Schnorr-like Proofs of Knowledge for Hidden Oil Subspaces in UOV

Zhiwei Wang

Nanjing University of Posts and Telecommunications
zhwwang@njupt.edu.cn

Abstract. A UOV public key hides a distinguished linear subspace: the public-coordinate image of the central oil-coordinate subspace. In central coordinates, the homogeneous quadratic part of each UOV polynomial contains no oil-oil monomials. Consequently, for every honestly generated UOV public key, each public homogeneous quadratic form vanishes when restricted to this hidden oil subspace. We formalize the hidden oil-subspace relation and construct a Schnorr-like Sigma protocol that proves knowledge of such a subspace without revealing it. The witness consists of matrices B, W satisfying $WB = I_o$ and $Q_k(Bz) \equiv 0$ for all public quadratic forms Q_k. The prover masks the witness linearly and responds to a challenge c with $Z = A + cB$, $Y = E + cW$, yielding a protocol with computational 3-special soundness and computational honest-verifier zero knowledge. We prove a generic uniqueness theorem: in the random homogeneous UOV model, the hidden oil subspace is, with overwhelming probability, the unique o-dimensional common zero subspace. The uniqueness bound is explicit and negligible for UOV-type parameters. This strengthens the interpretation of the protocol as proving knowledge of the unique hidden oil subspace, rather than merely some oil-like subspace. We provide empirical observations on small parameters and discuss applications and limitations.

Keywords. Unbalanced Oil and Vinegar; hidden oil subspace; multivariate cryptography; zero-knowledge proofs; proofs of knowledge; Sigma protocols; Schnorr-like protocols; homomorphic commitments; generic uniqueness

1 Introduction

The Unbalanced Oil and Vinegar (UOV) signature scheme is one of the most prominent families of multivariate quadratic signature schemes. Its central map is defined over variables divided into vinegar variables and oil variables. The defining structural constraint is that the homogeneous quadratic part of each central polynomial contains no oil-oil monomials. This structural property is hidden in the public key by an affine change of variables. As a result, every honestly generated UOV public key hides a distinguished linear subspace: the public-coordinate image of the central oil-coordinate subspace. We call this subspace the hidden oil subspace. If Q_k denotes the homogeneous quadratic part

of the k-th public polynomial, then Q_k vanishes identically when restricted to the hidden oil subspace. This observation suggests a natural proof-of-knowledge problem: given a UOV public key, prove knowledge of an o-dimensional subspace on which all public homogeneous quadratic forms vanish, without revealing the subspace.

1.1 The Hidden Oil-Subspace Relation

We formalize this problem through the hidden oil-subspace relation. A witness is a pair (B, W), where $B \in \mathbb{F}_q^{n\times o}$ is a basis matrix for the hidden oil subspace and $W \in \mathbb{F}_q^{o\times n}$ is a left inverse. The relation is

$$WB = I_o, \qquad Q_k(Bz) \equiv 0 \text{ for all } k.$$

The equation $WB = I_o$ certifies that B has full column rank, while the identities $Q_k(Bz) \equiv 0$ express the public-coordinate form of the UOV oil-oil vanishing property.

1.2 Our Approach

We construct a Schnorr-like Sigma protocol for the hidden oil-subspace relation. The prover masks the witness linearly with random matrices A, E, receives a challenge c, and responds with $Z = A + cB$, $Y = E + cW$. The verifier checks identities involving YZ and the restrictions $Q_k(Zz)$. Since the relation contains quadratic equations, the verification conditions are quadratic in the challenge. Consequently, extraction requires three accepting transcripts with the same first message and distinct challenges. The resulting protocol satisfies computational 3-special soundness and computational honest-verifier zero knowledge.

1.3 Generic Uniqueness Theorem

A central question is whether the hidden oil-subspace relation captures a meaningful structural property or merely a weak algebraic coincidence. We address this through a generic uniqueness theorem: in the random homogeneous UOV model, the hidden oil subspace is, with overwhelming probability, the unique o-dimensional common zero subspace. We prove an explicit bound on the probability of additional competing subspaces by classifying them according to their vinegar-projection rank and applying a union bound over all types. For UOV-type parameters, the bound is negligible, confirming that the hidden oil subspace is generically unique. This result significantly strengthens the interpretation of our protocol: proving knowledge of an oil-subspace witness is, in the random UOV model, essentially equivalent to proving knowledge of the unique hidden oil subspace.

1.4 Structural Strength and Limitations

We further analyze the structural strength through random-MQ baseline counting, generic uniqueness bounds, local rigidity via linearized deformation equations, and empirical observations on small parameters. The relation captures a rare structural feature that distinguishes honest UOV keys from random MQ maps. However, our protocol proves knowledge of an oil-subspace witness, not knowledge of a complete UOV signing key or signing capability. Under the random homogeneous UOV model, this witness is generically unique. Our uniqueness theorem does not cover adversarially constructed keys or non-homogeneous variants. The protocol does not prove knowledge of the full affine transformation or signing capability. The commitment layer requires careful instantiation for post-quantum security. The protocol's practical value depends on applications where structural certification is sufficient.

1.5 Our Contributions

Our main contributions are as follows. First, we formalize the hidden oil-subspace relation and isolate the associated search problem. Second, we prove a generic uniqueness theorem showing that in the random homogeneous UOV model, the hidden oil subspace is, with overwhelming probability, the unique o-dimensional common zero subspace. The bound is explicit and negligible for UOV-type parameters. Third, we construct a Schnorr-like Sigma protocol with computational 3-special soundness and computational honest-verifier zero knowledge. Fourth, we analyze the structural strength of the relation through random-MQ counting arguments, generic uniqueness, and local rigidity via linearized deformation equations. We provide empirical observations on small parameters. Fifth, we carefully delineate the scope and limitations of our results, distinguishing what the protocol does and does not prove, and identifying remaining open problems.

1.6 Related Work

Oil-and-vinegar signatures were introduced by Patarin [1], and the unbalanced variant was proposed by Kipnis, Patarin, and Goubin [2]. UOV and its descendants remain central examples of multivariate quadratic signature schemes. Rainbow [4] generalized the oil-and-vinegar paradigm to multiple layers, while later cryptanalysis [5] illustrated the fragility of additional algebraic structure. More recent systems such as MAYO [6] continue to build on oil-vinegar ideas in post-quantum signature design. A central theme in multivariate cryptography is the recovery or certification of hidden algebraic structure from public quadratic maps. Classical examples include the Kipnis–Shamir attack on balanced oil-and-vinegar [3] and later structural attacks on Rainbow-type systems. Our work does not present a new practical attack on standard UOV. Instead, it isolates one hidden structural component—the oil subspace—as a proof-of-knowledge witness.

Generic zero-knowledge systems can prove knowledge of solutions to systems of multivariate quadratic equations, for example through arithmetic-circuit reductions or MPC-in-the-head techniques [7,8,9,10]. Such approaches are flexible

but typically do not exploit the specific structure of the UOV oil subspace. Our protocol is specialized to a narrower relation and therefore follows a more direct algebraic route. Schnorr's identification protocol [11] is the prototypical Sigma protocol with linear masking and special soundness. The broader Sigma-protocol literature [12,13] provides the conceptual framework for commit-challenge-response proofs of knowledge. Since the relation studied here is quadratic, our protocol has 3-special soundness rather than ordinary 2-special soundness. Non-interactive transformations such as Fiat–Shamir [14], especially in post-quantum settings [15,16,17,18], require additional care.

Our protocol relies on homomorphic commitments supporting linear combinations of witness-dependent values. Pedersen commitments [19] provide a simple classical baseline, while post-quantum instantiations may be sought from lattice-based commitments [20,21,22] and related linear-homomorphic tools. The hidden oil-subspace relation can also be viewed as a question about common zero subspaces of systems of quadratic forms over finite fields. The random-MQ baseline and uniqueness analysis therefore connect our relation to classical counting arguments on Grassmannians, finite fields, and finite geometry [23,24].

1.7 Scope and Non-Goals

The protocol proves knowledge of a hidden oil-subspace witness. Under the random homogeneous UOV model, this witness is generically unique. The protocol does not prove knowledge of the full affine transformation or complete central map, ability to sign arbitrary messages, correctness of a particular UOV signature, or use of a declared vinegar in a signing execution. The fixed-vinegar leakage result discussed later is not a break of standard UOV. It requires repeated use of the same internal vinegar, which ordinary UOV signing avoids by fresh sampling.

1.8 Organization

Section 2 reviews notation, UOV structure, and homomorphic commitments. Section 3 studies the hidden oil-subspace relation, proves the generic uniqueness theorem, analyzes local rigidity, and presents empirical observations. Section 4 presents the Schnorr-like Sigma protocol and its security properties. Section 5 explains the fixed-vinegar motivation. Section 6 discusses commitment instantiation and efficiency. Section 7 summarizes applications and limitations. Section 8 concludes.

2 Preliminaries

2.1 Notation

Let q be a prime power and $\mathbb{F}_q$ the finite field with q elements. Let v be the number of vinegar variables, o the number of oil variables, and $n = v + o$. Vectors are written in lowercase bold or ordinary lowercase letters, and matrices in

uppercase letters. For a matrix $B \in \mathbb{F}_q^{n\times o}$, $\mathrm{Col}(B)$ denotes its column space. For a polynomial p, we write Q for its homogeneous quadratic part. For a polynomial in formal variables $z = (z_1, \ldots, z_o)$, $\mathsf{Coeff}_z(\cdot)$ denotes its coefficient vector under a fixed monomial ordering. $x \xleftarrow{\$} S$ denotes uniform sampling from a finite set S. The Gaussian binomial coefficient $\begin{bmatrix} n \\ k \end{bmatrix}_q$ counts the number of k-dimensional subspaces of $\mathbb{F}_q^n$:

$$\begin{bmatrix} n \\ k \end{bmatrix}_q = \frac{(q^n-1)(q^{n-1}-1)\cdots(q^{n-k+1}-1)}{(q^k-1)(q^{k-1}-1)\cdots(q-1)}.$$

Asymptotically, $\begin{bmatrix} n \\ k \end{bmatrix}_q \approx q^{k(n-k)}$ for fixed k and large n.

2.2 UOV Signatures

We recall the UOV syntax only to fix notation. A UOV central map $F = (f^{(1)}, \ldots, f^{(o)}) : \mathbb{F}_q^n \to \mathbb{F}_q^o$ with parameters (q, v, o) and $n = v + o$ has components

$$f^{(k)}(x) = \sum_{1\le i\le j\le n} \alpha_{ij}^{(k)} x_i x_j + \sum_{i=1}^{n} \beta_i^{(k)} x_i + \gamma^{(k)},$$

subject to the UOV constraint $\alpha_{ij}^{(k)} = 0$ whenever $v < i \le j \le n$. We write $x = (a\|u)$ where $a \in \mathbb{F}_q^v$ (vinegar) and $u \in \mathbb{F}_q^o$ (oil). A UOV key is generated by sampling a central UOV map F and an invertible affine transformation T. The public key is $P = F \circ T$, and the secret key contains F and T^{-1}.

