Scilight

Applied Mathematics and Statistics
https://www.sciltp.com/journals/ams



Article

Check for updates

Identifiability of Propagation Delay from Finite Time-Transfer Records

Angelina Hintsanen

Nexus Laboratory, Helsinki, Finland; angelina@nexus-ailab.org

How To Cite: Hintsanen, A. Identifiability of Propagation Delay from Finite Time-Transfer Records. *Applied Mathematics and Statistics* **2026**, *3*(2), 23. https://doi.org/10.53941/ams.2026.100023

Received: 25 August 2026
Revised: 14 September 2026
Accepted: 23 October 2026
Published: 27 October 2026

Abstract: Objectives: Finite time-transfer records can preserve several parameter decompositions that assign different values to propagation delay. This study determines which delay functionals a record identifies and which closure equations make the full delay vector unique. Methods: The finite record is written as $y = A\tau + B\kappa$. Record-preserving perturbations are projected onto the delay coordinates to form $\Pi_\tau(\ker J)$. Closure is analysed through its induced action on delay-changing equivalence classes, and a projected Moore–Penrose inverse quantifies sensitivity to closure error. Results: A linear delay functional is identifiable exactly when it annihilates the projected null space. Its identifiable subspace has codimension $\dim \Pi_\tau(\ker J)$. Full delay identification under closure is equivalent to $\Pi_\tau(\ker J \cap \ker G) = \{0\}$, with effective closure rank computed after zero-delay null directions are removed. The criterion is realised in a standard two-way transfer model, verified by synthetic operator tests, and traced through raw and residualised GNSS/RINEX records. Conclusion: Propagation delay has a unique record-based meaning when the target functional is invariant on the record-equivalence class. Closure can supply the information that creates this invariance; stable reconstruction alone does not establish it.

Keywords: structural identifiability; inverse problems; time transfer; propagation delay; closure constraints; GNSS

1. Introduction

A finite timing record may be compatible with more than one propagation latency. Recorded observations combine propagation delay with clock offsets, receiver and transmitter delays, cable or fibre delays, modem delays, calibration constants, environmental corrections, and synchronisation conventions. Several joint parameter vectors can therefore produce the same record and assign different values to the intended delay. The record has then failed to specify a unique latency measurand, so an ordinary uncertainty calculation has no unique record-based target.

Write a finite or linearised record as

$$y = A\tau + B\kappa,$$

where τ contains propagation-latency components and κ contains clock, hardware, calibration, and convention parameters. A perturbation $(\delta\tau, \delta\kappa)$ belongs to the record-preserving null space when it leaves the observations fixed. Its latency projection can remain nonzero. The projected null space $\Pi_\tau(\ker J)$ is the object that decides whether the finite record identifies the intended latency functional.

The analysis computes the record-preserving null space, projects it onto the latency coordinates, and tests the intended latency functional on that projection. A nonzero value along any projected direction proves non-identifiability from the record. Calibration, loop, gauge, or synchronisation closure succeeds when it removes every direction that changes the target. In some architectures, closure supplies the information that gives the latency a unique inverse-problem meaning.

Copyright: © 2026 by the authors. This is an open access article under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).
Publisher's Note: Scilight stays neutral about jurisdictional claims in published maps and institutional affiliations.

Structural uniqueness is assessed before conditioning and uncertainty propagation. Once a latency functional is invariant across the record-equivalence class, a GUM-style budget can quantify uncertainty for that specified measurand [1, 2]. When the functional varies across the class, a regularisation penalty, prior, convention, or calibration equation selects a value. Stability of that selected value does not establish record-based identification.

Structural identifiability asks whether parameters are uniquely determined by an input–output map before numerical estimation [3–8]. For a finite linear model, the kernel criterion for a target functional is the classical estimable-function condition [9, 10]. Applied to the latency block, it gives the identifiable subspace $\Pi_\tau(\ker J)^\perp$ and its codimension. The contribution of this paper is the closure analysis on the quotient of record-null directions by zero-latency directions. Proposition 2 and Theorem 2 give the necessary and sufficient effective-rank condition. Proposition 3 gives the condition for a specified ambiguity class, and Corollary 2 bounds sensitivity to closure error through the projected Moore–Penrose inverse [11, 12].

Classical inverse-problem theory supplies the treatment of nonuniqueness, rank deficiency, ill-posedness, and regularised solution selection [13–18]. The present question concerns a specified component: the full parameter vector may remain nonunique even when one latency functional is identifiable. The projected null space isolates that component-level question and exposes the closure information relevant to it.

Time-transfer practice supplies the physical setting. Reviews, standards, and operational methods for GNSS common-view, precise point positioning, two-way satellite time and frequency transfer, RINEX records, clock products, calibrated receiver or modem delays, and traceability to UTC show that timing records are interpreted through modelling and calibration layers [19–28]. An ideal reciprocal two-way exchange separates clock offset from symmetric path delay. Directional equipment delays and conventions add unknowns to the same equations. GNSS/RINEX records provide a public-data setting for receiver-clock and calibration gauge freedom. These examples illustrate the projected-null-space mechanism; the proof rests on the operator statements.

