Secret Sharing Despite Non-Extractable Randomness

Divesh Aggarwal
National University of Singapore
dcsdiva@nus.edu.sg

November 3, 2026

Abstract

Bosley and Dodis (TCC 2007) proved that a class of imperfect randomness sources that suffices for information-theoretic encryption of messages longer than the logarithm of the key length must admit deterministic extraction of nearly uniform bits, and asked whether the same is true of two-out-of-two secret sharing. Aggarwal, Chung, Obremski, and Ribeiro (TCC 2022) gave a positive answer for schemes that are in addition leakage-resilient or non-malleable; the original question remained open.

We answer it in the negative. For every $b \geq 2$ we construct a class of flat sources on $\{0,1\}^{b+2}$, each of min-entropy at least b, together with a perfectly correct and perfectly private two-out-of-two scheme that shares a b-bit secret into two $(b+1)$-bit shares using any source in the class, and we show that every deterministic map from $\{0,1\}^{b+2}$ to $\{0,1\}$ is at statistical distance at least $1/6$ from uniform on some source in the class. This implies that sharing a long enough secret does not require extractable randomness. The proof is a short parity argument. The idea of the construction extends to other thresholds: for two-out-of-n and r-out-of-r sharing we give perfect schemes with $(b+1)$-bit shares, private for classes of flat sources from which no single bit can be extracted with error below $1/(4n-2)$ and $1/(2(2^r-1))$, respectively.

1 Introduction

Uniform randomness is the most fundamental resource in information-theoretic cryptography. A uniform one-time pad gives perfect encryption, and splitting a secret into a uniform one-time pad and the padded secret gives perfect two-out-of-two secret sharing [Sha49, Sha79, Bla79]. In contrast, physical sources of randomness are typically not uniform, and their exact distribution is usually unknown; the most one can assume is that the distribution belongs to some well-defined *class* of sources, for instance the class of all sources with a given min-entropy. The cleanest way to use such a source is to extract uniform randomness from it: if a single deterministic function maps every source in the class to a nearly uniform string, every cryptographic task requiring uniform randomness becomes possible. A basic question, going back to [MP91, DS02, DOPS04], is whether extraction is necessary:

If a class of sources suffices for a given cryptographic task, must it be extractable?

The answer depends on the task. Entropy alone does not suffice: deterministic extraction fails for sources with almost full min-entropy [SV84, CG88], and such sources do not support information-theoretic privacy either [MP91, DOPS04]. For carefully chosen classes, on the other hand, privacy is possible without extraction when the message is a single bit: Dodis and Spencer separated one-bit encryption from extraction [DS02], and Dodis, Pietrzak, and Przydatek separated one-bit secret sharing from one-bit encryption [DPP06].

Encryption implies extraction. Bosley and Dodis [BD07] showed that these separations only work if the message length is small. If a class of distributions on d-bit keys permits perfect encryption of b-bit messages, then for every $\varepsilon > 0$ with

$$b > \log d + 2\log(1/\varepsilon)$$

the class admits a deterministic extractor with $\lfloor b - 2\log(1/\varepsilon) \rfloor$ output bits and error ε [BD07, Theorem 1]. They complemented this result with non-extractable classes that support encryption of $b \le \log d - \log\log d - 2$ bits. Hence, as soon as the message is sufficiently longer than the logarithm of the key length, *encryption requires extractable randomness.*

Encryption yields two-out-of-two secret sharing (give the key to one party and the ciphertext to the other), but not conversely [DPP06], so secret sharing is the weaker and more basic primitive. Bosley and Dodis left open whether it obeys the same principle [BD07, Section 5]:

Does two-out-of-two secret sharing of long secrets require extractable randomness?

Aggarwal, Chung, Obremski, and Ribeiro [ACOR22] revisited the question fifteen years later. They proved that the answer is *yes* when the scheme additionally resists one bit of local leakage from each share, or is non-malleable (unless d is exponentially large); see also [Chu24]. They posed the question for general thresholds and explicitly left ordinary secret sharing open.

1.1 Our results

We resolve the question in the negative, in a strong form: the scheme is perfect, the secret is only two bits shorter than the source string, every source is flat with high min-entropy, and not even one bit can be extracted with error below a fixed constant.

A *source class* is a fixed family of distributions on the dealer's random input, which we call the *key*. A single sharing scheme must be private whichever distribution in the class the key is drawn from, and a single deterministic extractor must work for every distribution in the class. Extraction error is measured by statistical distance from uniform. Formal definitions are in Section 2.

Theorem 1.1 (Main result)**.** *For every integer $b \ge 2$ there are a source class $\mathcal{S}_b$ on $\{0,1\}^{b+2}$ and a perfectly correct two-out-of-two secret sharing scheme for b-bit messages that is perfectly private for $\mathcal{S}_b$, such that:*

1. *each share has $b+1$ bits;*
2. *every distribution in $\mathcal{S}_b$ is uniform on a set of size between 2^b and 2^{b+1}, and hence has min-entropy at least b;*
3. *for every $\ell \ge 1$ and every function $\mathrm{Ext} : \{0,1\}^{b+2} \to \{0,1\}^\ell$,*

$$\max_{P \in \mathcal{S}_b} \mathrm{TV}(\mathrm{Ext}(K), U_\ell) \ge \frac{1}{6}, \qquad K \sim P.$$

Sharing and reconstruction use a constant number of operations in $\mathbb{F}_{2^b}$.

The extractor in item 3 is arbitrary: it may be computationally unbounded and may depend on the whole key. The key has $d = b + 2$ bits, so the secret is far longer than $\log d$, which is the regime in which encryption forces extraction [BD07]. The constant $1/6$ cannot be improved for this class: a suitable function of the key has error at most $1/6$ on every source in $\mathcal{S}_b$ (Theorem 4.1 with $n = 2$).

Other thresholds. Aggarwal, Chung, Obremski, and Ribeiro [ACOR22] also asked the question for t-out-of-n sharing. Their random-less reductions turn any t-out-of-n scheme with $t \geq 2$ into a two-out-of-two scheme for the same source class, so a separation for any threshold gives one for two-out-of-two, and two-out-of-two is the easiest case in which to separate sharing from extraction. The reductions do not go in the other direction (see the end of Section 5), so Theorem 1.1 by itself says nothing about other thresholds. We extend it to two-out-of-n sharing for every n, and to r-out-of-r sharing for every r.