2.3 Quadratic Forms and Polarization

For a homogeneous quadratic form $Q : \mathbb{F}_q^n \to \mathbb{F}_q$, its polar form is $\mathsf{Pol}_Q(x, y) = Q(x + y) - Q(x) - Q(y)$. For $A, B \in \mathbb{F}_q^{n\times o}$, $c \in \mathbb{F}_q$, and formal variables z,

$$Q((A + cB)z) = Q(Az) + c\mathsf{Pol}_Q(Az, Bz) + c^2 Q(Bz).$$

2.4 Linearly Homomorphic Commitments

The protocol is stated in an abstract linearly homomorphic commitment model. All message spaces used by the protocol are finite-dimensional $\mathbb{F}_q$-vector spaces. Matrices and coefficient vectors are encoded as vectors over $\mathbb{F}_q$ in the natural way. The abstraction is intentionally separated from any concrete instantiation: the security proof of the protocol uses only correctness, computational hiding, and a binding property that is stable under public linear combinations of commitments.

Let M be a finite-dimensional $\mathbb{F}_q$-vector space. A linearly homomorphic commitment scheme for M is a tuple

$$\Pi_{\mathsf{com}} = (\mathsf{Setup}, \mathsf{Com}, \mathsf{Open}, \mathsf{Verify}, \mathsf{LinComb}, \mathsf{LinRand})$$

with the following syntax. The setup algorithm $\mathsf{pp} \leftarrow \mathsf{Setup}(1^\lambda, M)$ outputs public parameters for message space M. The commitment algorithm $C = \mathsf{Com}(\mathsf{pp}, x; r)$ commits to a message $x \in M$ using randomness r. The opening algorithm outputs $\pi = \mathsf{Open}(\mathsf{pp}, C, x, r)$, and the verification algorithm returns $\mathsf{Verify}(\mathsf{pp}, C, x, \pi) \in \{0, 1\}$. The public algorithm $\mathsf{LinComb}$ derives commitments to linear combinations. Given commitments $C_1, \ldots, C_t$ and scalars $\alpha_1, \ldots, \alpha_t \in \mathbb{F}_q$, it outputs

$$C = \mathsf{LinComb}(\mathsf{pp}; C_1, \ldots, C_t; \alpha_1, \ldots, \alpha_t).$$

The algorithm $\mathsf{LinRand}$ derives the corresponding opening randomness from the individual randomnesses:

$$r = \mathsf{LinRand}(r_1, \ldots, r_t; \alpha_1, \ldots, \alpha_t).$$

Definition 1 (Correctness). *A linearly homomorphic commitment scheme is correct if the following two conditions hold. First, for every $x \in M$, every valid randomness r, and $C = \mathsf{Com}(\mathsf{pp}, x; r)$, we have $\mathsf{Verify}(\mathsf{pp}, C, x, \mathsf{Open}(\mathsf{pp}, C, x, r)) = 1$. Second, for honestly generated commitments $C_i = \mathsf{Com}(\mathsf{pp}, x_i; r_i)$ $(i = 1, \ldots, t)$ and scalars $\alpha_1, \ldots, \alpha_t \in \mathbb{F}_q$, the derived commitment $C = \mathsf{LinComb}(\mathsf{pp}; C_1, \ldots, C_t; \alpha_1, \ldots, \alpha_t)$ opens correctly to $x = \sum_{i=1}^{t} \alpha_i x_i$ using derived randomness $r = \mathsf{LinRand}(r_1, \ldots, r_t; \alpha_1, \ldots, \alpha_t)$.*

Definition 2 (Computational hiding). *A linearly homomorphic commitment scheme is computationally hiding if, for every probabilistic polynomial-time adversary $\mathcal{A}$, every message space M used by the protocol, and every pair of equal-length messages $x_0, x_1 \in M$, the distributions $\mathsf{Com}(\mathsf{pp}, x_0; r_0)$ and $\mathsf{Com}(\mathsf{pp}, x_1; r_1)$ are computationally indistinguishable, where $\mathsf{pp} \leftarrow \mathsf{Setup}(1^\lambda, M)$ and r_0, r_1 are sampled according to the commitment randomness distribution. Equivalently, the usual left-or-right commitment-hiding experiment has advantage at most $\mathsf{negl}(\lambda)$.*

Definition 3 (Binding). *A commitment scheme is computationally binding if, except with negligible probability over $\mathsf{pp} \leftarrow \mathsf{Setup}(1^\lambda, M)$, no probabilistic polynomial-time adversary can output a commitment C and two valid openings (x, π), (x', π') such that $x \neq x'$ and $\mathsf{Verify}(\mathsf{pp}, C, x, \pi) = \mathsf{Verify}(\mathsf{pp}, C, x', \pi') = 1$.*

For protocols that open homomorphically derived commitments, ordinary binding of base commitments is not by itself a convenient formulation. We therefore use the following linear consistency property.

Definition 4 (Linear binding). *A linearly homomorphic commitment scheme is linearly binding if, except with negligible probability over $\mathsf{pp} \leftarrow \mathsf{Setup}(1^\lambda, M)$, no probabilistic polynomial-time adversary can output commitments $C_1, \ldots, C_t$, scalars $\alpha_1, \ldots, \alpha_t \in \mathbb{F}_q$, and two valid openings of the derived commitment $C = \mathsf{LinComb}(\mathsf{pp}; C_1, \ldots, C_t; \alpha_1, \ldots, \alpha_t)$ to two distinct messages $x \neq x'$. Equivalently, once base commitments $C_1, \ldots, C_t$ are fixed, every homomorphically derived commitment is binding to a unique message. If the base commitments are validly openable to messages $x_1, \ldots, x_t$, then the derived commitment must be bound to $\sum_{i=1}^{t} \alpha_i x_i$.*

The protocol uses commitments to several message spaces, including matrix spaces such as $\mathbb{F}_q^{n\times o}$, $\mathbb{F}_q^{o\times n}$, $\mathbb{F}_q^{o\times o}$, and coefficient-vector spaces such as $\mathbb{F}_q^{o(o+1)/2}$. This can be handled in one of two equivalent ways. One may instantiate the commitment scheme separately for each finite-dimensional $\mathbb{F}_q$-vector space, or one may embed all protocol messages into a single sufficiently large $\mathbb{F}_q$-vector space with fixed tags identifying the block type. The security proof only requires that linear combinations are applied within compatible message spaces and that the resulting derived commitments satisfy the linear binding property.

The soundness proof in Section 4 uses linear binding to ensure that, for a fixed first message, all accepting responses to different challenges are consistent with one fixed set of committed linear blocks. The honest-verifier zero-knowledge proof uses only computational hiding and does not require equivocation, trapdoors, or programmable commitments.

Classical Pedersen-style commitments provide a simple intuition for linear homomorphism over suitable scalar fields under the discrete-logarithm assumption. However, when q is not the scalar field of the commitment group, for example $q = 2^8$, an explicit encoding is needed to preserve the intended $\mathbb{F}_q$-linear operations. For post-quantum security, lattice-based homomorphic or vector commitments are natural candidates, but concrete parameter selection and a full instantiation-level analysis are outside the scope of this work. The results below should therefore be read as security statements in the abstract linearly homomorphic commitment model defined here.

3 The Hidden Oil-Subspace Relation

3.1 Hidden Oil Subspaces in Public Coordinates

Let the secret affine transformation be $T^{-1}(x) = L^{-1}x + t$ with $L \in \mathrm{GL}_n(\mathbb{F}_q)$, $t \in \mathbb{F}_q^n$. The central oil-coordinate subspace is $\{0\}^v \times \mathbb{F}_q^o$; its image under L^{-1} is

$$\mathcal{O} = \{L^{-1}(0\|y) : y \in \mathbb{F}_q^o\} \subseteq \mathbb{F}_q^n,$$

called the hidden oil subspace. The translation t does not affect the homogeneous quadratic part.

3.2 Definition of the Relation

Let $P = (p_1, \ldots, p_m)$ be a public MQ map and Q_k the homogeneous quadratic part of p_k. Define

$$\mathcal{R}_{\mathrm{OS}} = \{(P; (B, W)) : B \in \mathbb{F}_q^{n\times o}, W \in \mathbb{F}_q^{o\times n}, WB = I_o, Q_k(Bz) \equiv 0 \ \forall k\},$$

where $z = (z_1, \ldots, z_o)$ is a vector of formal variables. B is a basis of an o-dimensional subspace; $WB = I_o$ certifies full column rank; $Q_k(Bz) \equiv 0$ states that all public quadratic forms vanish on $\mathrm{Col}(B)$.

Lemma 1 (Correctness for honest UOV keys). *For every honestly generated UOV public key $P = F \circ T$, there exists a witness (B, W) such that $(P; (B, W)) \in \mathcal{R}_{\mathrm{OS}}$.*

Proof. Let $\mathcal{O}$ be the hidden oil subspace. Choose any basis B of $\mathcal{O}$; B has full rank so a left inverse W exists. For any $Bz \in \mathcal{O}$ there exists y with $Bz = L^{-1}(0\|y)$. In central coordinates this is a pure oil direction, and the central UOV forms have no oil-oil terms, hence $Q_k(Bz) \equiv 0$.

3.3 The Oil-Subspace Search Problem

The relation $\mathcal{R}_{\mathrm{OS}}$ naturally defines a search problem that isolates the structural task of finding a common zero subspace for the public quadratic forms. Given a public MQ map $P = (p_1, \ldots, p_m)$ over $\mathbb{F}_q$ with homogeneous quadratic parts $Q_1, \ldots, Q_m$, the hidden oil-subspace search problem (HOSP) asks to find a matrix $B \in \mathbb{F}_q^{n \times o}$ of rank o such that $Q_k(Bz) \equiv 0$ for all $k = 1, \ldots, m$. The witness relation $\mathcal{R}_{\mathrm{OS}}$ is the proof-of-knowledge version of this search problem, augmented with an explicit left inverse W satisfying $WB = I_o$. For honest UOV keys, HOSP always has at least one solution. For a random MQ map, existence is a strong algebraic condition.

3.4 Random-MQ Baseline

Let $\mathrm{Gr}(o, n)$ denote the Grassmannian of o-dimensional subspaces of $\mathbb{F}_q^n$. For a fixed subspace $S \in \mathrm{Gr}(o, n)$, the restriction of a homogeneous quadratic form Q to S is a homogeneous quadratic form in o variables, which has $N_o = o(o+1)/2$ coefficients. After choosing a basis of S, the restriction $Q|_S$ has $o(o+1)/2$ coefficients that are uniformly distributed over $\mathbb{F}_q$. They are all zero with probability $q^{-o(o+1)/2}$. Independence gives the m-form case. The number of o-dimensional subspaces is $\begin{bmatrix} n \\ o \end{bmatrix}_q$. For each fixed S, the probability that all m forms vanish on S is $q^{-mo(o+1)/2}$. Linearity of expectation yields the expected number of common zero subspaces:

$$\mathbb{E}[X] = \begin{bmatrix} n \\ o \end{bmatrix}_q q^{-mo(o+1)/2} \approx q^{o(n-o)-mo(o+1)/2}.$$

For UOV-type parameters, the exponent $o(n-o) - mo(o+1)/2$ is typically strongly negative. Thus a random MQ map is extremely unlikely to admit an o-dimensional common zero subspace. By contrast, an honestly generated UOV public key always admits at least one such subspace by Lemma 1. Hence this property sharply separates honest UOV structure from the random-MQ baseline at the level of existence. For example, for $n = 156$, $o = 44$, $m = 44$ we have $o(n-o) = 44 \cdot 112 = 4928$, while $mo(o+1)/2 = 44 \cdot 990 = 43560$. The exponent is $4928 - 43560 = -38632$, so $\mathbb{E}[X] \approx q^{-38632}$, which is negligible.

