TRAVESS: Traceable Threshold VSS with Publicly Verifiable Shares and Tracing Evidence

Wenxing Duan[1][0009-0008-6338-7050], Huiran Zheng[2][0009-0009-5771-3240], and Man Ho Au[1][0000-0003-2068-9530]

[1] The Hong Kong Polytechnic University, HK
[2] University of Cambridge, UK
wenxing.duan@connect.polyu.hk, hz507@cam.ac.uk, mhaau@polyu.edu.hk

Abstract. Secret sharing protects a secret against coalitions below the reconstruction threshold, but it does not stop participants from selling their shares. Traceable secret sharing (TSS) holds participants accountable for such leakage: when a coalition packages its shares into a leaked program, called a *reconstruction box*, that recovers the secret given enough additional shares, a tracer can identify the traitors by querying this box on specially crafted *artificial shares*. Meanwhile, by allowing participants to verify that their shares are well-formed, verifiable secret sharing (VSS) addresses a different threat: a dishonest dealer distributing inconsistent shares. Boneh, Partap, and Rotem (CRYPTO 2024) posed the combination of TSS and VSS as an open problem, observing that share verification allows a reconstruction box to detect and reject the artificial shares used by their tracing procedures. We present TRAVESS, a threshold construction of traceable verifiable secret sharing. To our knowledge, it is the first construction to combine public share verification and reconstruction consistency with secrecy, tracer secrecy, non-imputability, and black-box traceability in the framework of Boneh et al. Our key idea is to use a trapdoor commitment scheme to enable the tracer to generate artificial shares that pass the same public verification as genuine shares. The tracer gains this capability without learning the secret, while the dealer remains bound to consistent shares. TRAVESS supports imperfect reconstruction boxes and produces evidence that anyone can verify using public information alone. We establish these guarantees in an ideal-functionality hybrid model.

1 Introduction

Secret sharing encodes a secret into multiple shares held by different participants. In a threshold scheme, a dealer generates n shares and privately distributes one to each participant; any t participants can subsequently reconstruct the secret by combining their shares. A typical application is storing a wallet's private key across several servers, avoiding the single point of failure created by storing the entire key on one server [7]. The *secrecy*, or *privacy*, requirement ensures that any coalition holding fewer than t shares learns nothing about the secret.

Ordinary secret sharing assumes that the dealer distributes shares correctly, but provides no mechanism for recipients to check their consistency. A malicious dealer may therefore distribute shares that cause different groups of participants to reconstruct different secrets. Verifiable secret sharing (VSS) [10,23] addresses this problem through

share verification and *reconstruction consistency*: all qualifying collections of accepted shares reconstruct the same secret, even when the dealer is malicious. We focus on schemes in which shares can be verified using public information.

Even when shares are distributed consistently, a malicious coalition may package and sell the shares in its possession. Goyal, Song, and Srinivasan (CRYPTO 2021) [16] introduced traceable secret sharing (TSS) to hold participants accountable for such leakage. Boneh, Partap, and Rotem (CRYPTO 2024) [7] subsequently introduced a black-box tracing framework and gave efficient constructions.

The need for black-box tracing arises from how a coalition holding $f < t$ shares may package its shares for sale. Rather than exposing raw shares directly, the coalition may embed them into a reconstruction program (modeled by Boneh et al. as a *reconstruction box* $\mathcal{R}$) that recovers the secret when supplied with $t - f$ additional shares. In building $\mathcal{R}$, the coalition may combine its shares, obfuscate the program code, or apply other unknown post-processing that preserves this reconstruction functionality. Because a tracer cannot be assumed to reverse these transformations even when $\mathcal{R}$'s code is available, it must interact with $\mathcal{R}$ solely by submitting inputs and observing outputs. In addition to secrecy, their framework requires the following properties:

- **Traceability.** Given black-box access to a reconstruction box, the tracer can identify the participants whose shares contributed to it and produce evidence of their involvement. This requirement also covers imperfect boxes that succeed with sufficient probability on fresh honest inputs.
- **Tracer secrecy.** The secret remains hidden from the tracer even when it holds the tracing material together with fewer than t participant shares.
- **Non-imputability.** A malicious tracer cannot fabricate evidence that passes verification and falsely implicates an honest participant in a leak.

The tracing procedures of Boneh et al. query the reconstruction box on carefully chosen artificial shares and use its responses to identify the contributors. This approach does not directly extend to VSS: the artificial shares need not pass verification, so a box can reject tracing queries while continuing to reconstruct correctly on genuine shares. Boneh et al. [7, Sec. 5] therefore explicitly posed the construction of traceable VSS as an open problem. In this work, we ask:

Can we combine VSS with all the security guarantees of black-box TSS, while making both shares and tracing evidence verifiable from public information alone?

We address this question with a *traceable verifiable secret sharing* (TVSS) scheme that retains all the TSS security guarantees and allows both shares and tracing evidence to be verified using public information alone.

1.1 Our Results

We present TRAVESS, a threshold TVSS construction based on Shamir secret sharing with three field elements per participant share (Section 5.6). Our formal definition in Section 4 captures all four security properties of Boneh et al.'s TSS framework [7]:

secrecy, tracer secrecy, non-imputability, and black-box traceability, alongside public share verification, reconstruction consistency, and public verification of tracing evidence. To our knowledge, TRAVESS is the first construction to achieve all of these guarantees simultaneously, resolving Boneh et al.'s open problem in the ideal committed-receiver oblivious linear evaluation ($\mathcal{F}_{\text{CROLE}}$) hybrid model. For the concrete instantiation in Section 5.6, security in this hybrid model relies on the discrete-logarithm assumption and a correct, position-binding vector commitment (Definition 6).

Moreover, TRAVESS supports black-box tracing of *imperfect* reconstruction boxes. Rather than requiring a box to follow the prescribed reconstruction algorithm or succeed on every valid input, our tracing procedure handles any adversarially constructed ϵ-*good box* [7], which is a box whose probability of recovering the secret is at least ϵ, averaged over fresh honest shares and the box's internal randomness. Following Boneh et al., we consider stateless boxes that take $t - f$ additional shares (for a static coalition of size $f < t$) and use fresh independent coins for each response without retaining memory of earlier calls. Given a tracing parameter δ, our tracing algorithm makes $O(f\delta^{-4} + \lambda\delta^{-2})$ oracle queries to the box and runs in polynomial time (Section 5.7 and Appendix A).

Our two main contributions are as follows.

- **Public share verification and reconstruction consistency.** TRAVESS lets recipients verify their shares using public verification information and ensures that any two qualifying collections of verified shares reconstruct to the same secret. To reconcile this guarantee with black-box tracing, the tracer holds a commitment trapdoor that remains unknown to the dealer. This trapdoor, together with privately computed masked coefficients, allows the tracer to generate artificial shares that pass the same public verification as genuine shares when querying a reconstruction box, while a malicious dealer facing an honest tracer remains bound to consistent shares prior to tracing. Conversely, when the dealer is honest, granting this capability to a potentially malicious tracer preserves both tracer secrecy and non-imputability, so that the tracer learns nothing about the secret and cannot falsely accuse an honest participant of leaking its share.
- **Publicly verifiable tracing evidence.** TRAVESS produces evidence that anyone can verify from the public record, without the private tracing key or access to the reconstruction box. We bind the evaluation points recovered by tracing to registered participant identities through openings of a vector commitment fixed during sharing (Section 5.5). Only the accused positions are explicitly opened. Non-imputability ensures that a malicious tracer cannot use this verification procedure to substantiate a false accusation against an honest participant.

TSS-PV [20] combines public share verification with black-box tracing, but provides neither tracer secrecy nor oracle-free evidence verification. Accountable VSS [4] instead combines public share verification with publicly verifiable transcript-based accountability, but does not provide black-box traceability. Neither provides VSS together with all the TSS security guarantees in Boneh et al.'s framework and tracing evidence verifiable without box access. Section 2 and Table 1 give the detailed comparison.

1.2 Technical Overview

We develop TRAVESS in stages, starting from Pedersen VSS. We first enable the tracer to generate queries that pass public share verification, then protect honest participants against false accusations by the tracer. These extensions leave the dealer able to distribute inconsistent shares, so we next revise initialization to restore reconstruction consistency. Finally, we make tracing evidence publicly verifiable through openings at the accused table positions.

1.2.1 Starting from Pedersen VSS. To make shares publicly verifiable without exposing the secret, we build on Pedersen VSS [23] (Section 5.1). To share a secret $s \in \mathbb{F}_q$, the dealer samples a Shamir polynomial P with constant term $a_0 = s$, together with an independent blinding polynomial R, both of degree at most $t-1$:

$$P(X) = \sum_{j=0}^{t-1} a_j X^j, \qquad R(X) = \sum_{j=0}^{t-1} b_j X^j.$$

Using public generators g, h, the dealer publishes a Pedersen commitment $C_j = g^{a_j} h^{b_j}$ for each coefficient pair. Participant P_i is assigned an evaluation point x_i and receives the share

$$\mathsf{sh}_i = (x_i, y_i, \rho_i) = \big(x_i, P(x_i), R(x_i)\big).$$

Anyone can verify this share against the public commitments by checking the verification equation:

$$g^{y_i} h^{\rho_i} = \prod_{j=0}^{t-1} C_j^{x_i^j},$$

which checks the pair (y_i, ρ_i) against the public coefficient commitments. Any t verified shares at distinct evaluation points reconstruct $s = P(0)$ by Lagrange interpolation of their y-coordinates. However, because the evaluation points x_i are public and no tracing material is generated, Pedersen VSS provides no mechanism to identify which participants contributed shares to a leaked reconstruction box.

1.2.2 Adding traceability. To trace a reconstruction box $\mathcal{R}$ built by $f < t$ corrupt participants, we first recall how black-box tracing works in Boneh et al.'s Shamir-based TSS [7, Sec. 3.2]. Instead of using public coordinates, the dealer assigns each participant P_i a private random evaluation point $x_i \in \mathbb{F}_q^*$ and records the ordered list of registered points $\mathbf{x} = (x_1, \ldots, x_n)$ in the tracer's key. To explain the tracing idea, first consider a box that combines its f embedded coalition shares with $t-f$ additional shares by interpolation. The tracer probes $\mathcal{R}$ using paired queries: in each round, it queries $\mathcal{R}$ twice on the same $t-f$ input shares, except that a single input share at point x' has its value perturbed from y' to $y'+\Delta$. When both responses agree with reconstruction and all combined evaluation points are distinct, linearity of Lagrange interpolation ensures that the output difference reveals the Lagrange coefficient of x', which is a rational function of the coalition's hidden evaluation points. With sufficiently many such rounds across

random points x', Reed–Solomon list decoding recovers a list containing the degree-f polynomial whose roots are precisely the coalition's private points. Imperfect boxes may produce erroneous or missing samples; list decoding accommodates these errors, and candidate roots are matched against the registered list $\mathbf{x} = (x_1, \ldots, x_n)$ to filter the candidates.

Directly applying this procedure to Pedersen VSS fails because changing a share's reconstruction value without adjusting its verification randomness invalidates the verification equation above, allowing the box to reject the query. Moreover, tracer secrecy prohibits giving $P(X)$ to the tracer. We overcome this obstacle using the trapdoor equivocation property of Pedersen commitments (Sections 3.6.2 and 5.2). Suppose $h = g^{\alpha}$, where $\alpha \in \mathbb{F}_q^*$ is the commitment trapdoor. Observe that a share (x, y, ρ) satisfies this verification equation if and only if

$$y + \alpha\rho = D(x),$$

where

$$D(X) = P(X) + \alpha R(X), \qquad d_j = a_j + \alpha b_j$$

define a masked polynomial and its coefficients. In our preliminary construction, the dealer gives the tracer the trapdoor α, the coefficient vector $\mathbf{d} = (d_0, \ldots, d_{t-1})$ of $D(X)$, and the registered point list $\mathbf{x} = (x_1, \ldots, x_n)$. Because $R(X)$ is uniform and independent, $D(X)$ completely hides $P(X)$ and the secret s from the tracer. At the same time, for any chosen point x and any chosen value y, the tracer can compute

$$\rho = \alpha^{-1}\big(D(x) - y\big)$$

so that the artificial share (x, y, ρ) passes public verification. In each paired query, the tracer simply perturbs

$$(x', y', \rho') \quad \text{into} \quad \big(x', y' + \Delta, \rho' - \Delta/\alpha\big),$$

both of which pass verification. An ϵ-good box is guaranteed to recover the original secret with probability at least ϵ on fresh honest inputs, but this does not directly guarantee success on artificial tracing shares. To bridge this gap, we prove bidirectional sharing (Definition 17). Averaged over a uniformly sampled secret and the sharing randomness, the joint experiment using honest inputs and the original secret is statistically close to that using artificial valid inputs and their reconstructed value, while retaining the coalition's view and tracing material. This property connects success on honest inputs to tracing on artificial valid inputs, yielding the stated guarantee for imperfect (ϵ-good) boxes.

1.2.3 Adding non-imputability. In the preliminary scheme above, the tracing key contains the plaintext list $\mathbf{x} = (x_1, \ldots, x_n)$ of all registered points so that the tracer can match recovered roots to participant identities. However, this also enables a malicious tracer to falsely accuse any honest participant P_i simply by outputting x_i from its key, without ever querying a reconstruction box. Valid tracing evidence must therefore require secret information held by the participant that is not contained in the tracing key.

Following Boneh et al. [7, Sec. 3.3], we replace the plaintext points in the tracing and verification keys with their one-way images $u_i = \mathsf{OW}(x_i)$ under a one-way function OW, giving the tracer the image list $\mathbf{u} = (u_1, \ldots, u_n)$ instead of $\mathbf{x}$ (Section 5.3). After recovering the coalition's evaluation points as roots from the box, an honest tracer applies OW to each root to find the matching entries in $\mathbf{u}$, and outputs the plaintext roots themselves as the tracing evidence. Verification checks that the one-way images of the submitted roots match the accused participants' entries in $\mathbf{u}$. By the one-wayness of OW, a malicious tracer cannot invert u_i to fabricate a valid preimage x_i for an honest participant.

1.2.4 Restoring reconstruction consistency. A major vulnerability remains in this intermediate construction: because the dealer knows the trapdoor α, it can equivocate its own Pedersen commitments. The very capability that lets the tracer craft artificial shares also allows a malicious dealer to distribute shares that pass public verification yet interpolate to different secrets for different groups of participants. To restore reconstruction consistency, we move trapdoor generation from the dealer to the tracer (Section 5.4). During initialization, the tracer samples $\alpha \xleftarrow{\$} \mathbb{F}_q^*$ and sends $h = g^\alpha$ to the dealer. Because α is hidden from the dealer, the discrete-logarithm assumption binds the dealer to a single polynomial $P(X)$ under the published commitments $C_j = g^{a_j} h^{b_j}$. This creates a new challenge: the tracer still needs the coefficients $d_j = a_j + \alpha b_j$ of $D(X) = P(X) + \alpha R(X)$, which must be computed without revealing α to the dealer or exposing the dealer's private coefficients (a_j, b_j) to the tracer. We solve this using committed-receiver oblivious linear evaluation ($\mathcal{F}_{\text{CROLE}}$) [1, Sec. 2.3], defined in Definition 7, which binds the tracer's input α to the public element $h = g^\alpha$ across all calls and delivers only the combined coefficients $d_j = a_j + \alpha b_j$ to the tracer. The tracer then checks

$$g^{d_j} = C_j \qquad (0 \le j < t)$$

against the public commitments. This preserves both tracing and tracer secrecy while guaranteeing reconstruction consistency against a malicious dealer (with an honest tracer and before tracing queries or outputs are exposed).

1.2.5 Making tracing evidence publicly verifiable. Even after the steps above, verifying an accusation still requires the identity table $\mathbf{u} = (u_1, \ldots, u_n)$ that links participant indices to their one-way images $u_i = \mathsf{OW}(x_i)$. Publishing this table directly during sharing would disclose all registered one-way images and their association with participant identities, including entries unrelated to an accusation. Instead, during sharing, the dealer publishes a position-binding vector commitment c_{tvk} to the ordered table $\mathbf{u}$ (Definition 6), and sends $\mathbf{u}$ along with the opening material privately to the tracer (Section 5.5). After tracing, the tracer outputs the recovered roots together with vector-commitment openings for only the accused positions. Anyone can then verify against the public commitment c_{tvk} that the one-way images of the disclosed roots occupy the accused positions in the committed table, without needing the full table $\mathbf{u}$ or access to the reconstruction box. Position binding prevents a malicious tracer from opening a position to a different identity's image. With the optional statistical selective-opening hiding

property, the commitment and disclosed openings also protect unopened table entries against additional leakage.

These steps yield TRAVESS, which combines public share verification and reconstruction consistency with black-box tracing while preserving secrecy, tracer secrecy, and non-imputability. The tracer can construct queries that pass the same verification as genuine shares, and the recovered points can be linked to participant identities through openings of the public commitment. This allows anyone to verify the resulting tracing evidence using public information, without the private tracing key or access to the reconstruction box.

2 Related Works

2.1 Verifiable secret sharing.

Verifiable secret sharing (VSS), introduced by Chor et al. [10], addresses the consistency of shares distributed by a potentially malicious dealer. Feldman [13] developed a practical non-interactive scheme in which the dealer publishes the coefficients of a sharing polynomial in the exponent, allowing recipients to check their shares against common public information. Pedersen [23] combined polynomial sharing with perfectly hiding commitments, preserving information-theoretic privacy of the secret while providing computational guarantees against inconsistent sharing. Each share includes an evaluation of an independent blinding polynomial, which enables verification against the published coefficient commitments.

Beyond the verification algorithm itself, VSS requires a security condition connecting local acceptance to consistent reconstruction. Bellare, Dai, and Rogaway [6, Appendix A.1] formalized this connection through an experiment requiring two authorized collections of verified shares under common public information to reconstruct to the same non-failure value. Subsequent work has explored alternative mechanisms for efficient share verification. Atapoor et al. [2] constructed non-interactive VSS using zero-knowledge proofs on secret-shared data, including instantiations suitable for post-quantum applications. Baghery [3] presented a unified framework for computational VSS in the honest-majority setting, using a random oracle and commitments without requiring homomorphic structure.

Publicly verifiable secret sharing (PVSS), studied by Stadler [25] and Schoenmakers [24], extends verification to external observers, who can check the correctness of share distribution without learning the private shares. This distribution-level guarantee differs from verifying an individual share using public verification information.

2.2 Traceable secret sharing.

Goyal, Song, and Srinivasan [16] gave the first TSS construction based on secure two-party computation, with shares quadratic in the secret length. Boneh, Partap, and Rotem [7] provided efficient constructions based on Shamir and Blakley secret sharing, with shares twice the size of the secret. Their Shamir-based construction uses private random evaluation points and recovers the points associated with the leaking coalition from the box's responses. This construction provides the starting point for our tracing procedure.

Subsequent work explored other algebraic constructions and more expressive access structures. Hoffmann [17] developed a traceable variant of Mignotte secret sharing based on the Chinese remainder theorem, while Leslie and Dutta [19] constructed traceable schemes from additive variants of Asmuth–Bloom secret sharing. Farràs and Guiot [12] extended the study to general access structures, giving an information-theoretic construction using anonymous secret sharing in a model that hides the access structure from participants.

2.3 Stronger tracing notions.

Goyal, Jain, and Partap [15] revisited the tracing definitions and identified useful reconstruction boxes not covered by earlier guarantees, including boxes that require more than the minimum number of additional shares or reveal only partial information about the secret. They introduced strong traceability and $\emptyset$-strong traceability. The latter restricts the box's designated input identities to be disjoint from the corrupted set, whereas the former permits overlap. They gave constructions for access structures represented by monotone circuits and established limitations on combining strong traceability with security against malicious tracers. Their $\emptyset$-strong construction relies on indistinguishability obfuscation, posing a significant obstacle to practical instantiation.

Farràs and Guiot [11, Theorem 3.1] subsequently eliminated this assumption with a fully information-theoretic construction achieving $\emptyset$-strong traceability for general access structures. However, its share size is $O(d \cdot \mathrm{poly}(n, \log m))$, where d is the DNF size of the access structure and m is the size of the secret domain. For a t-out-of-n threshold structure, this representation contains $\binom{n}{t}$ minimal authorized sets and is exponential when t is close to $n/2$. They also introduced hidden mildly strong traceability, a relaxation motivated by the limitations of strong traceability, and gave a separate information-theoretic construction for this notion.

Our work retains the more restricted $t - f$-input threshold framework of Boneh et al. [7] and studies its integration with public share verification, tracer secrecy, and publicly verifiable tracing evidence.

2.4 Traceable verifiable secret sharing.

Two directly relevant approaches are TSS-PV and accountable VSS, which provide different forms of accountability.

Luong et al. [20] proposed TSS-PV, which combines public share verification with black-box tracing. Its tracing procedure requires participants to disclose their shares and therefore does not preserve tracer secrecy. Public verification of tracing evidence also requires oracle access to the same reconstruction box. The main guarantees assume perfect reconstruction boxes; an extension to pointwise ϵ-good boxes, which succeed with probability at least ϵ on every valid input, establishes non-imputability only under honest share distribution, rather than against a malicious dealer as in the perfect-box setting.

Baghery et al. [4] proposed accountable VSS, which identifies contributors from reconstruction transcripts containing NIZK proofs. Its accountability guarantee applies to contributions accepted by the prescribed verification procedure and does not cover

reconstruction boxes built outside that protocol. It therefore provides transcript-based accountability rather than black-box traceability for arbitrary reconstruction programs.

