formal random-oracle model to appeal to here: [1]'s Section 3.4 is prose, reading "it is instructive to compare… the heuristic 'dream' bound… we claim", and the random oracle appears in no definition, lemma or theorem of that paper. Its formal Definition 2.1 has Ext a plain function and X a distribution quantified over, with no oracle anywhere. So the question of how many oracle queries the source may make has no answer there, not because the answer is zero but because the object is a distribution rather than an algorithm.

It is also why that comparison needs no salt while `main.tex` does. Once the distinguisher is unbounded the estimate $T/2^{\mu}$ is vacuous: in the game of `main.tex` every party holds all KD entries of the table, so T is effectively KD and $T/2^{\mu} = KD\epsilon$ exceeds 1 for every interesting ϵ. Unsalted extraction against an unbounded distinguisher simply fails, and this note's own arithmetic agrees: at $K = 1$ the bound (2) reads $c(\sqrt{\epsilon R} + \sqrt{\log_2 D})$, which is vacuous for every $D \geq 2$. The salt is what rescues it, and the salt is what the question is about. Summarising the two models:

	[1], Section 3.4	`main.tex`
construction	$H(X)$, unsalted	$H(sd, x)$, salted
source	a distribution, independent of H; makes no queries	an algorithm reading all of H; unbounded
distinguisher	running time $\leq T$	unbounded; holds all of H
bound	$\epsilon\, 2^{-L}$	$c(\sqrt{\epsilon R} + \sqrt{\log_2 D/K})$

[1] Boaz Barak, Yevgeniy Dodis, Hugo Krawczyk, Olivier Pereira, Krzysztof Pietrzak, François-Xavier Standaert and Yu Yu. Leftover Hash Lemma, Revisited. In *Advances in Cryptology – CRYPTO 2011*, LNCS 6841, pages 1–20. Springer, 2011. Cryptology ePrint Archive, Report 2011/088, https://eprint.iacr.org/2011/088. (Read at first hand, and quoted in Section 2: the seed-length and entropy-loss limitations are the two bullets of its abstract, developed on page 1; the generalised leftover hash lemma is its Theorem 3.1, applied in Theorems 3.3–3.7; expand-then-extract is its Section 4; the unsalted random-oracle comparison is its Section 3.4, which is prose and carries no theorem.)

[2] Douglas R. Stinson. Universal hashing and authentication codes. *Designs, Codes and Cryptography*, 4(4):369–380, 1994. (Cited by [1] as the source of that seed-length bound, and quoted in the proof of Proposition 3 for the ε-almost-universal lower bound $K \geq D(R-1)/(\varepsilon R(D-R) + R^2 - D)$. *Not* read at first hand for this note: the attribution is [1]'s

and the formula is taken from a secondary restatement. Its special case $K \geq D/R$ at $\varepsilon = 1/R$ was checked independently here — it is attained by $h(k,(a,b)) = ak + b$ over $\mathbb{F}_q$, and a randomised search over $R \leq 5$, $K \leq 9$ found no family beating it.)

[3] Salil P. Vadhan. *Pseudorandomness.* Foundations and Trends in Theoretical Computer Science. Chapter 6, Randomness Extractors. (Theorem 6.17 is (2) in extractor parameters; the remark following Theorem 6.18 is the seed length of universal hashing. Read at first hand.)

[4] Jaikumar Radhakrishnan and Amnon Ta-Shma. Bounds for dispersers, extractors, and depth-two superconcentrators. *SIAM Journal on Discrete Mathematics*, 13(1):2–24, 2000. (The optimum $d \geq \log(n-k) + 2\log(1/\epsilon) - O(1)$ that (2) meets up to log $\frac{n}{n-k}$, since it pays $\log n$ where the optimum pays $\log(n-k)$; that is 0.2 bits in the first row of Section 3. Read at first hand.)