Theorem 1.2 (see Theorems 4.1 and 5.1)**.**

1. *For every $n \geq 2$ and every b with $2^b > n$, there are a source class on a key set of size $(2n-1)2^b$ and a perfectly correct two-out-of-n scheme for b-bit messages, with $(b+1)$-bit shares, that is perfectly private for this class, such that every function* Ext *to $\{0,1\}$ has statistical error at least $1/(4n-2)$ on some distribution in the class.*

2. *For every $r \geq 2$ and every $b \geq r$, there are a source class on a key set of size $(2^r-1)2^{(r-1)b}$ and a perfectly correct r-out-of-r scheme for b-bit messages, with $(b+1)$-bit shares, that is perfectly private for this class, such that every function* Ext *to $\{0,1\}$ has statistical error at least $1/(2(2^r-1))$ on some distribution in the class.*

In both cases every distribution in the class is flat, with min-entropy at least $\log|\mathcal{K}| - \log 3$, where $\mathcal{K}$ is the key set, and the bound on the error is tight for the class.

1.2 Technical overview

For the purpose of the technical overview, we focus on two-out-of-two sharing.

Why the argument for encryption does not transfer. The proof of Bosley and Dodis considers, for each key k, the distribution of the ciphertext of a uniform message under k. Decryptability makes each of these 2^d distributions flat with min-entropy b, and so a random function f extracts from all these distributions simultaneously when $b > \log d$ and hence from $\mathrm{Enc}(U_b, K)$ where K is a valid key source; privacy lets one replace the uniform b-bit message U_b by the fixed message 0^b, and $k \mapsto f(\mathrm{Enc}(0^b, k))$ is the extractor.

Let us see which part of this idea carries over to two-out-of-two secret sharing, and where it breaks down. Let the two shares of a b-bit message m be $\mathrm{Shr}_1(m, K)$ and $\mathrm{Shr}_2(m, K)$. Given $\mathrm{Shr}_2(U_b, K)$, the share $\mathrm{Shr}_1(U_b, K)$ determines U_b, which is independent of $\mathrm{Shr}_2(U_b, K)$ by privacy; hence $\mathrm{Shr}_1(U_b, K)$ has min-entropy at least b. By privacy, $\mathrm{Shr}_1(0^b, K)$ has the same distribution as $\mathrm{Shr}_1(U_b, K)$, so it too has min-entropy at least b. To extract from it with a single function, as Bosley and Dodis do, one would write its distribution, for every source in the class, as a convex combination of a fixed and "small" family of high-min-entropy sources. For encryption this family is $\{\mathrm{Enc}(U_b, k)\}_k$; for secret sharing we see no analogous family in general. Unfortunately, any attempt to try to do this has to restrict the share size or the key size to be nearly b.

For secret sharing, the object that is private (one share) and the object that determines the message (both shares) are different. A single share, as a function of the message, need not be injective, and may even be constant.

Our counter-example exploits exactly this difference as we will explain below.

The three-key gadget. Let $V = \mathbb{F}_{2^b}$ (which has the same additive structure as $\mathbb{F}_2^b$) and let $\alpha \notin \{0,1\}$ be any fixed element of $\mathbb{F}_{2^b}$. For every index $i \in V$ there are three keys. The *central*

key C_i shares m as $(i+m, \alpha \cdot i + m)$, where addition and multiplication are in $\mathbb{F}_{2^b}$; each share is a one-time-padded copy of m, and together they determine m. The *left* key L_i gives the first party the same share $i+m$ and gives the second party a dummy $\bar{i}$, a tagged copy of the index i. The tagged values $\{\bar{i} : i \in \mathbb{F}_{2^b}\}$ are distinct from the untagged values in $\mathbb{F}_{2^b}$. The *right* key R_i is symmetric, and gives the first party the share $\bar{i}$, and the second party the share $\alpha \cdot i + m$. The sources in our class are the uniform distributions on sets of keys that contain, for each i, either the single key C_i or the pair $\{L_i, R_i\}$; there are 2^{2^b} of them.

Let us first see why this scheme is a perfect 2-out-of-2 secret sharing scheme. Clearly, given either (i) $i+m$ and $\alpha \cdot i + m$, or (ii) $i+m$ and $\bar{i}$, or (iii) $\bar{i}$ and $\alpha \cdot i + m$, it is straightforward to recover m.

Now we argue that neither share reveals any information about the message. Consider a source K in the class, uniform on a set Z of keys. Let $T \subseteq \mathbb{F}_{2^b}$ be the set of indices i with $C_i \in Z$, so that $\mathbb{F}_{2^b} \setminus T$ is the set of indices i with $L_i, R_i \in Z$.

Then, for any fixed message m, the first share is uniformly distributed over the set

$$\{i+m : i \in T\} \ \cup \ \{i+m : i \in \mathbb{F}_{2^b} \setminus T\} \ \cup \ \{\bar{i} : i \in \mathbb{F}_{2^b} \setminus T\} \ ,$$

where the three sets correspond to center keys, left keys, and right keys. This implies that the first share is uniform over $\mathbb{F}_{2^b} \cup \{\bar{i} : i \in \mathbb{F}_{2^b} \setminus T\}$, and hence is independent of the message m.

The argument for why the second share is independent of the message m is identical.

Parity defeats every extractor. We now explain why there is no deterministic extractor that extracts even a single bit from all such randomness sources. For convenience, we will view a candidate one-bit extractor as a labelling h of the keys by $\{-1, +1\}$.

Notice that a valid randomness source K is random over some set of keys Z, and this set Z contains, for each index i, either C_i, or both L_i and R_i. Thus, for h to be an extractor, the output of h must be (almost) unbiased regardless of whether Z contains C_i or both L_i and R_i for any index i. In other words, we want

$$\frac{\sum_{k \in Z} h(k)}{|Z|}$$

to be small.

We will argue that for any given h, there exists a choice of Z for which the above mentioned bias is a constant.

At index i the two allowed choices contribute $h(C_i)$ or $h(L_i) + h(R_i)$ to the signed sum. Their difference $h(C_i) - h(L_i) - h(R_i)$ is a sum of three ± 1 values, and is hence odd, i.e. one of $\{+1, +3, -1, -3\}$.

Let Q contain C_i for every index i with $h(C_i) - h(L_i) - h(R_i) > 0$, and both L_i and R_i for every other index, and let Q^c be the set of the remaining keys, which makes the other choice at every index. Both Q and Q^c are valid choices of Z, and

$$\sum_{k \in Q} h(k) - \sum_{k \in Q^c} h(k) = \sum_i |h(C_i) - h(L_i) - h(R_i)| \geq 2^b, \qquad |Q| + |Q^c| = 3 \cdot 2^b.$$

So one of the two sources has normalized signed sum at least $1/3$ in absolute value, that is, statistical error at least $1/6$.