3.5 Random Homogeneous UOV Model

We introduce a precise random model for the uniqueness analysis.

Definition 5 (Random homogeneous UOV model). *Let $V = \mathbb{F}_q^v$, $O = \mathbb{F}_q^o$, and write elements of $V \oplus O$ as (a, u). A homogeneous UOV-type quadratic form is a homogeneous quadratic polynomial of the form*

$$Q(a,u) = \sum_{1 \le i \le j \le v} \alpha_{ij} a_i a_j + \sum_{1 \le i \le v,\ 1 \le j \le o} \beta_{ij} a_i u_j,$$

with no monomials $u_i u_j$. The coefficients α_{ij} and β_{ij} are elements of $\mathbb{F}_q$. A random homogeneous UOV map with m components is a tuple $(Q_1, \ldots, Q_m)$ where the Q_k's are sampled independently and uniformly from this vector space of UOV-type homogeneous quadratic forms.

This definition is coefficient-based and is valid in all characteristics, including characteristic 2. The affine and linear terms of a UOV public key are irrelevant for the hidden oil-subspace relation, which depends only on the homogeneous quadratic parts. A public-coordinate instance is obtained by applying an invertible linear change of variables $L \in \mathrm{GL}_{v+o}(\mathbb{F}_q)$. Since such a change of variables induces a bijection on the Grassmannian, uniqueness in central coordinates is equivalent to uniqueness in public coordinates.

3.6 Classification of Competing Subspaces

Work in central coordinates and write vectors as $(a, u) \in \mathbb{F}_q^v \times \mathbb{F}_q^o$. The standard oil subspace is $\mathcal{O}_0 = \{0\}^v \times \mathbb{F}_q^o$. Any o-dimensional subspace $S \subseteq \mathbb{F}_q^v \oplus \mathbb{F}_q^o$ can be classified by its vinegar projection rank. Define the projection $\pi_V : \mathbb{F}_q^v \oplus \mathbb{F}_q^o \to \mathbb{F}_q^v$. For any o-dimensional subspace S, let $r = \dim \pi_V(S)$. If $r = 0$, then $S \subseteq \mathcal{O}_0$. Since $\dim S = o = \dim \mathcal{O}_0$, we have $S = \mathcal{O}_0$. Therefore, any competing subspace $S \neq \mathcal{O}_0$ must satisfy $r \geq 1$.

3.7 Parametrization of r-Type Subspaces

For fixed $r \in \{1, \ldots, o\}$, an r-type subspace S with $\dim \pi_V(S) = r$ can be described by an r-dimensional vinegar subspace $A \subseteq \mathbb{F}_q^v$, an $(o-r)$-dimensional oil subspace $K \subseteq \mathbb{F}_q^o$, and a linear map $\phi : A \to \mathbb{F}_q^o / K$. The subspace S is then the graph $S = \{(a, \tilde{\phi}(a) + k) : a \in A, k \in K\}$, where $\tilde{\phi} : A \to \mathbb{F}_q^o$ is any lift of ϕ. The number of such subspaces is

$$N_{\mathrm{sub}}(r) = \begin{bmatrix} v \\ r \end{bmatrix}_q \begin{bmatrix} o \\ r \end{bmatrix}_q q^{r^2}.$$

Asymptotically, $N_{\mathrm{sub}}(r) \approx q^{r(v-r)+r(o-r)+r^2} = q^{r(n-r)}$.

3.8 Restriction Dimension for r-Type Subspaces

Lemma 2 (Rank of the UOV restriction map). *Let $S \subseteq \mathbb{F}_q^v \oplus \mathbb{F}_q^o$ be an o-dimensional subspace with $r = \dim \pi_V(S)$. Let Q_{UOV} be the vector space of homogeneous quadratic forms on $\mathbb{F}_q^v \oplus \mathbb{F}_q^o$ with no oil-oil terms. Then the restriction map $\mathrm{Res}_S : Q_{\mathrm{UOV}} \to \mathrm{Quad}(S)$ has rank*

$$\mathrm{rank}(\mathrm{Res}_S) = r(o-r) + \frac{r(r+1)}{2}.$$

Proof. Using the parametrization from the previous subsection, we can identify S with $A \oplus K$ where $\dim A = r$ and $\dim K = o - r$. A UOV-type quadratic form has the structure $Q(a,u) = Q_{vv}(a) + Q_{vo}(a,u)$, where Q_{vv} contains vinegar-vinegar terms and Q_{vo} contains vinegar-oil terms, but there are no oil-oil terms. The restriction $Q|_S$ can be written as $Q|_S(a,k) = Q_{vv}(a) + Q_{vo}(a, \tilde{\phi}(a)) + Q_{vo}(a,k)$, where $a \in A$ and $k \in K$. The coefficient space of $Q|_S$ decomposes into bilinear terms in $A \times K$ coming from $Q_{vo}(a,k)$ with dimension $r(o-r)$, and quadratic terms in A coming from $Q_{vv}(a) + Q_{vo}(a, \tilde{\phi}(a))$ with dimension $r(r+1)/2$. Any desired bilinear coefficient in $A \times K$ can be realized by choosing appropriate Q_{vo} coefficients on the corresponding vinegar-oil pairs. Any desired quadratic coefficient in A can be realized by choosing appropriate Q_{vv} and Q_{vo} coefficients. These two types of coefficients are independent. Therefore, the image dimension is $\dim \mathrm{Bil}(A,K) + \dim \mathrm{Quad}(A) = r(o-r) + r(r+1)/2$.

Define $N_r = r(o-r) + r(r+1)/2 = ro - r(r-1)/2$. For a fixed r-type subspace S and a uniformly random UOV-type homogeneous quadratic form Q, the restriction $Q|_S$ is determined by N_r independent uniformly random coefficients. The probability that all are zero is q^{-N_r}. For m independent random UOV-type forms, the probability that all vanish on S is q^{-mN_r}.

3.9 Generic Uniqueness Theorem

Theorem 1 (Generic uniqueness bound for random UOV maps). *Let $V = \mathbb{F}_q^v$, $O = \mathbb{F}_q^o$, and let $Q_1, \ldots, Q_m$ be independent random UOV-type homogeneous quadratic forms on $V \oplus O$ according to Definition 5. Let $\mathcal{O}_0 = \{0\} \oplus O$. Then*

$$\Pr[\exists S \in \mathrm{Gr}(o, V \oplus O), S \neq \mathcal{O}_0, Q_k|_S \equiv 0 \ \forall k] \leq \sum_{r=1}^{o} \begin{bmatrix} v \\ r \end{bmatrix}_q \begin{bmatrix} o \\ r \end{bmatrix}_q q^{r^2} q^{-mN_r}.$$

In particular, if the right-hand side is negligible, then $\mathcal{O}_0$ is the unique o-dimensional common zero subspace with overwhelming probability.

Proof. Let $S \neq \mathcal{O}_0$ be an o-dimensional subspace, and let $r = \dim \pi_V(S)$. Since $S \neq \mathcal{O}_0$, we have $r \geq 1$. For fixed r, each such S can be uniquely represented by a triple (A, K, ϕ) where $A \in \mathrm{Gr}(r, V)$, $K \in \mathrm{Gr}(o-r, O)$, $\phi : A \to O/K$. The number of such subspaces is $\begin{bmatrix} v \\ r \end{bmatrix}_q \begin{bmatrix} o \\ r \end{bmatrix}_q q^{r^2}$. By Lemma 2, for a fixed r-type subspace S, $\Pr[Q_k|_S \equiv 0 \ \forall k] = q^{-mN_r}$. Applying the union bound over all r-type subspaces for each $r \in \{1, \ldots, o\}$ gives the inequality.

Table 1. Dominant exponent in the generic uniqueness bound. The worst case is typically $r = 1$. For UOV-type parameters, the bound is negligible.

q	(v, o, m)	worst r	exponent E_1	conclusion
251	(8,4,4)	1	-5	small but visible
251	(12,6,6)	1	-19	negligible for toy
251	(16,8,8)	1	-41	negligible
256	(112,44,44)	1	-1781	overwhelming uniqueness

Lemma 3 (Invariance under linear change of variables). *If $\mathcal{O}_0$ is the unique o-dimensional common zero subspace for the central quadratic forms, then $L^{-1}\mathcal{O}_0$ is the unique o-dimensional common zero subspace for the public quadratic forms.*

Proof. An invertible linear map L^{-1} induces a bijection on the Grassmannian $\mathrm{Gr}(o, n)$. If Q_k are the central forms and $\tilde{Q}_k$ are the public forms, then $\tilde{Q}_k(x) = Q_k(Lx)$. Thus $\tilde{Q}_k|_S \equiv 0$ if and only if $Q_k|_{LS} \equiv 0$. Uniqueness is preserved under this bijection.

3.10 Parameter Analysis

We now analyze the bound for concrete UOV-type parameters. For each r, the exponent in the bound is approximately

$$E_r = r(n-r) - m\left(ro - \frac{r(r-1)}{2}\right) = rn - r^2 - mro + \frac{mr(r-1)}{2}.$$

For typical UOV parameters, the dominant term is usually $r = 1$, which gives $E_1 = n-1-mo$. Table 1 shows that for all tested parameter sets, the exponent is strongly negative. For the standard UOV parameter set $(v, o, m) = (112, 44, 44)$ over $\mathbb{F}_{256}$, the probability of an additional competing subspace is bounded by approximately q^{-1781}, which is overwhelmingly negligible.

The generic uniqueness theorem applies to the random homogeneous UOV model, where central quadratic forms have random vinegar-vinegar and vinegar-oil coefficients. It does not cover adversarially constructed public keys with deliberately planted additional subspaces, non-homogeneous or affine variants where the affine and linear terms may introduce additional structure, or structured or low-rank coefficient distributions. Nevertheless, for honestly generated UOV keys under the standard random-coefficient model, the theorem provides strong evidence that the hidden oil subspace is generically unique.

3.11 Local Rigidity and the Linearized Benchmark

The generic uniqueness theorem provides a global probabilistic guarantee. We now complement it with a local rigidity analysis via linearized deformation equations. Work in central coordinates and write vectors as $(a, u) \in \mathbb{F}_q^v \times \mathbb{F}_q^o$. The

standard oil subspace is $\mathcal{O}_0 = \{0\}^v \times \mathbb{F}_q^o$. Any nearby o-dimensional subspace that is transverse to the vinegar coordinate space can be represented as the graph of a linear map $R : \mathbb{F}_q^o \to \mathbb{F}_q^v$, $S_R = \{(Ru, u) : u \in \mathbb{F}_q^o\}$. The undeformed oil subspace corresponds to $R = 0$. Because the central UOV forms contain no oil-oil terms, $R = 0$ is indeed a solution of the vanishing equations. Write the homogeneous quadratic part of the k-th central polynomial as $Q_k(a, u) = Q_k^{vv}(a) + Q_k^{vo}(a, u)$, where Q_k^{vv} contains the vinegar-vinegar terms and Q_k^{vo} contains the vinegar-oil terms. The condition that all central quadratic forms vanish on S_R is $Q_k(Ru, u) \equiv 0$ for all k. The first-order deformation map is therefore $L : \mathbb{F}_q^{v\times o} \to (\mathbb{F}_q^{o(o+1)/2})^m$, $R \mapsto (\mathsf{Coeff}_u(Q_k^{vo}(Ru, u)))_{k=1}^m$. Choose the standard basis of $\mathbb{F}_q^{v\times o}$ and the coefficient basis of $(\mathbb{F}_q^{o(o+1)/2})^m$. Let $M_{\text{lin}} \in \mathbb{F}_q^{mo(o+1)/2\times vo}$ be the matrix representing L in these bases. The honest oil subspace is first-order rigid if and only if $\text{rank}(M_{\text{lin}}) = vo$.