Table 1 summarizes the guarantees and model restrictions. Under the initialization and role assumptions stated below the table, TRAVESS combines all the desired VSS and TSS security guarantees with publicly verifiable tracing evidence. It also relaxes the success requirement imposed by the pointwise extension of TSS-PV. A pointwise ϵ-good box must reconstruct correctly with probability at least ϵ on every valid input. Following Boneh et al. [7, Definition 4], our average-case condition requires only that the success probability, averaged over fresh additional honest shares and the box's randomness, be at least ϵ for a fixed sharing instance. Thus, a box may fail systematically on some inputs while succeeding sufficiently often overall. Our average-case formulation therefore imposes a weaker success requirement on reconstruction boxes than TSS-PV's pointwise formulation, extending the tracing guarantee to a broader class of boxes.

Table 1. Share verification and accountability: guarantees and model restrictions.

Property	**TSS-PV** [20]	**Acc-VSS** [4]	**Ours**
Black-box traceability	Yes	No	**Yes**
Tracer secrecy	No	Not covered	**Yes**
Oracle-free evidence verification[A]	No	Yes	**Yes**
Imperfect-box support	Pointwise only[B]	N/A	**Average-case**[C]
Public share verification	Yes	Yes	Yes
Non-imputability[D]	Yes	Yes	Yes

"Not covered" means the guarantee is not established; N/A denotes a transcript-based model. TRAVESS uses interactive dealer–tracer initialization in the ideal $\mathcal{F}_{\text{CROLE}}$-hybrid model. Its experiments do not jointly corrupt the dealer and tracer (Table 2); the remaining assumptions appear in Section 5.7.

[A] Public information and the supplied evidence suffice for verification, without the tracing key or reconstruction-box access.

[B] For every valid input, the box reconstructs correctly with probability at least ϵ over its own randomness.

[C] For a fixed sharing instance, the probability of recovering the secret is at least ϵ, averaged over fresh additional honest shares and the box's randomness.

[D] For TRAVESS, non-imputability assumes an honest dealer and permits a malicious tracer holding all non-target shares, before external target disclosures (Definition 14). Reconstruction consistency instead assumes an honest tracer and holds before tracing queries or outputs are exposed.

3 Preliminary

3.1 Notation

The security parameter is λ, and PPT denotes probabilistic polynomial time. We write $\xleftarrow{\$}$ for uniform sampling and $\leftarrow$ for assignment or a randomized algorithm's output. Each occurrence of $\mathsf{negl}(\lambda)$ denotes a negligible function that may differ between occurrences and may depend on the adversary and the polynomially bounded parameters; $\mathsf{nonnegl}(\lambda)$ denotes a non-negligible function. For discrete distributions X, Y,

$$\Delta(X,Y) := \frac{1}{2}\sum_{z} |\Pr[X=z] - \Pr[Y=z]|$$

is their statistical distance. For distribution ensembles, $\overset{\mathrm{s}}{\approx}$ denotes negligible statistical distance, and $\overset{\mathrm{c}}{\approx}$ denotes negligible distinguishing advantage for every PPT distinguisher. We write $X \overset{\mathrm{p}}{=} Y$ for equality in distribution (perfect indistinguishability): $\Pr[X=z] = \Pr[Y=z]$ for every z. Probabilities include all random choices unless stated otherwise.

Write $[n] = \{1,\dots,n\}$. In this section and Section 4, use T for a corrupted coalition, J for a reconstruction subset, and I for an accusation. Participants are indexed by $i \in [n]$, and subvectors are ordered by increasing index. We write $\|$ for tuple concatenation. Reconstruction ignores tuple order but retains multiplicities; braces used for such inputs denote the same collection, without removing repetitions. For prime q, we identify $\mathbb{F}_q$ with $\mathbb{Z}_q$ and put $\mathbb{F}_q^* = \mathbb{F}_q \setminus \{0\}$. Algorithm names are local to their indicated primitive.

3.2 Threshold Secret Sharing

We specialize the secret-sharing definition of [5, Secs. 2 and 3.1] to threshold access structures.

Definition 1 (Threshold Secret Sharing). Let $1 \le t \le n$ and let $\mathcal{S}$ be a finite secret space with $|\mathcal{S}| \ge 2$. A t-out-of-n scheme $\mathsf{SS} = (\mathsf{Share}, \mu)$ consists of a distribution μ on a finite randomness space $\mathcal{R}$ and a map

$$\mathsf{Share} : \mathcal{S} \times \mathcal{R} \longrightarrow \prod_{i=1}^{n} \mathcal{SH}_i.$$

The dealer samples $r \sim \mu$, computes $\mathsf{Share}(s;r) = (\mathsf{sh}_1,\dots,\mathsf{sh}_n)$, and privately gives sh_i to P_i. For $J \subseteq [n]$, write $\mathsf{Share}(s;r)_J$ for the corresponding subvector. The requirements are:

- **Correctness.** For each $J \subseteq [n]$ with $|J| \ge t$, there is a function $\mathsf{Rec}_J : \prod_{i\in J} \mathcal{SH}_i \to \mathcal{S}$ such that, for every $s \in \mathcal{S}$,

$$\Pr_{r\sim\mu}[\mathsf{Rec}_J(\mathsf{Share}(s;r)_J) = s] = 1.$$

- **Perfect privacy.** For every $T \subseteq [n]$ with $|T| < t$ and every $s_0, s_1 \in \mathcal{S}$,

$$\mathsf{Share}(s_0;r)_T \overset{\mathrm{p}}{=} \mathsf{Share}(s_1;r)_T, \qquad r \sim \mu.$$

Hereafter, for simplicity, we write $\mathsf{Share}(s)$ for $\mathsf{Share}(s;r)$ with $r \sim \mu$, and similarly omit explicit random-coin arguments for other randomized algorithms unless needed.

3.3 Non-Interactive Verifiable Secret Sharing

In non-interactive VSS, the dealer broadcasts verification information and privately distributes shares; recipients verify locally. Following the verification requirements of Pedersen [23, Definition 4.1 and Lemma 4.2] and the consistency formulation of [6, Appendix A.1, Figure 10], we separate dealer initialization from single-share generation to match the extension in Section 4. We also separate reconstruction on admissible inputs from the security condition that its result is consistent across such inputs.

Definition 2 (Non-Interactive Verifiable Threshold Secret Sharing). For polynomially bounded $2 \le t \le n$, the scheme specifies a PPT procedure that generates public parameters pp on input $(1^\lambda, n, t)$. In security experiments, the challenger generates pp before the adversary chooses its inputs; probabilities include this randomness unless pp is fixed. These parameters specify a finite secret space $\mathcal{S}$, with $|\mathcal{S}| \ge 2$, and a share space $\mathcal{SH}$. A scheme

$$\mathsf{VSS} = (\mathsf{Gen}, \mathsf{Share}, \mathsf{VerifyShare}, \mathsf{Rec})$$

has four polynomial-time interfaces:

- $\mathsf{Gen}(\mathsf{pp}, s) \to (\mathsf{svk}, \sigma)$ is randomized dealer-local initialization for $s \in \mathcal{S}$. It produces public verification information svk and private state σ, or returns $\perp$.
- $\mathsf{Share}(\mathsf{pp}, \sigma) \to \mathsf{sh}$ generates one share using fresh randomness and leaves σ unchanged. All shares of an instance use the state fixed by its initialization.[3]
- $\mathsf{VerifyShare}(\mathsf{pp}, \mathsf{svk}, \mathsf{sh}) \to \{0, 1\}$ is a public deterministic check. It needs no interaction and rejects malformed keys and shares.
- $\mathsf{Rec}(\mathsf{pp}, \mathsf{svk}, \mathbf{q}) \to \mathcal{S} \cup \{\perp\}$ is deterministic and invariant under permutations of its input tuple.

Let $\mathsf{id} : \mathcal{SH} \to \mathcal{I}$ be an efficiently computable identifier map. A tuple $\mathbf{q}$ is *admissible* if it contains exactly t shares that pass $\mathsf{VerifyShare}$ and have pairwise distinct identifiers. As an explicit interface requirement, Rec returns an element of $\mathcal{S}$ on every admissible tuple and returns $\perp$ otherwise.

The dealer runs Gen, broadcasts svk on success, and privately distributes shares generated by Share according to the scheme's polynomial-time distribution rule. A completed honest distribution gives each P_i a share sh_i to verify locally, with distinct identifiers across $\mathbf{sh} = (\mathsf{sh}_1, \ldots, \mathsf{sh}_n)$. On abort, previously broadcast information and delivered shares remain available; missing outputs are denoted by $\perp$.

The scheme satisfies the following properties, for every polynomially bounded choice of n, t:

1. **Honest correctness.** For every valid input secret, except with negligible probability over parameter generation and distribution, the execution completes, every delivered share is accepted, and

$$\mathsf{Rec}(\mathsf{pp}, \mathsf{svk}, (\mathsf{sh}_i)_{i \in J}) = s \quad \text{for every } J \subseteq [n] \text{ with } |J| = t.$$

   Aborts count as failures.

[3] For the fixed-coordinate Pedersen baseline in Section 5.1, one may write $\mathsf{Share}(\mathsf{pp}, \sigma, i)$ to select the assigned point x_i. The randomized-coordinate TVSS construction uses the index-free interface displayed here.

2. **Privacy.** For every PPT adversary $\mathcal{A}$, generate pp as above and give it to $\mathcal{A}$, which chooses valid secrets s_0, s_1 and a coalition $T \subseteq [n]$ with $|T| < t$ before initialization. For $b \in \{0,1\}$, let V_b denote its joint view of the honest distribution of s_b, including pp, its choices and random tape, the public transcript (including svk), and the shares delivered to T, with observable abort information and prefixes retained as above; privacy requires
$$V_0 \overset{\mathrm{c}}{\approx} V_1.$$
Statistical and perfect privacy replace $\overset{\mathrm{c}}{\approx}$ by $\overset{\mathrm{s}}{\approx}$ and $\overset{\mathrm{p}}{=}$, respectively.
3. **Reconstruction consistency.** For every PPT adversary $\mathcal{A}$, generate pp as above and run $(\mathsf{svk}, \mathbf{q}_0, \mathbf{q}_1) \leftarrow \mathcal{A}(\mathsf{pp})$. We require
$$\Pr\left[\begin{array}{l}\mathbf{q}_0, \mathbf{q}_1 \text{ are admissible under } (\mathsf{pp}, \mathsf{svk}) \\ \wedge\ \mathsf{Rec}(\mathsf{pp}, \mathsf{svk}, \mathbf{q}_0) \neq \mathsf{Rec}(\mathsf{pp}, \mathsf{svk}, \mathbf{q}_1)\end{array}\right] \leq \mathsf{negl}(\lambda).$$

The interface excludes failure on admissible inputs independently of the consistency requirement. For Shamir-based reconstruction, distinct evaluation points make interpolation well defined. This requirement does not force a malicious dealer to deliver enough admissible shares: share identifiers need not identify holders, and different recipients may receive the same identifier. Complaint resolution and a common disqualification decision [14, Section 4] are outside this definition.

3.4 One-Way Functions

We use the standard inversion requirement [18, Sec. 8.1.1] on the parameter-dependent field domain needed by our construction.

Definition 3 (One-Way Function). For polynomially bounded n, t, suppose the prescribed public-parameter generation procedure on $(1^\lambda, n, t)$ produces pp specifying a prime q and a deterministic polynomial-time map $\mathsf{OW} : \mathbb{F}_q^* \to \mathcal{U}$. For any PPT adversary $\mathcal{A}$, sample
$$x \xleftarrow{\$} \mathbb{F}_q^*, \qquad x' \leftarrow \mathcal{A}(1^\lambda, \mathsf{pp}, \mathsf{OW}(x)).$$
The function family is one-way under this parameter-generation procedure if, for every such $\mathcal{A}$,
$$\Pr[x' \in \mathbb{F}_q^* \ \wedge\ \mathsf{OW}(x') = \mathsf{OW}(x)] \leq \mathsf{negl}(\lambda).$$
The probability includes parameter generation; inversion may return any preimage.

3.5 Discrete-Logarithm Assumption

Following [18, Sec. 9.3.2, Definition 9.63], let $\mathsf{GroupGen}(1^\lambda)$ output $(\mathbb{G}, q, g)$, where $\mathbb{G} = \langle g \rangle$ has prime order q and q has λ bits. The generator is PPT; group elements have unique bit-string encodings, and membership testing and group operations take polynomial time. For the concrete map $\mathsf{OW}(x) = g^x$, we use canonical fixed-length encodings in $\mathcal{U} = \{0,1\}^{\ell(\lambda)}$, where ℓ is a fixed polynomial large enough for all generated group elements; the vector commitment uses this same message space independently of group generation.

Definition 4 (Discrete-Logarithm Assumption). For every PPT adversary $\mathcal{A}$, sample

$$(\mathbb{G}, q, g) \leftarrow \mathsf{GroupGen}(1^\lambda), \qquad h \xleftarrow{\$} \mathbb{G}, \qquad x \leftarrow \mathcal{A}(\mathbb{G}, q, g, h).$$

The discrete-logarithm problem is hard relative to GroupGen if, for every such adversary,

$$\Pr[x \in \mathbb{F}_q \ \wedge \ g^x = h] \leq \mathsf{negl}(\lambda).$$

We use this assumption for the group-generation algorithm in our construction.

3.6 Commitment Schemes

We use Pedersen commitments for share verification and a separate vector commitment for public tracing.

3.6.1 Non-Interactive Commitments.

We follow [18, Sec. 6.6.5, Definition 6.13], allowing parameter-dependent spaces.

Definition 5 (Non-Interactive Commitment). A scheme has two algorithms, $\mathsf{COM} = (\mathsf{Gen}, \mathsf{Com})$. The PPT setup $\mathsf{Gen}(1^\lambda)$ samples $\mathsf{pp}_{\mathsf{com}}$, specifying message, randomness, and commitment spaces $\mathcal{M}_{\mathsf{pp}_{\mathsf{com}}}$, $\mathcal{R}_{\mathsf{pp}_{\mathsf{com}}}$, and $\mathcal{C}_{\mathsf{pp}_{\mathsf{com}}}$. The deterministic polynomial-time map $\mathsf{Com}(\mathsf{pp}_{\mathsf{com}}, m; r)$ outputs a commitment. An honest commitment uses uniform r; an opening (m, r) is accepted exactly when it lies in the specified spaces and recomputes the commitment. We require:

- **Hiding.** Sample $\mathsf{pp}_{\mathsf{com}} \leftarrow \mathsf{Gen}(1^\lambda)$, let $(m_0, m_1, \mathsf{state}) \leftarrow \mathcal{A}_1(\mathsf{pp}_{\mathsf{com}})$ contain valid messages and private working memory state passed from $\mathcal{A}_1$ to $\mathcal{A}_2$, and sample $b \xleftarrow{\$} \{0,1\}$ and $r \xleftarrow{\$} \mathcal{R}_{\mathsf{pp}_{\mathsf{com}}}$. Set $C = \mathsf{Com}(\mathsf{pp}_{\mathsf{com}}, m_b; r)$. For every PPT adversary $(\mathcal{A}_1, \mathcal{A}_2)$, we require

$$\Pr[\mathcal{A}_2(\mathsf{state}, C) = b] \leq \frac{1}{2} + \mathsf{negl}(\lambda).$$

- **Binding.** Sample $\mathsf{pp}_{\mathsf{com}} \leftarrow \mathsf{Gen}(1^\lambda)$ and $(m_0, r_0, m_1, r_1) \leftarrow \mathcal{A}(\mathsf{pp}_{\mathsf{com}})$, and set $C_b = \mathsf{Com}(\mathsf{pp}_{\mathsf{com}}, m_b; r_b)$ for $b \in \{0,1\}$. For every PPT adversary $\mathcal{A}$, we require

$$\Pr[m_0 \neq m_1 \ \wedge \ C_0 = C_1] \leq \mathsf{negl}(\lambda).$$

  Invalid outputs are rejected.

3.6.2 Pedersen Commitments.

For Pedersen commitments [23, Sec. 3, Theorem 3.1], sample

$$(\mathbb{G}, q, g) \leftarrow \mathsf{GroupGen}(1^\lambda), \qquad h \xleftarrow{\$} \mathbb{G} \setminus \{1\},$$

and publish $\mathsf{pp}_{\mathsf{com}} = (\mathbb{G}, q, g, h)$. Setup outputs no private key and does not provide the committer with $\log_g h$. Messages and randomness lie in $\mathbb{F}_q$, and

$$\mathsf{Com}(\mathsf{pp}_{\mathsf{com}}, m; r) := g^m h^r.$$

Uniform r makes this commitment perfectly hiding. Two openings to distinct messages reveal $\log_g h$, giving computational binding under Definition 4.

Pedersen commitments also support equivocation when the trapdoor $\alpha = \log_g h$ is known [21]. To generate parameters with this trapdoor, sample $\alpha \xleftarrow{\$} \mathbb{F}_q^*$ and set $h = g^\alpha$; the public parameters have the same distribution as above. Given α and an opening (m, r), one can open the same commitment to any $m' \in \mathbb{F}_q$ by computing

$$r' = r + (m - m')\alpha^{-1}, \qquad g^{m'} h^{r'} = g^m h^r,$$

as in [9, Sec. 2.4, p. 233]. Binding applies to committers without the trapdoor.

3.6.3 Vector Commitments. We use the vector-commitment syntax of [22, Sec. 3.1, Definitions 1–3] and the position-binding requirement of [8, Sec. 2, Definition 2]. We also record statistical selective-opening hiding as an optional privacy property. Succinctness, updates, and extraction are not required.

Definition 6 (Vector Commitment). For polynomially bounded n, a vector commitment has four algorithms:

- $\mathsf{VC.Setup}(1^\lambda, n) \to \mathsf{pp}_{\mathrm{vc}}$ specifies a message space $\mathcal{U}$ and length n.
- $\mathsf{VC.Commit}(\mathsf{pp}_{\mathrm{vc}}, \mathbf{u}) \to (c, \mathsf{aux}_{\mathrm{vc}})$ commits to $\mathbf{u} \in \mathcal{U}^n$ and retains private opening state $\mathsf{aux}_{\mathrm{vc}}$.
- $\mathsf{VC.Open}(\mathsf{pp}_{\mathrm{vc}}, \mathsf{aux}_{\mathrm{vc}}, i) \to \omega_i$ opens position $i \in [n]$.
- $\mathsf{VC.Verify}(\mathsf{pp}_{\mathrm{vc}}, c, i, u, \omega) \to \{0, 1\}$ checks an opening to u.

Setup and commitment are PPT; opening and verification are deterministic polynomial-time algorithms, with all opening randomness fixed in $\mathsf{aux}_{\mathrm{vc}}$. Verification rejects malformed inputs and out-of-range indices. The TVSS security results require correctness and position binding. We additionally define an optional hiding property:

- **Correctness.** Fix any honestly generated $\mathsf{pp}_{\mathrm{vc}}$, vector $\mathbf{u} \in \mathcal{U}^n$, and position $i \in [n]$. Sample $(c, \mathsf{aux}_{\mathrm{vc}}) \leftarrow \mathsf{VC.Commit}(\mathsf{pp}_{\mathrm{vc}}, \mathbf{u})$. Compute the opening $\omega_i = \mathsf{VC.Open}(\mathsf{pp}_{\mathrm{vc}}, \mathsf{aux}_{\mathrm{vc}}, i)$. Then
$$\Pr[\mathsf{VC.Verify}(\mathsf{pp}_{\mathrm{vc}}, c, i, u_i, \omega_i) = 1] = 1.$$
- **Position binding.** Sample $\mathsf{pp}_{\mathrm{vc}} \leftarrow \mathsf{VC.Setup}(1^\lambda, n)$. Let $(c, i, u, u', \omega, \omega')$ be the output of $\mathcal{A}(\mathsf{pp}_{\mathrm{vc}})$. For every PPT adversary $\mathcal{A}$, we require
$$\Pr\begin{bmatrix} i \in [n] \ \wedge \ u \neq u' \\ \wedge \ \mathsf{VC.Verify}(\mathsf{pp}_{\mathrm{vc}}, c, i, u, \omega) = 1 \\ \wedge \ \mathsf{VC.Verify}(\mathsf{pp}_{\mathrm{vc}}, c, i, u', \omega') = 1 \end{bmatrix} \leq \mathsf{negl}(\lambda).$$
The adversary may retain the state of commitments it generates.
- **Statistical selective-opening hiding (optional).** For $K \subseteq [n]$, let
$$\mathsf{View}_{\mathsf{pp}_{\mathrm{vc}}}(\mathbf{u}, K) := \big(c, (\omega_i)_{i \in K}\big),$$

where $(c, \mathsf{aux}_{\mathrm{vc}}) \leftarrow \mathsf{VC.Commit}(\mathsf{pp}_{\mathrm{vc}}, \mathbf{u})$ and $\omega_i = \mathsf{VC.Open}(\mathsf{pp}_{\mathrm{vc}}, \mathsf{aux}_{\mathrm{vc}}, i)$. For every honestly generated $\mathsf{pp}_{\mathrm{vc}}$, every $K \subseteq [n]$, and all $\mathbf{u}_0, \mathbf{u}_1 \in \mathcal{U}^n$ agreeing on K, we require

$$\Delta\big(\mathsf{View}_{\mathsf{pp}_{\mathrm{vc}}}(\mathbf{u}_0, K), \mathsf{View}_{\mathsf{pp}_{\mathrm{vc}}}(\mathbf{u}_1, K)\big) \le \mathsf{negl}(\lambda),$$

with a negligible bound uniform over these choices.