2. Methods

2.1. Forward Operator, Record Equivalence, and Identifiability

2.1.1. Operator Model and Notation

Let $y \in \mathbb{R}^m$ denote a finite vector of timing-transfer observations. Write the unknown parameter vector as

$$\theta = (\tau, \kappa) \in \Theta_\tau \times \Theta_\kappa,$$

where $\tau \in \mathbb{R}^p$ denotes propagation-latency components of interest and $\kappa \in \mathbb{R}^q$ collects clock offsets, instrumental delays, calibration constants, and other nuisance parameters. A finite transfer architecture is represented by an observation operator

$$F : \Theta_\tau \times \Theta_\kappa \to \mathbb{R}^m, \qquad y = F(\tau, \kappa) + \varepsilon,$$

with ε denoting measurement noise. In the linearised case around a nominal point θ_0,

$$\delta y = J_\tau \delta\tau + J_\kappa \delta\kappa + \varepsilon = J\delta\theta + \varepsilon, \qquad J = [J_\tau \; J_\kappa] \in \mathbb{R}^{m\times(p+q)}.$$

Throughout, $L(\tau) = \ell^T\tau$ denotes a linear latency functional with coefficient vector $\ell \in \mathbb{R}^p$, and $\Pi_\tau : \mathbb{R}^{p+q} \to \mathbb{R}^p$ denotes the natural projection onto the latency block.

Definition 1 (Record equivalence). *Two parameter values* (τ, κ) *and* (τ', κ') *are* record-equivalent *if* $F(\tau, \kappa) = F(\tau', \kappa')$. *The equivalence class is*

$$[\tau, \kappa]_F = \{(\tau', \kappa') : F(\tau', \kappa') = F(\tau, \kappa)\}.$$

Definition 2 (Structural identifiability). *A latency functional* $L(\tau)$ *is* structurally identifiable *from records generated by* F *if*

$$F(\tau, \kappa) = F(\tau', \kappa') \quad \Longrightarrow \quad L(\tau) = L(\tau')$$

for all admissible (τ, κ) *and* (τ', κ').

Definition 3 (Closure constraint). *A* closure constraint *is a linear condition* $G(\delta\tau, \delta\kappa) = 0$, *equivalently* $G_\tau\delta\tau + G_\kappa\delta\kappa = 0$, *that restricts the admissible perturbation set. The constrained null space is*

$$\mathcal{N}_G = \ker J \cap \ker G.$$

2.1.2. Physical Meaning of Nuisance Parameters and Closure

In the notation $y = A\tau + B\kappa$, the nuisance vector κ represents physical and conventional components that enter the finite record together with propagation latency: local clock offsets, receiver and transmitter hardware delays, modem delays, antenna-terminal delays, cable and fibre delays, timestamping offsets, reference-clock distribution delays, atmospheric or environmental correction parameters, and synchronisation or gauge conventions. Which of these terms is treated as known, estimated, calibrated, or fixed by convention determines the actual measurand realised by the record model.

A closure constraint is therefore the mathematical representation of metrological traceability information. It may come from an independently calibrated delay, a loop closure, a reciprocal two-way assumption, a hardware calibration campaign, a clock product, a common-clock experiment, or a stated synchronisation convention. In operator form, augmenting $y = A\tau + B\kappa$ by $G(\tau, \kappa) = g$ restricts the admissible record-equivalence class. Closure is effective for a latency measurand only when it removes the latency-changing component of the record-preserving null space.

This interpretation prevents a common category error. Calibration is often described as reducing the uncertainty of an already-defined quantity. In the setting studied here, calibration or closure may be required before the latency measurand is uniquely defined. If $L(\tau)$ changes along a record-preserving direction, then no uncertainty budget for L is complete until the closure information selecting a representative of that equivalence class has been stated.

2.1.3. Identifiable Latency Measurands

The target of identification is a linear functional of the latency vector. The following criterion tests its invariance on the projected record-null space.

Theorem 1 (Kernel criterion for local latency identifiability). *In the linearised model $\delta y = J\delta\theta$, a linear latency functional $L(\delta\tau) = \ell^T\delta\tau$ is locally identifiable if and only if*

$$\ell \perp \mathcal{N}_\tau, \qquad \mathcal{N}_\tau := \Pi_\tau(\ker J),$$

that is, $\ell^T v = 0$ for every $v \in \mathcal{N}_\tau$.

Proof. Let $\mathcal{N}_\tau = \Pi_\tau(\ker J)$. If L is locally identifiable, then for every $(v, w) \in \ker J$ the record-equivalent perturbation $(\delta\tau, \delta\kappa) = (v, w)$ satisfies $\delta y = 0$; identifiability requires $\ell^T v = 0$. Hence $\ell \perp \mathcal{N}_\tau$.