1.3 Related work

The separation of Dodis, Pietrzak, and Przydatek [DPP06] shares a single bit; its entropy amplification enlarges the randomness space while retaining that message space. The small-message

counterexamples for encryption in [BD07] yield sharing counterexamples as well, so the long-secret regime is essential to our claim. Randomness requirements of other primitives have been studied too: Santha–Vazirani and min-entropy sources fail already for one-bit encryption [MP91, DOPS04], computational assumptions can partially substitute for extraction [KLR09], and in [ACM+14], the authors prove impossibility results for encryption, commitment and zero knowledge under efficient online tampering with the random tape.

Dodis and Yao [DY15, Definition 9, Theorem 4(c), and Remark 2] derive a different, weak extraction consequence from secret sharing. Their extractor may output a failure symbol, and the lower bound on its success probability is required under the uniform input distribution rather than under every allowed source. This does not give the deterministic extraction considered here. The positive results of [ACOR22] require leakage resilience or non-malleability; Remark 3.4 shows the one-bit leakage attack on our scheme that their theorem predicts.

Yao and Li [YL25] state an extraction implication for ordinary threshold secret sharing. Their Proposition 2 extends the encryption reduction by applying privacy to the complete tuple of shares, whereas ordinary sharing hides only unauthorized subsets and the full tuple determines the message. Section A identifies the step that fails, and shows that the proposition itself is false for sharing.

Organization. Section 2 contains definitions. Section 3 proves Theorem 1.1. Section 4 considers two-out-of-n sharing and Section 5 extends the result to r-out-of-r secret sharing.

2 Definitions

All logarithms are base two. For a nonempty finite set A, write U_A for the uniform distribution on A, and abbreviate $U_{\{0,1\}^\ell}$ by U_ℓ. For distributions P, P' on the same finite set, the total variational distance (also called statistical distance, or statistical error) is defined as

$$\mathrm{TV}(P,P') = \frac{1}{2}\sum_x |P(x) - P'(x)| = \max_A |P(A) - P'(A)|.$$

For a random variable X we write $\mathrm{TV}(X,P')$ for the distance between the distribution of X and P'. The min-entropy of P is $H_\infty(P) = -\log \max_x P(x)$. A *flat* distribution is uniform on a nonempty subset of its domain.

Definition 2.1 (Threshold sharing from a source class)**.** Let $\mathcal{K}$ be a finite key set, $\mathcal{M}$ a message set, and $\mathcal{S}$ a class of distributions on $\mathcal{K}$. A *t-out-of-n sharing scheme* Π consists of a deterministic map

$$\mathrm{Shr} = (\mathrm{Shr}_1, \ldots, \mathrm{Shr}_n) : \mathcal{M} \times \mathcal{K} \to \mathcal{C}_1 \times \cdots \times \mathcal{C}_n$$

together with a reconstruction map for every set of at least t parties. The dealer takes one sample K from an unknown $P \in \mathcal{S}$ and outputs $\mathrm{Shr}(m,K)$; there is no other randomness. For $J \subseteq [n]$ write Shr_J for the shares of the parties in J.

- Π is *perfectly correct* if, for every J with $|J| \geq t$, all $m \in \mathcal{M}$, and all $k \in \mathcal{K}$, the reconstruction map Rec_J satisfies
$$\mathrm{Rec}_J(\mathrm{Shr}_J(m,k)) = m.$$
- Π is *δ-private for $\mathcal{S}$* if for every $P \in \mathcal{S}$, every J with $|J| < t$, and all $m, m' \in \mathcal{M}$,
$$\mathrm{TV}(\mathrm{Shr}_J(m,K), \mathrm{Shr}_J(m',K)) \leq \delta, \qquad K \sim P.$$
It is *perfectly private* if $\delta = 0$.

For $t = n = 2$ we write $F = \mathrm{Shr}_1$ and $G = \mathrm{Shr}_2$.

The distribution of a share may depend on P; it must not depend on the message.

Definition 2.2 (Extraction error)**.** For $\mathrm{Ext} : \mathcal{K} \to [D]$ and a class $\mathcal{S}$ of distributions on $\mathcal{K}$, let

$$\mathrm{err}(\mathrm{Ext}, \mathcal{S}) = \sup_{P \in \mathcal{S}} \mathrm{TV}\big(\mathrm{Ext}(K), U_{[D]}\big), \qquad K \sim P.$$

The class is (ℓ, ε)-*extractable* if some $\mathrm{Ext} : \mathcal{K} \to \{0,1\}^\ell$ has $\mathrm{err}(\mathrm{Ext}, \mathcal{S}) \le \varepsilon$.

This is the formulation of [BD07, ACOR22]. Composing an ℓ-bit extractor with the projection to its first bit does not increase the error, so a lower bound on one-bit error applies to every output length.

3 The construction

Let $b \ge 2$, $V = \mathbb{F}_{2^b}$ and $B = 2^b$. Fix an $\alpha \in \mathbb{F}_{2^b} \setminus \{0, 1\}$.

The key set is $\mathcal{K} = \{L, R, C\} \times V$, and we write L_i, R_i, C_i for $(L, i), (R, i), (C, i)$. Both share alphabets are $\{0,1\} \times V$; we write x for $(0, x)$ and $\bar{x}$ for $(1, x)$, and call the latter *tagged*. The bar is a formal tag, not an operation on V. The message set is $\mathcal{M} = V$, and the scheme Π_b shares $m \in V$ as follows:

key k	$F(m,k)$	$G(m,k)$
L_i	$i + m$	$\bar{i}$
R_i	$\bar{i}$	$\alpha \cdot i + m$
C_i	$i + m$	$\alpha \cdot i + m$

(1)

For every $T \subseteq V$ define

$$\Omega_T = \{C_i : i \in T\} \cup \{L_i, R_i : i \notin T\}, \qquad Z_T = |\Omega_T| = 2B - |T|, \qquad P_T = U_{\Omega_T}. \tag{2}$$

The source class is $\mathcal{S}_b = \{P_T : T \subseteq V\}$. It is fixed once and for all, independently of the message and of any candidate extractor. Since $B \le Z_T \le 2B$, every P_T is flat with min-entropy at least b.