For fixed $\mathbf{u}_0, \mathbf{u}_1$, the last condition also covers adaptive queries restricted to positions where they agree: the interaction is a function of the view opening all such positions. Repeated queries return the same opening.

We commit to $\mathbf{u} = (\mathsf{OW}(x_1), \ldots, \mathsf{OW}(x_n))$, where x_i is P_i's private evaluation point. Setup is independent of α; position binding must hold even against a tracer retaining α and $\mathsf{aux}_{\mathrm{vc}}$. The optional hiding property concerns leakage from the commitment and selectively disclosed openings; it is not used in the TVSS security proofs.

3.7 Committed-Receiver OLE

We use the committed-receiver oblivious linear evaluation functionality of [1, Sec. 2.3], in the notation of the preceding commitment scheme.

Definition 7 (Committed-Receiver OLE). The functionality $\mathcal{F}_{\mathrm{CROLE}}$ takes common input $(\mathbb{G}, q, g, h)$, where $\mathbb{G} = \langle g \rangle$ has prime order q and $h \in \mathbb{G}$. The sender S supplies $(a, b) \in \mathbb{F}_q^2$, and the receiver R supplies $\alpha \in \mathbb{F}_q$. Only the receiver obtains

$$d = \begin{cases} b\alpha + a, & g^{\alpha} = h, \\ \bot, & \text{otherwise.} \end{cases}$$

The sender receives no output.

A realization must be maliciously secure with abort: for every PPT adversary corrupting either party, there is a PPT simulator such that real and ideal executions are computationally indistinguishable, with the corrupted party allowed to abort. Our construction invokes this functionality t times sequentially under the same authenticated common input $(\mathbb{G}, q, g, h)$. A concrete realization must support sequential composition within the TVSS experiments, including correlated inputs and auxiliary information from earlier calls; standalone single-invocation security alone does not suffice.

4 Traceable Verifiable Secret Sharing with Publicly Verifiable Tracing Evidence (TVSS)

4.1 System and Adversarial Model

A sharing instance involves a dealer holding $s \in \mathcal{S}$, a tracer, and participants $P_1, \ldots, P_n$. We use the public-parameter generation convention and staged VSS interfaces of Definition 2, with polynomially bounded $2 \le t \le n$. Public parameters are generated by the challenger before the secret is chosen. In our construction, $\mathsf{pp} = (\mathbb{G}, q, g, n, t, \mathsf{OW}, \mathsf{pp}_{\mathrm{vc}})$;

the tracer generates h during instance initialization, so h is not part of this initial setup (Section 5.6). The participant ordering is fixed before initialization; public outputs enter a common authenticated record, and private messages use private authenticated channels. The final record is fixed before tracing.

Participant corruptions are static: the adversary chooses a set $T \subseteq [n]$ after seeing pp and before initialization. We reserve J for a reconstruction subset and I for an accusation.

Table 2 gives the dealer and tracer roles in each security experiment: honest parties follow the protocol, whereas malicious parties are controlled by the adversary.

Table 2. Dealer and tracer roles in the security experiments.

Property	Dealer	Tracer
Reconstruction consistency	Malicious	Honest
Tracer secrecy; non-imputability	Honest	Malicious
Secrecy; traceability	Honest	Honest

The three rows address attacks by the dealer, the tracer, and corrupted participants, respectively. No experiment jointly corrupts the dealer and tracer. Secrecy and tracer secrecy expose fewer than t shares; non-imputability exposes all non-target shares. Tracing uses the stateless-box model of Section 4.7.

Unless stated otherwise, each security requirement quantifies over every PPT adversary and polynomially bounded parameter choice. We use ϵ for box goodness and δ for the tracing parameter. All experiment probabilities are unconditional. Invalid adversarial choices give output 0 unless the experiment specifies otherwise.

The adversary's view includes its own randomness and all messages and shares received during the execution, including those received before an abort and any observable abort information. Public outputs and exposed shares are delivered as they become available; grouped arrows in the figures abbreviate these deliveries. Missing outputs are denoted by $\perp$.

4.2 Syntax and Correctness

We extend the single-share VSS interfaces of Definition 2 with the tracing interfaces of [7, Sec. 2.1]. The resulting primitive is traceable verifiable secret sharing (TVSS) with publicly verifiable tracing evidence.

Definition 8 (TVSS with Publicly Verifiable Tracing Evidence). A t-out-of-n TVSS scheme consists of

$$\Pi = (\mathsf{Gen}, \mathsf{Share}, \mathsf{VerifyShare}, \mathsf{Rec}, \mathsf{Trace}, \mathsf{VerifyTrace}).$$

It retains the $\mathsf{VerifyShare}$ and Rec interfaces, the identifier map, and the admissibility and reconstruction conventions of Definition 2. Its remaining interfaces are:

- $\mathsf{Gen}(\mathsf{pp}, s) \rightarrow (\mathsf{svk}, \sigma, \mathsf{tk}_0)$ extends VSS initialization to an interactive protocol between the dealer and tracer. The common input is pp; only the dealer supplies

s. The outputs are public svk, dealer-private σ, and tracer-private tk_0. Either party may reject locally.
- $\mathsf{Share}(\mathsf{pp}, \sigma) \to (\mathsf{sh}, \mathsf{tk}', \mathsf{tvk}')$ extends single-share generation with private tracing-key and verification-key contributions. It uses fresh randomness, leaves σ unchanged, and requires no tracer interaction.[4]
- $\mathsf{Trace}^{\mathcal{R}}(\mathsf{pp}, \mathsf{svk}, \mathsf{tvk}, \mathsf{tk}, 1^{\lceil 1/\delta \rceil}) \to (I, \pi)$ or $\perp$ is randomized, with tracing parameter $\delta \in (0, 1)$. Given black-box access to a stateless reconstruction box $\mathcal{R}$ as defined in Section 4.7, it accuses a set $I \subseteq [n]$ with $1 \le |I| < t$ and supplies evidence π. Its running time, counting each oracle call as one step, is polynomial in $\lambda, n, 1/\delta$.
- $\mathsf{VerifyTrace}(\mathsf{pp}, \mathsf{svk}, \mathsf{tvk}, I, \pi) \to \{0, 1\}$ is deterministic and uses only the public record and supplied evidence. It rejects malformed evidence or an accusation outside the preceding range. Sets of indices are encoded in increasing order.

Initialization and share generation are PPT; the deterministic algorithms run in polynomial time in their input lengths. In security experiments, $1/\delta$ is polynomially bounded. The scheme also specifies the admission and key-assembly rules below, and the auxiliary valid-share distribution of Section 4.7.

After completing its local initialization, the dealer uses the distribution rule of Definition 2 to obtain the admitted outputs $(\mathsf{sh}_i, \mathsf{tk}_i, \mathsf{tvk}_i)$ for $i \in [n]$. The participant receives sh_i; the dealer retains the key contributions. A polynomial-time local assembly rule, including its randomness and checks, produces public tvk and private completion material for the tracer. Using the latter and tk_0, the tracer forms tk and performs its local checks. These contributions need not themselves be public verification keys.

The final record is $\mathsf{pub} = (\mathsf{pp}, \mathsf{svk}, \mathsf{tvk})$ with the fixed participant ordering. Let acc denote the event that the dealer and tracer finish their prescribed stages without rejection, this record is fixed, and every participant receives a share accepted by $\mathsf{VerifyShare}$. This is an experiment event, not a protocol message. The dealer continues according to its own local state; neither initialization nor key assembly discloses the tracer's private checks or acceptance decision to it. No reply from the tracer is required after delivery of the completion material.

The extension inherits the VSS interfaces, rather than the arbitrary-key experiment for ordinary VSS: its tracer-assisted initialization determines which records a malicious dealer can establish. Section 4.4 specifies the corresponding consistency experiment.

Definition 9 (Correctness). A scheme is correct if it satisfies the following two conditions.

1. **Sharing and reconstruction.** For every λ, every admissible n, t, every pp that the prescribed generation procedure can output, every $s \in \mathcal{S}$, and every fixed $J \subseteq [n]$ with $|J| = t$, honest sharing satisfies

$$\Pr[\mathsf{acc} \wedge \mathsf{Rec}(\mathsf{pp}, \mathsf{svk}, \mathbf{sh}_J) = s] \ge 1 - \mathsf{negl}(\lambda).$$

   The negligible bound is uniform over these inputs. The probability is over sharing randomness with pp fixed; acc includes acceptance of all delivered shares.

[4] As in [7, Secs. 2.1–2.2], the centralized interface specifies a joint output distribution; it does not preclude a distributed realization.

2. **Tracing-evidence validity.** Fix any PPT box $\mathcal{R}$ and polynomially bounded parameters $n, t, 1/\delta$. On every accepted honest instance, if

$$o \leftarrow \mathsf{Trace}^{\mathcal{R}}(\mathsf{pp}, \mathsf{svk}, \mathsf{tvk}, \mathsf{tk}, 1^{\lceil 1/\delta \rceil}),$$

then

$$\Pr[o = (I, \pi) \ \wedge \ \mathsf{VerifyTrace}(\mathsf{pp}, \mathsf{svk}, \mathsf{tvk}, I, \pi) = 0] \leq \mathsf{negl}(\lambda).$$

The bound is uniform over accepted instances. The probability is over tracing and box randomness. The event is false for $o = \perp$; no goodness assumption is imposed.

Correctness is perfect if both errors are zero.

Evidence validity concerns verification of the output. The separate properties of non-imputability and traceability protect against false accusations and require recovery of a good box's coalition, respectively.

4.3 Secrecy

Definition 10 (Secrecy). For $b \in \{0, 1\}$, define V_b as in the privacy clause of Definition 2, using an honest TVSS dealer and tracer to share s_b. The view includes all public messages, the record pub, and the shares delivered to T, following the online and abort conventions of Section 4.1; tracer-private information remains internal. Secrecy requires

$$V_0 \overset{c}{\approx} V_1.$$

For use in the proofs, $\mathsf{Exp}^{\mathsf{sec}}_{\mathcal{A}}(b)$ denotes $\mathcal{A}$'s output bit after completion or abort in this experiment.

4.4 Public Share Verifiability and Consistency

Definition 11 (Public Share Verifiability). Public share verifiability is the public $\mathsf{VerifyShare}$ interface inherited from Definition 2, together with honest-share acceptance in Definition 9.

Definition 12 (Reconstruction Consistency). Generate public parameters pp as in Definition 2 and let $\mathcal{A}(\mathsf{pp})$ act as the dealer against an honest tracer, controlling share generation, delivery, and key assembly. It then outputs two tuples $\mathbf{q}_0, \mathbf{q}_1$. The scheme specifies publicly checkable conditions, satisfied by every honestly completed sharing, for compatibility of the final record with the tracer's authenticated public initialization messages; these conditions do not depend on private initialization or key-assembly checks. The experiment $\mathsf{Exp}^{\mathsf{cons}}_{\mathcal{A},\Pi}(1^{\lambda}, n, t)$ returns 1 exactly when the final record is fixed and compatible with these messages, and the tuples are admissible under that record but reconstruct different secrets. Otherwise it returns 0. The dealer receives its full sharing view, with no disclosure of tracer acceptance and no later tracing queries or outputs. Consistency requires

$$\Pr[\mathsf{Exp}^{\mathsf{cons}}_{\mathcal{A},\Pi}(1^{\lambda}, n, t) = 1] \leq \mathsf{negl}(\lambda).$$

This guarantee does not force a malicious dealer to deliver usable shares to every participant: it may withhold shares or assign repeated identifiers. Such collections are not admissible. For admissible collections, reconstruction is defined by the inherited interface. Local verification does not establish which participant was assigned a share.

4.5 Tracer Secrecy

Following [7, Definitions 2 and 3], we strengthen secrecy to allow an actively malicious tracer. Figure 1 defines the experiment $\mathsf{Exp}^{\mathsf{tsec}}_{\mathcal{A}}(b)$ for fixed λ, n, t and $b \in \{0, 1\}$.

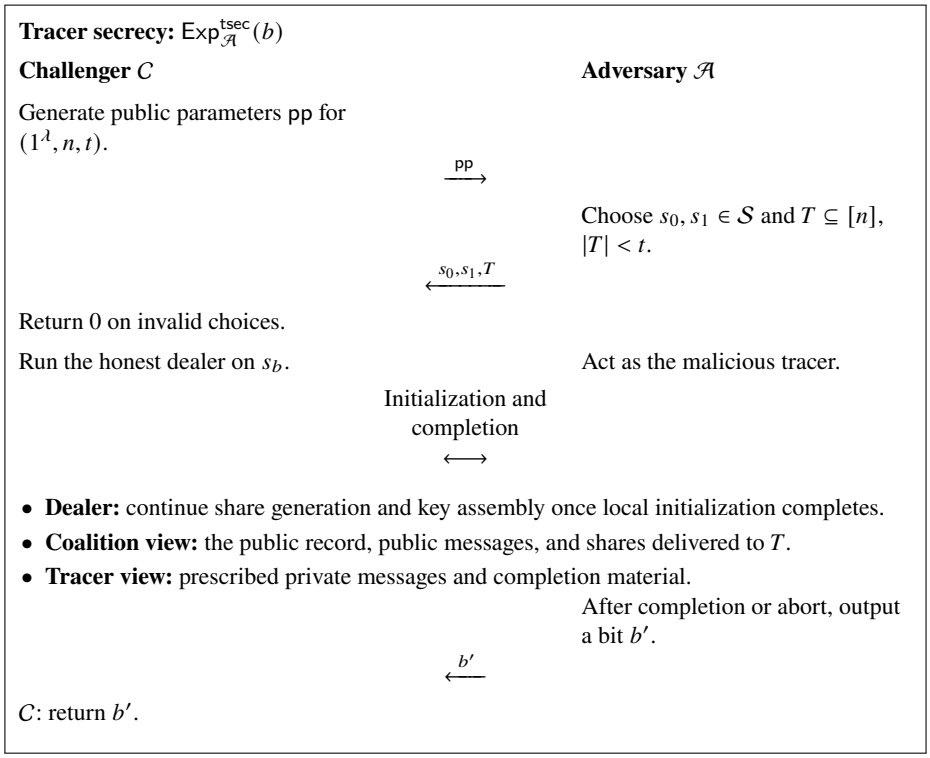


Fig. 1. The tracer-secrecy experiment for Definition 13. Deliveries and aborts follow the online-view convention of Section 4.1.

Definition 13 (Tracer Secrecy). In the experiment of Figure 1, tracer secrecy requires

$$\left|\Pr[\mathsf{Exp}^{\mathsf{tsec}}_{\mathcal{A}}(0) = 1] - \Pr[\mathsf{Exp}^{\mathsf{tsec}}_{\mathcal{A}}(1) = 1]\right| \leq \mathsf{negl}(\lambda).$$

Tracer secrecy implies secrecy (Definition 10) by executing an honest tracer and passing only the public view and coalition shares to the secrecy adversary. Neither experiment conditions on acceptance or supplies reconstruction or tracing oracles.

4.6 Non-Imputability

We adapt [7, Definition 6] to an honest dealer and an actively malicious tracer. In Figure 2, the adversary chooses the secret and target $i^\star$ before initialization and receives all $n - 1$ non-target shares. The target remains honest.

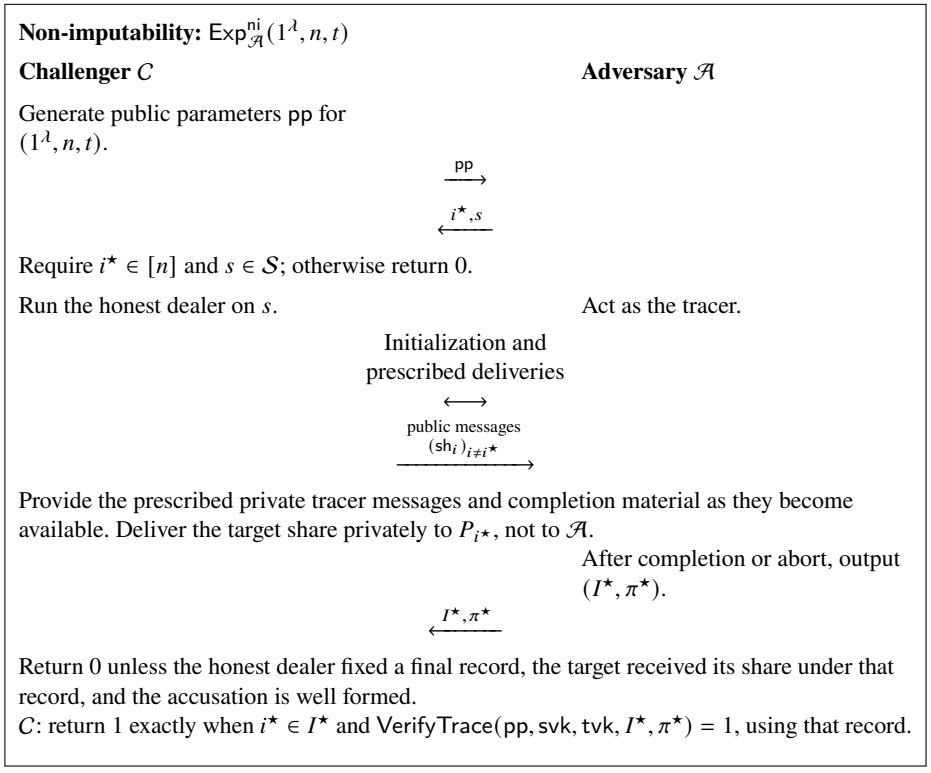


Fig. 2. The malicious tracer's own acceptance decision does not affect the winning condition. The secret and all non-target shares may be known.

Definition 14 (Non-Imputability). A scheme satisfies non-imputability if

$$\Pr[\mathsf{Exp}^{\mathsf{ni}}_{\mathcal{A}}(1^\lambda, n, t) = 1] \leq \mathsf{negl}(\lambda),$$

where the experiment is Figure 2.

The experiment supplies no previous target-share disclosure, target tracing evidence, or oracle using the target's share. Information obtained through the adversary's own deviations remains available to it. This is a single-instance guarantee before external target disclosures: public evidence can be copied and does not certify a new leakage event. Its probability is not conditioned on sharing acceptance.

4.7 Reconstruction Boxes

We adapt [7, Definitions 4 and 7]. Fix an accepted honest sharing of s, including the final keys, admitted shares $\mathbf{sh}$, and dealer state σ. Let $T \subseteq [n]$ be a coalition of size $1 \leq f < t$. A PPT adversary given the public sharing view and $\mathbf{sh}_T$ may produce any

PPT reconstruction box $\mathcal{R}$ with public input arity $t - f$. The box outputs an element of $\mathcal{S} \cup \{\perp\}$ and need not run Rec or verify its inputs.

The box is stateless: its embedded data, including any coalition shares, remain fixed, and each response depends only on these data, the current input, and fresh independent randomness. It retains no memory of earlier calls. The tracer has only black-box access: it supplies a tuple Q and receives $\mathcal{R}(Q)$.

The following two distributions supply additional shares under the original record:

- $\mathcal{H}_{\mathsf{pp},\sigma}$ is the share component of a fresh call to $\mathsf{Share}(\mathsf{pp}, \sigma)$. Discard the key contributions. No admission, resampling, or registration is performed.
- $\mathcal{D}_{\mathsf{pub},\mathsf{tk}}$ is specified by the scheme and supported on shares that pass verification. It is fixed before the box is chosen and efficiently samplable from $(\mathsf{pub}, \mathsf{tk})$, without σ. Its samples need not be honest shares of the original secret.

All additional samples are independent and leave the public record and registration unchanged. The box receives neither σ nor tk. Neither distribution removes point collisions.

We also require correctness when additional honest shares are sampled: for every secret and every $T \subseteq [n]$ with $|T| < t$, the tuple $\mathbf{sh}_T$ together with $t - |T|$ independent samples from $\mathcal{H}_{\mathsf{pp},\sigma}$ must reconstruct s except with negligible probability over honest sharing and sampling. Sharing aborts and reconstruction failures, including collisions, count as failures.

Definition 15 (Good Reconstruction Boxes). For $\epsilon : \mathbb{N} \to [0, 1]$, a box is $\epsilon(\lambda)$-good for the fixed instance and coalition T if

$$\Pr_{Q \leftarrow \mathcal{H}_{\mathsf{pp},\sigma}^{t-f}} [\mathcal{R}(Q) = s] \geq \epsilon(\lambda).$$

Here Q is an ordered tuple of additional shares. The probability also includes box randomness; the original sharing instance and σ remain fixed, as do the box and its embedded data.

Definition 16 (Universally-Good Reconstruction Boxes). For $\epsilon : \mathbb{N} \to [0, 1]$, sample $Q \leftarrow \mathcal{D}_{\mathsf{pub},\mathsf{tk}}^{t-f}$ and set

$$s_Q := \mathsf{Rec}(\mathsf{pp}, \mathsf{svk}, \mathbf{sh}_T \parallel Q).$$

A box is $\epsilon(\lambda)$-universally good for the fixed instance and coalition shares if

$$\Pr_{Q \leftarrow \mathcal{D}_{\mathsf{pub},\mathsf{tk}}^{t-f}} [\mathcal{R}(Q) = s_Q \ \wedge \ s_Q \in \mathcal{S}] \geq \epsilon(\lambda).$$

The probability also includes box randomness, with $\mathsf{pub}, \mathsf{tk}, \mathbf{sh}_T$ fixed, together with the box and its embedded data. A reconstruction failure makes the event false, even if the box also returns $\perp$; sampling is not conditioned on successful reconstruction.