Conversely, suppose $\ell \perp \mathcal{N}_\tau$. Take any two admissible parameter pairs with $J(\delta\tau, \delta\kappa)^T = 0$, so $(\delta\tau, \delta\kappa) \in \ker J$. Then $\delta\tau \in \mathcal{N}_\tau$, so $\ell^T\delta\tau = 0$, giving $L(\tau_1) = L(\tau_2)$. Local identifiability follows.

□

The following proposition is the algebraic version of a metrological statement: a latency measurand is determined by a finite record only if it has the same value for all parameter decompositions that produce that record. Directions in $\Pi_\tau(\ker J)$ are latency changes that the record cannot distinguish. Identifiable latency measurands are precisely the linear combinations that are insensitive to those changes.

Proposition 1 (Identifiable latency subspace). *The set of locally identifiable linear latency functionals is the subspace*

$$\mathcal{I} = \mathcal{N}_\tau^\perp = \{\ell \in \mathbb{R}^p : \ell^T v = 0 \ \textit{for all} \ \ v \in \mathcal{N}_\tau\},$$

with

$$\dim(\mathcal{I}) = p - \dim(\mathcal{N}_\tau).$$

The codimension of the identifiable subspace is therefore the dimension of the latency-projection of the null space. An explicit basis for $\mathcal{I}$ is obtained from any orthonormal basis of $(\mathcal{N}_\tau)^\perp$ in $\mathbb{R}^p$, computable via the SVD of J restricted to the latency block.

Proof. By Theorem 1, ℓ is identifiable iff $\ell \in \mathcal{N}_\tau^\perp$. The dimension formula follows from the identity $\dim(\mathcal{N}_\tau^\perp) = p - \dim(\mathcal{N}_\tau)$ in finite dimensions. For the computational claim: partition the SVD $J = U\Sigma V^T$ and extract the right singular vectors corresponding to zero singular values; project their τ-components to obtain a basis for $\mathcal{N}_\tau$. Its orthogonal complement in $\mathbb{R}^p$ is $\mathcal{I}$.

□

Remark 1 (Special cases). *If* $\ker J \cap (\mathbb{R}^p \times \{0\}) = \{0\}$*—that is, no null vector has nonzero latency component and zero nuisance component—then* $\mathcal{N}_\tau$ *may still be nontrivial because nuisance directions can mix into latency through the null space. The identifiable subspace can be strictly smaller than* $\mathbb{R}^p$ *even when* J_τ *alone has full column rank.*

2.1.4. Closure Rank Condition

Closure restricts the record-preserving directions. Its effect on the latency projection determines the remaining identifiability condition.

Proposition 2 (Closure-rank condition for identifiability). *Let*

$$\mathcal{N} = \ker J, \qquad P = \Pi_\tau|_{\mathcal{N}} : \mathcal{N} \to \mathbb{R}^p, \qquad R = G|_{\mathcal{N}} : \mathcal{N} \to \mathbb{R}^r.$$

Under the closure constraint $G(\delta\tau, \delta\kappa) = 0$*, a latency functional* $L(\delta\tau) = \ell^T \delta\tau$ *is identifiable if and only if*

$$\ell \perp P(\ker R).$$

Full latency identifiability is achieved if and only if

$$P(\ker R) = \{0\}, \qquad \textit{equivalently} \qquad \ker R \subseteq \ker P.$$

To express the effective rank of the closure information, quotient out null directions that do not change latency. Define

$$\widetilde{R} : \mathcal{N}/\ker P \longrightarrow \mathbb{R}^r / R(\ker P), \qquad \widetilde{R}([n]) = [Rn].$$

This map is well defined because replacing n *by* $n + k$ *with* $k \in \ker P$ *changes* Rn *only by an element of* $R(\ker P)$*. Full latency identifiability is equivalent to injectivity of* $\widetilde{R}$*. Since*

$$\dim(\mathcal{N}/\ker P) = \dim P(\mathcal{N}) = \dim(\mathcal{N}_\tau),$$

full closure requires the effective closure rank

$$\operatorname{rank} \widetilde{R} = \dim(\mathcal{N}_\tau).$$

Thus closure information that is spent only on zero-latency null directions does not remove latency ambiguity; only the rank remaining after quotienting by $R(\ker P)$ *is structurally effective for the latency measurand. A direction* $k \in \ker P$ *satisfies* $Pk = 0$*. Moving along* k *changes nuisance or calibration terms and leaves the latency vector unchanged. A closure constraint that removes only these directions reduces nuisance-parameter ambiguity. It supplies no information that separates record-equivalent solutions with different latency vectors.*

Proof. Closure restricts admissible record-preserving perturbations from $\mathcal{N}$ to $\mathcal{N} \cap \ker G = \ker R$. A functional $L(\delta\tau) = \ell^T \delta\tau$ is identifiable under closure exactly when it vanishes on every latency component generated by such perturbations, i.e., when $\ell \perp P(\ker R)$. Full latency identifiability requires every constrained null perturbation to have zero latency projection, which is $P(\ker R) = \{0\}$, or equivalently $\ker R \subseteq \ker P$.