3.1 Correctness, privacy, and non-extractability

Lemma 3.1 (Correctness)**.** *There is a map* Rec *with* $\mathrm{Rec}(F(m,k), G(m,k)) = m$ *for all* $m \in V$ *and* $k \in \mathcal{K}$.

Proof. The pattern of tags identifies the branch. If only the second share is tagged, the shares are $(x, \bar{i})$ with $x = i + m$, and Rec outputs $x + i$. If only the first share is tagged, the shares are $(\bar{i}, y)$ with $y = \alpha \cdot i + m$, and Rec outputs $y + \alpha \cdot i$. If neither share is tagged, the shares are $(x, y) = (i + m, \alpha \cdot i + m)$, so $y + \alpha \cdot x = (\alpha + 1) \cdot m$ and Rec outputs $(\alpha \cdot x + y)/(\alpha + 1)$. □

Lemma 3.2 (Privacy)**.** *For every* $T \subseteq V$ *and every* $m \in V$*, if* $K \sim P_T$ *then each of* $F(m, K)$ *and* $G(m, K)$ *is uniform on*

$$V \ \cup \ \{\bar{i} : i \notin T\}. \tag{3}$$

In particular Π_b *is perfectly private for* $\mathcal{S}_b$.

Proof. Fix T and m. For each i, exactly one of L_i, C_i lies in Ω_T, and either key has first share $i+m$. As $i \mapsto i+m$ is a permutation of V, each untagged value has probability exactly $1/Z_T$. The tagged first shares arise from the keys R_i with $i \notin T$; each value $\bar{i}$ with $i \notin T$ has probability $1/Z_T$.

For the second share the argument is the same: exactly one of R_i, C_i lies in Ω_T, either key has second share $\alpha \cdot i + m$, and $i \mapsto \alpha \cdot i + m$ is a permutation because $\alpha \neq 0$; the tagged values come from L_i, $i \notin T$. The set (3), on which each share is uniform over, does not depend on m. □

Lemma 3.3 (No bit can be extracted)**.** *For every* $\mathrm{Ext} : \mathcal{K} \to \{0,1\}$ *there is* $T \subseteq V$ *such that* Ext *has statistical error at least* $1/6$ *on* P_T *or on* P_{T^c}*, where* $T^c = V \setminus T$*.*

Proof. Let $h(k) = (-1)^{\mathrm{Ext}(k)}$ and, for $T \subseteq V$, let

$$H_T := \sum_{k \in \Omega_T} h(k) = \sum_{i \in T} h(C_i) + \sum_{i \in T^c} \Big(h(L_i) + h(R_i) \Big) .$$

For $K \sim P_T$ the distance of $\mathrm{Ext}(K)$ from U_1 equals $|H_T|/(2Z_T)$. For each $i \in V$ the integer

$$v_i = h(C_i) - h(L_i) - h(R_i)$$

is a sum of three odd numbers and is therefore odd; in particular $|v_i| \geq 1$. Let $T = \{i : v_i > 0\}$. Then, $T^c = \{i : v_i < 0\}$. In particular, for each $i \in T$, v_i is either 1 or 3, and for each i in T^c, v_i is either -1, or -3. Then

$$H_T - H_{T^c} = \sum_{i \in T} v_i - \sum_{i \notin T} v_i = \sum_{i \in V} |v_i| \geq B, \qquad Z_T + Z_{T^c} = 3B.$$

Consequently

$$\max \left\{ \frac{|H_T|}{Z_T}, \frac{|H_{T^c}|}{Z_{T^c}} \right\} \geq \frac{|H_T| + |H_{T^c}|}{Z_T + Z_{T^c}} \geq \frac{H_T - H_{T^c}}{3B} \geq \frac{1}{3}, \tag{4}$$

and the source corresponding to the maximum (Ω_T or Ω_{T^c}) has error at least $1/6$. □

3.2 Bit strings and running time

Proof of Theorem 1.1. Encode a key by two branch bits followed by the b-bit index, and a share by its tag bit followed by b bits. Three of the four branches are L, R, C. To define the scheme on all of $\{0,1\}^{b+2}$, let the fourth branch consist of keys $\bar{C}_i$, tagged copies of the central keys, with

$$\mathrm{Shr}(m, \bar{C}_i) = (\overline{i+m}, \overline{\alpha \cdot i + m});$$

on two tagged shares, Rec removes the tags and applies the rule of the central branch. Correctness on $\mathcal{K}$ is Lemma 3.1. Every P_T assigns probability zero to the fourth branch, so Lemma 3.2 gives privacy. The set T in Lemma 3.3 depends only on the restriction of the extractor to $\mathcal{K}$, so the lemma applies verbatim to functions on $\{0,1\}^{b+2}$, and the remark after Definition 2.2 extends it to every output length. Items 1 and 2 are immediate.

The running time for sharing and reconstruction algorithms is immediate from the construction. □

Remark 3.4 (Consistency with the leakage-resilient case)**.** By [ACOR22], a class that supports a two-out-of-two scheme resisting one bit of leakage from each share is extractable, so Π_b must be vulnerable to such leakage. It is. Write $\mathrm{lsb}(z)$ for the least significant bit of $z \in V$ in a fixed basis over $\mathbb{F}_2$, which is $\mathbb{F}_2$-linear in z. For untagged shares x and y of the two parties let $f(x) = \mathrm{lsb}(\alpha \cdot x)$

and $g(y) = \mathrm{lsb}(y)$ (on tagged shares f and g may be arbitrary). Under the source P_V, which is uniform on the central keys, the two leaked bits satisfy

$$f(i+m) + g(\alpha \cdot i + m) = \mathrm{lsb}\big((\alpha+1)\cdot m\big) \qquad \text{for every } i,$$

which is a nontrivial linear functional of the secret. Hence the leakage distributions of two secrets on which this functional differs have disjoint supports.

4 Two-out-of-n sharing

We now extend the construction of Section 3 from two to $n \geq 2$ parties, keeping the threshold at two. The left and right keys become n *side keys*, one for each party, and the central key is repeated $n-1$ times.