Universal goodness is an average agreement requirement for the specified distribution. The reconstructed value s_Q may differ from the original secret, and verification alone does not determine the sampling distribution.

The following distributional condition adapts [7, Definition 9] to the public record and tracing key. It is the hypothesis used below to pass from universal to ordinary traceability.

Definition 17 (Bidirectional Sharing). Fix a PPT coalition selector $\mathcal{A}_0$. On input $(1^\lambda, n, t)$, each sampler separately performs the following common initialization:

- Generate public parameters pp and run $(T, \mathsf{state}) \leftarrow \mathcal{A}_0(\mathsf{pp})$, where state is the selector's private working memory and random tape.
- Return $\perp$ on malformed output, $T \not\subseteq [n]$, or $|T| \notin \{1, \ldots, t-1\}$. Otherwise set $f = |T|$.
- Sample $z \xleftarrow{\$} \mathcal{S}$ and complete honest sharing of z. Return $\perp$ unless sharing is accepted.

Let V contain $\mathsf{pp}, T, \mathsf{state}$, the final public record, and the ordered public messages and coalition shares. Retain the dealer state σ and tracing key tk internally, and continue as follows:

Forward $\mathsf{Fwd}_{\mathcal{A}_0}$

- Sample $Q \leftarrow \mathcal{H}^{t-f}_{\mathsf{pp},\sigma}$.
- Return (V, tk, z, Q).

Backward $\mathsf{Bwd}_{\mathcal{A}_0}$

- Sample $Q \leftarrow \mathcal{D}^{t-f}_{\mathsf{pub},\mathsf{tk}}$.
- Return (V, tk, s_Q, Q).

Here $s_Q = \mathsf{Rec}(\mathsf{pp}, \mathsf{svk}, \mathbf{sh}_T \parallel Q)$ as in Definition 16; reconstruction failure retains the tuple $(V, \mathsf{tk}, \perp, Q)$. The scheme is bidirectional if, for every such selector,

$$\Delta\Big(\mathsf{Fwd}_{\mathcal{A}_0}(1^\lambda, n, t), \mathsf{Bwd}_{\mathcal{A}_0}(1^\lambda, n, t)\Big) \le \mathsf{negl}(\lambda).$$

The comparison includes parameter generation and honest-sharing randomness, without conditioning on success.

The tracing key is included to preserve the tracing inputs, without disclosing it to the adversary. The comparison excludes the dealer state and the backward sampler's original secret; it need not hold for each fixed dealer state.

4.8 Publicly Verifiable Traceability

We adapt [7, Definitions 5 and 8]. Tracing must recover the entire coalition and produce publicly accepted evidence under the original record. Let $\epsilon, \delta : \mathbb{N} \to (0, 1)$ be efficiently computable with polynomially bounded reciprocals.

For fixed pp, an independent uniform guess of the secret succeeds with probability $1/|\mathcal{S}|$. Since the secret space is specified by pp, define the average guessing baseline

$$p_{\mathsf{guess}}(\lambda, n, t) := \mathbb{E}_{\mathsf{pp}}\left[\frac{1}{|\mathcal{S}(\mathsf{pp})|}\right],$$

where the expectation is over the prescribed public-parameter generation and $\mathcal{S}(\mathsf{pp})$ makes the parameter dependence explicit. If this size is fixed across parameter choices, the expression reduces to $1/|\mathcal{S}|$, the guessing baseline in [7, Definition 5].

Figure 3 defines $\mathsf{Exp}^{\mathsf{trace}}_{\mathcal{A},\epsilon,\delta}$ and $\mathsf{Exp}^{\mathsf{utrace}}_{\mathcal{A},\epsilon}$. The universal experiment uses universal goodness and sets the tracing parameter to $\delta = \epsilon$. Inside the figure, ϵ, δ abbreviate their values at λ; $(1^\lambda, n, t)$ is understood.

Ordinary and universal tracing experiments

Challenger C **Adversary** $\mathcal{A}$

Generate public parameters pp for $(1^\lambda, n, t)$.

$\xrightarrow{\mathsf{pp}}$

$\xleftarrow{T}$

Require $T \subseteq [n]$ and $1 \le f := |T| < t$; otherwise return 0.

Sample $s \xleftarrow{\$} \mathcal{S}$; run the honest dealer and tracer.

$\xrightarrow{\text{public messages } \mathbf{sh}_T}$

Complete sharing with online deliveries. Return 0 unless accepted. The challenger stores $(\mathsf{pub}, \mathsf{tk}, \mathbf{sh}, \sigma)$; reveal neither tk nor σ.

Using its sharing view, output a stateless box $\mathcal{R}$ with fixed embedded data.

$\xleftarrow{\mathcal{R}}$

Require a well-formed box with input arity $t - f$. Ordinary tracing requires ϵ-goodness (Definition 15); universal tracing requires ϵ-universal goodness (Definition 16). Return 0 if the applicable requirement fails.

C runs $o \leftarrow \mathsf{Trace}^{\mathcal{R}}(\mathsf{pp}, \mathsf{svk}, \mathsf{tvk}, \mathsf{tk}, 1^{\lceil 1/\delta \rceil})$ with black-box access to $\mathcal{R}$, setting $\delta = \epsilon$ in the universal experiment.

The box retains no memory of earlier calls and uses fresh independent coins for each response.

Return 0 if $o = (I, \pi)$, $I = T$, and $\mathsf{VerifyTrace}(\mathsf{pp}, \mathsf{svk}, \mathsf{tvk}, I, \pi) = 1$. Otherwise return 1, including when $o = \perp$.

Fig. 3. A common failure experiment for both tracing notions. The goodness requirements are mathematical predicates, not efficient tests.

Definition 18 (Publicly Verifiable Universal Traceability). For a negligible guessing baseline p_{guess}, a scheme has publicly verifiable universal traceability at threshold ϵ if

$$\Pr[\mathsf{Exp}^{\mathsf{utrace}}_{\mathcal{A},\epsilon}(1^\lambda, n, t) = 1] \le p_{\mathsf{guess}}(\lambda, n, t) + \mathsf{negl}(\lambda).$$

Definition 19 (Publicly Verifiable Traceability). A scheme has publicly verifiable traceability with parameters ϵ, δ if

$$\epsilon(\lambda)\ \Pr[\mathsf{Exp}^{\mathsf{trace}}_{\mathcal{A},\epsilon,\delta}(1^\lambda, n, t) = 1] \le \delta(\lambda) + p_{\mathsf{guess}}(\lambda, n, t) + \mathsf{negl}(\lambda).$$

The factor ϵ and guessing baseline account for boxes that ignore their inputs and guess a secret: such a box can be good on a small fraction of instances without identifying a coalition. The additive δ controls the tracing budget; it is not itself a bound

on the unweighted failure probability. Both experiments use the original authenticated record and unconditional probabilities, including parameter generation, sharing, adversary, tracer, and box randomness. Tracing receives only its declared inputs and oracle access.

The following lemma adapts the distribution-switching argument of [7, Lemma 1] to Definition 17 and our quantitative tracing bound.

Lemma 1. Let Π be bidirectional. For every PPT adversary $\mathcal{A}$ and parameters n, t, ϵ, δ as above,

$$\epsilon(\lambda) \cdot \Pr\left[\mathsf{Exp}^{\mathsf{trace}}_{\mathcal{A},\epsilon,\delta}(1^\lambda, n, t) = 1\right] \leq \delta(\lambda) + \Pr\left[\mathsf{Exp}^{\mathsf{utrace}}_{\mathcal{A},\delta}(1^\lambda, n, t) = 1\right] + \mathsf{negl}(\lambda).$$

Consequently, publicly verifiable universal traceability at threshold $\delta(\lambda)$ implies publicly verifiable traceability with parameters $\epsilon(\lambda), \delta(\lambda)$.

Proof. Fix λ, n, t and omit these arguments from the experiment names. Continue $\mathcal{A}$ from its sharing view to obtain $\mathcal{R}$, and run $\mathsf{Trace}^{\mathcal{R}}$ with parameter $\delta(\lambda)$. Let F_δ be the event that tracing fails to return exactly T with accepted evidence. Let E be the event that an additional test of $\mathcal{R}$ on Q returns the displayed secret and that this secret is not $\perp$. The test uses randomness independent of tracing; the box is stateless, so prior tracing queries do not change the test's distribution. Set both events to false on a $\perp$ sampler output or malformed box.

First condition on a complete honest sharing instance, including the dealer state σ, and the box with its fixed embedded data; whenever the box is $\epsilon(\lambda)$-good, the independent forward test succeeds with probability at least $\epsilon(\lambda)$. Averaging over these conditioned data gives the first inequality below. For the second inequality, the test and tracing form the same randomized postprocessing of (V, tk, z, Q) under both samplers, using no dealer state or goodness test, so bidirectionality gives

$$\begin{aligned}\epsilon(\lambda)\Pr\left[\mathsf{Exp}^{\mathsf{trace}}_{\mathcal{A},\epsilon,\delta} = 1\right] &\leq \Pr_{\mathsf{Fwd}}[E \wedge F_\delta] \\ &\leq \Pr_{\mathsf{Bwd}}[E \wedge F_\delta] + \mathsf{negl}(\lambda).\end{aligned}$$

For a fixed $(V, \mathsf{tk}, \mathcal{R})$, let p_{u} be the backward test's success probability, set to zero on a $\perp$ sampler output or malformed box. Independence of the test and tracing yields

$$\begin{aligned}\Pr_{\mathsf{Bwd}}[E \wedge F_\delta] &= \mathbb{E}\left[p_{\mathrm{u}}\,\mathbf{1}_{F_\delta}\right] \\ &\leq \delta(\lambda) + \Pr[p_{\mathrm{u}} \geq \delta(\lambda) \wedge F_\delta].\end{aligned}$$

The backward sampler has the same marginal distribution of (V, tk) as honest sharing. Thus the last probability is exactly the failure probability in the universal experiment at threshold $\delta(\lambda)$, which uses the same tracing parameter. Combining the inequalities proves the bound. □

5 TRAVESS: A Shamir-Based Construction

We start from Pedersen VSS and add traceability and non-imputability in turn. The resulting basic scheme motivates two further extensions: reconstruction consistency and publicly verifiable tracing.

5.1 Pedersen VSS

We begin with Pedersen's verifiable secret-sharing scheme [23], which combines Shamir secret sharing with Pedersen commitments to the polynomial coefficients. The dealer shares the secret using a random polynomial and uses a second, independent polynomial to provide the randomness for share verification. The coefficient commitments allow each participant to check its share, while reconstruction proceeds by Lagrange interpolation.

Let $\mathbb{G} = \langle g \rangle$ be a cyclic group of prime order q, and let h be a second generator for which no party knows $\log_g h$. Fix $2 \le t \le n < q$. All polynomial operations are over $\mathbb{F}_q$. The public parameters contain the group description and g, h, q, n, t. We assume private authenticated channels for share delivery and a common authenticated public record for the sharing instance.

To share a secret $s \in \mathbb{F}_q$, the dealer sets $a_0 = s$, samples $a_1, \dots, a_{t-1}$ and $b_0, \dots, b_{t-1}$ independently and uniformly from $\mathbb{F}_q$, and defines

$$P(X) = \sum_{j=0}^{t-1} a_j X^j, \qquad R(X) = \sum_{j=0}^{t-1} b_j X^j.$$

The polynomial P carries the secret in its constant term, while R supplies the verification randomness. For each coefficient, the dealer computes a Pedersen commitment and publishes the resulting share-verification key:

$$C_j = g^{a_j} h^{b_j} \quad (0 \le j < t), \qquad \mathsf{svk} = (C_0, \dots, C_{t-1}).$$

Fix distinct public evaluation points $x_1, \dots, x_n \in \mathbb{F}_q^*$, with x_i assigned to participant P_i; for example, take $x_i = i$. Participant P_i privately receives

$$\mathsf{sh}_i = (x_i, y_i, \rho_i) = \big(x_i, P(x_i), R(x_i)\big).$$

Given a share (x, y, ρ), a participant checks that $x \in \mathbb{F}_q^*$ and $y, \rho \in \mathbb{F}_q$, and accepts exactly when

$$g^y h^\rho = \prod_{j=0}^{t-1} C_j^{x^j}. \tag{1}$$

The equation checks the supplied polynomial evaluation and its opening against the published coefficient commitments.

Given t shares, relabel them as (x_k, y_k, ρ_k) for $k = 1, \dots, t$. Reconstruction first verifies every share using Equation (1) and checks that the evaluation points are pairwise distinct. If any check fails, it returns $\perp$. Otherwise, it computes

$$s^* = \sum_{k=1}^{t} \lambda_k y_k, \qquad \lambda_k = \prod_{\substack{j=1 \\ j \ne k}}^{t} \frac{-x_j}{x_k - x_j}.$$

The verification coordinates ρ_k are used to validate the shares, while the secret is reconstructed from their evaluation points and y-coordinates.

Although this construction provides share verification and reconstruction, it has no mechanism for identifying the participants behind a reconstruction box. Its evaluation points are public, and it supplies no tracing key. Section 5.2 makes these points private and gives a designated tracer the material needed to generate valid query shares.

5.2 Traceable Pedersen VSS

To add traceability, we modify the choice of evaluation points and the generation of the commitment parameters. Sharing still uses the polynomials P, R and coefficient commitments above, and share verification and reconstruction retain the same equations.

Following the approach of Boneh, Partap, and Rotem [7], we replace the fixed public coordinates with randomly chosen evaluation points. Once P, R are fixed, each dealer-local call to Share samples $x \xleftarrow{\$} \mathbb{F}_q^*$ and outputs the share $(x, P(x), R(x))$ together with the key components $\mathsf{tk}_i = \mathsf{tvk}_i = x$. To register n participants, the dealer calls this algorithm repeatedly, discarding a candidate if its point repeats an already registered point, and assigning each retained output to the next participant. Registration aborts if fewer than n outputs have been admitted after $2n$ calls. The underlying Shamir share of participant P_i is now $(x_i, P(x_i))$, rather than $(i, P(i))$. The evaluation points are no longer public: each participant privately receives its own share, while their association with participant identities is recorded in the tracing material.

The tracer queries a reconstruction box with additional shares whose reconstruction values it must be able to vary. Because the box may verify these shares, each query must still satisfy Equation (1). We achieve this using the equivocation property of Pedersen commitments from Section 3.6.2, adapted so that the tracer does not need the dealer's original openings (a_j, b_j).

For each coefficient commitment $C_j = g^{a_j} h^{b_j}$, give the tracer the trapdoor $\alpha = \log_g h$ and the exponent $d_j = a_j + \alpha b_j$, so that $C_j = g^{d_j}$. It can then open C_j to any $m \in \mathbb{F}_q$ by computing

$$r^* := \mathsf{Equiv}(\mathsf{pp}_{\mathrm{com}}, \alpha, d_j, m) = (d_j - m)\alpha^{-1}, \qquad g^m h^{r^*} = C_j.$$

This is the usual equivocation formula $r^* = b_j + (a_j - m)\alpha^{-1}$ expressed using d_j instead of the original opening. For fixed public parameters, α, a_j, m and uniform $b_j \in \mathbb{F}_q$, the exponent d_j is uniform, giving

$$(C_j, r^*) \triangleq (\mathsf{Com}(\mathsf{pp}_{\mathrm{com}}, m; r), r), \qquad r \xleftarrow{\$} \mathbb{F}_q.$$

The values $(\alpha, \mathbf{d})$ therefore let the tracer compute new openings without receiving P, R separately. The trapdoor alone would not supply the exponents d_j; Section 5.4 explains how the tracer obtains them and checks $g^{d_j} = C_j$ when the dealer may be malicious.

In this basic variant, the dealer initializes the instance by sampling $\alpha \xleftarrow{\$} \mathbb{F}_q^*$, setting $h = g^\alpha$, and computing the coefficient commitments and $d_j = b_j\alpha + a_j$ for $0 \le j < t$. After the individual shares and key components have been generated, it privately delivers the assembled tracing key

$$\mathsf{tk} = (\alpha, \mathbf{d}, \mathbf{x}), \qquad \mathbf{d} = (d_0, \ldots, d_{t-1}), \qquad \mathbf{x} = (x_1, \ldots, x_n)$$

to the tracer. For checking tracing evidence, it privately supplies an authorized verifier with the original tracing-verification key

$$\mathsf{tvk}_{\mathrm{in}} = \mathbf{x} = (x_1, \ldots, x_n),$$

where entry x_i is associated with participant P_i. The share-verification key svk is fixed in the authenticated public record; the tracing key and the original tracing-verification key are supplied over private authenticated channels. After that, α is removed from the dealer's state. The subscript in distinguishes this private identity table from the public tracing-verification information to be introduced in Section 5.5.

To construct a query share, the tracer uses the coefficients $\mathbf{d}$ in tk to evaluate the combined polynomial

$$D(X) := \sum_{j=0}^{t-1} d_j X^j = P(X) + \alpha R(X).$$

The tracer evaluates $D(x)$ locally using $\mathbf{d}$, without obtaining $P(x)$ and $R(x)$ separately.

For a chosen point $x \in \mathbb{F}_q^*$ and value $y \in \mathbb{F}_q$, it computes

$$\rho = \alpha^{-1}\big(D(x) - y\big), \tag{2}$$

and forms the artificial share (x, y, ρ). The value y need not equal $P(x)$, and share verification does not require x to be a registered participant point.

In particular, given such a share and any $\Delta \in \mathbb{F}_q^*$, the tracer can form the pair

$$(x, y, \rho) \qquad \text{and} \qquad (x, y + \Delta, \rho - \Delta/\alpha).$$

The pair changes the special share's y-coordinate by Δ and adjusts its verification coordinate by $-\Delta/\alpha$. These are the input perturbations used by the tracing procedure.

A box containing $f < t$ genuine shares takes $t - f$ additional shares as input. We need two sampling rules to distinguish its ordinary reconstruction ability from its behavior on tracing queries. Additional honest shares test whether the box can recover the original secret, as required by ordinary goodness. Artificial valid shares give the tracer inputs it can generate from its own key, without the dealer's state; they define the test for universal goodness and supply the share construction used in tracing. The bidirectional-sharing condition of Section 4.7 connects these two tests in the traceability analysis.

We instantiate the two sampling rules of Section 4.7 using the public sharing record $\mathsf{pub} = (\mathsf{pp}, \mathsf{svk})$ at this stage; the original identity table remains private. Given the fixed dealer state $\sigma = (P, R)$, obtain k additional honest shares by making k independent calls to $\mathsf{Share}(\mathsf{pp}, \sigma)$ and retaining only the share outputs. For reconstruction-box tests, we take $k = t - f$. Each call samples a new $x'_j \xleftarrow{\$} \mathbb{F}_q^*$ and returns $(x'_j, P(x'_j), R(x'_j))$. For the additional valid-share distribution $\mathcal{D}_{\mathsf{pub},\mathsf{tk}}$, sample $x \xleftarrow{\$} \mathbb{F}_q^*$ and $y \xleftarrow{\$} \mathbb{F}_q$ independently and output

$$\big(x, y, \alpha^{-1}(D(x) - y)\big).$$

Repeated samples from this distribution are independent, as required by Definition 16. The additional Share calls discard their key components and perform no registration.

Neither sampling rule changes the public record or rejects repeated points; collisions within the test inputs or with coalition points are counted as reconstruction failures. The paired queries below use the artificial-share construction with distinct-point sampling and a controlled perturbation of one reconstruction value.

Our procedure builds on the paired-query and polynomial-recovery approach of Boneh, Partap, and Rotem [7, Sec. 3.2]. We work in the stateless-box model of Section 4.7. The public input arity $t-f$ determines the coalition size f. Write T for the coalition and $\xi_1, \ldots, \xi_f$ for the points of its fixed genuine shares; their association with the box is unknown to the tracer. Fix integer parameters N, C with $\sqrt{fN} < C \le N \le q-1$.

In each of N rounds, the tracer samples a special point x' uniformly from the nonzero points not previously used as special points. The other $t-f-1$ points, $x_{f+1}, \ldots, x_{t-1}$, are sampled uniformly without replacement from $\mathbb{F}_q^* \setminus \{x'\}$. All y-coordinates are sampled independently and uniformly from $\mathbb{F}_q$. It completes each share using Equation (2). Both query tuples list the ordinary shares in the order $f+1, \ldots, t-1$, followed by the special share. They use (x', y', ρ') and $(x', y'+\Delta, \rho'-\Delta/\alpha)$ as their last entries, respectively, for $\Delta \xleftarrow{\$} \mathbb{F}_q^*$. Let v, v' denote the responses. If either response is not a field element or $v = v'$, the tracer treats the round as an erasure. Otherwise, it records (x', z), where

$$z = \frac{\Delta}{v' - v} \prod_{i=f+1}^{t-1} \frac{x_i}{x_i - x'}. \tag{3}$$

Empty products are one. The polynomial targeted by tracing is

$$H_T(X) = \prod_{k=1}^{f} \frac{\xi_k - X}{\xi_k}. \tag{4}$$

Its roots identify the coalition points. The recorded samples may contain errors.