To interpret this structurally, we compare M_{lin} with a random linear map of the same dimensions. Let $M = m \cdot o(o + 1)/2$, $N = vo$, and let $U \xleftarrow{\$} \mathbb{F}_q^{M\times N}$ be a uniformly random matrix, with $M \geq N$. The columns of U are $u_1, \ldots, u_N \in \mathbb{F}_q^M$. The matrix has full column rank if and only if these columns are linearly independent. The first column u_1 is uniformly random in $\mathbb{F}_q^M$. The probability that $u_1 \neq 0$ is $1 - q^{-M}$. Conditioned on $u_1 \neq 0$, the column u_2 is uniformly random and independent of u_1. The span of u_1 is a 1-dimensional subspace of $\mathbb{F}_q^M$, which contains q vectors. The probability that u_2 lies outside this span is $1 - q^{1-M}$. Conditioned on $u_1, \ldots, u_j$ being linearly independent, they span a j-dimensional subspace containing q^j vectors. The probability that u_{j+1} lies outside this span is $1-q^{j-M}$. Multiplying these probabilities gives $\Pr[\text{rank}(U) = N] = \prod_{j=0}^{N-1}(1 - q^{j-M})$. The probability of rank deficiency is $\Pr[\text{rank}(U) < N] = 1 - \prod_{j=0}^{N-1}(1 - q^{j-M})$. Using the inequality $1 - \prod_i(1 - a_i) \leq \sum_i a_i$ for $a_i \in [0, 1]$, we obtain $\Pr[\text{rank}(U) < N] \leq \sum_{j=0}^{N-1} q^{j-M}$. This is a geometric series: $\sum_{j=0}^{N-1} q^{j-M} = q^{-M} \sum_{j=0}^{N-1} q^j = q^{-M} \cdot \frac{q^N-1}{q-1} < \frac{q^{N-M}}{q-1}$. For $M \gg N$, this bound is exponentially small in $M - N$.

For randomly generated UOV central maps, the matrix M_{lin} often behaves experimentally like a random matrix in $\mathbb{F}_q^{M\times N}$ with $M = m\cdot o(o+1)/2$, $N = vo$. Accordingly, when $M \gg N$, we expect first-order rigidity with overwhelming probability. For the parameter choice $n = 156$, $v = 112$, $o = 44$, $m = 44$, we obtain $N = vo = 112 \cdot 44 = 4928$, $M = m \cdot o(o + 1)/2 = 44 \cdot 990 = 43560$. Thus $M - N = 38632$, and the benchmark bound gives $\Pr[\text{rank}(U) < N] < q^{-38632}/(q - 1)$. For $q = 256$, this failure probability is negligible. Table 2 shows that the linearized system is already overdetermined for small toy parameters and dramatically overdetermined for UOV-type parameters.

3.12 Empirical Observations

We conducted experiments to support three claims suggested by the theory: first, the linearized rigidity matrix M_{lin} should have full column rank with high probability for random UOV-type instances; second, randomized search should not

Table 2. Dimension parameters for the linearized-rigidity benchmark. The overdetermination gap $\Delta = M - N$ is large for all parameter sets.

q	(v, o, m)	$N = vo$	$M = m \cdot o(o+1)/2$	$\Delta = M - N$
251	(8,4,4)	32	40	8
251	(12,6,6)	72	126	54
251	(16,8,8)	128	288	160
256	(112,44,44)	4928	43560	38632

Table 3. Empirical rank tests for the linearized matrix M_{lin}. All tested random UOV instances achieved full column rank. The final column is the random-matrix upper bound $q^{N-M}/(q-1)$.

q	(v, o, m)	trials	$M \times N$	full-rank count	benchmark failure bound
251	(6,3,3)	200	18×18	199	$< 1/(250)$
251	(8,4,4)	200	40×32	200	$< 251^{-8}/250$
251	(12,6,6)	100	126×72	100	$< 251^{-54}/250$
251	(16,8,8)	50	288×128	50	$< 251^{-160}/250$

find additional common zero subspaces in parameter ranges where the uniqueness bound is already small; third, in the faulty fixed-vinegar scenario, signature differences generated from a single vinegar coset should recover the hidden oil subspace once the corresponding oil differences span $\mathbb{F}_q^o$.

All experiments were implemented in Python using SageMath for finite-field arithmetic and linear algebra. The experiments are intended as sanity checks for the algebraic analysis, not as evidence of cryptographic security. Exhaustive search over the Grassmannian is feasible only for very small parameters, since $\|\operatorname{Gr}(o,n)\| = \begin{bmatrix} n \\ o \end{bmatrix}_q \approx q^{o(n-o)}$. For larger parameters we rely on randomized sampling of candidate subspaces.

Linearized rigidity experiment. For each parameter set (q, v, o, m), we generated independent random homogeneous UOV central maps according to Definition 5. For each instance we constructed the linearized matrix $M_{\text{lin}} \in \mathbb{F}_q^{mo(o+1)/2 \times vo}$ defined in Section 3 and computed its rank by Gaussian elimination over $\mathbb{F}_q$. We compare the observed rank-deficiency rate with the random-matrix benchmark

$$\Pr[\operatorname{rank}(U) < N] \leq \frac{q^{N-M}}{q-1},$$

where $U \xleftarrow{\$} \mathbb{F}_q^{M \times N}$, $M = mo(o+1)/2$, $N = vo$.

The borderline case $(v, o, m) = (6, 3, 3)$ has $M = N = 18$, so full rank is not overwhelmingly likely; the observed one rank-deficient instance among 200 trials is consistent with this. In the overdetermined cases, all tested instances had full column rank, matching the prediction that local deformations should be rare when $M - N$ is large.

Table 4. Randomized search for additional common zero subspaces. No additional subspaces were found in the tested random UOV instances.

q	(v, o, m)	instances	candidates per instance	extra subspaces found
251	(6,3,3)	50	10^4	0
251	(8,4,4)	50	10^4	0
251	(12,6,6)	20	10^4	0

Table 5. Fixed-vinegar recovery sanity check. The recovered difference space reaches dimension o once the sampled oil differences span $\mathbb{F}_q^o$.

q	(v, o)	trials	average samples to rank o	failures
251	(8,4)	100	4.02	0
251	(12,6)	100	6.03	0
251	(16,8)	50	8.04	0

Randomized uniqueness search. We next tested whether randomized search finds additional o-dimensional subspaces $S \neq \mathcal{O}_0$ satisfying $Q_k|_S \equiv 0$ for all k. For each generated instance, we sampled random full-rank matrices $B' \in \mathbb{F}_q^{n\times o}$, rejected those whose column space coincided with the known oil subspace, and tested whether $Q_k(B'z) \equiv 0$ for all k. This search is not exhaustive except in the smallest cases, but it is useful as a sanity check against accidental additional subspaces.

These results are consistent with Theorem 1. They should not be interpreted as a proof of uniqueness, since the number of sampled candidate subspaces is negligible compared with the full Grassmannian for larger parameters.

Fixed-vinegar recovery sanity check. Finally, we simulated the faulty fixed-vinegar scenario described in Section 5. For each instance, we sampled a random invertible linear map L, fixed a vinegar value a_0, sampled oil values $u_i \in \mathbb{F}_q^o$, and formed public-coordinate signatures $\sigma_i = L^{-1}(a_0\|u_i) + t$. We then computed the rank of the difference space $D_t = \{\sigma_i - \sigma_0 : 1 \le i \le t\}$. As predicted, $\dim D_t$ matched the rank of the oil-difference matrix $[u_1 - u_0 \ \cdots \ u_t - u_0]$, and reached o once the oil differences spanned $\mathbb{F}_q^o$.

This experiment does not model the full UOV signing distribution. It only checks the linear-algebraic claim that signatures generated from the same vinegar coset reveal vectors in the hidden oil subspace through differences. The result is therefore a validation of the motivation in Section 5, not an attack experiment against standard UOV.

Summary. The experiments support the qualitative predictions of the theoretical analysis: the linearized system is full rank in overdetermined regimes, randomized search does not reveal additional common zero subspaces in tested instances, and fixed-vinegar differences recover the hidden oil subspace in the expected linear-algebraic sense. The experiments remain limited to small parameters and do not replace the generic uniqueness theorem.

3.13 Consequences for Proof-of-Knowledge Interpretation

The generic uniqueness theorem significantly strengthens the interpretation of our protocol. In the random UOV model of Theorem 1, except with probability bounded by $\sum_{r=1}^{o} \begin{bmatrix} v \\ r \end{bmatrix}_q \begin{bmatrix} o \\ r \end{bmatrix}_q q^{r^2} q^{-mN_r}$, any witness (B, W) extracted by the special-soundness extractor satisfies $\mathrm{Col}(B) = \mathcal{O}_0$ in central coordinates, and hence corresponds to the unique hidden oil subspace in public coordinates. By Theorem 1, with overwhelming probability there is only one o-dimensional common zero subspace. By the soundness of the protocol, the extractor outputs a valid witness $(B, W) \in \mathcal{R}_{\mathrm{OS}}$. Therefore $\mathrm{Col}(B)$ is an o-dimensional common zero subspace, which must be $\mathcal{O}_0$ except with the stated probability. Lemma 3 extends this to public coordinates. This shows that, in the random UOV model, proving knowledge of an oil-subspace witness is essentially equivalent to proving knowledge of the unique hidden oil subspace.

3.14 Relation Strength

The relation $\mathcal{R}_{\mathrm{OS}}$ is intentionally weaker than the full UOV secret-key relation. A witness (B, W) certifies knowledge of an o-dimensional subspace on which all public homogeneous quadratic forms vanish. It does not specify the complete affine transformation T^{-1}, the full central map F, or an efficient signing algorithm. Nevertheless, $\mathcal{R}_{\mathrm{OS}}$ is not an arbitrary weak algebraic condition. Lemma 1 shows that every honestly generated UOV public key contains at least one valid witness. The random-MQ baseline shows that the expected number of such witnesses for a random MQ system is negligible in UOV-type parameter regimes. Theorem 1 shows that for random UOV-type quadratic maps, the hidden oil subspace is generically unique. The local-rigidity analysis further indicates that, near the honest oil subspace, nontrivial deformations should be rare when the linearized system is strongly overdetermined. Thus, the relation $\mathcal{R}_{\mathrm{OS}}$ should be interpreted as a proof of knowledge of a rare structural feature that distinguishes honest UOV keys from random MQ maps, and which is generically unique in the random UOV model.