4.1 The construction

Let $n \geq 2$ and let b satisfy $2^b > n$. Let $V = \mathbb{F}_{2^b}$ and $B = 2^b$, and fix distinct nonzero elements $\alpha_1, \ldots, \alpha_n \in V$, which exist because $2^b - 1 \geq n$. For every index $i \in V$ there are $2n-1$ keys: $n-1$ *central* keys $C_i^{(1)}, \ldots, C_i^{(n-1)}$ and n *side* keys $X_i^{(1)}, \ldots, X_i^{(n)}$. The key set is

$$\mathcal{K}_n = \{C_i^{(c)} : c \in [n-1],\ i \in V\} \cup \{X_i^{(j)} : j \in [n],\ i \in V\}, \qquad |\mathcal{K}_n| = (2n-1)B.$$

As in Section 3, the message set is V, every share alphabet is $\{0,1\} \times V$, and we write x for the untagged value $(0,x)$ and $\bar{x}$ for the tagged value $(1,x)$. The scheme $\Pi_{b,n}$ gives party $j \in [n]$ the share

$$\mathrm{Shr}_j(m, C_i^{(c)}) = \alpha_j i + m, \qquad \mathrm{Shr}_j(m, X_i^{(j')}) = \begin{cases} \bar{i} & \text{if } j = j', \\ \alpha_j i + m & \text{if } j \neq j'. \end{cases} \tag{5}$$

Thus every central key gives every party its *true share* $\alpha_j i + m$, and the side key $X_i^{(j')}$ gives party j' the dummy $\bar{i}$ and every other party its true share. The $n-1$ central keys at an index give identical shares; they differ only as keys. For $n = 2$, with $\alpha_1 = 1$ and $\alpha_2 = \alpha$, this is the scheme Π_b of (1), with $C_i = C_i^{(1)}$, $L_i = X_i^{(2)}$, and $R_i = X_i^{(1)}$.

For every $T \subseteq V$ define

$$\begin{gathered} \Omega_T = \{C_i^{(c)} : i \in T,\ c \in [n-1]\} \cup \{X_i^{(j)} : i \notin T,\ j \in [n]\}, \\ Z_T = |\Omega_T| = nB - |T|, \qquad P_T = U_{\Omega_T}. \end{gathered} \tag{6}$$

The source class is $\mathcal{S}_{b,n} = \{P_T : T \subseteq V\}$. As in Section 3, at every index a source contains either all central keys or all side keys, and for $n = 2$ these are the sources of (2). Since $(n-1)B \leq Z_T \leq nB$, every P_T is flat with min-entropy at least $b + \log(n-1)$.

The number of central keys is dictated by privacy. Fix a party j. At an index $i \notin T$, party j receives its true share $\alpha_j i + m$ from the $n-1$ side keys $X_i^{(j')}$ with $j' \neq j$, and a dummy from $X_i^{(j)}$. The $n-1$ central keys make it receive its true share equally often at an index $i \in T$, and therefore every untagged value is equally likely, whatever T and m are.

Theorem 4.1. *Let $n \geq 2$ and let b satisfy $2^b > n$. The scheme $\Pi_{b,n}$ is a perfectly correct two-out-of-n scheme for b-bit messages that is perfectly private for $\mathcal{S}_{b,n}$, and:*

1. *each share has $b+1$ bits, and the key set $\mathcal{K}_n$ has $(2n-1)2^b$ elements;*

2. *every distribution in $\mathcal{S}_{b,n}$ is uniform on a set of between $(n-1)2^b$ and $n2^b$ keys, and hence has min-entropy at least $b + \log(n-1)$;*

3. *for every $\ell \geq 1$ and every function $\mathrm{Ext} : \mathcal{K}_n \to \{0,1\}^\ell$,*

$$\max_{P \in \mathcal{S}_{b,n}} \mathrm{TV}(\mathrm{Ext}(K), U_\ell) \;\geq\; \frac{1}{4n-2}, \qquad K \sim P;$$

4. *for $\ell = 1$, some function of the index i alone attains equality in item 3.*

For $n = 2$ the scheme is Π_b and the bound is the $1/6$ of Theorem 1.1; item 4 shows that this bound is exact for $\mathcal{S}_b$.

4.2 Correctness, privacy, and non-extractability

Lemma 4.2 (Correctness)**.** *For any two distinct parties j, j' there is a map $\mathrm{Rec}_{\{j,j'\}}$ such that*

$$\mathrm{Rec}_{\{j,j'\}}(\mathrm{Shr}_j(m,k), \mathrm{Shr}_{j'}(m,k)) = m \qquad \text{for all } m \in V \text{ and } k \in \mathcal{K}_n.$$

Proof. No key gives a dummy to two parties, so at least one of the two shares is untagged, and, as in Lemma 3.1, the pattern of tags identifies the case. If party j holds the dummy $\bar{i}$ and party j' holds $y = \alpha_{j'} i + m$, then $\mathrm{Rec}_{\{j,j'\}}$ outputs $y + \alpha_{j'} i$; the case in which party j' holds the dummy is symmetric. If neither share is tagged, the shares are $x = \alpha_j i + m$ and $y = \alpha_{j'} i + m$, so $x + y = (\alpha_j + \alpha_{j'})i$ with $\alpha_j \neq \alpha_{j'}$, and $\mathrm{Rec}_{\{j,j'\}}$ outputs $x + \alpha_j(x+y)/(\alpha_j + \alpha_{j'})$. □

Lemma 4.3 (Privacy)**.** *For every $T \subseteq V$, every $m \in V$, and every party j, if $K \sim P_T$ then $\mathrm{Shr}_j(m, K)$ takes each untagged value $x \in V$ with probability $(n-1)/Z_T$ and each tagged value $\bar{i}$ with $i \notin T$ with probability $1/Z_T$. In particular $\Pi_{b,n}$ is perfectly private for $\mathcal{S}_{b,n}$.*

Proof. Fix T, m, and j. For each $i \in V$, exactly $n-1$ keys of Ω_T at index i give party j the share $\alpha_j i + m$: the central keys $C_i^{(1)}, \ldots, C_i^{(n-1)}$ if $i \in T$, and the side keys $X_i^{(j')}$ with $j' \neq j$ if $i \notin T$. Since $\alpha_j \neq 0$, the map $i \mapsto \alpha_j i + m$ is a permutation of V, so each untagged value has probability exactly $(n-1)/Z_T$. The tagged shares of party j come from the keys $X_i^{(j)}$ with $i \notin T$, one for each such i. This distribution does not depend on m. □

For $n = 2$ the distribution is uniform, as in Lemma 3.2; for $n \geq 3$ it is not uniform, but it still does not depend on the message.