The quotient map $\widetilde{R}$ records only closure information that distinguishes latency-changing equivalence classes. If $\widetilde{R}$ is injective and $n \in \ker R$, then $\widetilde{R}([n]) = 0$, so $[n] = 0$ and hence $n \in \ker P$; therefore $\ker R \subseteq \ker P$. Conversely, if $\ker R \subseteq \ker P$ and $\widetilde{R}([n]) = 0$, then $Rn = Rk$ for some $k \in \ker P$. Hence $n - k \in \ker R$, so $n - k \in \ker P$, and therefore $[n] = 0$. Thus $\widetilde{R}$ is injective. The rank condition follows from rank-nullity on the finite-dimensional quotient space. □

2.1.5. Minimum Effective Closure Rank

The relevant closure rank is measured on the latency-changing ambiguity quotient, after zero-latency null directions are removed.

Proposition 3 (Minimum effective closure rank for ambiguity removal). *Let* $\mathcal{N} = \ker J$*, let* $P = \Pi_\tau|_{\mathcal{N}}$*, and let* $d = \dim P(\mathcal{N})$*. A closure matrix* G *achieves full latency identifiability if and only if the effective quotient map*

$$\widetilde{R} : \mathcal{N}/\ker P \to \mathbb{R}^r / R(\ker P), \qquad \widetilde{R}([n]) = [Gn],$$

has rank d. The effective quantity is the rank of the closure information that remains after closure action on zero-latency null directions has been quotiented out. The raw number of closure rows does not determine that rank. Rows acting only on $\ker P$ *can constrain nuisance or calibration components. Their induced action on* $\mathcal{N}/\ker P$ *is zero, so they cannot remove a latency-changing equivalence class.*

A specified latency-changing equivalence class $[n_*] \in \mathcal{N}/\ker P$ *is removed by* G *if and only if*

$$Gn_* \notin G(\ker P).$$

For a single scalar closure row $g^T n = 0$*, this reduces to the condition that the row vanishes on the zero-latency null space and acts nontrivially on the specified latency-changing class:*

$$g^T k = 0 \quad \textit{for all } k \in \ker P, \qquad g^T n_* \neq 0.$$

If $\ker P = \{0\}$*, the condition reduces to* $g^T n_* \neq 0$*.*

Proof. The rank statement is Proposition 2. For a specified class $[n_*]$, the class survives closure precisely when there is some $k \in \ker P$ such that $n_* - k \in \ker G$, equivalently $Gn_* = Gk$. Thus the class is removed exactly when $Gn_* \notin G(\ker P)$. For a scalar row, $g(\ker P)$ is either $\{0\}$ or all of $\mathbb{R}$. A scalar row can therefore remove the specified class only when it is zero on $\ker P$ and nonzero on n_*. If $\ker P = \{0\}$, there is no zero-latency null subspace to quotient out, so the condition becomes $g^T n_* \neq 0$. □

Corollary 1 (Common-mode ambiguity removal by a single constraint). *Suppose the latency-changing ambiguity quotient* $\mathcal{N}/\ker P$ *is one-dimensional and is represented by* $n_* = (v_*, w_*)$ *with* $v_* \neq 0$*. A scalar closure constraint* $g^T n = 0$ *suffices to make every latency functional identifiable if the row annihilates zero-latency null directions and acts nontrivially on the latency-changing class:*

$$g^T k = 0 \quad (k \in \ker P), \qquad g^T n_* \neq 0.$$

In architectures where $\ker P = \{0\}$*, this reduces to the simpler nonzero-action condition* $g^T n_* \neq 0$*.*

2.1.6. Stability under Approximate Closure

An approximate closure equation can change the reconstructed latency. The next bound controls this change on the identifiable subspace.

Corollary 2 (Stability under approximate closure). *Let*

$$M_G = \begin{pmatrix} J \\ G \end{pmatrix}$$

denote the augmented finite operator and let $\eta = (\delta y, \delta g)$ *be the joint observation–closure residual. Choose the orthogonal complement*

$$\mathcal{S} = (\ker M_G)^{\perp}$$

and denote the restriction of M_G *to* $\mathcal{S}$ *by* $M_{G,\mathcal{S}}$*. An SVD of* M_G *supplies* $\mathcal{S}$*: its basis consists of the right singular vectors associated with singular values above the stated numerical-rank tolerance. The restricted operator is injective, and* $\sigma_{\min}(M_{G,\mathcal{S}}) = \sigma^{+}_{\min}(M_G)$*, the smallest nonzero singular value of* M_G*. If* $\sigma_{\min}(M_{G,\mathcal{S}}) > 0$*, then the identifiable latency perturbation obeys*

$$\|\Pi_\tau \delta\theta_{\mathrm{id}}\| \leq \|\Pi_\tau M^{\dagger}_{G,\mathcal{S}}\| \, \|\eta\| \leq \frac{\|\Pi_\tau\|}{\sigma_{\min}(M_{G,\mathcal{S}})} \, \|\eta\|.$$

In the special case where the record residual is zero, and only closure is uncertain, $\|\eta\| = \|\delta g\| \leq \varepsilon$*, giving the corresponding closure-error bound. This choice fixes the Euclidean minimum-norm identifiable component. The bound then depends on the nonzero singular spectrum of* M_G*, with no arbitrary complement left to select. A non-orthogonal complement changes the norm of the restricted inverse and can change the numerical constant in the bound.*

Proof. On the identifiable complement $\mathcal{S}$ the restricted augmented map $M_{G,\mathcal{S}}$ is injective and has Moore–Penrose

inverse with operator norm $1/\sigma_{\min}(M_{G,\mathcal{S}})$; see, e.g., standard generalized-inverse and finite-dimensional inverse-problem perturbation estimates in [11, 12, 15, 16]. Applying this inverse to η and then projecting onto latency coordinates gives the first inequality. The singular-value bound gives the second.