4 A Schnorr-like Sigma Protocol

We now construct a public-coin three-move protocol for proving knowledge of a hidden oil-subspace witness. The protocol follows the Schnorr paradigm of linear masking, but the relation is quadratic. Consequently, extraction requires three accepting transcripts with the same first message and distinct challenges. Throughout this section, the public input is $P = (p_1, \ldots, p_m)$, and $Q_1, \ldots, Q_m$ denote the homogeneous quadratic parts of the public polynomials. The witness is a pair $(B, W) \in \mathbb{F}_q^{n \times o} \times \mathbb{F}_q^{o \times n}$ satisfying $WB = I_o$ and $Q_k(Bz) \equiv 0$ for all $k = 1, \ldots, m$, where $z = (z_1, \ldots, z_o)$ is a vector of formal variables.

4.1 Commitment Interface

We use the linearly homomorphic commitment interface defined in Section 2.4. All committed messages are encoded as elements of finite-dimensional $\mathbb{F}_q$-vector spaces. The protocol uses four committed blocks:

$$X_1 = (B, W), \qquad X_2 = (A, E), \qquad X_3 = (M_0, M_1), \qquad X_4 = (U_{1,0}, U_{1,1}, \ldots, U_{m,0}, U_{m,1}).$$

For each challenge c, the verifier derives commitments to the linear combinations

$$Z = A + cB, \quad Y = E + cW, \quad M = M_0 + cM_1, \quad U_k = U_{k,0} + cU_{k,1}.$$

Correctness of the commitment scheme ensures that honest derived openings verify. Computational hiding is used in the honest-verifier zero-knowledge proof. Linear binding, in the sense of Definition 4, is used in the special-soundness proof to guarantee that accepting responses to different challenges are consistent with one fixed set of committed blocks. The commitment layer is the only computational assumption used by the protocol itself. The algebraic extraction argument is information-theoretic once the committed linear decompositions are fixed by binding.

4.2 Protocol Description

Let $\mathcal{C} \subseteq \mathbb{F}_q$ be the challenge set, with $|\mathcal{C}| \geq 3$. The public input is $P = (p_1, \ldots, p_m)$, together with the homogeneous quadratic forms $Q_1, \ldots, Q_m$. The prover knows (B, W) such that $WB = I_o$ and $Q_k(Bz) \equiv 0$ for all $k = 1, \ldots, m$.

In the commitment phase, the prover samples $A \xleftarrow{\$} \mathbb{F}_q^{n \times o}$, $E \xleftarrow{\$} \mathbb{F}_q^{o \times n}$. It computes $M_0 := EA$, $M_1 := EB + WA$, and, for each $k = 1, \ldots, m$, $U_{k,0} := \mathsf{Coeff}_z(Q_k(Az))$, $U_{k,1} := \mathsf{Coeff}_z(\mathsf{Pol}_{Q_k}(Az, Bz))$. Define four committed blocks: $X_1 := (B, W)$, $X_2 := (A, E)$, $X_3 := (M_0, M_1)$, $X_4 := (U_{1,0}, U_{1,1}, \ldots, U_{m,0}, U_{m,1})$. The prover samples commitment randomnesses r_1, r_2, r_3, r_4 and sends $C_j := \mathsf{Com}(\mathsf{pp}, X_j; r_j)$ for $j = 1, 2, 3, 4$. The first message is $\mathsf{com} := (C_1, C_2, C_3, C_4)$.

In the challenge phase, the verifier samples $c \xleftarrow{\$} \mathcal{C}$ and sends c to the prover.

In the response phase, the prover computes $Z := A + cB$, $Y := E + cW$, $M := M_0 + cM_1$, and, for each $k = 1, \ldots, m$, $U_k := U_{k,0} + cU_{k,1}$. Using the homomorphic commitment operations, the prover derives commitments to $Z, Y, M, U_1, \ldots, U_m$ from $C_1, \ldots, C_4$. More explicitly, these are commitments to the following linear combinations of the committed blocks: $Z = A + cB$, $Y = E + cW$, $M = M_0 + cM_1$, $U_k = U_{k,0} + cU_{k,1}$. The prover sends $\mathsf{rsp} := (Z, Y, M, \{U_k\}_{k=1}^m, \pi_Z, \pi_Y, \pi_M, \{\pi_{U_k}\}_{k=1}^m)$, where $\pi_Z, \pi_Y, \pi_M, \pi_{U_k}$ are valid openings of the corresponding derived commitments.

In the verification phase, the verifier first checks all opening proofs: $\mathsf{Verify}(\mathsf{pp}, C_Z^{(c)}, Z, \pi_Z) = 1$, $\mathsf{Verify}(\mathsf{pp}, C_Y^{(c)}, Y, \pi_Y) = 1$, $\mathsf{Verify}(\mathsf{pp}, C_M^{(c)}, M, \pi_M) = 1$, and, for all $k = 1, \ldots, m$, $\mathsf{Verify}(\mathsf{pp}, C_{U_k}^{(c)}, U_k, \pi_{U_k}) = 1$. Here $C_Z^{(c)}, C_Y^{(c)}, C_M^{(c)}, C_{U_k}^{(c)}$ denote the commitments obtained from $C_1, \ldots, C_4$ by the public linear-homomorphic derivation corresponding to challenge c. The verifier then checks the matrix identity $YZ = M + c^2 I_o$, and, for every $k = 1, \ldots, m$, the quadratic-form identity $\mathsf{Coeff}_z(Q_k(Zz)) = U_k$. The verifier accepts if and only if all checks pass.

4.3 Algebraic Identities

The protocol is based on two elementary identities. First, for matrices A, B, E, W of compatible dimensions and any $c \in \mathbb{F}_q$, $(E + cW)(A + cB) = EA + c(EB + WA) + c^2WB$. Second, for every homogeneous quadratic form Q_k, $Q_k((A + cB)z) = Q_k(Az) + c\mathsf{Pol}_{Q_k}(Az, Bz) + c^2Q_k(Bz)$. Thus, if $WB = I_o$ and $Q_k(Bz) \equiv 0$, then $(E + cW)(A + cB) = M_0 + cM_1 + c^2I_o$ and $\mathsf{Coeff}_z(Q_k((A + cB)z)) = U_{k,0} + cU_{k,1}$. These are exactly the identities checked by the verifier.

4.4 Completeness

Theorem 2 (Completeness). *If $(P; (B, W)) \in \mathcal{R}_{\mathrm{OS}}$, then an honest prover following the protocol is accepted by an honest verifier with probability* 1*.*

Proof. Let the prover follow the protocol honestly. By construction, $Z = A + cB$ and $Y = E + cW$. Therefore, $YZ = (E + cW)(A + cB) = EA + c(EB + WA) + c^2WB$. Since $M_0 = EA$, $M_1 = EB + WA$, $WB = I_o$, we obtain $YZ = M_0 + cM_1 + c^2I_o = M + c^2I_o$. Thus the verifier's matrix identity holds. For the quadratic checks, fix $k \in \{1, \ldots, m\}$. Since Q_k is homogeneous quadratic, $Q_k((A+cB)z) = Q_k(Az) + c\mathsf{Pol}_{Q_k}(Az, Bz) + c^2Q_k(Bz)$. Because (B, W) is a valid witness, $Q_k(Bz) \equiv 0$. Hence $\mathsf{Coeff}_z(Q_k(Zz)) = \mathsf{Coeff}_z(Q_k(Az)) + c\mathsf{Coeff}_z(\mathsf{Pol}_{Q_k}(Az, Bz))$. By definition of $U_{k,0}$ and $U_{k,1}$, $\mathsf{Coeff}_z(Q_k(Zz)) = U_{k,0} + cU_{k,1} = U_k$. Thus all quadratic checks pass. Finally, all derived commitments are opened with correctly derived randomness, so all opening verifications succeed by correctness of the commitment scheme. Therefore the verifier accepts with probability 1.

4.5 Computational 3-Special Soundness

We next prove that three accepting transcripts with the same first message and distinct challenges allow extraction of a valid witness.

Theorem 3 (Computational 3-special soundness). *Assume that the commitment scheme satisfies linear binding in the sense of Definition 4. Given three accepting transcripts $(\mathsf{com}, c_i, \mathsf{rsp}_i)$, $i = 1, 2, 3$, with the same first message com and pairwise distinct challenges $c_1, c_2, c_3 \in \mathcal{C}$, there exists a polynomial-time extractor that, except with negligible probability, outputs a witness (B, W) such that $(P; (B, W)) \in \mathcal{R}_{\mathrm{OS}}$.*

Proof. Let the three accepting responses contain $(Z_i, Y_i, M_i, \{U_{k,i}\}_{k=1}^m)$, $i = 1, 2, 3$. By linear binding of the commitment scheme, the commitments in the common first message $\mathsf{com} = (C_1, C_2, C_3, C_4)$ are bound to unique messages, except with negligible probability. Denote these messages by $X_1 = (B, W)$, $X_2 = (A, E)$, $X_3 = (M_0, M_1)$, $X_4 = (U_{1,0}, U_{1,1}, \ldots, U_{m,0}, U_{m,1})$. Again by linear binding, any valid opening of a homomorphically derived commitment for challenge c_i must open to the corresponding linear combination of these fixed committed blocks. Hence every accepting response satisfies $Z_i = A + c_iB$, $Y_i =$

$E + c_i W$, $M_i = M_0 + c_i M_1$, and $U_{k,i} = U_{k,0} + c_i U_{k,1}$ for every $k = 1, \ldots, m$, except with negligible probability.

The extractor computes $B := (c_2 - c_1)^{-1}(Z_2 - Z_1)$, $W := (c_2 - c_1)^{-1}(Y_2 - Y_1)$. This is well-defined because $c_1 \neq c_2$. Indeed, $Z_2 - Z_1 = (c_2 - c_1)B$ and $Y_2 - Y_1 = (c_2 - c_1)W$. It remains to prove that the extracted (B, W) satisfies the witness conditions.

Define the matrix-valued polynomial $F(c) := (E + cW)(A + cB) - (M_0 + cM_1) - c^2 I_o$. This polynomial has degree at most 2 in c. Since each transcript is accepting, the verifier's matrix check gives $Y_i Z_i = M_i + c_i^2 I_o$. Using the committed linear decompositions, $Y_i = E + c_i W$, $Z_i = A + c_i B$, $M_i = M_0 + c_i M_1$, we get $F(c_i) = 0$ for $i = 1, 2, 3$. Since $F(c)$ is a degree-at-most-2 polynomial over $\mathbb{F}_q$ and vanishes at three distinct points c_1, c_2, c_3, it is identically zero: $F(c) \equiv 0$. The coefficient of c^2 in $F(c)$ is $WB - I_o$. Therefore $WB = I_o$.