Lemma 4.4 (No bit can be extracted)**.** *For every $\mathrm{Ext} : \mathcal{K}_n \to \{0,1\}$ there is $T \subseteq V$ such that Ext has statistical error at least $1/(4n-2)$ on P_T or on P_{T^c}. Conversely, some Ext that depends only on the index i has error at most $1/(4n-2)$ on every source in $\mathcal{S}_{b,n}$, with equality for some source.*

Proof. The first statement is proved exactly as Lemma 3.3, with the $n-1$ central keys in place of C_i and the n side keys in place of L_i, R_i. Let $h(k) = (-1)^{\mathrm{Ext}(k)}$ and $H_T = \sum_{k \in \Omega_T} h(k)$, so that Ext has error $|H_T|/(2Z_T)$ on P_T. For each $i \in V$ the integer

$$v_i = \sum_{c \in [n-1]} h(C_i^{(c)}) - \sum_{j \in [n]} h(X_i^{(j)})$$

is a sum of $2n-1$ odd numbers and is therefore odd; in particular $|v_i| \geq 1$. With $T = \{i : v_i > 0\}$ we get, as in the proof of Lemma 3.3,

$$H_T - H_{T^c} = \sum_{i \in V} |v_i| \geq B, \qquad Z_T + Z_{T^c} = (2n-1)B,$$

and (4) becomes $\max\{|H_T|/Z_T, |H_{T^c}|/Z_{T^c}\} \geq 1/(2n-1)$. So Ext has error at least $1/(4n-2)$ on P_T or on P_{T^c}.

For the second statement, let $g : V \to \{0,1\}$ be balanced and let $\mathrm{Ext}(k) = g(i)$ for every key k at index i. Fix T. Under P_T, index i carries w_i keys, where $w_i = n-1$ if $i \in T$ and $w_i = n$ otherwise. With $A_\beta = \sum_{i:\, g(i)=\beta} w_i$ we have $\Pr[\mathrm{Ext}(K) = 1] = A_1/(A_0 + A_1)$, and both A_0 and A_1 lie between $(n-1)B/2$ and $nB/2$. Hence $\Pr[\mathrm{Ext}(K) = 1]$ lies between $(n-1)/(2n-1)$ and $n/(2n-1)$, so the error is at most $1/(4n-2)$, with equality for $T = g^{-1}(0)$. □

Lemmas 4.2 to 4.4 prove Theorem 4.1: a set of more than two parties reconstructs from any two of its members, and the remark after Definition 2.2 extends Lemma 4.4 to ℓ output bits.

5 r-out-of-r secret sharing

We now extend our result to r-out-of-r secret sharing. Fix $r \geq 2$ and $b \geq r$. Let $V = \mathbb{F}_{2^b}$, $B = 2^b$, and $N = B^{r-1}$, and fix $\gamma_1, \ldots, \gamma_r \in V$ linearly independent over $\mathbb{F}_2$. Such $\gamma_1, \ldots, \gamma_r$ exist because $b \geq r$. The index is now a vector $u = (u_1, \ldots, u_{r-1}) \in V^{r-1}$, and the *pad* of party j is

$$p_j(u) = u_j \quad (j < r), \qquad p_r(u) = u_1 + \cdots + u_{r-1}.$$

We use two properties of the pads:

(P1) $p_1(u) + \cdots + p_r(u) = 0$ for every u;

(P2) for every set J of $r-1$ parties, the map $p_J(u) = (p_j(u))_{j \in J}$ is a bijection from V^{r-1} to V^J.

A key is a pair (u, S) with $u \in V^{r-1}$ and $\emptyset \neq S \subseteq [r]$; let $\mathcal{K}^{(r)}$ be the set of keys. The message set is V, and the scheme $\Pi_b^{(r)}$ shares $m \in V$ as

$$\mathrm{Shr}_j(m, (u, S)) = \begin{cases} p_j(u) + \gamma_j m & \text{if } j \in S, \\ \overline{p_j(u)} & \text{if } j \notin S. \end{cases} \tag{7}$$

We say that the key (u, S) *activates* the parties in S. Thus an active party receives its pad plus a multiple of the secret, and an inactive party receives its pad, tagged, as a dummy. For a share x or $\bar{x}$ we call $x \in V$ its *value*.

Let $\mathcal{O}_r$ be the family of subsets of $[r]$ of odd size and $\mathcal{E}_r$ the family of nonempty subsets of even size, so that $|\mathcal{O}_r| = 2^{r-1}$ and $|\mathcal{E}_r| = 2^{r-1} - 1$. For every $T \subseteq V^{r-1}$ define

$$\begin{gathered} \Omega_T = \{(u, S) : u \in T,\ S \in \mathcal{O}_r\} \cup \{(u, S) : u \notin T,\ S \in \mathcal{E}_r\}, \\ Z_T = |\Omega_T| = (2^{r-1} - 1)N + |T|, \qquad P_T = U_{\Omega_T}. \end{gathered} \tag{8}$$

The source class is $\mathcal{S}_b^{(r)} = \{P_T : T \subseteq V^{r-1}\}$. At every index a source contains either the 2^{r-1} keys with an odd set or the $2^{r-1} - 1$ keys with a nonempty even set, just as in Section 3 it contains either C_i or L_i, R_i.

For $r = 2$ we have $\mathcal{O}_2 = \{\{1\}, \{2\}\}$ and $\mathcal{E}_2 = \{\{1, 2\}\}$, and both pads equal u. Take $\gamma_1 = 1$ and $\gamma_2 = \alpha^{-1}$, which are linearly independent over $\mathbb{F}_2$ because $\alpha \notin \{0, 1\}$. After multiplying the untagged shares of party 2 by α, the keys $(u, \{1\})$, $(u, \{2\})$, and $(u, \{1, 2\})$ become L_u, R_u, and C_u of (1), and P_T becomes the source $P_{V \setminus T}$ of Section 3. For $r = 3$ the seven keys at an index are the three singletons and $\{1, 2, 3\}$, used at indices in T, and the three pairs, used at indices outside T.