Let M be the number of recorded samples. If $M < C$, tracing returns $\perp$. Otherwise, it uses Reed–Solomon list decoding to recover all polynomials of degree at most f agreeing with at least C samples. The tracer factors the candidates and retains those of degree exactly f with f distinct roots all belonging to the registered vector $\mathbf{x}$ in tk. If none remain, it returns $\perp$; otherwise, it selects a candidate by a fixed ordering. It orders the matched participant indices as $i_1 < \cdots < i_f$ and orders the recovered roots $w_1, \ldots, w_f$ accordingly. The tracer submits the accusation and evidence

$$I = (i_1, \ldots, i_f), \qquad \pi = (w_1, \ldots, w_f), \qquad w_k = x_{i_k} \quad (k \in [f]).$$

Given the original authenticated key $\mathsf{tvk}_{\mathrm{in}} = \mathbf{x}$, an authorized verifier checks that the table contains n distinct nonzero field elements, that $1 \le |I| < t$, and that the indices are strictly increasing and belong to $[n]$. It rejects malformed evidence or evidence not containing exactly $|I|$ distinct nonzero field elements. It then accepts if and only if

$$w_k = x_{i_k} \qquad \text{for every } k \in [|I|].$$

This defines the preliminary verifier $\mathsf{VerifyTrace}_0(\mathsf{pp}, \mathsf{svk}, \mathsf{tvk}_{\mathrm{in}}, I, \pi)$. The check certifies the association between the submitted roots and the registered identities.

However, the tracing key contains every registered point. A malicious tracer can therefore submit any participant's point and obtain an accepting accusation without querying a reconstruction box. The current evidence offers no protection against such false accusations. In Section 5.3, we replace the points in the tracing and verification keys with their one-way images, so that an accusation must supply a preimage that the keys do not disclose.

5.3 Traceable Pedersen VSS with Non-Imputability

Following the approach of Boneh, Partap, and Rotem [7, Sec. 3.3], we replace the cleartext points in both keys with their one-way images. Let $\mathsf{OW} : \mathbb{F}_q^* \to \mathcal{U}$ be a one-way function for uniformly sampled inputs, with its description included in the public parameters. Each call to Share now outputs $\mathsf{tk}_i = \mathsf{tvk}_i = u_i$, where

$$u_i = \mathsf{OW}(x_i) \quad (i \in [n]), \qquad \mathbf{u} = (u_1, \ldots, u_n).$$

Formal registration requires both the evaluation points and their images to be distinct: the dealer discards a candidate output and calls Share again if its point or image repeats a previously registered entry. This rejection rule belongs to registration, not to the single-share algorithm. It gives an unambiguous association between images and participant indices without assuming that OW is injective.

The dealer privately delivers

$$\mathsf{tk} = (\alpha, \mathbf{d}, \mathbf{u})$$

to the tracer and privately supplies an authorized verifier with the original authenticated key $\mathsf{tvk}_{\mathsf{in}} = \mathbf{u}$, with entry u_i associated with participant P_i. The vector $\mathbf{u}$ replaces $\mathbf{x}$ in both keys and is not published; the cleartext evaluation points are neither published nor additionally supplied to the tracer. Participant P_i still receives $(x_i, P(x_i), R(x_i))$. Share verification, reconstruction, and the generation of artificial shares using $(\alpha, \mathbf{d})$ are unchanged.

The tracer performs the same paired queries and list decoding as in the basic construction. For each degree-f candidate with f distinct nonzero roots $w_1, \ldots, w_f$, it computes their images and looks them up in $\mathbf{u}$. It retains the candidate only if the images match f distinct registered entries. If no candidate remains, it returns $\perp$; otherwise, it selects one by the fixed ordering used in the basic construction. It orders the matched participant indices as $i_1 < \cdots < i_f$ and orders the roots accordingly, outputting

$$I = (i_1, \ldots, i_f), \qquad \pi = (w_1, \ldots, w_f), \qquad \mathsf{OW}(w_k) = u_{i_k} \quad (k \in [f]).$$

The verifier $\mathsf{VerifyTrace}_0(\mathsf{pp}, \mathsf{svk}, \mathsf{tvk}_{\mathsf{in}}, I, \pi)$ first checks that $\mathsf{tvk}_{\mathsf{in}}$ contains n well-formed distinct images, that $1 \le |I| < t$, and that the indices in I are strictly increasing and lie in $[n]$. It rejects malformed evidence or evidence not containing exactly $|I|$ distinct elements of $\mathbb{F}_q^*$. It accepts if and only if

$$\mathsf{OW}(w_k) = u_{i_k} \qquad \text{for every } k \in [|I|]. \tag{5}$$

The verifier uses the original authenticated identity table. Since successful inversion means finding any preimage, the check does not require a root to equal the originally sampled point beyond satisfying this equation.

A remaining problem is that the dealer still generates α and can therefore equivocate its coefficient commitments. It can issue verification-valid shares that reconstruct to different secrets, and erasing α afterwards does not prevent this deviation. Section 5.4 moves trapdoor generation to the tracer and uses committed-receiver OLE to supply the tracing coefficients privately to it, keeping α from the dealer.

5.4 Reconstruction Consistency

To keep the trapdoor from the dealer, we generate h in the interactive initialization $\mathsf{Gen}(\mathsf{pp}, s)$, after the group and other public parameters have been fixed. The tracer samples $\alpha \xleftarrow{\$} \mathbb{F}_q^*$ and sends $h = g^\alpha$ to the dealer, which rejects unless $h \in \mathbb{G} \setminus \{1\}$. Using its private input s, the dealer samples P, R, forms the coefficient commitments, and fixes

$$\mathsf{svk} = (h, C_0, \ldots, C_{t-1})$$

in the common authenticated record. Public record compatibility in Definition 12 requires the h in this key to equal the authenticated value initially sent by the tracer, independently of the tracer's private check result. Every $\mathcal{F}_{\mathrm{CROLE}}$ call must use the same h in its common input, and the tracer rejects locally on any mismatch. Share verification and reconstruction obtain h from this key, reject malformed keys or invalid group elements, and otherwise use the same verification and interpolation rules as before.

The tracer still needs $d_j = b_j\alpha + a_j$ to generate artificial shares. During Gen, we obtain these coefficients using $\mathcal{F}_{\mathrm{CROLE}}$ from Definition 7, with the malicious security and abort guarantee required in Section 3.7. After fixing svk, the dealer and tracer make t sequential calls, one for each $j = 0, \ldots, t-1$. In call j, the dealer is the sender with input (a_j, b_j), and the tracer is the receiver with input α. Every call uses the same common input $(\mathbb{G}, q, g, h)$ and returns only to the tracer

$$d_j = b_j\alpha + a_j \in \mathbb{F}_q.$$

The dealer receives no functionality output and retains $\sigma = (P, R)$ for its subsequent local Share calls. Those calls generate the participant shares and identity-table components as in Section 5.3; the dealer no longer generates α or directly supplies $\mathbf{d}$.

After the calls, the tracer rejects locally if any output is $\perp$ or is not in $\mathbb{F}_q$, or if any check

$$g^{d_j} = C_j \qquad (0 \le j < t) \tag{6}$$

fails. This check is necessary because $\mathcal{F}_{\mathrm{CROLE}}$ constrains the receiver's input but does not bind the sender's inputs to the published coefficient commitments. The tracer retains the initialization output $\mathsf{tk}_0 = (\alpha, \mathbf{d})$ and its local check result privately, without sending an acceptance message. The table supplied after the individual Share calls completes the tracing key as $\mathsf{tk} = (\alpha, \mathbf{d}, \mathbf{u})$. The tracer uses $D(X) = \sum_{j=0}^{t-1} d_j X^j$ and Equation (2) to generate artificial shares. The tracing and original evidence-verification procedures remain unchanged.

With an honest tracer, α remains private, and the dealer receives neither a $\mathcal{F}_{\text{CROLE}}$ output nor the tracer's local check results. Under the fixed verification key, Lagrange interpolation of any t verified shares at distinct points yields an opening of C_0 to the reconstructed secret. Two successful reconstructions of different secrets would therefore give openings of the same Pedersen commitment to different messages, contradicting its binding under the discrete-logarithm assumption as discussed in Section 3.6.2. This gives reconstruction consistency in the sharing-stage view of Definition 12, before tracing queries or outputs are exposed.

Tracing evidence is still verified using the private identity table $\mathsf{tvk}_{\text{in}} = \mathbf{u}$, so only a verifier holding that table can check an accusation. Publishing the entire table during sharing would disclose all registered one-way images and their association with participant identities, including entries unrelated to an accusation. Section 5.5 enables public verification through a commitment to the ordered images and openings at the accused positions.

5.5 Publicly Verifiable Tracing

We use the vector commitment of Definition 6 to bind the private identity table to a public commitment. The dealer publishes a commitment to the ordered table during sharing, and the tracer later opens only the positions named in an accusation. The recovered roots remain public, so these openings let anyone check their association with the committed identities. Statistical selective-opening hiding is an optional additional property that limits leakage about unopened entries from the commitment and disclosed openings; the TVSS security proofs require only correctness and position binding from the vector commitment.

Before the secret is chosen and Gen is run, the public parameters pp include

$$\mathsf{pp}_{\text{vc}} \leftarrow \mathsf{VC.Setup}(1^{\lambda}, n),$$

with message space $\mathcal{U}$. These parameters are generated independently of the tracing trapdoor α. After registering n outputs of Share, the dealer locally assembles their private key components into the ordered image table $\mathbf{u}$ and commits to it:

$$(c_{\text{tvk}}, \mathsf{aux}_{\text{vc}}) \leftarrow \mathsf{VC.Commit}(\mathsf{pp}_{\text{vc}}, \mathbf{u}).$$

It fixes c_{tvk} alongside svk and the participant ordering in the common authenticated record, and privately sends $(\mathbf{u}, \mathsf{aux}_{\text{vc}})$ to the tracer. The tracer checks that the table contains n well-formed, distinct images. For each position $i \in [n]$, it computes an opening $\omega_i = \mathsf{VC.Open}(\mathsf{pp}_{\text{vc}}, \mathsf{aux}_{\text{vc}}, i)$. It checks that $\mathsf{VC.Verify}(\mathsf{pp}_{\text{vc}}, c_{\text{tvk}}, i, u_i, \omega_i) = 1$. Malformed material, a failed opening computation, or a failed check causes local rejection. Together with the initialization of Section 5.4, an accepted execution gives

$$\mathsf{tvk} = c_{\text{tvk}}, \qquad \mathsf{tk} = (\alpha, \mathbf{d}, \mathbf{u}, \mathsf{aux}_{\text{vc}}).$$

The share-verification key remains $\mathsf{svk} = (h, C_0, \ldots, C_{t-1})$. Participant shares and their verification and reconstruction rules are unchanged. The table and opening state are supplied privately to the tracer and are not explicitly published; public tracing verification uses tvk in place of the original private verification key.

The tracer uses the same paired queries, list decoding, and matching against $\mathbf{u}$ as in Section 5.3. If this procedure fails, it returns $\perp$. Otherwise, let $I = (i_1, \ldots, i_f)$ with $i_1 < \cdots < i_f$ be the matched indices, and let $w_1, \ldots, w_f$ be the corresponding roots, satisfying $\mathsf{OW}(w_k) = u_{i_k}$. For each $k \in [f]$, the tracer computes $\omega_{i_k} = \mathsf{VC.Open}(\mathsf{pp}_{\mathrm{vc}}, \mathsf{aux}_{\mathrm{vc}}, i_k)$ and outputs (I, π), where

$$\pi = (\pi_0, \pi_{\mathrm{tr}}), \qquad \pi_0 = (w_1, \ldots, w_f), \qquad \pi_{\mathrm{tr}} = (\omega_{i_1}, \ldots, \omega_{i_f}).$$

Thus the original root evidence is accompanied by an opening for each accused position. The roots and their one-way images are public; only the accused table positions are explicitly opened.

The algorithm $\mathsf{VerifyTrace}(\mathsf{pp}, \mathsf{svk}, \mathsf{tvk}, I, \pi)$ uses the parameters and commitment fixed in the sharing record. It rejects malformed inputs and requires $1 \le f = |I| < t$, strictly increasing indices in $[n]$, and exactly f roots and f openings. It checks that the roots are distinct elements of $\mathbb{F}_q^*$ and that their images are well-formed, distinct elements of $\mathcal{U}$. It then accepts exactly when

$$\mathsf{VC.Verify}\big(\mathsf{pp}_{\mathrm{vc}}, c_{\mathrm{tvk}}, i_k, \mathsf{OW}(w_k), \omega_{i_k}\big) = 1 \qquad \text{for every } k \in [f]. \tag{7}$$

Verification requires only the public record and the submitted evidence. It checks the association between the disclosed roots and registered identities without access to the private table or the reconstruction box.

5.6 Complete Construction

We collect the preceding components into the final Shamir-based construction, using the six interfaces of Definition 8. The description below fixes the sharing order, final key formats, and tracing parameters. The verification equations and paired-query procedure are those specified above.

Public parameters. For $(1^\lambda, n, t)$, the public parameters are honestly generated before the secret is chosen. They contain $(\mathbb{G}, q, g) \leftarrow \mathsf{GroupGen}(1^\lambda)$ as in Definition 4, with $q > n$, the description of $\mathsf{OW} : \mathbb{F}_q^* \to \mathcal{U}$, and independently sampled parameters $\mathsf{pp}_{\mathrm{vc}} \leftarrow \mathsf{VC.Setup}(1^\lambda, n)$ with message space $\mathcal{U}$. Thus

$$\mathsf{pp} = (\mathbb{G}, q, g, n, t, \mathsf{OW}, \mathsf{pp}_{\mathrm{vc}}), \qquad \mathcal{S} = \mathbb{F}_q.$$

The element h is generated during Gen, independently of $\mathsf{pp}_{\mathrm{vc}}$.

A concrete instantiation is $\mathsf{OW}(x) = g^x$, using the unique group-element encoding as the vector-commitment message. This map is injective on $\mathbb{F}_q^*$ and one-way under Definition 4. Its collision probability on two independent uniform inputs is $1/(q-1)$, supplying the uniform honest-completion bound of Lemma 2; see Appendix A.1. The generic construction retains an arbitrary OW satisfying the stated one-wayness and completion hypotheses.

Instance generation. On common input pp and the dealer's private input $s \in \mathbb{F}_q$, the dealer and tracer execute $\mathsf{Gen}(\mathsf{pp}, s)$ as follows, using the channels and authenticated record of Section 4.1.

1. The tracer samples $\alpha \xleftarrow{\$} \mathbb{F}_q^*$ and sends $h = g^\alpha$ to the dealer. The dealer rejects unless $h \in \mathbb{G} \setminus \{1\}$.
2. The dealer samples P, R with $P(0) = s$ as in Section 5.1 and computes $C_j = g^{a_j} h^{b_j}$ for $0 \le j < t$. Using the previously fixed participant ordering, it publishes the share-verification key
$$\mathsf{svk} = (h, C_0, \ldots, C_{t-1})$$
in the common authenticated record. Public record compatibility requires the h in this key to equal the authenticated value sent by the tracer in Step 1, independently of the tracer's private check result. Every $\mathcal{F}_{\mathrm{CROLE}}$ call must use the same h in its common input, and the tracer rejects locally on any mismatch.
3. The parties make t sequential $\mathcal{F}_{\mathrm{CROLE}}$ calls, all with common input $(\mathbb{G}, q, g, h)$. For $j = 0, \ldots, t-1$, the dealer supplies (a_j, b_j) and the tracer supplies α; only the tracer receives $d_j = b_j\alpha + a_j$. The calls use the functionality and abort guarantee of Section 3.7.
4. The tracer checks that every $d_j \in \mathbb{F}_q$ satisfies Equation (6). It rejects locally on a missing or malformed value, a functionality abort, or a failed check. Otherwise, it retains $\mathsf{tk}_0 = (\alpha, \mathbf{d})$. These checks and their outcome remain private; the dealer receives no functionality output or acceptance message.

The outputs are the public key svk, the dealer's private state $\sigma = (P, R)$, and the tracer's private initialization material tk_0 when its checks succeed. The state is bound to $s = P(0)$. The dealer proceeds to share generation after completing its own initialization steps; it neither waits for nor learns the tracer's local acceptance decision.

Single-share generation. The dealer-local algorithm $\mathsf{Share}(\mathsf{pp}, \sigma)$ parses $\sigma = (P, R)$, samples $x \xleftarrow{\$} \mathbb{F}_q^*$, and computes $u = \mathsf{OW}(x)$. It outputs

$$\big((x, P(x), R(x)), u, u\big),$$

where the first component is the share and the last two are its tracing-key and tracing-verification-key components. The share identifier is its evaluation point: $\mathsf{id}((x, y, \rho)) = x$, with identifier space $\mathcal{I} = \mathbb{F}_q^*$. The algorithm leaves σ unchanged and uses independent randomness on every call. The key components remain private until they are processed by the registration stage below.

Registration and key assembly. For each participant P_i in the fixed order, the dealer calls $\mathsf{Share}(\mathsf{pp}, \sigma)$. It discards the output and repeats the call if its point or image coincides with a previously registered entry. Otherwise, it records $u_i = \mathsf{tk}_i = \mathsf{tvk}_i$ and privately delivers the retained share sh_i to P_i. The participant checks that share using $\mathsf{VerifyShare}$.

Registration makes at most $2n$ calls in total; if fewer than n outputs have been admitted, the dealer aborts without publishing a final tvk. Otherwise, it locally computes

$$(c_{\mathrm{tvk}}, \mathsf{aux}_{\mathrm{vc}}) \leftarrow \mathsf{VC.Commit}(\mathsf{pp}_{\mathrm{vc}}, (u_1, \ldots, u_n)),$$

and publishes $\mathsf{tvk} = c_{\mathrm{tvk}}$ in the authenticated record. It sends the pair $(\mathbf{u}, \mathsf{aux}_{\mathrm{vc}})$ privately to the tracer. The tracer checks that the table contains n well-formed distinct images and that every position opens correctly to its supplied image, as specified in Section 5.5. It rejects locally on missing or malformed material or a failed check. If these checks and its earlier initialization checks pass, its final key is $\mathsf{tk} = (\alpha, \mathbf{d}, \mathbf{u}, \mathsf{aux}_{\mathrm{vc}})$. No check result or acceptance message is sent to the dealer.

Key assembly is a local stage of the sharing execution, followed by private delivery to the tracer; it requires no further interactive key generation. The final public record is $\mathsf{pub} = (\mathsf{pp}, \mathsf{svk}, \mathsf{tvk})$ with the fixed participant ordering. The shares are private to their recipients and the dealer; the final tracing key is private to the tracer. Acceptance in Definition 8 covers initialization, share delivery, and key assembly, and does not require an additional protocol message.

Share verification. The algorithm $\mathsf{VerifyShare}$ parses svk as $(h, C_0, \ldots, C_{t-1})$ and sh as (x, y, ρ). It rejects malformed inputs, requires $h \in \mathbb{G} \setminus \{1\}$, $C_j \in \mathbb{G}$, $x \in \mathbb{F}_q^*$, and $y, \rho \in \mathbb{F}_q$, and accepts exactly when Equation (1) holds. Verification does not consult the identity table or require a registered evaluation point.

Reconstruction. The algorithm Rec requires exactly t shares, retaining multiplicities. It returns $\perp$ if any share fails $\mathsf{VerifyShare}$ or if two evaluation points coincide. Otherwise, it returns the Lagrange interpolation of the y-coordinates at zero, as specified in Section 5.1.

Tracing. Given $(\mathsf{pp}, \mathsf{svk}, \mathsf{tvk}, \mathsf{tk}, 1^{\lceil 1/\delta \rceil})$ and access to $\mathcal{R}$, Trace obtains f from the box's input arity $t - f$ and returns $\perp$ unless $1 \le f < t$. Let $m = \lceil 1/\delta \rceil$, as supplied in unary, and recover λ as the bit length of q. Set

$$N_0 = 64 f m^4 + 16(\lambda + 1)m^2, \qquad N = \min\{N_0, q - 1\},$$
$$C = \max\left\{ \left\lfloor \frac{N}{4m^2} \right\rfloor, \left\lfloor \sqrt{fN} \right\rfloor + 1 \right\}.$$

These integers satisfy $\sqrt{fN} < C \le N \le q - 1$ on every admissible input. With these parameters, tracing executes the ordered paired queries and list decoding of Section 5.2, using $(\alpha, \mathbf{d})$ to form artificial shares.[5] The running-time bounds treat exact uniform sampling from finite fields and integer intervals as primitive operations. Sampling without

[5] Following [7, footnote 7], we implement list decoding and candidate factorization in strict polynomial time using standard time limits and independent retries. The all-input decoding bound is justified in Appendix A.9. Uniformly over valid inputs, these routines return complete, exact results except with negligible probability of failure; on failure, Trace returns $\perp$ and retains no candidates. We absorb this failure probability into the negligible terms in the security bounds.

replacement maps a uniform rank among the remaining field elements through a sorted list of excluded points.

Candidate filtering uses the one-way-image matching rule of Section 5.3, with the same fixed candidate ordering. On success, tracing orders the matched indices and roots and produces $\pi = (\pi_0, \pi_{\mathrm{tr}})$ by opening the accused positions as in Section 5.5; otherwise, it returns $\perp$. Appendix A.9 proves the parameter bounds and the resulting tracing guarantee.

Tracing verification. The algorithm $\mathsf{VerifyTrace}(\mathsf{pp}, \mathsf{svk}, \mathsf{tvk}, I, \pi)$ applies all format, index, root, and image checks of Section 5.5. It accepts exactly when every submitted opening satisfies Equation (7) under the commitment tvk in the original record.

Additional-share sampling. The construction uses the two sampling rules of Section 5.2 under the final record pub. To obtain k additional honest shares, make k independent calls to $\mathsf{Share}(\mathsf{pp}, \sigma)$ for the fixed state $\sigma = (P, R)$ and discard their key components. For reconstruction-box tests, we take $k = t - f$. These calls perform no registration or rejection sampling. The distribution $\mathcal{D}_{\mathsf{pub},\mathsf{tk}}$ independently samples $x \xleftarrow{\$} \mathbb{F}_q^*$ and $y \xleftarrow{\$} \mathbb{F}_q$, returning $(x, y, \alpha^{-1}(D(x) - y))$ with $D(X) = \sum_{j=0}^{t-1} d_j X^j$. Both sampling rules retain point collisions and leave the public record and participant registration unchanged, as required by the experiments of Section 4.7.