□

Corollary 3 (Regularisation selects a representative, not an identified latency). *If $L(\tau) = \ell^T\tau$ is not structurally identifiable, then there exist admissible record-equivalent parameter values (τ_1, κ_1) and (τ_2, κ_2) such that*

$$J(\tau_1, \kappa_1)^T = J(\tau_2, \kappa_2)^T, \qquad L(\tau_1) \neq L(\tau_2).$$

Any deterministic estimator that depends only on the finite record must return a single value for that record and therefore cannot equal both latency values. Regularisation can select a stable representative of the record-equivalence class. It cannot make a non-invariant latency functional identifiable from the record alone.

Proof. Non-identifiability means there is a record-preserving perturbation with nonzero value under L, yielding the two displayed admissible parameter values. Since the finite record is the same for both, any deterministic map from records to estimates returns the same output for both parameter values. A Tikhonov estimator such as $\hat{\theta}_\lambda = \arg\min_\theta\{\|y - J\theta\|^2 + \lambda\|\theta\|^2\}$ selects a particular representative determined by the penalty and parameterisation; in the usual rank-deficient linear limit this tends toward a minimum-norm generalized-inverse representative. A different penalty or convention can select a different representative. This is regularised selection within an equivalence class, not structural uniqueness of L on that class [14–17].

□

2.1.7. Admissible Null-Direction Range in the Linear Box

Admissible parameter bounds restrict how far a record-equivalent perturbation can move along a null direction.

Proposition 4 (Admissible range along a latency-changing null direction). *Let $n_* = (v_*, w_*) \in \ker J$ be a latency-changing null direction and let the admissible parameter set be the box $[\tau_{\min}, \tau_{\max}] \times [\kappa_{\min}, \kappa_{\max}]$. For a particular solution (τ_0, κ_0), admissible motion along n_* is the interval*

$$I(n_*; \theta_0) = \{a \in \mathbb{R} : \theta_0 + an_* \in [\tau_{\min}, \tau_{\max}] \times [\kappa_{\min}, \kappa_{\max}]\}.$$

For any scalar latency component τ_j, the range of that component along this record-equivalent family is

$$\Delta\tau_j = |v_{*,j}| \operatorname{length} I(n_*; \theta_0),$$

with

$$\operatorname{length} I(n_*; \theta_0) \leq \min\left(\min_{v_{*,j}\neq 0} \frac{\tau_{j,\max} - \tau_{j,\min}}{|v_{*,j}|}, \min_{w_{*,i}\neq 0} \frac{\kappa_{i,\max} - \kappa_{i,\min}}{|w_{*,i}|}\right).$$

Thus the reported grid-family count S_R should be interpreted as a discretised diagnostic of admissible range over a chosen grid, not as a canonical invariant volume of J.

Proof. The equivalence class through a particular solution contains the affine line $\theta_0 + an_*$. Intersecting this line with the admissible box gives an interval in the scalar parameter a, obtained by intersecting the componentwise admissible intervals. The latency component τ_j varies by $|v_{*,j}|$ times the interval length, giving the formula and the bound.

□

2.2. Finite Time-Transfer Record Model

2.2.1. Mathematical Assumptions on the Operator Family

The following assumptions define the operator class studied in this paper.

Assumption 1 (Bounded linear observation operator). *The observation operator F is bounded and linear: $F(\tau, \kappa) = J_\tau\tau + J_\kappa\kappa$, with $J = [J_\tau \; J_\kappa] \in \mathbb{R}^{m\times(p+q)}$. Nonlinear forward models are admitted provided they are Fréchet-differentiable at the nominal point, in which case all results apply to the linearisation.*

Section 4.3 states what this local result establishes for a nonlinear forward model.