Theorem 5.1 (r-out-of-r)**.** *Let $r \geq 2$ and $b \geq r$. The scheme $\Pi_b^{(r)}$ is a perfectly correct r-out-of-r scheme for b-bit messages that is perfectly private for $\mathcal{S}_b^{(r)}$, and:*

1. *each share has $b+1$ bits, and the key set $\mathcal{K}^{(r)}$ has $(2^r-1)2^{(r-1)b}$ elements;*
2. *every distribution in $\mathcal{S}_b^{(r)}$ is uniform on a set of between $(2^{r-1}-1)N$ and $2^{r-1}N$ keys, and hence has min-entropy at least $(r-1)b + \log(2^{r-1}-1)$;*
3. *for every $\ell \geq 1$ and every function $\mathrm{Ext} : \mathcal{K}^{(r)} \to \{0,1\}^\ell$,*

$$\max_{P \in \mathcal{S}_b^{(r)}} \mathrm{TV}(\mathrm{Ext}(K), U_\ell) \;\geq\; \frac{1}{2(2^r-1)}, \qquad K \sim P;$$

4. *for $\ell = 1$, some function of the index u alone attains equality in item 3.*

For $r = 2$ the bound is $1/6$, as in Theorem 1.1, and for $r = 3$ it is $1/14$.

Lemma 5.2 (Correctness)**.** *For every $m \in V$ and every key (u, S), the values $z_1, \ldots, z_r$ of the r shares satisfy*

$$z_1 + \cdots + z_r = \Big(\sum_{j \in S} \gamma_j\Big) m,$$

and the coefficient of m is nonzero. Hence the r shares determine m.

Proof. By (7), $z_j = p_j(u) + \gamma_j m$ for $j \in S$ and $z_j = p_j(u)$ for $j \notin S$, so the identity follows from (P1). The coefficient $\sum_{j \in S} \gamma_j$ is nonzero because S is nonempty and the γ_j are linearly independent over $\mathbb{F}_2$. The tags reveal S, so Rec computes this coefficient and divides by it. □

Lemma 5.3 (Privacy)**.** *Let J be a set of $r-1$ parties. Record the shares of J as the pair (R, z), where $R = S \cap J$ is the tag pattern and $z \in V^J$ is the vector of values. For every $T \subseteq V^{r-1}$ and every $m \in V$, if $K \sim P_T$ then (R, z) is uniform on the set*

$$\big\{(R, z) : \emptyset \neq R \subseteq J,\ z \in V^J\big\} \;\cup\; \big\{(\emptyset, z) : z \in p_J(T)\big\}, \tag{9}$$

which has size Z_T and does not depend on m. In particular $\Pi_b^{(r)}$ is perfectly private for $\mathcal{S}_b^{(r)}$.

Proof. Fix T and m, and let j_0 be the party outside J. Fix a nonempty $R \subseteq J$. A key (u, S) shows J the tag pattern R, that is, $S \cap J = R$, exactly when $S = R$ or $S = R \cup \{j_0\}$. These two sets are nonempty and have different parities, so one of them lies in $\mathcal{O}_r$ and the other in $\mathcal{E}_r$. Hence Ω_T contains exactly one key at index u that shows J the pattern R, whether or not $u \in T$. That key gives J the values $p_J(u) + (\mathbf{1}[j \in R]\,\gamma_j m)_{j \in J}$, which by (P2) run over V^J exactly once as u runs over V^{r-1}. For the empty pattern, the key must activate $\{j_0\}$, which is odd; so Ω_T contains such a key at index u exactly when $u \in T$, and it gives J the values $p_J(u)$, which do not involve m. Altogether, every element of (9) is produced by exactly one key of Ω_T. A coalition of fewer than $r-1$ parties sees a projection of the shares of a coalition of $r-1$ parties, and so its view does not depend on m either. □

Lemma 5.4 (No bit can be extracted)**.** *For every $\mathrm{Ext} : \mathcal{K}^{(r)} \to \{0,1\}$ there is $T \subseteq V^{r-1}$ such that* Ext *has statistical error at least $1/(2(2^r-1))$ on P_T or on P_{T^c}. Conversely, some* Ext *that depends only on the index u has error at most $1/(2(2^r-1))$ on every source in $\mathcal{S}_b^{(r)}$, with equality for some source.*

Proof. This is the proof of Lemma 4.4, with the 2^{r-1} keys (u, S), $S \in \mathcal{O}_r$, in place of the central keys at an index and the $2^{r-1}-1$ keys (u, S), $S \in \mathcal{E}_r$, in place of the side keys. Now

$$v_u = \sum_{S \in \mathcal{O}_r} h(u, S) - \sum_{S \in \mathcal{E}_r} h(u, S)$$

is a sum of $2^r - 1$ odd numbers, hence odd, and $Z_T + Z_{T^c} = (2^r - 1)N$, which gives the first statement. For the second, under every P_T each index carries $2^{r-1} - 1$ or 2^{r-1} keys, and the computation in the proof of Lemma 4.4, with $n-1$ and n replaced by these two numbers, shows that a balanced function g of u has error at most $1/(2(2^r - 1))$, with equality for $T = g^{-1}(1)$. □

Combining Lemmas 5.2 to 5.4, we prove Theorem 5.1; the remark after Definition 2.2 extends Lemma 5.4 to ℓ output bits.

Relation to the reductions of [ACOR22]. A (t, n, t', n')-distribution design hands each of n new parties the shares of a set of old parties of a (t', n')-scheme, so that any t new parties together hold the shares of at least t' old parties, and any $t-1$ new parties hold the shares of fewer than t'. It converts every (t', n')-scheme into a (t, n)-scheme with the same keys, and hence private for the same source class. For $t = n = 2$ such a design exists whenever $2 \le t' \le n'$: give the first new party the shares of old parties $1, \ldots, t'-1$ and the second the share of old party t'. Consequently a separation for any (t', n') implies one for $(2, 2)$. Designs do not go in the other direction, so neither Theorem 4.1 for $n \ge 3$ nor Theorem 5.1 for $r \ge 3$ can be derived from Theorem 1.1 in this way. A $(2, 3, 2, 2)$-design would give three new parties subsets of $\{1, 2\}$ of size at most one whose pairwise unions all equal $\{1, 2\}$, and no such subsets exist; and in an $(r, r, 2, 2)$-design with $r \ge 3$, some $r-1$ of the r subsets would already cover $\{1, 2\}$.

Acknowledgments

This section is primarily an acknowledgment of defeat. Divesh Aggarwal (the person whose name appears as an author on this paper, henceforth DA) pursued this problem for more than ten years as a curiosity question, and discussed it with many people, including Yevgeniy Dodis, Maciej Obremski, Mingyuan Wang, Eldon Chung, Tomasz Kazana, Arka Rai Choudhuri, and Siyao Guo.

Throughout this period the conjecture was that every source that supports two-out-of-two secret sharing is extractable, with the lingering thought that an extractor, if it existed, could not be very complicated, so that the failure to find one hinted at a counterexample. All attempts to prove the conjecture or to find a counterexample failed.