5.7 Security Analysis

We assume discrete-logarithm hardness in $\mathbb{G}$ and one-wayness of OW on uniform inputs from $\mathbb{F}_q^*$. The vector commitment satisfies the correctness and position-binding requirements of Definition 6; its optional statistical selective-opening hiding property is not needed for the security results below. We use the $\mathcal{F}_{\mathrm{CROLE}}$ functionality of Section 3.7, whose realization must be maliciously secure with abort.

Throughout, n, t are polynomially bounded with $2 \le t \le n$, and the public parameters are honestly generated. We assume that $p_{\mathsf{guess}}(\lambda, n, t) = \mathbb{E}_{\mathsf{pp}}[1/q]$ is negligible. The functions ϵ, δ are as in Section 4.8, with polynomially bounded reciprocals.

The statements follow the order of the properties in Section 4. Appendix A provides their proofs in the same order.

Lemma 2 (Correctness and public share verifiability). Suppose an honest sharing execution, including Gen, registered Share calls, and key assembly, takes polynomial time and fails to be accepted with probability at most $\eta_{\mathrm{sh}}(\lambda)$, uniformly over the admissible public parameters and secrets, where η_{sh} is negligible. Then the construction satisfies correctness (Definition 9) and public share verifiability (Definition 11), as well as the additional honest-share requirement of Section 4.7. In particular, for every accepted honest sharing of s and every $J \subseteq [n]$ with $|J| = t$,

$$\mathsf{VerifyShare}(\mathsf{pp}, \mathsf{svk}, \mathsf{sh}_i) = 1 \quad (i \in [n]), \qquad \mathsf{Rec}(\mathsf{pp}, \mathsf{svk}, \{\mathsf{sh}_i\}_{i \in J}) = s.$$

For every reconstruction box in the prescribed model, each non-$\perp$ tracing output under that instance is accepted by $\mathsf{VerifyTrace}$.

Correctness consists of sharing and reconstruction correctness, proved in Appendix A.1, and tracing-evidence validity, proved in Appendix A.2. Appendix A.1 also establishes the additional honest-share requirement. **Public share verifiability** is proved in Appendix A.3.

Lemma 3 (Reconstruction consistency). Under the discrete-logarithm assumption of Definition 4, the construction satisfies reconstruction consistency in the $\mathcal{F}_{\mathrm{CROLE}}$-hybrid model. For every PPT adversary $\mathcal{A}$ in Definition 12, there exists a negligible function ν_{cons} such that

$$\Pr\left[\mathsf{Exp}^{\mathsf{cons}}_{\mathcal{A},\mathsf{TRAVESS}}(1^\lambda, n, t) = 1\right] \le \nu_{\mathrm{cons}}(\lambda).$$

The reduction simulates the dealer's full sharing view and extracts from any two admissible tuples reconstructing distinct secrets, even when the tracer privately rejects; see Appendix A.4.

We next address a malicious tracer, first considering secrecy and then protection against false accusations.

Lemma 4 (Tracer secrecy and secrecy). In the $\mathcal{F}_{\mathrm{CROLE}}$-hybrid model, for every PPT adversary $\mathcal{A}$ in Definition 13,

$$\Pr[\mathsf{Exp}^{\mathsf{tsec}}_{\mathcal{A}}(0) = 1] = \Pr[\mathsf{Exp}^{\mathsf{tsec}}_{\mathcal{A}}(1) = 1].$$

The equality includes the adversary's complete interactive view and observable aborts. Consequently, the construction also satisfies secrecy in Definition 10.

The distribution argument and the implication to secrecy are given in Appendices A.5 and A.6, respectively.

Lemma 5 (Non-imputability). Assume that OW is one-way on uniformly sampled inputs in $\mathbb{F}_q^*$ under the public-parameter distribution, and that the vector commitment is position binding. Then the construction satisfies non-imputability in the $\mathcal{F}_{\mathrm{CROLE}}$-hybrid model: for every PPT adversary $\mathcal{A}$ in Definition 14, there exists a negligible function ν_{ni} such that

$$\Pr[\mathsf{Exp}^{\mathsf{ni}}_{\mathcal{A}}(1^\lambda, n, t) = 1] \le \nu_{\mathrm{ni}}(\lambda).$$

The reduction permits the adversary to choose the secret and receive all non-target shares; see Appendix A.7.

We now turn from protection against a malicious tracer to the guarantees of honest tracing. Bidirectional sharing connects the two reconstruction-box distributions.

Lemma 6 (Bidirectional sharing). In the $\mathcal{F}_{\mathrm{CROLE}}$-hybrid model, for every PPT coalition selector $\mathcal{A}_0$ of Definition 17,

$$\Delta\Big(\mathsf{Fwd}_{\mathcal{A}_0}(1^\lambda, n, t), \mathsf{Bwd}_{\mathcal{A}_0}(1^\lambda, n, t)\Big) \le \mathbb{E}_{\mathsf{pp}}\left[\frac{\binom{t}{2}}{q-1}\right].$$

In particular, the construction is bidirectional when p_{guess} is negligible.

The comparison retains the complete sharing view and $\mathsf{tk} = (\alpha, \mathbf{d}, \mathbf{u}, \mathsf{aux}_{\mathrm{vc}})$. Its proof is given in Appendix A.8.

Theorem 1 (Publicly verifiable universal traceability). Fix a PPT adversary $\mathcal{A}$ and use the tracing parameter rule of Section 5.6, with $m = \lceil 1/\epsilon(\lambda) \rceil$. In the experiment of Definition 18, let G_ϵ be the event that sharing is accepted and the adversary outputs a well-formed $\epsilon(\lambda)$-universally good box for a valid coalition T.

Let L count the tracing rounds in which the supplied points avoid the coalition points and both responses equal reconstruction from the coalition shares and the supplied inputs. Let B be the event that a candidate returned by the algebraic routines passes the image-matching filter but identifies an index set different from T; on algebraic failure, the returned list is empty. Then

$$\begin{aligned}\Pr[\mathsf{Exp}^{\mathsf{utrace}}_{\mathcal{A},\epsilon}(1^\lambda, n, t) = 1] \\ \le \Pr[G_\epsilon \wedge L < C] + \Pr[G_\epsilon \wedge B] + \mathsf{negl}(\lambda).\end{aligned}$$

Moreover, for the non-imputability adversary $\mathcal{B}_{\mathrm{ni}}$ constructed in Appendix A.9,

$$\begin{aligned}\Pr[G_\epsilon \wedge L < C] + \Pr[G_\epsilon \wedge B] \\ \le e^{-\lambda-1} + \eta_{\mathrm{fld}}(\lambda) + n\nu_{\mathrm{ni},\mathcal{B}_{\mathrm{ni}}}(\lambda),\end{aligned} \tag{8}$$

where $\nu_{\mathrm{ni},\mathcal{B}_{\mathrm{ni}}}$ is the negligible bound from Lemma 5, and η_{fld} is the probability that $q - 1 < 2m^2(N_0 + \binom{t-f}{2})$ for the selected valid coalition. For $\overline{Q}$ in Equation (38),

$$\eta_{\mathrm{fld}}(\lambda) \le \Pr_{\mathsf{pp}}[q - 1 < \overline{Q}] \le (\overline{Q} + 1)\, p_{\mathsf{guess}}(\lambda, n, t).$$

It is negligible, and the exceptional event is empty for all sufficiently large λ because q has λ bits. Consequently, the construction satisfies publicly verifiable universal traceability. All probabilities are over the full experiment, without conditioning on acceptance or goodness.

Appendix A.9 proves the failure-event bound and Equation (8) for the stated parameter rule, accounting for the full public record, independent box randomness, and point collisions.

Universal traceability and bidirectional sharing together yield the following guarantee for ordinary good reconstruction boxes.

Corollary 1 (Publicly verifiable traceability). Instantiate Trace with the parameter rule of Section 5.6. Then, in the $\mathcal{F}_{\mathrm{CROLE}}$-hybrid model, the construction satisfies Definition 19: for every PPT adversary $\mathcal{A}$, there exists a negligible function ν such that

$$\begin{aligned}\epsilon(\lambda) \cdot \Pr[\mathsf{Exp}^{\mathsf{trace}}_{\mathcal{A},\epsilon,\delta}(1^\lambda, n, t) = 1] \\ \le \delta(\lambda) + p_{\mathsf{guess}}(\lambda, n, t) + \nu(\lambda).\end{aligned}$$

This follows from Lemmas 6 and 1 and Theorem 1; see Appendix A.10.

The construction of Section 5.6, with the stated tracing parameters, satisfies the security guarantees proved above in the $\mathcal{F}_{\mathrm{CROLE}}$-hybrid model. The overall TRAVESS claim also requires the uniform completion hypothesis in Lemma 2. Appendix A.1 establishes this hypothesis for the injective instantiation $\mathsf{OW}(x) = g^x$.

References

1. Adjedj, M., Blokh, C., Couteau, G., Galansky, A., Joux, A., Makriyannis, N.: Two-round 2PC ECDSA at the cost of 1 OLE: Applications to embedded cryptocurrency wallets. In: Bai, S., Persichetti, E. (eds.) PKC 2026: 29th International Conference on Theory and Practice of Public Key Cryptography, Part II. Lecture Notes in Computer Science, vol. 16552, pp. 355–386. Springer, Cham, Switzerland, West Palm Beach, FL, USA (May 25–28, 2026). https://doi.org/10.1007/978-3-032-26734-4_12
2. Atapoor, S., Baghery, K., Cozzo, D., Pedersen, R.: VSS from distributed ZK proofs and applications. In: Guo, J., Steinfeld, R. (eds.) Advances in Cryptology – ASIACRYPT 2023, Part I. Lecture Notes in Computer Science, vol. 14438, pp. 405–440. Springer, Singapore, Singapore, Guangzhou, China (Dec 4–8, 2023). https://doi.org/10.1007/978-981-99-8721-4_13
3. Baghery, K.: Π: A unified framework for computational verifiable secret sharing. In: Jager, T., Pan, J. (eds.) PKC 2025: 28th International Conference on Theory and Practice of Public Key Cryptography, Part IV. Lecture Notes in Computer Science, vol. 15677, pp. 110–142. Springer, Cham, Switzerland, Røros, Norway (May 12–15, 2025). https://doi.org/10.1007/978-3-031-91829-2_4
4. Baghery, K., Ebrahimi, E., Mirzamohammadi, O., Sedaghat, M.: Accountable verifiable secret sharing and applications. Cryptology ePrint Archive, Report 2025/318 (2025), https://eprint.iacr.org/2025/318
5. Beimel, A.: Secret-sharing schemes: A survey. In: Chee, Y.M., Guo, Z., Ling, S., Shao, F., Tang, Y., Wang, H., Xing, C. (eds.) Coding and Cryptology. pp. 11–46. Springer Berlin Heidelberg, Berlin, Heidelberg (2011). https://doi.org/10.1007/978-3-642-20901-7_2
6. Bellare, M., Dai, W., Rogaway, P.: Reimagining secret sharing: Creating a safer and more versatile primitive by adding authenticity, correcting errors, and reducing randomness requirements. Proceedings on Privacy Enhancing Technologies **2020**(4), 461–490 (Oct 2020). https://doi.org/10.2478/popets-2020-0082
7. Boneh, D., Partap, A., Rotem, L.: Traceable secret sharing: Strong security and efficient constructions. In: Reyzin, L., Stebila, D. (eds.) Advances in Cryptology – CRYPTO 2024, Part V. Lecture Notes in Computer Science, vol. 14924, pp. 221–256. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 18–22, 2024). https://doi.org/10.1007/978-3-031-68388-6_9
8. Catalano, D., Fiore, D.: Vector commitments and their applications. In: Kurosawa, K., Hanaoka, G. (eds.) PKC 2013: 16th International Conference on Theory and Practice of Public Key Cryptography. Lecture Notes in Computer Science, vol. 7778, pp. 55–72. Springer Berlin Heidelberg, Germany, Nara, Japan (Feb 26 – Mar 1, 2013). https://doi.org/10.1007/978-3-642-36362-7_5
9. Chin, J.J., Anada, H., Tan, S.Y.: Reset-secure identity-based identification schemes without pairings. In: Au, M.H., Miyaji, A. (eds.) ProvSec 2015: 9th International Conference on Provable Security. Lecture Notes in Computer Science, vol. 9451, pp. 227–246. Springer, Cham, Switzerland, Kanazawa, Japan (Nov 24–26, 2015). https://doi.org/10.1007/978-3-319-26059-4_13
10. Chor, B., Goldwasser, S., Micali, S., Awerbuch, B.: Verifiable secret sharing and achieving simultaneity in the presence of faults (extended abstract). In: 26th Annual Symposium on Foundations of Computer Science. pp. 383–395. IEEE Computer Society Press, Portland, Oregon (Oct 21–23, 1985). https://doi.org/10.1109/SFCS.1985.64
11. Farràs, O., Guiot, M.: Information-theoretic strong traceable secret sharing schemes. Cryptology ePrint Archive, Paper 2026/435 (2026), https://eprint.iacr.org/2026/435

12. Farràs, O., Guiot, M.: Traceable secret sharing schemes for general access structures. In: Daemen, J., Thomé, E. (eds.) Advances in Cryptology – EUROCRYPT 2026, Part III. Lecture Notes in Computer Science, vol. 16543, pp. 3–32. Springer, Cham, Switzerland, Rome, Italy (May 10–14, 2026). https://doi.org/10.1007/978-3-032-25324-8_1
13. Feldman, P.: A practical scheme for non-interactive verifiable secret sharing. In: 28th Annual Symposium on Foundations of Computer Science. pp. 427–437. IEEE Computer Society Press, Los Angeles, CA, USA (Oct 12–14, 1987). https://doi.org/10.1109/SFCS.1987.4
14. Gennaro, R., Jarecki, S., Krawczyk, H., Rabin, T.: Secure distributed key generation for discrete-log based cryptosystems. In: Stern, J. (ed.) Advances in Cryptology – EUROCRYPT'99. Lecture Notes in Computer Science, vol. 1592, pp. 295–310. Springer Berlin Heidelberg, Germany, Prague, Czech Republic (May 2–6, 1999). https://doi.org/10.1007/3-540-48910-X_21
15. Goyal, V., Jain, A., Partap, A.: Traceable secret sharing revisited. In: Daemen, J., Thomé, E. (eds.) Advances in Cryptology – EUROCRYPT 2026, Part III. Lecture Notes in Computer Science, vol. 16543, pp. 33–63. Springer, Cham, Switzerland, Rome, Italy (May 10–14, 2026). https://doi.org/10.1007/978-3-032-25324-8_2
16. Goyal, V., Song, Y., Srinivasan, A.: Traceable secret sharing and applications. In: Malkin, T., Peikert, C. (eds.) Advances in Cryptology – CRYPTO 2021, Part III. Lecture Notes in Computer Science, vol. 12827, pp. 718–747. Springer, Cham, Switzerland, Virtual Event (Aug 16–20, 2021). https://doi.org/10.1007/978-3-030-84252-9_24
17. Hoffmann, C.: Traceable secret sharing based on the chinese remainder theorem. Cryptology ePrint Archive, Report 2024/811 (2024), https://eprint.iacr.org/2024/811
18. Katz, J., Lindell, Y.: Introduction to Modern Cryptography. Chapman and Hall, CRC Press, third edn. (2014)
19. Leslie, M., Dutta, R.: A traceable threshold asmuth–bloom secret sharing scheme. Cryptology ePrint Archive, Report 2025/1561 (2025), https://eprint.iacr.org/2025/1561
20. Luong, D.A., Park, J.H., Lee, C., Kim, H.: TSS-PV: Traceable secret sharing with public verifiability. Cryptology ePrint Archive, Report 2025/2261 (2025), https://eprint.iacr.org/2025/2261
21. MacKenzie, P.D., Yang, K.: On simulation-sound trapdoor commitments. In: Cachin, C., Camenisch, J. (eds.) Advances in Cryptology – EUROCRYPT 2004. Lecture Notes in Computer Science, vol. 3027, pp. 382–400. Springer Berlin Heidelberg, Germany, Interlaken, Switzerland (May 2–6, 2004). https://doi.org/10.1007/978-3-540-24676-3_23
22. Muth, P., Geihs, M., Arul, T., Buchmann, J., Katzenbeisser, S.: ELSA: Efficient long-term secure storage of large datasets (full version). EURASIP Journal on Information Security **2020**(1), 9 (2020). https://doi.org/10.1186/s13635-020-00108-9
23. Pedersen, T.P.: Non-interactive and information-theoretic secure verifiable secret sharing. In: Feigenbaum, J. (ed.) Advances in Cryptology – CRYPTO'91. Lecture Notes in Computer Science, vol. 576, pp. 129–140. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 11–15, 1992). https://doi.org/10.1007/3-540-46766-1_9
24. Schoenmakers, B.: A simple publicly verifiable secret sharing scheme and its application to electronic voting. In: Wiener, M.J. (ed.) Advances in Cryptology – CRYPTO'99. Lecture Notes in Computer Science, vol. 1666, pp. 148–164. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 15–19, 1999). https://doi.org/10.1007/3-540-48405-1_10
25. Stadler, M.: Publicly verifiable secret sharing. In: Maurer, U.M. (ed.) Advances in Cryptology – EUROCRYPT'96. Lecture Notes in Computer Science, vol. 1070, pp. 190–199. Springer Berlin Heidelberg, Germany, Saragossa, Spain (May 12–16, 1996). https://doi.org/10.1007/3-540-68339-9_17

A Security Proofs

We prove the claims of Section 5.7 for the final construction of Section 5.6 in the $\mathcal{F}_{\text{CROLE}}$-hybrid model. Throughout, $T \subseteq [n]$ is the corrupted coalition, $f = |T|$, J is a reconstruction subset, and I is an accusation. We retain the notation

$$P(X) = \sum_{j=0}^{t-1} a_j X^j, \qquad R(X) = \sum_{j=0}^{t-1} b_j X^j, \qquad D(X) = P(X) + \alpha R(X) = \sum_{j=0}^{t-1} d_j X^j.$$

In an honest execution, $h = g^{\alpha}$, $d_j = a_j + \alpha b_j$, and $C_j = g^{d_j}$. The same authenticated value h is used in svk and in every ideal call. A record that substitutes a different value for this h is not a record of the prescribed initialization.

All probabilities retain the online deliveries and aborts of Section 4.1. In particular, we do not condition a security experiment on acceptance. We sometimes first prove a bound for a fixed accepted instance and then average it, assigning zero to the relevant event on executions that do not reach that instance. We account explicitly for the registered-point sampler and fresh randomness on calls to the stateless box. Algebraic failure is included in the negligible terms, as specified in Section 5.6.

Appendix A.9 proves the quantitative bounds in Equation (8) for an explicit parameter rule. Honest correctness for an arbitrary one-way family retains the uniform completion hypothesis stated in Lemma 2. We give its exact role and a concrete sufficient instantiation below.

A.1 Correctness

Fix pp, a secret $s \in \mathbb{F}_q$, and an accepted honest execution. For every delivered share $\mathsf{sh}_i = (x_i, y_i, \rho_i) = (x_i, P(x_i), R(x_i))$,

$$\prod_{j=0}^{t-1} C_j^{x_i^j} = g^{\sum_j a_j x_i^j} h^{\sum_j b_j x_i^j} = g^{P(x_i)} h^{R(x_i)} = g^{y_i} h^{\rho_i}. \tag{9}$$

All format and membership checks also pass. Registration ensures that the n admitted points are distinct and nonzero. For any $J \subseteq [n]$ with $|J| = t$, interpolation of the evaluations $\{(x_i, P(x_i))\}_{i \in J}$ therefore gives $P(0) = s$. This is a deterministic assertion for every accepted honest instance, simultaneously for all such J. Under the completion hypothesis of Lemma 2,

$$\Pr[\mathsf{acc} \wedge \mathsf{Rec}(\mathsf{pp}, \mathsf{svk}, \mathbf{sh}_J) = s] \geq 1 - \eta_{\text{sh}}(\lambda).$$

This proves the sharing and reconstruction clause of Definition 9.

Fix T of size $0 \leq f < t$, put $k = t - f$, and first fix an accepted honest instance. Each of the k additional calls to Share uses the same $\sigma = (P, R)$ and independently samples a uniform point in $\mathbb{F}_q^*$. No point is rejected or registered. The probability that an additional point equals one of the f coalition points, or that two additional points coincide, is at most

$$\frac{fk + \binom{k}{2}}{q-1} = \frac{\binom{t}{2} - \binom{f}{2}}{q-1} \leq \frac{\binom{t}{2}}{q-1}. \tag{10}$$

Outside this event, all t shares verify and interpolate the same P, so reconstruction returns s. Including sharing failures, the unconditional failure probability is at most

$$\eta_{\mathrm{sh}}(\lambda) + \mathbb{E}_{\mathsf{pp}}\left[\frac{\binom{t}{2}}{q-1}\right] \le \eta_{\mathrm{sh}}(\lambda) + 2\binom{t}{2} p_{\mathsf{guess}}(\lambda, n, t), \tag{11}$$

which is negligible. This proves the additional honest-share requirement in Section 4.7, with collisions counted as failures.