Fix $k \in \{1, \ldots, m\}$. Define the coefficient-vector-valued polynomial $G_k(c) := \mathsf{Coeff}_z(Q_k((A + cB)z)) - (U_{k,0} + cU_{k,1})$. Again, $G_k(c)$ has degree at most 2 in c, because Q_k is homogeneous quadratic. Since the i-th transcript is accepting, the verifier's quadratic check gives $\mathsf{Coeff}_z(Q_k(Z_i z)) = U_{k,i}$. Using $Z_i = A + c_i B$, $U_{k,i} = U_{k,0} + c_i U_{k,1}$, we obtain $G_k(c_i) = 0$ for $i = 1, 2, 3$. Hence $G_k(c) \equiv 0$. Expanding the quadratic form, $Q_k((A + cB)z) = Q_k(Az) + c\mathsf{Pol}_{Q_k}(Az, Bz) + c^2 Q_k(Bz)$. Therefore the coefficient of c^2 in $G_k(c)$ is $\mathsf{Coeff}_z(Q_k(Bz))$. Since $G_k(c) \equiv 0$, all its coefficients vanish, and hence $\mathsf{Coeff}_z(Q_k(Bz)) = 0$. Equivalently, $Q_k(Bz) \equiv 0$. Since this holds for every $k = 1, \ldots, m$, and since we have already shown $WB = I_o$, the extracted pair (B, W) satisfies $(P; (B, W)) \in \mathcal{R}_{\mathrm{OS}}$. The extractor runs in polynomial time and fails only if the binding property of the commitment scheme is violated.

4.6 Computational Honest-Verifier Zero Knowledge

We now prove honest-verifier zero knowledge. The proof uses only the hiding property of the commitment scheme and does not require equivocation.

Theorem 4 (Computational honest-verifier zero knowledge). *Assume that the commitment scheme is computationally hiding in the sense of Definition 2 for all message spaces used by the protocol. Then the protocol is computational honest-verifier zero knowledge.*

Proof. We construct a simulator $\mathcal{S}$ that, given the public input P and a challenge $c \in \mathcal{C}$, produces a transcript computationally indistinguishable from a real honest-verifier transcript conditioned on challenge c. On input P and c, the simulator proceeds as follows. First sample $B^\star \xleftarrow{\$} \mathbb{F}_q^{n \times o}$, $W^\star \xleftarrow{\$} \mathbb{F}_q^{o \times n}$, and independently sample $Z \xleftarrow{\$} \mathbb{F}_q^{n \times o}$, $Y \xleftarrow{\$} \mathbb{F}_q^{o \times n}$. Define $A^\star := Z - cB^\star$, $E^\star := Y - cW^\star$. Next sample $M_1^\star \xleftarrow{\$} \mathbb{F}_q^{o \times o}$. Set $M := YZ - c^2 I_o$, and define $M_0^\star := M - cM_1^\star$. For each $k = 1, \ldots, m$, sample a random coefficient vector $U_{k,1}^\star \xleftarrow{\$} \mathbb{F}_q^{o(o+1)/2}$. Set $U_k := \mathsf{Coeff}_z(Q_k(Zz))$, and define $U_{k,0}^\star := U_k - cU_{k,1}^\star$. The simulator now forms the blocks $X_1^\star := (B^\star, W^\star)$, $X_2^\star := (A^\star, E^\star)$, $X_3^\star := (M_0^\star, M_1^\star)$,

$X_4^\star := (U_{1,0}^\star, U_{1,1}^\star, \ldots, U_{m,0}^\star, U_{m,1}^\star)$, commits to them honestly, and derives the response commitments for challenge c. It outputs the first message $\mathsf{com}^\star = (C_1^\star, C_2^\star, C_3^\star, C_4^\star)$ and the response $\mathsf{rsp}^\star = (Z, Y, M, \{U_k\}_{k=1}^m, \pi_Z, \pi_Y, \pi_M, \{\pi_{U_k}\}_{k=1}^m)$, where the opening proofs are computed using the derived randomness.

By construction, $M = YZ - c^2 I_o$, so $YZ = M + c^2 I_o$. Also, for every k, $U_k = \mathsf{Coeff}_z(Q_k(Zz))$. Thus all algebraic verification checks pass. The opening proofs are valid by correctness of the commitment scheme. We compare the simulated transcript with a real transcript generated using an honest witness (B, W). In a real transcript conditioned on challenge c, the prover samples A, E uniformly and sends $Z = A + cB$, $Y = E + cW$. Since A and E are uniform and independent, Z and Y are also uniform and independent over $\mathbb{F}_q^{n\times o} \times \mathbb{F}_q^{o\times n}$. Thus the simulator's choice of uniform Z, Y has the same distribution. Furthermore, in a real accepting transcript, $M = YZ - c^2 I_o$ and $U_k = \mathsf{Coeff}_z(Q_k(Zz))$. These values are therefore distributed identically in the real and simulated transcripts.

It remains to compare the first-message commitments. In the real protocol, these commitments hide $(B, W), (A, E), (M_0, M_1), (U_{k,0}, U_{k,1})_{k=1}^m$. In the simulation, they hide $(B^\star, W^\star), (A^\star, E^\star), (M_0^\star, M_1^\star), (U_{k,0}^\star, U_{k,1}^\star)_{k=1}^m$. The derived response values Z, Y, M, U_k are distributed as in a real transcript, and all remaining differences appear only inside commitments. By computational hiding of the commitment scheme, replacing the real committed blocks by the simulated committed blocks changes the transcript distribution only negligibly. Therefore the simulator's output is computationally indistinguishable from a real honest-verifier transcript conditioned on challenge c. This proves computational honest-verifier zero knowledge.

4.7 Knowledge Error

Proposition 1 (Knowledge error). *Let $|\mathcal{C}| = \kappa$. Under the linear binding assumption of Definition 4, the knowledge error of the protocol is at most $2/\kappa$.*

Proof. Fix a first message com. If there exist accepting responses for three pairwise distinct challenges $c_1, c_2, c_3 \in \mathcal{C}$, then by Theorem 3 one can extract a valid witness (B, W). Therefore, for any first message from which no valid witness can be extracted, there can be accepting responses for at most two challenges in $\mathcal{C}$. Since the honest verifier samples the challenge uniformly from $\mathcal{C}$, the probability that such a first message is accepted is at most $2/|\mathcal{C}| = 2/\kappa$.

4.8 Soundness Amplification

Sequential repetition reduces the knowledge error in the standard way. If the protocol is repeated independently t times with fresh commitments and fresh challenges, then the knowledge error is at most $(2/|\mathcal{C}|)^t$. Sequential repetition preserves completeness and computational honest-verifier zero knowledge.

4.9 Fiat–Shamir Transform

The protocol can be made non-interactive by applying the Fiat–Shamir transform in the random oracle model. The challenge is computed as $c = H(\mathsf{pp}, P, \mathsf{com})$, or, for repeated executions, as a vector of challenges derived from the hash. A detailed proof of Fiat–Shamir security, especially in the quantum random oracle model, depends on the exact commitment instantiation and is outside the scope of this work. We therefore treat the interactive protocol as the primary object of analysis.

5 Motivation from Fixed-Vinegar Leakage

We now describe a concrete scenario that motivates the hidden oil-subspace relation. Standard UOV signing samples a fresh random vinegar value for each signature. If, contrary to the standard signing procedure, the same internal vinegar value is reused across many signatures, then the corresponding public signatures lie in a single affine coset of the hidden oil subspace. Taking differences of such signatures removes the affine offset and yields vectors in the hidden oil subspace.

This observation is not an attack on standard UOV, since standard UOV signing uses freshly sampled vinegar values. It is included only to illustrate that the hidden oil subspace is a meaningful structural secret.

5.1 Fixed Vinegar and an Affine Oil Coset

Let the secret affine transformation be

$$T^{-1}(x) = L^{-1}x + t,$$

where $L \in_n (\mathbb{F}_q)$ and $t \in \mathbb{F}_q^n$. In central coordinates, write an input as $(a\|u) \in \mathbb{F}_q^v \times \mathbb{F}_q^o$, where a is the vinegar part and u is the oil part.

Suppose that the signer incorrectly reuses the same vinegar value $a_0 \in \mathbb{F}_q^v$ in several signing executions. Then the corresponding central-coordinate preimages have the form $(a_0\|u_i)$, where the oil values $u_i \in \mathbb{F}_q^o$ depend on the message targets and on the solutions of the resulting linear systems. The public signatures are

$$\sigma_i = T^{-1}(a_0\|u_i) = L^{-1}(a_0\|u_i) + t.$$

Define the affine offset

$$x_0 := L^{-1}(a_0\|0) + t.$$

Then

$$\sigma_i - x_0 = L^{-1}(0\|u_i).$$

Therefore all signatures generated with the same vinegar value a_0 lie in the affine coset $x_0 + \mathcal{O}$, where

$$\mathcal{O} := L^{-1}(\{0\}^v \times \mathbb{F}_q^o)$$

is the hidden oil subspace.

5.2 Recovering Oil-Subspace Vectors from Signature Differences

An adversary does not know T^{-1} and cannot compute $T^{-1}(a_0\|u)$ on chosen oil values u. However, the adversary may observe public signatures. If two observed signatures σ_i, σ_j were generated with the same internal vinegar value a_0, then

$$\sigma_i - \sigma_j = \big(L^{-1}(a_0\|u_i) + t\big) - \big(L^{-1}(a_0\|u_j) + t\big) = L^{-1}(0\|u_i - u_j).$$

Hence

$$\sigma_i - \sigma_j \in \mathcal{O}.$$

Thus repeated use of the same vinegar value gives the adversary public vectors inside the hidden oil subspace.

Fix one reference signature σ_0. From signatures $\sigma_0, \sigma_1, \ldots, \sigma_t$ generated with the same vinegar value, the adversary can form

$$D_t := \{\sigma_i - \sigma_0 : 1 \le i \le t\}.$$

Then

$$D_t \subseteq \mathcal{O}.$$

Moreover, $\sigma_i - \sigma_0 = L^{-1}(0\|u_i - u_0)$, so the dimension of D_t is equal to the dimension of the span of the oil differences $u_i - u_0$ for $i = 1, \ldots, t$. Therefore, if these oil differences span $\mathbb{F}_q^o$, then $D_t = \mathcal{O}$. In that case, the adversary recovers the hidden oil subspace from public signature differences.

5.3 Spanning Probability

If the oil differences $u_i - u_0$ behave like independent uniformly random vectors in $\mathbb{F}_q^o$, then the probability that they span $\mathbb{F}_q^o$ is

$$\prod_{j=0}^{o-1}(1 - q^{j-t})$$

for $t \ge o$. Equivalently, after slightly more than o signatures generated with the same vinegar value, one expects the observed differences to span the hidden oil subspace with high probability.

This heuristic spanning estimate depends on the distribution of the oil solutions obtained during signing. It should not be interpreted as a statement about standard UOV security. In standard UOV, the vinegar value is freshly sampled for every signature, and signatures do not lie in a common affine oil coset.

5.4 Implications for the Hidden Oil-Subspace Relation

The fixed-vinegar scenario illustrates that the hidden oil subspace is a genuine structural secret. If enough signatures from a single vinegar coset are observed, then the subspace can be recovered by taking differences. Once a basis $B \in \mathbb{F}_q^{n\times o}$

of the recovered subspace is known, a left inverse $W \in \mathbb{F}_q^{o \times n}$ can be computed by Gaussian elimination, yielding a witness candidate for the hidden oil-subspace relation.

This does not mean that the protocol of Section 4 proves non-reuse of vinegar, nor does it prove signing capability. Rather, the example motivates why knowledge of the hidden oil subspace is a meaningful structural claim: it is a part of the UOV secret structure that can become exposed under faulty signing behavior.

6 Efficiency and Commitment Instantiation

6.1 Commitment Requirements

The protocol requires a linearly homomorphic commitment scheme satisfying correctness, computational hiding, and linear binding as defined in Section 2.4. The message spaces must support encoding of matrices and coefficient vectors over $\mathbb{F}_q$. The commitment scheme must allow public linear combinations of commitments and efficient opening proofs for derived commitments.