In [ACOR22], DA and coauthors considered a relaxation in which the scheme is additionally leakage-resilient or non-malleable, and a straightforward argument showed that every source supporting such a scheme is extractable.

Between June and August 2026, DA made several attempts to find an extractor with ChatGPT 5.5 Pro and ChatGPT 5.6 Pro/Ultra. In most of these attempts the model was fairly convinced that a deterministic extractor exists, and it proposed constructions that worked in every case except a special one.

In September 2026, the same question was posed to ChatGPT 6.0, which, after several attempts at a proof, stated that the conjecture is false and found a simple counterexample. Subsequent prompts generalized and cleaned up the counterexample, with Claude Fable 5.1 joining the effort. The paper was written with the assistance of Claude Fable 5.1; DA has verified everything to the best of his ability and takes full responsibility for the text.

References

[ACM+14] Per Austrin, Kai-Min Chung, Mohammad Mahmoody, Rafael Pass, and Karn Seth. On the impossibility of cryptography with tamperable randomness. In *Advances in Cryptology—CRYPTO 2014, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 462–479. Springer, 2014.

[ACOR22] Divesh Aggarwal, Eldon Chung, Maciej Obremski, and João Ribeiro. On secret sharing, randomness, and random-less reductions for secret sharing. In *Theory of Cryptography (TCC 2022), Part I*, volume 13747 of *Lecture Notes in Computer Science*, pages 327–354. Springer, 2022.

[BD07] Carl Bosley and Yevgeniy Dodis. Does privacy require true randomness? In *Theory of Cryptography (TCC 2007)*, volume 4392 of *Lecture Notes in Computer Science*, pages 1–20. Springer, 2007.

[Bla79] George R. Blakley. Safeguarding cryptographic keys. In *Proceedings of the National Computer Conference*, volume 48 of *AFIPS Conference Proceedings*, pages 313–317, 1979.

[CG88] Benny Chor and Oded Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM Journal on Computing*, 17(2):230–261, 1988.

[Chu24] Eldon Chung. *Thinking Inside the Box: Privacy Against Stronger Adversaries*. PhD thesis, National University of Singapore, 2024. arXiv:2406.16313.

[DOPS04] Yevgeniy Dodis, Shien Jin Ong, Manoj Prabhakaran, and Amit Sahai. On the (im)possibility of cryptography with imperfect randomness. In *45th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 196–205. IEEE Computer Society, 2004.

[DPP06] Yevgeniy Dodis, Krzysztof Pietrzak, and Bartosz Przydatek. Separating sources for encryption and secret sharing. In *Theory of Cryptography (TCC 2006)*, volume 3876 of *Lecture Notes in Computer Science*, pages 601–616. Springer, 2006.

[DS02] Yevgeniy Dodis and Joel Spencer. On the (non)universality of the one-time pad. In *43rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 376–385. IEEE Computer Society, 2002.

[DY15] Yevgeniy Dodis and Yanqing Yao. Privacy with imperfect randomness. In *Advances in Cryptology—CRYPTO 2015, Part II*, volume 9216 of *Lecture Notes in Computer Science*, pages 463–482. Springer, 2015.

[KLR09] Yael Tauman Kalai, Xin Li, and Anup Rao. 2-source extractors under computational assumptions and cryptography with defective randomness. In *50th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 617–626. IEEE Computer Society, 2009.

[MP91] James L. McInnes and Benny Pinkas. On the impossibility of private key cryptography with weakly random keys. In *Advances in Cryptology—CRYPTO '90*, volume 537 of *Lecture Notes in Computer Science*, pages 421–435. Springer, 1991.

[Sha49] Claude E. Shannon. Communication theory of secrecy systems. *Bell System Technical Journal*, 28(4):656–715, 1949.

[Sha79] Adi Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.

[SV84] Miklós Santha and Umesh V. Vazirani. Generating quasi-random sequences from slightly-random sources. In *25th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 434–440. IEEE Computer Society, 1984.

[YL25] Yanqing Yao and Zhoujun Li. Novel improvements and extensions of the extractable results about (leakage-resilient) privacy schemes with imperfect randomness. *Journal of Information Security and Applications*, 89:104008, 2025.

A Comparison with a statement of Yao and Li

This appendix concerns only the implication for ordinary secret sharing in [YL25]; statement numbers and hypotheses follow the published journal version, against which this appendix has been checked. We record the comparison because Theorem 1.1 is incompatible with that implication, and we indicate the step at which the argument departs from the encryption case.

Their Definition 4 requires privacy for unauthorized subsets. Definition 5 sets $\psi(m,k) = \mathrm{Shr}(m,k)$, the entire tuple of shares, and calls the scheme perfectly binding when $\psi(m,k)$ is injective in m for each fixed k. Perfect reconstruction implies this condition. They define $D_k = \psi(U_b, k)$ and $\mathcal{Z} = \{D_k : k \in \{0,1\}^d\}$, denoted $\mathcal{D}$ there. Proposition 2 states that an ℓ-bit extractor f for $\mathcal{Z}$ with error ε yields an extractor for the original class of key sources with error $\delta + \varepsilon$, namely

$$\widetilde{\mathrm{Ext}}(k) = f(\psi(0^b, k)). \tag{10}$$

The proof substitutes ψ for encryption in [BD07, Lemma 1]. That lemma needs the full output on a fixed message to be close to the full output on a uniform message. For sharing, perfect reconstruction gives instead

$$\mathrm{TV}\big(\mathrm{Shr}(0^b, K), \mathrm{Shr}(U_b, K)\big) = 1 - 2^{-b} \tag{11}$$

for every key distribution, because the distributions of the full tuple under different messages have disjoint supports. Concretely, for $1 \le \ell \le b$ let f reconstruct the message from the tuple and output its first ℓ bits, with arbitrary values outside the range of ψ. Then $f(D_k) = U_\ell$ for every k, while (10) is the constant 0^ℓ. The error is in the statement, not only in the proof: the scheme of Theorem 1.1 is, in their terminology, $(\mathcal{S}_b, \delta)$-secure for every $\delta > 0$, and with this f the family $\mathcal{Z}$ is $(1,0)$-extractable, so Proposition 2 would make $\mathcal{S}_b$ $(1,\delta)$-extractable for every $\delta > 0$, which Theorem 1.1 excludes for $\delta < 1/6$. Theorems 5 and 8 of [YL25] combine Proposition 2 with extraction statements for $\mathcal{Z}$ and fail for threshold sharing in the same way.