For fixed pp, define

$$\kappa(\mathsf{pp}) = \Pr_{X,X' \xleftarrow{\$} \mathbb{F}_q^*}[\mathsf{OW}(X) = \mathsf{OW}(X')], \qquad \epsilon_{\mathrm{coll}}(\lambda) = \mathbb{E}_{\mathsf{pp}}[\kappa(\mathsf{pp})]. \tag{12}$$

If the first n candidates have distinct images, then they also have distinct points, all are admitted, and registration finishes within its $2n$-call cap. Honest ideal initialization succeeds, and correctness of the vector commitment makes every honest key-assembly check succeed. Consequently, uniformly over secrets,

$$\Pr[\neg\mathsf{acc} \mid \mathsf{pp}, s] \le \min\left\{1, \binom{n}{2}\kappa(\mathsf{pp})\right\}. \tag{13}$$

The execution takes polynomial time: there are t ideal calls, at most $2n$ share-generation calls, and polynomial-time commitment operations.

One-wayness makes ϵ_{coll} negligible. Indeed, an inverter given $(\mathsf{pp}, \mathsf{OW}(X))$ can sample independent uniform $X' \in \mathbb{F}_q^*$ and return X'; its success probability is exactly ϵ_{coll}. Thus the assumptions already imply negligible honest failure *averaged over parameter generation*. They do not, by themselves, imply a negligible bound on $\kappa(\mathsf{pp})$ for every individual pp in the support. This is why Lemma 2 states its uniform completion hypothesis explicitly.

A sufficient condition for that hypothesis is a uniform negligible bound on $\kappa(\mathsf{pp})$. In particular, if OW is injective, then $\kappa(\mathsf{pp}) = 1/(q-1)$. Since q has λ bits,

$$\Pr[\neg\mathsf{acc} \mid \mathsf{pp}, s] \le \frac{\binom{n}{2}}{q-1} \le \frac{\binom{n}{2}}{2^{\lambda-1}-1} \tag{14}$$

for $\lambda \ge 2$, a uniform negligible bound for polynomially bounded n. One concrete admissible choice is $\mathsf{OW}(x) = g^x$ on $\mathbb{F}_q^*$, using canonical group-element encodings in the vector-commitment message space $\mathcal{U}$. It is injective, and inversion of a uniform non-identity image solves discrete logarithm on that distribution; the excluded identity has probability $1/q$ in Definition 4. Thus this choice discharges the completion hypothesis under the same discrete-logarithm assumption. For a generic one-way family, the hypothesis of Lemma 2 is not implied by average-case one-wayness. □

A.2 Validity of Tracing Evidence

Fix an accepted honest instance and any box with its fixed embedded data, without any goodness requirement. Every candidate retained by the tracing filter has degree exactly

f and f distinct roots in $\mathbb{F}_q^*$. Its root images match f distinct positions of the honest table $\mathbf{u} = (u_1, \ldots, u_n)$. After ordering these positions as $i_1 < \cdots < i_f$ and their roots accordingly, the tracer has

$$\mathsf{OW}(w_k) = u_{i_k} \qquad (k \in [f]).$$

It outputs $I = (i_1, \ldots, i_f)$ and

$$\pi = (\pi_0, \pi_{\mathrm{tr}}), \qquad \pi_0 = (w_1, \ldots, w_f), \qquad \pi_{\mathrm{tr}} = (\omega_{i_1}, \ldots, \omega_{i_f}),$$

where $\omega_i = \mathsf{VC.Open}(\mathsf{pp}_{\mathrm{vc}}, \mathsf{aux}_{\mathrm{vc}}, i)$. Vector-commitment correctness gives

$$\mathsf{VC.Verify}(\mathsf{pp}_{\mathrm{vc}}, c_{\mathrm{tvk}}, i_k, \mathsf{OW}(w_k), \omega_{i_k}) = 1 \quad \text{for all } k \in [f].$$

All remaining format, order, range, and distinctness checks follow from the filtering and output rules. Thus every non-$\perp$ output verifies, with probability one. This remains true with the algebraic implementation specified in Section 5.6, which returns $\perp$ on failure and checks candidate roots before use. The conclusion establishes the evidence-validity clause of Lemma 2 and Definition 9. It does not assert that the accused set is the true coalition; that assertion is proved separately. □

A.3 Public Share Verifiability

The algorithm $\mathsf{VerifyShare}$ parses svk and the candidate share, performs the specified group and field membership checks, and tests Equation (1) using only public inputs. These operations are deterministic and polynomial time, and require no tracing key, identity-table entry, or interaction. Equation (9) proves acceptance of every delivered honest share. Together with Appendix A.1, this proves Definition 11 and completes Lemma 2. □

A.4 Reconstruction Consistency

Proof of Lemma 3. For $b \in \{0, 1\}$, write an admissible tuple as

$$\mathbf{q}_b = \big((x_{b,k}, y_{b,k}, \rho_{b,k})\big)_{k=1}^{t}.$$

Define its interpolation coefficients and reconstructed values by

$$\lambda_{b,k} = \prod_{\substack{\ell \in [t] \\ \ell \neq k}} \frac{-x_{b,\ell}}{x_{b,k} - x_{b,\ell}}, \qquad s_b^* = \sum_{k=1}^{t} \lambda_{b,k} y_{b,k}, \qquad r_b^* = \sum_{k=1}^{t} \lambda_{b,k} \rho_{b,k}.$$

For each $0 \le j < t$, interpolation of X^j at zero gives

$$\sum_{k=1}^{t} \lambda_{b,k} x_{b,k}^{j} = \begin{cases} 1, & j = 0, \\ 0, & 1 \le j < t. \end{cases}$$

Exponentiating and multiplying the verified share equations yields

$$g^{s_b^*} h^{r_b^*} = \prod_{j=0}^{t-1} C_j^{\sum_k \lambda_{b,k} x_{b,k}^j} = C_0. \tag{15}$$

If $s_0^* \neq s_1^*$, then $r_1^* - r_0^* \neq 0$, because otherwise Equation (15) and the order of g would imply $s_0^* = s_1^*$. Therefore the two admissible tuples reveal

$$\log_g h = (s_0^* - s_1^*)(r_1^* - r_0^*)^{-1}. \tag{16}$$

We now specify a reduction that includes the dealer's complete sharing-stage view. Given a discrete-logarithm challenge $(\mathbb{G}, q, g, h)$ with $h \neq 1$, the reduction samples the remaining public parameters with their prescribed conditional distribution, runs the malicious dealer, and uses the challenge h as the honest tracer's first authenticated message. It simulates the sequential ideal calls from the dealer's perspective. The dealer, as sender, receives no functionality output; the receiver's input and outputs are private. Consequently, this simulation does not need to compute $\alpha = \log_g h$ or any d_j. Adversarial abort choices and public message scheduling are processed exactly as prescribed.

There are no acceptance messages to simulate: the tracer's initialization checks and key-assembly checks are private, and the dealer continues according to its own local state. Subsequent share generation and commitment assembly are dealer-local; the reduction records their outgoing messages. Thus the entire view on which the dealer's final tuples depend is distributed exactly as in Definition 12, even when the actual tracer would reject. The reduction does not need to decide that rejection.

On the two output tuples, the reduction checks public record compatibility and admissibility using the same authenticated h. Whenever they reconstruct distinct secrets it outputs Equation (16). Every winning execution of the consistency experiment is such an extraction event, including executions in which the tracer privately rejects. On an identity challenge the reduction instead outputs 0 immediately. If p_{cons} denotes the consistency success probability, the reduction's discrete-logarithm success probability is at least

$$\mathbb{E}_{\mathbb{G},q,g}\left[\left(1 - \frac{1}{q}\right) \Pr[\mathsf{Exp}^{\mathsf{cons}}_{\mathcal{A},\mathsf{TRAVESS}} = 1 \mid \mathbb{G}, q, g]\right] \geq \tfrac{1}{2} p_{\text{cons}}.$$

Here the conditional probability still averages the remaining public parameters and the sharing execution, and $q > n \geq 2$. It follows that $p_{\text{cons}} \leq 2\,\mathrm{Adv}_{\mathrm{DL}}$ for the constructed reduction. Definition 4 makes this negligible. The proof uses neither knowledge of the dealer's coefficients nor any post-sharing tracing information. □

A.5 Tracer Secrecy

Proof of the equality in Lemma 4. We give a measure-preserving bijection between the two honest-dealer executions. Fix pp, the adversary's random tape, its valid choices s_0, s_1, and its static coalition T, with $|T| < t$. If the adversary supplies an invalid or missing initial h, the dealer rejects before sending secret-dependent information and the two views coincide. Otherwise there is a unique $\alpha \in \mathbb{F}_q^*$ such that $h = g^\alpha$. We use this

value only to define the bijection; neither the adversary nor a simulator has to compute it.

Presample a length-$2n$ stream of candidate points and the vector-commitment coins. Couple these choices identically in the two experiments. Apply the registration rule to the stream even if the interactive execution would abort earlier. Let $T' \subseteq T$ be the coalition positions that would receive admitted points before the cap. This presampling does not change the distribution, since candidate points are independent of the sharing polynomials. Every coalition share actually delivered belongs to one of these positions. Define

$$\Delta P(X) = (s_1 - s_0) \prod_{i \in T'} \frac{X - x_i}{-x_i}, \qquad P_1 = P_0 + \Delta P, \quad R_1 = R_0 - \alpha^{-1} \Delta P. \tag{17}$$

An empty product is one. We have $\deg \Delta P \le |T'| < t$, $\Delta P(0) = s_1 - s_0$, and $\Delta P(x_i) = 0$ for every $i \in T'$. Hence the map sends the sampling space

$$\{P_0 : \deg P_0 < t,\ P_0(0) = s_0\} \times \{R_0 : \deg R_0 < t\}$$

bijectively onto its counterpart with constant term s_1. Its inverse subtracts ΔP from P_1 and adds $\alpha^{-1}\Delta P$ to R_1. Thus it preserves the uniform distribution on the prescribed coefficient spaces, also after averaging over the presampled points.

The map preserves $D = P + \alpha R$, and therefore every C_j. It also preserves both evaluations $P(x_i), R(x_i)$ for all exposed coalition shares. The candidate stream, admission decisions, table $\mathbf{u}$, commitment c_{tvk}, and complete opening state $\mathsf{aux}_{\mathrm{vc}}$ remain identical. In particular, secrecy does not rely on hiding $\mathbf{u}$ or $\mathsf{aux}_{\mathrm{vc}}$ from the tracer.

For completeness, consider an arbitrary ideal-call interaction. Inductively suppose the adversary has received identical prefixes in the two executions. It then makes the same next choice, including the same abort decision and receiver input α'. Missing or malformed inputs and functionality aborts give the same transition or terminated prefix in both executions. If $g^{\alpha'} \neq h$, the ideal output is $\perp$ in both executions. If $g^{\alpha'} = h$, prime order and canonical field inputs imply $\alpha' = \alpha$, and the output in call j is the same coefficient d_j of the preserved polynomial D. This proves equality of the interaction prefixes by induction. All subsequent public and prescribed private deliveries are also equal, as are the coalition shares and observable aborts. Unexposed honest shares may differ; their contents are not part of the adversary's view and all still pass verification.

The bijection therefore establishes equality in distribution of the entire malicious tracer's view, without conditioning on any acceptance event. Averaging over pp and the adversary's randomness proves

$$\Pr[\mathsf{Exp}^{\mathsf{tsec}}_{\mathcal{A}}(0) = 1] = \Pr[\mathsf{Exp}^{\mathsf{tsec}}_{\mathcal{A}}(1) = 1].$$

This is perfect tracer secrecy in the ideal model. □

A.6 Secrecy

Given an adversary for Definition 10, construct an adversary for Definition 13 that runs an honest tracer internally. It forwards the secrecy adversary's initial secrets and coalition, and then relays exactly the public messages and coalition shares prescribed by its

experiment, in their original online order. It retains tracer-private material internally. On completion or abort it outputs the secrecy adversary's bit. The resulting subadversary view is exactly its view in Definition 10. The equality proved in Appendix A.5 therefore gives equality of its two output distributions. This proves the secrecy part of Lemma 4, with perfect secrecy in the $\mathcal{F}_{\mathrm{CROLE}}$-hybrid model. □

A.7 Non-Imputability

Proof of Lemma 5. We first remove the dependency introduced by registration, with an explicit bound that includes abort prefixes. Let Game_0 be Definition 14. In Game_1, the dealer uses exactly one independent candidate per participant and aborts at the first repeated image, rather than resampling. It otherwise keeps the same initialization, delivery order, commitment assembly, and public verifier.

Couple the two games on the same independent candidate stream and other coins. Their entire observable histories are identical unless two of the first n candidate images coincide. Distinct images imply distinct points, so there is no separate point-collision event. Consequently, for any event in the adversary's full view, its probabilities differ by at most

$$\eta_{\mathrm{reg}}(\lambda) := \binom{n}{2}\epsilon_{\mathrm{coll}}(\lambda). \tag{18}$$

The same coupling applies if the tracer aborts during initialization or after only a prefix of deliveries. By Appendix A.1, η_{reg} is negligible under one-wayness.

Work in Game_1. In a winning output, let w and ω^* be the root and opening at the target position $i^\star$. Write $u_{i^\star}$ for its honestly committed image. There are two exhaustive cases.

Case 1: $\mathsf{OW}(w) \neq u_{i^\star}$. A position-binding adversary embeds its honestly generated $\mathsf{pp}_{\mathrm{vc}}$ in the public parameters, independently generates the other parameters, and simulates the honest dealer and the malicious tracer's ideal interaction. It knows all the points in this simulation and produces the genuine commitment and opening state. The adversary's winning output gives an accepted opening to $\mathsf{OW}(w)$ at $i^\star$, whereas commitment correctness gives an accepted opening

$$\omega = \mathsf{VC.Open}(\mathsf{pp}_{\mathrm{vc}}, \mathsf{aux}_{\mathrm{vc}}, i^\star)$$

to the different message $u_{i^\star}$ at the same position of the same c_{tvk}. These two openings break position binding. The simulation is efficient even when the malicious tracer chooses h without revealing its logarithm: the simulator knows (a_j, b_j) and can evaluate each ideal response on the supplied receiver input. Let ν_{bind} be this reduction's position-binding advantage. The probability of Case 1 is at most ν_{bind}.

Case 2: $\mathsf{OW}(w) = u_{i^\star}$. Consider an inverter given

$$(1^\lambda, \mathsf{pp}, u^*), \qquad u^* = \mathsf{OW}(X), \quad X \xleftarrow{\$} \mathbb{F}_q^*.$$

It gives only pp to the non-imputability adversary and obtains the latter's target $i^\star$ and chosen secret s. These choices are made before initialization and do not reveal the challenge image to the adversary. The inverter samples P, R exactly as the honest dealer would for s. On the adversary's initial h, it performs the dealer's membership test, computes the genuine coefficient commitments, and simulates the sequential ideal calls. For receiver input α', it returns $a_j + b_j\alpha'$ precisely when $g^{\alpha'} = h$, and returns $\perp$ otherwise. All abort choices and ensuing deliveries follow the original rules.

For every $i \neq i^\star$, it samples an independent uniform point $x_i \in \mathbb{F}_q^*$, computes $u_i = \mathsf{OW}(x_i)$, and can compute the entire non-target share. At the target position it sets $u_{i^\star} = u^*$. It detects image collisions among these known images and aborts at the first one, exactly as Game_1 prescribes. The target's conceptual share is $(X, P(X), R(X))$. Its point is independent and uniform, and this share always verifies. Its private delivery is never visible to the adversary and elicits no prescribed message to it, so the inverter does not need to compute its contents. This remains true when all $n-1$ other shares reveal P, R.

If registration completes, the inverter computes $\mathsf{VC.Commit}(\mathsf{pp}_{\mathrm{vc}}, \mathbf{u})$ on the fully known image vector. It publishes the genuine c_{tvk} and sends the genuine $(\mathbf{u}, \mathsf{aux}_{\mathrm{vc}})$ to the malicious tracer. Thus the simulation includes the real commitment, its entire opening state, all non-target shares, and all online prefixes; it does not replace any auxiliary information by a hiding hybrid. Its view is distributed exactly as Game_1, with the unknown challenge preimage occupying the target slot. If Case 2 occurs, the inverter outputs w, which is a valid preimage of u^*. Let ν_{ow} be its inversion advantage. This bounds the probability of Case 2.

Combining the game switch and both cases gives the explicit bound

$$\Pr[\mathsf{Exp}^{\mathsf{ni}}_{\mathcal{A}}(1^\lambda, n, t) = 1] \leq \eta_{\mathrm{reg}}(\lambda) + \nu_{\mathrm{bind}}(\lambda) + \nu_{\mathrm{ow}}(\lambda) =: \nu_{\mathrm{ni},\mathcal{A}}(\lambda). \tag{19}$$

Every term is negligible. The reduction uses the experiment's absence of prior target-share disclosure, prior target evidence, or an oracle supplied with the target share. It does not depend on the malicious tracer's acceptance decision, on secrecy of the chosen secret, or on the number of exposed non-target shares. This proves Lemma 5. □

A.8 Bidirectional Sharing

Proof of Lemma 6. Fix pp and a valid output (T, state) of the selector, including its random tape, and put $k = t - f$. Presample the registration candidate stream, the honest $\alpha \in \mathbb{F}_q^*$, and the commitment coins. These choices are independent of the uniformly sampled secret z and the polynomial coefficients. Invalid selector outputs cause both samplers to return $\perp$. If the presampled stream does not complete registration, both also return $\perp$. All other honest checks succeed in the ideal model.

On a stream that completes, the registered points and the image vector are fixed. Since z is uniform, P is uniform over all degree-at-most-$t-1$ polynomials, with no restriction on its constant coefficient, and R is an independent uniform polynomial in the same space. For fixed nonzero α, the transformation

$$(P, R) \longmapsto (P, D = P + \alpha R)$$

is bijective. Hence P and D are independent uniform polynomials. Condition on D and on the f coalition evaluations $\{P(x_i)\}_{i\in T}$. Then P is uniform on the affine space

$$\mathcal{P}_T = \{\widetilde{P} \in \mathbb{F}_q[X] : \deg \widetilde{P} < t,\ \widetilde{P}(x_i) = P(x_i)\ (i \in T)\}, \tag{20}$$

which has dimension k and cardinality q^k. The coalition's R-evaluations supply no further constraint, because $R(x_i) = \alpha^{-1}(D(x_i) - P(x_i))$.

These conditioned quantities determine the entire retained view V and tracing key

$$\mathsf{tk} = (\alpha, \mathbf{d}, \mathbf{u}, \mathsf{aux}_{\mathrm{vc}}).$$

Indeed, the public coefficient commitments depend on D, all ideal receiver outputs are its coefficients, coalition shares have been fixed, and the identity commitment and its state depend only on the registered images and commitment coins. In particular, the dealer state and the backward sampler's original secret are not retained in the output.

Couple the two samplers to use identical additional points $x'_1, \ldots, x'_k$, independently uniform in $\mathbb{F}_q^*$. Let $\mathsf{GoodPts}$ mean that these points are distinct and avoid the coalition points. On this event, the map

$$\widetilde{P} \in \mathcal{P}_T \longmapsto (\widetilde{P}(x'_1), \ldots, \widetilde{P}(x'_k)) \in \mathbb{F}_q^k \tag{21}$$

is bijective. To see injectivity, a difference of two such polynomials has $f + k = t$ distinct roots and degree less than t, and is therefore zero; equal finite cardinalities then give surjectivity.

Accordingly, in the forward sampler the additional y-coordinates are independent uniform field elements, conditional on all the retained data and these points. Conversely, the backward sampler's independent uniform y-coordinates determine a unique $P^* \in \mathcal{P}_T$ with that same uniform distribution. In either case, all additional verification coordinates are

$$\rho'_j = \alpha^{-1}(D(x'_j) - P^*(x'_j)),$$

and the displayed secret is $P^*(0)$. In the forward experiment this is the original z; in the backward experiment it is the nonfailure reconstruction s_Q. Thus the joint output distributions are identical on $\mathsf{GoodPts}$ while retaining (V, tk) exactly.

Both samplers use the same extra-point distribution. Equation (10) gives

$$\Pr[\neg\mathsf{GoodPts}] \le \frac{f(t-f) + \binom{t-f}{2}}{q-1} \le \frac{\binom{t}{2}}{q-1}.$$

A coupling can therefore make the outputs equal except with this probability. There is no factor of two: we couple the same bad-point event in the two samplers. Averaging all fixed choices proves

$$\Delta(\mathsf{Fwd}_{\mathcal{A}_0}(1^\lambda, n, t), \mathsf{Bwd}_{\mathcal{A}_0}(1^\lambda, n, t)) \le \mathbb{E}_{\mathsf{pp}}\left[\frac{\binom{t}{2}}{q-1}\right] =: \epsilon_{\mathsf{bd}}(\lambda). \tag{22}$$

This is negligible because $1/(q-1) \le 2/q$ and t is polynomially bounded. Notice that this proof needs no assumption that sharing completes with high probability: both samplers already give the same $\perp$ output on a registration failure. This proves Lemma 6. □

A.9 Publicly Verifiable Universal Traceability

We specify the tracing parameters on the algorithm's actual inputs, then prove the quantitative bounds of Theorem 1. The analysis follows the paired-query and polynomial-recovery approach of [7, Sec. 3.2], with the verification coordinate and public record retained throughout. The false-candidate bound uses Lemma 5 rather than an independence claim about unopened evaluation points.