The security proof is modular in the commitment layer. Completeness uses correctness, honest-verifier zero knowledge uses computational hiding, and 3-special soundness uses linear binding. No equivocation trapdoor or programmability is required for the interactive protocol.

Classical Pedersen commitments provide a useful baseline for understanding the homomorphic structure, but care is needed when the UOV field $\mathbb{F}_q$ is not the same as the scalar field of the commitment group. For post-quantum security, lattice-based vector or homomorphic commitments are natural candidates. A full instantiation requires specifying the encoding of $\mathbb{F}_q$-vectors, the commitment randomness space, the derived-opening procedure, and concrete parameters ensuring hiding and binding. This work treats the commitment scheme as an abstract primitive and leaves optimized post-quantum instantiation to future work.

6.2 Communication Complexity

The first message consists of four commitments C_1, C_2, C_3, C_4. The response consists of matrices $Z \in \mathbb{F}_q^{n \times o}$, $Y \in \mathbb{F}_q^{o \times n}$, $M \in \mathbb{F}_q^{o \times o}$, and m coefficient vectors $U_k \in \mathbb{F}_q^{o(o+1)/2}$, together with opening proofs. The total algebraic response size is

$$2no + o^2 + m \cdot \frac{o(o+1)}{2}$$

field elements per round. For UOV-type parameters with $n = 156$, $o = 44$, $m = 44$, this gives

$$2 \cdot 156 \cdot 44 + 44^2 + 44 \cdot \frac{44 \cdot 45}{2} = 13728 + 1936 + 43560 = 59224$$

field elements. Over $\mathbb{F}_{256}$, this is approximately 59 KB of raw field data per round, before accounting for commitment and opening overhead. Commitment overhead depends on the concrete scheme; a naive componentwise commitment would multiply this size significantly, while more efficient vector commitments can reduce the overhead.

If the challenge set is $\mathcal{C} = \mathbb{F}_{256}$, the single-round knowledge error is $2/256 = 2^{-7}$. To achieve knowledge error approximately 2^{-128}, the protocol must be repeated sequentially $t = \lceil 128/7 \rceil = 19$ times. The total raw response data for 19 rounds is approximately $19 \cdot 59224 \approx 1.13 \times 10^6$ bytes (about 1.08 MiB) over $\mathbb{F}_{256}$, before commitment overhead. Sequential repetition also multiplies the number of commitments and opening proofs. The protocol is therefore not competitive with optimized post-quantum signature schemes in terms of communication efficiency. Its value lies in the structural certification it provides, not in raw efficiency.

6.3 Computational Complexity

We now estimate the algebraic cost of one protocol round, ignoring the cost of the commitment scheme. The dominant algebraic operation is the computation or verification of coefficient vectors of restricted quadratic forms.

Let $Q_k(x) = \sum_{1 \le i \le j \le n} q_{ij}^{(k)} x_i x_j$ be the homogeneous quadratic part of the k-th public polynomial. Given $Z \in \mathbb{F}_q^{n \times o}$, the verifier must compute the coefficient vector of $Q_k(Zz)$ as a quadratic form in $z \in \mathbb{F}_q^o$. If Q_k is represented densely, this can be done by forming the restricted quadratic form $Z^T H_k Z$, where H_k is a matrix representation of Q_k. With straightforward matrix multiplication, this costs $O(n^2 o + n o^2)$ field operations per quadratic form. Therefore, for m public quadratic forms, the verifier's quadratic-checking cost is $O\big(m(n^2 o + n o^2)\big)$ field operations per round in the dense representation.

Equivalently, if the public quadratic forms are stored as lists of nonzero monomials, then the cost can be expressed as $O\left(\sum_{k=1}^{m} \mathrm{nnz}(Q_k)\, o^2\right)$, where $\mathrm{nnz}(Q_k)$ denotes the number of nonzero quadratic monomials in Q_k. For dense public keys, $\mathrm{nnz}(Q_k) \approx \frac{n(n+1)}{2}$, which gives the same order as the dense estimate above.

The prover must compute, for each k, $U_{k,0} = \mathsf{Coeff}_z(Q_k(Az))$ and $U_{k,1} = \mathsf{Coeff}_z(\mathsf{Pol}_{Q_k}(Az, Bz))$. For dense public quadratic forms, computing $U_{k,0}$ has the same order as computing $Q_k(Zz)$, namely $O(n^2 o + n o^2)$ per k. Computing the polar term $U_{k,1}$ can similarly be done from restricted bilinear products involving A and B, with the same asymptotic cost. Hence the prover's algebraic cost is $O\big(m(n^2 o + n o^2)\big)$ field operations per round, up to small constant factors.

The matrix identity check $YZ = M + c^2 I_o$ costs $O(o^2 n)$ field operations and is smaller than the quadratic-form checks for dense public keys. The additional linear operations for forming $Z = A + cB$, $Y = E + cW$, $M = M_0 + cM_1$, $U_k = U_{k,0} + cU_{k,1}$ are lower order compared with the restriction of dense quadratic forms.

Thus, in the dense public-key setting, the per-round algebraic complexity of both prover and verifier is dominated by $O\big(m(n^2 o + n o^2)\big)$ field operations, plus the cost of producing and verifying commitment openings. The commitment

overhead is instantiation-dependent and may dominate the overall runtime for concrete implementations.

7 Applications and Limitations

The protocol should be understood as a zero-knowledge structural well-formedness proof for UOV-like MQ public keys. Its purpose is not to authenticate a signer or to prove signing capability. Rather, it certifies possession of the hidden oil-subspace invariant:

$$\exists (B, W): \qquad WB = I_o, \qquad Q_k(Bz) \equiv 0 \quad \forall k.$$

This invariant is a genuine part of the UOV secret structure. It is always present in honestly generated UOV keys, while random MQ maps almost never contain such an o-dimensional common zero subspace for UOV-type parameters. Moreover, in the random homogeneous UOV model, Theorem 1 shows that this subspace is generically unique.

The applications below are therefore best interpreted as uses of the protocol as a structural admissibility proof. The verifier learns that the submitted public MQ map contains the core UOV oil-subspace invariant and that the prover knows it, but learns nothing about the subspace itself. This is weaker than proving full honest key generation or signing capability, but it is stronger than merely declaring that a public key was generated by a UOV-like algorithm.

7.1 Case A: Rogue-Key Filtering in MQ Public-Key Registration

Consider a public-key registry, certificate authority, or protocol coordinator that accepts public MQ maps claimed to be UOV-like. A malicious registrant may submit an arbitrary MQ map, a malformed key, or a key generated by a non-standard procedure, while still claiming that it belongs to a UOV-like family. If subsequent protocols rely on the existence of the UOV oil-subspace invariant, then accepting arbitrary MQ maps is undesirable.

Our protocol provides a zero-knowledge structural admissibility test for such registries. During registration, the registrant submits the public MQ map

$$P = (p_1, \ldots, p_m)$$

together with a proof of knowledge of matrices (B, W) satisfying

$$WB = I_o, \qquad Q_k(Bz) \equiv 0 \quad \text{for all } k = 1, \ldots, m.$$

The verifier learns that P contains an o-dimensional common zero subspace of the public homogeneous quadratic forms, and that the registrant knows such a subspace. The subspace itself remains hidden.

This gives a first-line filter against arbitrary MQ substitutions. For a random MQ map, the expected number of such subspaces is

$$\begin{bmatrix} n \\ o \end{bmatrix}_q q^{-mo(o+1)/2} \approx q^{o(n-o)-mo(o+1)/2},$$

which is negligible for UOV-type parameters. Thus a random MQ key is extremely unlikely to pass the structural test. For random homogeneous UOV keys, Theorem 1 further shows that the witness subspace is generically the canonical hidden oil subspace.

This use case should be interpreted carefully. The proof does not certify that the key was generated honestly, nor does it exclude adversarially constructed keys with deliberately planted oil-like structure. It certifies only the existence and possession of the core oil-subspace invariant. Nevertheless, in a registration setting where the first goal is to rule out arbitrary or malformed MQ keys that lack the UOV structural invariant, this proof provides a concrete zero-knowledge well-formedness check.

7.2 Case B: A Layered Verifiable-Generation and Audit Pipeline

Full certification of a UOV public key is a stronger statement than the relation studied in this work. A complete proof of honest UOV key generation could require proving knowledge of an affine transformation, a central UOV map, and possibly signing capability. Such a monolithic proof may be substantially more complex and less efficient than a proof of a single structural invariant.

The hidden oil-subspace proof can instead serve as the first layer in a verifiable-generation or audit pipeline. The certification process may be decomposed into several increasingly strong layers:

Layer 1: knowledge of an oil-subspace witness;
Layer 2: compatibility with a central UOV representation;
Layer 3: knowledge of a full signing key or signing capability.

This work addresses Layer 1. It proves knowledge of an o-dimensional subspace on which all public homogeneous quadratic forms vanish. This is the defining quadratic invariant of the UOV oil space.

Such a layered approach is useful in settings where different verifiers require different assurance levels. A lightweight registry may only need to check that a submitted MQ map is structurally UOV-like. A stronger certification authority may additionally require evidence of compatibility with a full UOV central-map description. A high-assurance deployment may require a separate proof of signing capability. In this modular view, the protocol of this paper supplies the structural layer and can be combined with stronger relations when needed.

This also clarifies the role of the generic uniqueness theorem. In the random homogeneous UOV model, passing the Layer-1 proof means proving knowledge of the unique hidden oil subspace except with negligible probability. Therefore, although Layer 1 is weaker than full key-generation certification, it is not an arbitrary weak condition: it captures a rare and generically unique structural feature of UOV-like keys.

7.3 Case C: Privacy-Preserving Proof of Structural Leakage

The fixed-vinegar discussion in Section 5 suggests a forensic use case. Suppose an implementation incident is suspected: for example, a faulty signer may have

reused internal vinegar values across multiple signatures. As explained in Section 5, signatures generated with the same vinegar value lie in a common affine coset of the hidden oil subspace, and their differences reveal vectors inside that subspace.

An analyst may recover a candidate hidden oil subspace from observed signature differences. Directly disclosing this subspace may be undesirable: the subspace is a sensitive structural secret, and publishing it could further weaken the affected key. At the same time, the analyst may need to convince a vendor, auditor, or coordinating party that a genuine structural leakage has occurred.

Our protocol provides a privacy-preserving proof of possession of the recovered structure. The analyst proves knowledge of (B, W) such that

$$WB = I_o, \qquad Q_k(Bz) \equiv 0 \quad \forall k,$$

without revealing B, W, or the recovered subspace $\mathrm{Col}(B)$. This can serve as a zero-knowledge demonstration that the analyst possesses a valid oil-subspace witness for the public key.

This application has a limited but concrete scope. The proof does not by itself show that vinegar reuse occurred, nor does it identify which signatures shared a vinegar value. It only proves possession of a valid common zero subspace. Nevertheless, in coordinated disclosure, incident response, or forensic audits, this can be useful: it allows one to demonstrate knowledge of a sensitive structural leakage without releasing the leaked structure.