Assumption 2 (Nontrivial record-equivalence structure). *The nontrivial case studied in the paper is the case in which the observation operator has record-preserving null directions,* $\ker J \neq \{0\}$*, and the latency ambiguity is governed by* $\mathcal{N}_\tau = \Pi_\tau(\ker J)$*. Full latency identifiability corresponds to the special case* $\mathcal{N}_\tau = \{0\}$*. The analysis does not require an underdetermined dimension count such as* $m < p+q$*; rank deficiency and the latency projection of the null space are the relevant conditions.*

Assumption 3 (Admissible parameter set). *The admissible set* $\Theta_\tau \times \Theta_\kappa \subseteq \mathbb{R}^{p+q}$ *is a convex compact box. Positivity constraints (e.g. non-negative propagation times) enter as box bounds.*

Assumption 4 (Closure constraints are linear). *Auxiliary closure constraints take the form* $G_\tau \tau + G_\kappa \kappa = g$ *for known matrices* $G_\tau \in \mathbb{R}^{r\times p}$*,* $G_\kappa \in \mathbb{R}^{r\times q}$ *and vector* $g \in \mathbb{R}^r$*. Nonlinear constraints are admitted after local linearisation.*

2.2.2. Discretisation of Integral Measurement-Chain Models

Physical timing chains are often modelled by integral operators. The following lemma connects those models to the finite linear operator studied here.

Lemma 1 (Discretisation of integral timing-chain operator). *Let the physical record be* $R_i(t) = \int K_i(t,s)\, P_\tau(E)(s)\, ds$ *for measurement- chain kernel* $K_i(t,s)$ *and propagated signal* $P_\tau(E)$*. Suppose* K_i *and* $P_\tau(E)$ *are square-integrable over* $[0,T]$*, and that sampling at epochs* $t_1,\dots,t_m$ *yields the finite record vector* $y_j = R_i(t_j)$*. If* P_τ *is affine in* (τ,κ)*, then the discrete map* $(\tau,\kappa) \mapsto y$ *is linear with Jacobian*

$$J_{j,\ell} = \int K_i(t_j,s)\, \frac{\partial P_\tau(E)(s)}{\partial \theta_\ell}\, ds,$$

and all results of Section 2.1 apply to J*.*

Proof. Linearity of the map $(\tau,\kappa) \mapsto y$ follows from the assumed affine dependence of P_τ on parameters. The Jacobian formula follows by differentiating under the integral sign, which is valid since $K_i \in L^2([0,T]^2)$ and the derivative is uniformly bounded on the admissible set.

□

Remark 2. *In practice, the integral kernel* $K_i(t,s)$ *represents detector response, bandwidth filtering, and timestamp extraction. Lemma 1 does not claim that finite sampling preserves every null direction of the continuous operator. Sampling can remove, merge, or introduce null directions. The theorem therefore analyzes the finite sampled operator* J *actually used to form the record; the continuous formulation supplies a modelling route to* J*, not an additional identifiability claim.*

2.2.3. Finite Light-Transfer Record Equation

In what follows, identifiability is used in the structural sense of [3–8, 18]. Let $y \in \mathbb{R}^m$ denote recorded timing observations, $\tau \in \mathbb{R}^p$ propagation intervals, and $\kappa \in \mathbb{R}^q$ endpoint timing terms: clock bias, hardware delay, synchronisation offsets, and calibration residuals.

A finite light-transfer protocol under Assumption 1 is modelled as

$$y = A\tau + B\kappa,$$

with $A \in \mathbb{R}^{m\times p}$ and $B \in \mathbb{R}^{m\times q}$. The notation (A,B) is the specialisation of (J_τ, J_κ) to the timing-record context.

Definition 4 (Finite light-transfer propagation identifiability). *Propagation* τ *is identifiable from* $y = A\tau + B\kappa$ *if* $A\tau_1 + B\kappa_1 = A\tau_2 + B\kappa_2 \Rightarrow \tau_1 = \tau_2$ *for all admissible* $(\tau_1,\kappa_1),(\tau_2,\kappa_2)$*. Non-identifiability occurs iff there exists* $(\Delta\tau,\Delta\kappa) \neq 0$ *with* $\Delta\tau \neq 0$ *and* $A\Delta\tau + B\Delta\kappa = 0$*.*

Definition 5 (Reconstruction equivalence class). *For record* y*, the admissible reconstruction set is*

$$\mathcal{A}(y) = \{(\tau,\kappa) : A\tau + B\kappa = y,\ (\tau,\kappa) \in \Theta_\tau \times \Theta_\kappa\}.$$

Equivalently, for any particular solution (τ_0,κ_0) *satisfying* $A\tau_0 + B\kappa_0 = y$*,*

$$\mathcal{A}(y) = \big[(\tau_0,\kappa_0) + \ker[A\ B]\big] \cap (\Theta_\tau \times \Theta_\kappa).$$

Elements of $\mathcal{A}(y)$ are propagation–calibration decompositions consistent with the same record.

The grid-family size $N_R = |\mathcal{A}(y)|$ (for a discretised admissible set) and its logarithm $S_R = \log N_R$ serve throughout as a *descriptive diagnostic* of equivalence class size over a fixed tested parameter grid. S_R depends on the grid resolution and the box bounds; it is not an intrinsic measure of the operator's null structure. The admissible-range characterisation is given by Proposition 4.