A.9.1 An explicit parameter rule. The interface supplies 1^m, where $m = \lceil 1/\delta \rceil$; it does not supply the exact real number δ. Infer f from the public box arity $t - f$, and return $\perp$ on invalid arities. Define the integers

$$\begin{aligned} N_0 &= 64fm^4 + 16(\lambda + 1)m^2, \\ C_{\text{main}} &= 16fm^2 + 4(\lambda + 1) = \frac{N_0}{4m^2}. \end{aligned} \tag{23}$$

Here λ can be recovered as the bit length of q. For a rule defined at every finite input size, use

$$\begin{aligned} N &= \min\{N_0, q - 1\}, \\ C &= \max\left\{ \left\lfloor \frac{N}{4m^2} \right\rfloor, \left\lfloor \sqrt{fN} \right\rfloor + 1 \right\}. \end{aligned} \tag{24}$$

The square-root floor is an integer square root. Because $q - 1 \geq n \geq t > f$ and $N_0 > f$, we have $N > f$, whence $\lfloor \sqrt{fN} \rfloor + 1 \leq N$. Therefore

$$\sqrt{fN} < C \leq N \leq q - 1 \tag{25}$$

on every admissible input. There are at most $2N_0$ oracle calls, giving a query budget $O(f\delta^{-4} + \lambda\delta^{-2})$.

For the main analysis put $k = t - f$ and consider

$$q - 1 \;\geq\; 2m^2 \left(N_0 + \binom{k}{2} \right). \tag{26}$$

This implies $N = N_0$. Moreover,

$$\frac{C_{\text{main}}^2}{fN_0} = \frac{N_0}{16fm^4} = 4 + \frac{\lambda + 1}{fm^2} > 4.$$

Thus $C_{\text{main}} > \sqrt{fN_0}$ and the all-input rule gives $C = C_{\text{main}}$.

A.9.2 Valid shares and the interpolation identity. On an accepted honest instance, for any $x \in \mathbb{F}_q^*$ and $y \in \mathbb{F}_q$ the tracer forms $\rho = \alpha^{-1}(D(x) - y)$. Then

$$g^y h^\rho = g^{y + \alpha\rho} = g^{D(x)} = \prod_{j=0}^{t-1} C_j^{x^j}. \tag{27}$$

The perturbed share $(x, y+\Delta, \rho-\Delta/\alpha)$ satisfies the same equation. Both inputs therefore pass public verification, irrespective of whether $y = P(x)$.

Write $\xi_1, \ldots, \xi_f$ for the coalition points. Define

$$H_T(X) = \prod_{r=1}^{f} \frac{\xi_r - X}{\xi_r}. \tag{28}$$

This is the polynomial of Equation (4); its degree is f and its roots are the f distinct nonzero coalition points.

Consider a round whose supplied points avoid the coalition points and whose two responses are correct reconstructions. The supplied points are already pairwise distinct by the tracing sampler. The special point x' has nonzero Lagrange coefficient

$$a(x') = \left(\prod_{r=1}^{f} \frac{\xi_r}{\xi_r - x'}\right)\left(\prod_{i=f+1}^{t-1} \frac{x_i}{x_i - x'}\right).$$

Hence $v' - v = \Delta a(x') \neq 0$. The round is never erased, and

$$z = \frac{\Delta}{v' - v} \prod_{i=f+1}^{t-1} \frac{x_i}{x_i - x'} = H_T(x'). \tag{29}$$

For $k = 1$, the products over ordinary supplied points are empty and equal one, so the same identity holds.

A.9.3 A lower bound for one paired round. Fix a complete accepted honest instance and an initialized ϵ-universally-good box $\mathcal{R}$ of arity $k = t - f$. The dealer state, coalition shares, public record, tracing key, and box state are fixed; only new query randomness and fresh independent box coins remain random. In the universal experiment $\delta = \epsilon$, so $m = \lceil 1/\epsilon \rceil$ and $1/m \leq \epsilon$.

First consider a reference distribution. Choose all k input points independently and uniformly in $\mathbb{F}_q^*$, with the last coordinate special in both ordered tuples. Choose the $k-1$ ordinary y-coordinates and two special values y_0, y_1 independently and uniformly in $\mathbb{F}_q$. Complete each share by Equation (27). Each marginal query is exactly $\mathcal{D}^k_{\mathsf{pub},\mathsf{tk}}$, even for an order-sensitive box.

Let W collect the k points and the $k-1$ ordinary values. Write $Q(W, y)$ for the resulting ordered tuple, and set

$$\begin{aligned} a_W(y) &= \Pr\left[\mathcal{R}(Q(W,y)) = \mathsf{Rec}(\mathsf{pp}, \mathsf{svk}, \mathbf{sh}_T \| Q(W,y)) \in \mathbb{F}_q\right], \\ p(W) &= \mathbb{E}_y[a_W(y)]. \end{aligned} \tag{30}$$

The probability in $a_W(y)$ is over one fresh box call. In particular, $a_W(y) = 0$ whenever points collide with each other or with the coalition; reconstruction failure never counts as success. Universal goodness gives $\mathbb{E}_W[p(W)] \geq \epsilon$. Conditioned on W, the special values and the two calls' coins are independent. Jensen's inequality implies

$$\Pr[\text{both reference calls succeed}] = \mathbb{E}_W[p(W)^2] \geq (\mathbb{E}_W[p(W)])^2 \geq \epsilon^2. \tag{31}$$

This independence statement concerns the reference distribution; we next account for the actual nonzero perturbation.

Before round ℓ, fix any preceding tracing history and let $S_{\ell-1}$ contain its $\ell-1$ earlier special points. The actual point distribution is the independent reference distribution conditioned on $x' \notin S_{\ell-1}$ and all k supplied points being distinct. Indeed, the conditional special point is uniform outside $S_{\ell-1}$, and the other ordered points are uniform without replacement outside $\{x'\}$, as prescribed in Section 5.2. A distribution has statistical distance $\Pr[\neg E]$ from its conditional distribution on an event E of positive probability. Here the conditioning costs at most

$$\frac{\ell - 1 + \binom{k}{2}}{q-1}. \tag{32}$$

This is uniform over the actual earlier special points and hence over the full history.

Choosing uniform y_0 and $\Delta \xleftarrow{\$} \mathbb{F}_q^*$ and setting $y_1 = y_0 + \Delta$ is the independent uniform pair conditioned on $y_0 \neq y_1$. Its statistical distance from the reference value pair is exactly $1/q$. Share completion and fresh box calls cannot increase statistical distance. Thus, for every preceding history $\mathcal{F}_{\ell-1}$,

$$\Pr[Z_\ell = 1 \mid \mathcal{F}_{\ell-1}] \geq \epsilon^2 - \frac{1}{q} - \frac{\ell - 1 + \binom{k}{2}}{q-1}, \tag{33}$$

where Z_ℓ indicates a round counted by L in Theorem 1. Coalition-point collisions require no further subtraction: they already count as failures in Equation (30).

Under Equation (26), $N = N_0$ and the loss in Equation (33) is at most $(N_0 + \binom{k}{2})/(q-1) \leq 1/(2m^2)$. Since $\epsilon^2 \geq 1/m^2$, every round has conditional success probability at least

$$p := \frac{1}{2m^2}. \tag{34}$$

A.9.4 Concentration with history-dependent sampling. The indicators $Z_1, \ldots, Z_N$ need not be independent. For every $u > 0$, however,

$$\mathbb{E}[e^{-uZ_\ell} \mid \mathcal{F}_{\ell-1}] = 1 - (1 - e^{-u}) \Pr[Z_\ell = 1 \mid \mathcal{F}_{\ell-1}] \leq 1 - p + pe^{-u}.$$

Iterating conditional expectation, with $L = \sum_\ell Z_\ell$,

$$\mathbb{E}[e^{-uL}] \leq (1 - p + pe^{-u})^N \leq \exp[-pN(1 - e^{-u})].$$

Taking $u = \log 2$ and applying Markov's inequality at $pN/2$,

$$\Pr[L \leq pN/2] \leq \exp[-(1 - \log 2)pN/2] \leq \exp[-pN/8].$$

Since $C_{\text{main}} = pN_0/2$, each fixed instance under consideration has

$$\Pr[L < C_{\text{main}}] \leq \exp\left[-\frac{N_0}{16m^2}\right] = e^{-4fm^2-\lambda-1} \leq e^{-\lambda-1}. \tag{35}$$

The bound is uniform over accepted universally-good instances satisfying Equation (26). It can therefore be averaged without dividing by the probability of acceptance or goodness. Statelessness is essential here: the embedded data stay fixed and each call uses fresh independent coins, so previous responses do not otherwise change future box behavior.

A.9.5 False candidates and non-imputability. For analysis, extend tracing to every accepted sharing with a well-formed box, without testing goodness. This extension agrees with the original experiment on G_ϵ. Define B on this extended execution as the event that some returned candidate passes the image-matching filter but identifies a set different from T. Extending execution outside the goodness event does not change any event intersected with G_ϵ. The reduction uses the same polynomial-time tracing implementation and failure convention specified in Section 5.6.

Construct a PPT adversary $\mathcal{B}_{\mathrm{ni}}$ for Definition 14 as follows. On pp, start the tracing adversary $\mathcal{A}$ and obtain its static coalition T; invalid choices are unsuccessful. Sample a target uniformly from $[n] \setminus T$ and a secret uniformly from $\mathbb{F}_q$, and submit both to the non-imputability challenger before initialization. Act as an honest tracer, sampling α and making the prescribed ideal calls and local checks. Forward to $\mathcal{A}$ only the public messages and shares in T, in the original online order. All these shares are available since the target is outside T. Ignore the additional non-target shares supplied by the challenger.

After an accepted sharing, the reduction has the genuine $\mathsf{tk} = (\alpha, \mathbf{d}, \mathbf{u}, \mathsf{aux}_{\mathrm{vc}})$. It obtains $\mathcal{A}$'s box and performs exactly the prescribed tracing queries, decoding, and filtering. It examines *all* retained candidates, rather than only the one selected by the fixed ordering. If some candidate accuses the target, it selects the first such candidate under that ordering and outputs its set I_H, its ordered roots, and the genuine VC openings. Appendix A.2 ensures that this is accepted. Otherwise the reduction outputs an unsuccessful accusation.

Couple all target choices to the same full honest sharing, execution of $\mathcal{A}$, and tracing execution. Before the last accusation-selection step, the simulated tracing execution is independent of the selected target: the latter lies outside T, the public record and tracing key are genuine, and the extra non-target shares are unused. This assertion concerns the simulated tracing execution, not the reduction's unused NI-only share disclosures, which can depend on the target. For each retained candidate H, let I_H be its index set and put

$$U = \bigcup_{H \text{ retained}} (I_H \setminus T),$$

with $U = \varnothing$ on invalid or aborted executions. Each retained set has size $f = |T|$, so B occurs exactly when $U \neq \varnothing$. Conditional on the coupled execution, the uniform target lies in U with probability $|U|/(n-f)$. Consequently,

$$\Pr[\mathsf{Exp}^{\mathsf{ni}}_{\mathcal{B}_{\mathrm{ni}}}(1^\lambda, n, t) = 1] = \mathbb{E}\left[\frac{|U|}{n-f}\right] \geq \frac{\Pr[B]}{n}, \tag{36}$$

$$\Pr[G_\epsilon \wedge B] \leq \Pr[B] \leq n\nu_{\mathrm{ni},\mathcal{B}_{\mathrm{ni}}}(\lambda). \tag{37}$$

The expectation is zero on invalid coalitions. Lemma 5 makes the last expression negligible.

This reduction never tests universal goodness. No target-share oracle is supplied: the box is locally constructed from $\mathcal{A}$'s permitted public view, private working memory, and $\mathbf{sh}_T$. Any information it can compute from that view is already permitted by the non-imputability experiment. The reduction preserves the real public commitment and all tracing auxiliary information. It does not assume that hidden points are independent of their images, the commitment, or the tracer's view. Statistical selective-opening

hiding need not be invoked for this bound. Distinct polynomial scalings or preimages that identify the same set T are harmless and are not counted by B. There is no circularity: Lemma 5 was established using one-wayness and position binding without either traceability property.

A.9.6 Accounting for small fields unconditionally. In the universal experiment use $m = \lceil 1/\epsilon \rceil$ and define the public upper bounds

$$\begin{aligned} \overline{N} &= 64(t-1)m^4 + 16(\lambda+1)m^2, \\ \overline{Q} &= 2m^2\left(\overline{N} + \binom{t}{2}\right). \end{aligned} \tag{38}$$

Let η_{fld} be the probability that Equation (26) fails for the chosen valid coalition. Since $N_0 \le \overline{N}$ and $\binom{k}{2} \le \binom{t}{2}$, this event implies $q < \overline{Q} + 1$, and hence

$$\eta_{\mathrm{fld}} \le \Pr[q < \overline{Q} + 1] \le (\overline{Q} + 1)\mathbb{E}_{\mathsf{pp}}[1/q] = (\overline{Q} + 1)p_{\mathsf{guess}}. \tag{39}$$

This is negligible: $\overline{Q}$ is polynomially bounded. Indeed the λ-bit requirement on q makes the event empty for all sufficiently large λ, for each admissible polynomial parameter family. The capped rule still specifies the algorithm at every finite input size.

Averaging Equation (35) and allowing the exceptional field-size event yields

$$\Pr[G_\epsilon \wedge \{L < C\}] \le e^{-\lambda-1} + \eta_{\mathrm{fld}} \le e^{-\lambda-1} + (\overline{Q} + 1)p_{\mathsf{guess}}. \tag{40}$$

Sharing failures require no added error term here: the tracing experiment returns zero on them, and they are outside G_ϵ.

A.9.7 Completing the universal-traceability proof. Let M count the recorded pairs, with distinct special points. If $L \ge C$, then $C \le M \le N$ and at least C pairs agree with H_T: a round counted by L is never erased. The Reed–Solomon list-decoding guarantee of [7, Theorem 1] applies because

$$C > \sqrt{fN} \ge \sqrt{fM}.$$

Whenever decoding is invoked, $C \le M \le N$ and integrality gives

$$\begin{aligned} C^2 - fM \ge C^2 - fN \ge 1, \qquad N \le N_0, \\ C - \sqrt{fM} = \frac{C^2 - fM}{C + \sqrt{fM}} \ge \frac{1}{2N}. \end{aligned}$$

Thus the gap above the decoding threshold has a polynomially bounded reciprocal, including when $N = q - 1 < N_0$. Together with the cited Reed–Solomon theorem, this gives a uniform polynomial running-time bound on every admissible input. On successful completion, the decoder returns all degree-at-most-f polynomials with at least C agreements, including H_T. The complete list has polynomial length, so factoring and

checking all candidates entails polynomially many subroutine invocations. The implementation of Section 5.6 has negligible algebraic failure probability, uniformly over the recorded inputs.

Suppose G_ϵ and $L \geq C$ hold and no algebraic failure occurs. The list contains H_T, which passes the filter because its roots are the true coalition points. On $\neg B$, every retained candidate identifies T. Thus the fixed selection ordering returns T, with accepted evidence by Appendix A.2. Therefore failure of the tracing algorithm is contained in the union of insufficient correct rounds, a false retained candidate, and algebraic failure:

$$\begin{aligned}\Pr[\mathsf{Exp}^{\mathsf{utrace}}_{\mathcal{A},\epsilon} = 1] \leq\ & \Pr[G_\epsilon \wedge \{L < C\}] \\ & + \Pr[G_\epsilon \wedge B] + \mathsf{negl}(\lambda).\end{aligned} \tag{41}$$

The quantitative bound is

$$\begin{aligned}&\Pr[G_\epsilon \wedge \{L < C\}] + \Pr[G_\epsilon \wedge B] \\ &\quad \leq e^{-\lambda-1} + (\overline{Q}+1)p_{\mathsf{guess}} + n\nu_{\mathrm{ni},\mathcal{B}_{\mathrm{ni}}}(\lambda).\end{aligned} \tag{42}$$

Including the negligible algebraic failure probability, we have

$$\begin{aligned}\Pr[\mathsf{Exp}^{\mathsf{utrace}}_{\mathcal{A},\epsilon}(1^\lambda, n, t) = 1] \leq\ & e^{-\lambda-1} + (\overline{Q}+1)p_{\mathsf{guess}} \\ & + n\nu_{\mathrm{ni},\mathcal{B}_{\mathrm{ni}}}(\lambda) + \mathsf{negl}(\lambda).\end{aligned} \tag{43}$$

Absorb the algebraic failure probability together with $e^{-\lambda-1}$, $\overline{Q}p_{\mathsf{guess}}$, and $n\nu_{\mathrm{ni},\mathcal{B}_{\mathrm{ni}}}(\lambda)$ into a negligible function $\nu_{\mathrm{tr}}(\lambda)$, with dependence on the adversary and polynomial parameter family allowed by Section 4.1. Consequently,

$$\Pr[\mathsf{Exp}^{\mathsf{utrace}}_{\mathcal{A},\epsilon}(1^\lambda, n, t) = 1] \leq p_{\mathsf{guess}}(\lambda, n, t) + \nu_{\mathrm{tr}}(\lambda). \tag{44}$$

This proves Definition 18 and Equation (8) using the concrete tracing parameters, without an additional assumption on the box's construction or its pointwise success behavior. □

A.10 Publicly Verifiable Traceability

We give the distribution-switching argument explicitly to justify the use of Lemma 1 with the retained tracing key. Fix ϵ, δ with polynomially bounded reciprocals, and use the tracing implementation above with input $1^{\lceil 1/\delta \rceil}$.

Use $\mathcal{A}$'s coalition-selection phase as the selector of Definition 17. From either sampler's output (V, tk, z, Q), continue $\mathcal{A}$ to obtain its box with its fixed embedded data. Run Trace at parameter δ, and let F_δ mean failure to return precisely T with accepted evidence. Independently sample fresh coins for one further box call on Q; let E mean that this call returns the displayed z and $z \neq \perp$. On a $\perp$ sampler output or malformed box set both events to false. These operations are a common randomized postprocessing kernel on the two sampler outputs; they use no dealer state or backward original secret.

For the forward distribution, first fix a complete honest instance, including its dealer state, and the box with its fixed embedded data. If the box is ϵ-good, the extra test succeeds with probability at least ϵ. The test sample and coins are independent of tracing, and the stateless box does not retain memory of tracing queries. Hence

$$\epsilon \Pr[\mathsf{Exp}^{\mathsf{trace}}_{\mathcal{A},\epsilon,\delta}(1^\lambda, n, t) = 1] \le \Pr_{\mathsf{Fwd}}[E \wedge F_\delta]. \tag{45}$$

Conditioning on the dealer state justifies this inequality; it does not disclose that state to the postprocessing kernel. By Lemma 6 and statistical-distance contraction,

$$\Pr_{\mathsf{Fwd}}[E \wedge F_\delta] \le \Pr_{\mathsf{Bwd}}[E \wedge F_\delta] + \epsilon_{\mathsf{bd}}(\lambda). \tag{46}$$

In the backward distribution, condition on $(V, \mathsf{tk}, \mathcal{R})$, including the box and its fixed embedded data, and let p_{u} be the independent test's success probability. This is exactly the universal agreement probability of Definition 16: the coalition shares are in V, and $Q \leftarrow \mathcal{D}^{t-f}_{\mathsf{pub},\mathsf{tk}}$. Let r_δ be the conditional tracing failure probability. The test and tracing are independent given these data, so

$$\begin{aligned} \Pr_{\mathsf{Bwd}}[E \wedge F_\delta] &= \mathbb{E}[p_{\mathsf{u}} r_\delta] \\ &\le \delta + \mathbb{E}[\mathbf{1}_{\{p_{\mathsf{u}} \ge \delta\}} r_\delta] \\ &= \delta + \Pr[\mathsf{Exp}^{\mathsf{utrace}}_{\mathcal{A},\delta}(1^\lambda, n, t) = 1]. \end{aligned} \tag{47}$$

The equality uses the honest-sharing marginal of (V, tk) in the backward sampler and the same tracing parameter δ in its universal experiment. The proof neither conditions the whole experiment on goodness nor needs to test goodness efficiently.

Combining the preceding three inequalities proves Lemma 1 with the explicit error from Lemma 6. Applying Equation (44) at threshold δ gives

$$\begin{aligned} \epsilon(\lambda) \Pr[\mathsf{Exp}^{\mathsf{trace}}_{\mathcal{A},\epsilon,\delta}(1^\lambda, n, t) = 1] \le \delta(\lambda) + p_{\mathsf{guess}}(\lambda, n, t) \\ + \nu_{\mathrm{tr},\delta}(\lambda) + \epsilon_{\mathsf{bd}}(\lambda). \end{aligned} \tag{48}$$

The last two terms are negligible, proving Corollary 1 and Definition 19. The factor ϵ is retained; the conclusion does not bound the unweighted failure probability by δ.

The proofs establish the security claims of Section 5.7 with explicit polynomial tracing parameters. For a generic OW family, correctness retains the uniform completion hypothesis of Lemma 2. Appendix A.1 establishes this hypothesis for $\mathsf{OW}(x) = g^x$. With that instantiation, discrete-logarithm hardness and a vector commitment satisfying the correctness and position-binding clauses of Definition 6 for the chosen encodings imply all the claimed TVSS properties in the ideal model.

All these statements concern ideal $\mathcal{F}_{\mathrm{CROLE}}$ calls. A maliciously secure concrete realization must be composed for the sequential invocations and the relevant outer experiment. An appropriate composition proof may transfer computational security; computational realization alone does not preserve Lemma 4's perfect equalities or Lemma 6's statistical sampler bound. No stronger real-protocol statistical claim is made. □