7.4 Limitations of the Structural Relation

The relation $\mathcal{R}_{\mathrm{OS}}$ is intentionally weaker than the full UOV secret-key relation. A witness (B, W) certifies knowledge of an o-dimensional subspace on which all public homogeneous quadratic forms vanish. It does not specify the complete affine transformation T^{-1}, the central UOV map F, or an efficient signing algorithm. Therefore, the protocol does not prove signing capability.

The protocol also does not prove honest key generation. An adversary may be able to construct public MQ maps that deliberately contain one or more oil-like subspaces while not being generated by the standard UOV key-generation algorithm. In such adversarial settings, a proof of membership in $\mathcal{R}_{\mathrm{OS}}$ should be interpreted only as a proof of knowledge of some valid common zero subspace, not necessarily as proof of knowledge of a unique canonical subspace.

The generic uniqueness theorem applies to the random homogeneous UOV model. It does not cover all non-homogeneous variants, structured coefficient distributions, or adversarially planted public keys. Thus the uniqueness interpretation is strongest when the key is assumed to come from the random homogeneous UOV model, or when the application only requires the structural admissibility property itself.

7.5 Practical Considerations

The protocol has substantial communication cost, especially after sequential repetition for small knowledge error. Its efficiency depends heavily on the chosen

homomorphic commitment scheme and on the cost of opening derived commitments. The present construction should therefore be viewed primarily as a structural proof-of-knowledge mechanism rather than as a compact authentication protocol.

A concrete deployment would require at least three additional ingredients: first, an efficient linearly homomorphic commitment instantiation over the required $\mathbb{F}_q$-vector spaces; second, a careful accounting of commitment and opening sizes; and third, if non-interactivity is desired, a Fiat–Shamir security analysis in the appropriate random-oracle or quantum-random-oracle model. Designing more efficient protocols for the same relation, and extending the relation toward full UOV secret-key knowledge, remain important directions for future work.

8 Conclusion

We have formalized the hidden oil-subspace relation for UOV public keys and constructed a Schnorr-like Sigma protocol for proving knowledge of a hidden oil-subspace witness. The protocol achieves computational 3-special soundness and computational honest-verifier zero knowledge. We proved a generic uniqueness theorem showing that in the random homogeneous UOV model, the hidden oil subspace is, with overwhelming probability, the unique o-dimensional common zero subspace. This result significantly strengthens the interpretation of the protocol: proving knowledge of an oil-subspace witness is essentially equivalent to proving knowledge of the unique hidden oil subspace. We analyzed the structural strength of the relation through random-MQ baseline counting, generic uniqueness bounds, local rigidity via linearized deformation equations, and empirical observations on small parameters. We carefully delineated the scope and limitations of our results, distinguishing what the protocol does and does not prove.

Several open problems remain. Extending the protocol to prove signing capability or full secret-key knowledge would significantly enhance its practical value. Constructing efficient post-quantum homomorphic commitments with linearly consistent binding is an important practical challenge. Analyzing the protocol's security under the Fiat–Shamir transform in the quantum random oracle model is left for future work. Extending the generic uniqueness theorem to non-homogeneous or adversarially constructed keys would broaden its applicability. Exploring applications in key registration, distributed key generation, and auditing protocols is a promising direction for future research.

References

1. J. Patarin. The oil and vinegar signature scheme. Presented at the Dagstuhl Workshop on Cryptography, 1997.
2. A. Kipnis, J. Patarin, and L. Goubin. Unbalanced oil and vinegar signature schemes. In *Advances in Cryptology – EUROCRYPT '99*, volume 1592 of *LNCS*, pages 206–222. Springer, 1999.

3. A. Kipnis and A. Shamir. Cryptanalysis of the oil and vinegar signature scheme. In *Advances in Cryptology – CRYPTO '98*, volume 1462 of *LNCS*, pages 257–266. Springer, 1998.
4. J. Ding and D. Schmidt. Rainbow, a new multivariable polynomial signature scheme. In *Applied Cryptography and Network Security – ACNS 2005*, volume 3531 of *LNCS*, pages 164–175. Springer, 2005.
5. W. Beullens. Breaking Rainbow takes a weekend on a laptop. In *Advances in Cryptology – CRYPTO 2022*, volume 13508 of *LNCS*, pages 464–479. Springer, 2022.
6. W. Beullens. MAYO: Practical post-quantum signatures from oil-and-vinegar maps. In *Selected Areas in Cryptography – SAC 2021*, volume 13203 of *LNCS*, pages 355–376. Springer, 2022.
7. Y. Ishai, E. Kushilevitz, R. Ostrovsky, and A. Sahai. Zero-knowledge from secure multiparty computation. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing – STOC 2007*, pages 21–30. ACM, 2007.
8. I. Giacomelli, J. Madsen, and C. Orlandi. ZKBoo: Faster zero-knowledge for Boolean circuits. In *USENIX Security Symposium 2016*, pages 1069–1083. USENIX Association, 2016.
9. J. Katz, V. Kolesnikov, and X. Wang. Improved non-interactive zero knowledge with applications to post-quantum signatures. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security – CCS 2018*, pages 525–537. ACM, 2018.
10. M. Chase, D. Derler, S. Goldfeder, C. Orlandi, S. Ramacher, C. Rechberger, D. Slamanig, and G. Zaverucha. Post-quantum zero-knowledge and signatures from symmetric-key primitives. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security – CCS 2017*, pages 1825–1842. ACM, 2017.
11. C. P. Schnorr. Efficient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, 1991.
12. R. Cramer, I. Damgård, and B. Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In *Advances in Cryptology – CRYPTO '94*, volume 839 of *LNCS*, pages 174–187. Springer, 1994.
13. I. Damgård. On Σ-protocols. Lecture notes, University of Aarhus, 2010.
14. A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *Advances in Cryptology – CRYPTO '86*, volume 263 of *LNCS*, pages 186–194. Springer, 1987.
15. D. Unruh. Non-interactive zero-knowledge proofs in the quantum random oracle model. In *Advances in Cryptology – EUROCRYPT 2015*, volume 9057 of *LNCS*, pages 755–784. Springer, 2015.
16. D. Unruh. Post-quantum security of Fiat–Shamir. In *Advances in Cryptology – ASIACRYPT 2017*, volume 10624 of *LNCS*, pages 65–95. Springer, 2017.
17. J. Don, S. Fehr, C. Majenz, and C. Schaffner. Security of the Fiat–Shamir transformation in the quantum random-oracle model. In *Advances in Cryptology – CRYPTO 2019*, volume 11693 of *LNCS*, pages 356–383. Springer, 2019.
18. Q. Liu and M. Zhandry. Revisiting post-quantum Fiat–Shamir. In *Advances in Cryptology – CRYPTO 2019*, volume 11693 of *LNCS*, pages 326–355. Springer, 2019.
19. T. P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In *Advances in Cryptology – CRYPTO '91*, volume 576 of *LNCS*, pages 129–140. Springer, 1992.

20. A. Kawachi, K. Tanaka, and K. Xagawa. Concurrently secure identification schemes based on the worst-case hardness of lattice problems. In *Advances in Cryptology – ASIACRYPT 2008*, volume 5350 of *LNCS*, pages 372–389. Springer, 2008.
21. C. Baum, I. Damgård, V. Lyubashevsky, S. Oechsner, and C. Peikert. More efficient commitments from structured lattice assumptions. In *Security and Cryptography for Networks – SCN 2018*, volume 11035 of *LNCS*, pages 368–385. Springer, 2018.
22. V. Lyubashevsky, N. K. Nguyen, and G. Seiler. Shorter lattice-based zero-knowledge proofs via one-time commitments. In *Public-Key Cryptography – PKC 2021*, volume 12710 of *LNCS*, pages 215–241. Springer, 2021.
23. R. Lidl and H. Niederreiter. *Finite Fields*, volume 20 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 2nd edition, 1997.
24. J. W. P. Hirschfeld. *Projective Geometries over Finite Fields*. Oxford University Press, 2nd edition, 1998.
25. J. Ding and B.-Y. Yang. Multivariate public key cryptography. In *Post-Quantum Cryptography*, pages 193–241. Springer, 2009.
26. C. Wolf and B. Preneel. Taxonomy of public key schemes based on the problem of multivariate quadratic equations. Cryptology ePrint Archive, Report 2005/077, 2005.
27. L. Bettale, J.-C. Faugère, and L. Perret. Hybrid approach for solving multivariate systems over finite fields. *Journal of Mathematical Cryptology*, 3(3):177–197, 2009.
28. J.-C. Faugère. A new efficient algorithm for computing Gröbner bases (F4). *Journal of Pure and Applied Algebra*, 139(1–3):61–88, 1999.
29. J.-C. Faugère. A new efficient algorithm for computing Gröbner bases without reduction to zero (F5). In *Proceedings of the 2002 International Symposium on Symbolic and Algebraic Computation – ISSAC 2002*, pages 75–83. ACM, 2002.
30. M. Bardet, J.-C. Faugère, and B. Salvy. On the complexity of Gröbner basis computation of semi-regular overdetermined algebraic equations. In *Proceedings of the International Conference on Polynomial System Solving – ICPSS 2004*, pages 71–74, 2004.
31. E. Thomae and C. Wolf. Solving underdetermined systems of multivariate quadratic equations revisited. In *Public Key Cryptography – PKC 2012*, volume 7293 of *LNCS*, pages 156–171. Springer, 2012.
32. A. Petzoldt, S. Bulygin, and J. Buchmann. Selecting parameters for the Rainbow signature scheme. In *Post-Quantum Cryptography – PQCrypto 2010*, volume 6061 of *LNCS*, pages 218–240. Springer, 2010.
33. A. Petzoldt, M.-S. Chen, B.-Y. Yang, C. Tao, and J. Ding. Design principles for HFEv- based multivariate signature schemes. In *Advances in Cryptology – ASIACRYPT 2015*, volume 9452 of *LNCS*, pages 311–334. Springer, 2015.
34. J. Ding, M.-S. Chen, A. Petzoldt, D. Schmidt, and B.-Y. Yang. *Rainbow*. NIST Post-Quantum Cryptography Standardization Round 3 Submission, 2020.
35. W. Beullens and B. Preneel. Field lifting for smaller UOV public keys. In *Progress in Cryptology – INDOCRYPT 2017*, volume 10698 of *LNCS*, pages 227–246. Springer, 2017.
36. K. Sakumoto, T. Shirai, and H. Hiwatari. Public-key identification schemes based on multivariate quadratic polynomials. In *Advances in Cryptology – CRYPTO 2011*, volume 6841 of *LNCS*, pages 706–723. Springer, 2011.
37. J. Stern. A new identification scheme based on syndrome decoding. In *Advances in Cryptology – CRYPTO '93*, volume 773 of *LNCS*, pages 13–21. Springer, 1994.
38. P. Véron. Improved identification schemes based on error-correcting codes. *Applicable Algebra in Engineering, Communication and Computing*, 8(1):57–69, 1997.

39. P.-L. Cayrel, R. Lindner, M. Rückert, and R. Silva. A lattice-based threshold ring signature scheme. In *Progress in Cryptology – LATINCRYPT 2010*, volume 6212 of *LNCS*, pages 255–272. Springer, 2010.
40. O. Goldreich. *Foundations of Cryptography: Volume 1, Basic Tools*. Cambridge University Press, 2001.
41. O. Goldreich. *Foundations of Cryptography: Volume 2, Basic Applications*. Cambridge University Press, 2004.