Remark 3 (Interpretation of the grid-family diagnostic). *$S_R = \log N_R$ is the logarithm of a finite count on a specified grid. Grid spacing and box bounds determine its value. Changing either can change S_R while the operator J, its null space, and its identifiable subspace remain fixed. The calculation contains no probability model and no continuous measure. Reading S_R as entropy or ambiguity volume would give it physical content absent from the calculation. Values can be compared only when the same grid construction and admissible bounds are used.*

2.3. Computational Implementation

The numerical checks were run with Python 3. NumPy supplied array operations, singular-value decompositions, matrix ranks, null-space calculations, and Moore–Penrose inverses. pandas handled tabular inputs and generated CSV outputs. Matplotlib produced the figures. The RINEX analyses used GeoRinex with xarray and netCDF4; unlzw3 decompressed Unix-compressed files, and requests retrieved public records. Numerical-rank tolerances and all fixed-grid bounds are recorded in the scripts and machine-readable summaries supplied in Supplementary Material S1 and the archived reproducibility package. The reported grid counts serve solely as deterministic bookkeeping outputs.

Scale parameter columns to comparable numerical magnitudes before computing ranks. For $K \in \{J, M_G\}$ with dimensions $a \times b$, use $\epsilon_{\mathrm{mach}} \max(a,b)\sigma_{\max}(K)$ as the baseline tolerance for exact double-precision matrices; for an estimated noisy K, set the threshold above an independent operator-perturbation estimate. Repeat the calculation across a documented tolerance range. A direction that changes null status across that range has unresolved identifiability at the available precision.

3. Results

3.1. Finite Transfer-Record Closure Theorem

The finite-record result is stated in terms of null directions that change latency while preserving the observations.

Theorem 2 (Finite transfer-record closure theorem). *Under Assumptions 1–4, consider the linear transfer-record model $y = A\tau + B\kappa$.*

Part (i) — records alone. *At interior points of the admissible set, or for operator-level structural identifiability on the affine parameter space, τ is identifiable from the finite record if and only if*

$$\ker[A\ B] \cap \{(\Delta\tau, \Delta\kappa) : \Delta\tau \neq 0\} = \varnothing.$$

For a bounded admissible box, the same condition is sufficient globally; necessity holds whenever a latency-changing null direction admits a nonzero admissible motion from the parameter point under consideration.

Part (ii) — records plus closure. *Under additional closure constraints $G_\tau\tau + G_\kappa\kappa = g$, τ is identifiable if and only if the augmented map*

$$M_G = \begin{pmatrix} A & B \\ G_\tau & G_\kappa \end{pmatrix}$$

has no null vector $(\Delta\tau, \Delta\kappa)$ with $\Delta\tau \neq 0$.

Part (iii) — local vs. affine-global. *For the linear model F (Assumption 1), the null-space condition is both local and global on the affine parameter space. Box constraints can truncate admissible motion, as noted in the admissibility remark below.*

Part (iv) — necessary and sufficient condition on G. *Closure constraints G achieve full latency identifiability if and only if*

$$\Pi_\tau(\ker J \cap \ker G) = \{0\},$$

where $\ker J = \ker[A\ B]$ (Proposition 2). Equivalently, if $P = \Pi_\tau|_{\ker J}$ and $R = G|_{\ker J}$, the quotient map

$$\widetilde{R} : \ker J / \ker P \to \mathbb{R}^r / R(\ker P), \qquad \widetilde{R}([n]) = [Rn],$$

must be injective, or equivalently must have effective rank $\dim P(\ker J)$.

Proof. **Part (i).** Suppose τ is identifiable. If there existed $(\Delta\tau, \Delta\kappa) \in \ker[A\ B]$ with $\Delta\tau \neq 0$, then for any admissible (τ, κ) such that $(\tau + \Delta\tau, \kappa + \Delta\kappa)$ is also admissible, $A(\tau + \Delta\tau) + B(\kappa + \Delta\kappa) = A\tau + B\kappa$, so two distinct τ-values produce the same record, contradicting identifiability.

Conversely, if $\ker[A\ B]$ has no vector with nonzero τ-component, then $A\tau_1 + B\kappa_1 = A\tau_2 + B\kappa_2$ implies $(\tau_1 - \tau_2, \kappa_1 - \kappa_2) \in \ker[A\ B]$, so $\tau_1 = \tau_2$.

Part (ii). With closure constraints, admissible perturbations satisfy $G_\tau\Delta\tau + G_\kappa\Delta\kappa = 0$ in addition to $A\Delta\tau + B\Delta\kappa = 0$. Admissible perturbations thus lie in $\ker M_G$. The same injectivity argument as Part (i) applies to M_G.

Part (iii). For linear F, the unconstrained solution set is $\mathcal{A}(y) = (\tau_0, \kappa_0) + \ker[A\ B]$. Identifiability on this affine set requires projection to a single τ-value. This is a null-space property independent of a purely local numerical assessment. When admissible box constraints are imposed, the same null-space condition remains sufficient; if a latency-changing null direction permits nonzero admissible motion from an interior or relative-interior point, it is also necessary at that point.

Part (iv). Follows directly from Proposition 2 with $J = [A\ B]$ and $G = (G_\tau, G_\kappa)$. The quotient-rank condition states exactly that the closure rows separate every nonzero latency-changing equivalence class in $\ker J / \ker P$.

□

Remark 4 (Admissibility interaction). *Box constraints such as non-negativity restrict the admissible affine equivalence class and leave the null-space criterion unchanged at interior points. If a latency-changing null vector exists, and the true parameter lies in the relative interior of the admissible box, then sufficiently small motion along at least one admissible sign of that direction produces another record-equivalent point with a different latency component. Boundary constraints can truncate the admissible range; they do not supply closure unless they remove every latency-changing null direction, consistent with Proposition 4.*

The theorem converts identifiability into a closure problem. Raw records determine τ only up to the equivalence class $\mathcal{A}(y)$ of Definition 5; closure removes equivalence by eliminating all latency-changing null directions. When $[A\ B]$ is non-injective on τ, calibration and protocol structure change the measurement-identifiability problem by removing admissible equivalence directions. Figure 1 illustrates the record map and closure geometry.

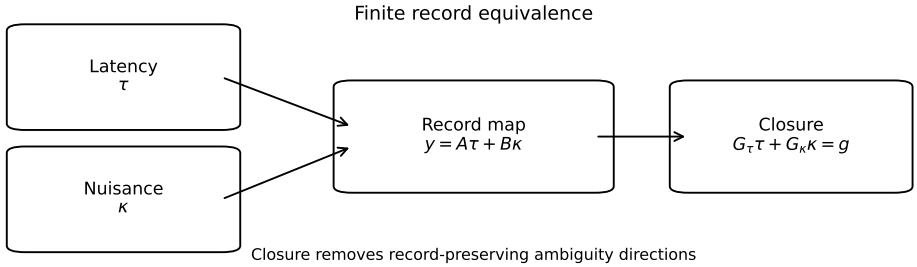


Figure 1. Schematic of the finite-record closure theorem. The observation map $y = A\tau + B\kappa$ sends joint parameter pairs to records. A null direction $A\Delta\tau + B\Delta\kappa = 0$ preserves the record and changes τ. Closure constraints $G_\tau\tau + G_\kappa\kappa = g$ augment the system; latency identifiability is restored when the augmented null space $\ker M_G$ has no vector with $\Delta\tau \neq 0$.

3.2. Corollaries and Specialisations

3.2.1. Two-Way Time-Transfer Non-Identifiability

The symmetric two-record model exhibits the parameter decompositions that preserve both timing records.

Corollary 4 (Two-way time-transfer non-identifiability). *Consider the symmetric two-record time-transfer model*

$$y_{AB} = \tau + \kappa + \delta_{AB}, \qquad y_{BA} = \tau - \kappa + \delta_{BA},$$

where τ is the reciprocal one-way propagation interval, κ is the clock offset between the two local time scales, and δ_{AB}, δ_{BA} are direction- dependent hardware, modem, antenna, cable, fibre, or timestamping delays. With

$$x = (\tau, \kappa, \delta_{AB}, \delta_{BA})^T,$$

the Jacobian is

$$J = \begin{pmatrix} 1 & 1 & 1 & 0 \\ 1 & -1 & 0 & 1 \end{pmatrix}, \qquad \operatorname{rank}(J) = 2, \quad \operatorname{nullity}(J) = 2.$$

A basis for the null space is

$$n_\tau = (1, 0, -1, -1)^T, \qquad n_\kappa = (0, 1, -1, 1)^T,$$

and direct multiplication gives $Jn_\tau = Jn_\kappa = 0$*. The projected latency null space is therefore*

$$\Pi_\tau(\ker J) = \operatorname{span}\{1\},$$

so the one-dimensional latency component is not identifiable from the two records alone. Explicitly, for any admissible a, b*,*

$$\tau' = \tau + a, \quad \kappa' = \kappa + b, \quad \delta'_{AB} = \delta_{AB} - a - b, \quad \delta'_{BA} = \delta_{BA} - a + b$$

preserves (y_{AB}, y_{BA}) *and shifts* τ *by* a*.*

Proof. The two rows of J are linearly independent, so rank–nullity gives $\dim \ker J = 2$. The displayed vectors satisfy $Jn_\tau = Jn_\kappa = 0$ and are linearly independent, hence form a basis of $\ker J$. Projection onto the latency coordinate maps n_τ to 1 and n_κ to 0. Thus $\Pi_\tau(\ker J) = \mathbb{R}$, and no nonzero linear functional of the single latency component is invariant on the record-equivalence class. The explicit transformation is the linear combination $an_\tau + bn_\kappa$ and follows by substitution.

□

3.2.2. Calibration Closure

Independent calibration constrains the direction-dependent delay terms in the preceding two-way model.

Corollary 5 (Calibration as closure). *In the two-way model above, exact independent knowledge of the direction-dependent delay terms* δ_{AB} *and* δ_{BA} *imposes two closure constraints whose combined action on* $\ker J$ *sets* $a = b = 0$*. The identifiable latency and clock offset are then*

$$\tau = \frac{(y_{AB} - \delta_{AB}) + (y_{BA} - \delta_{BA})}{2}, \qquad \kappa = \frac{(y_{AB} - \delta_{AB}) - (y_{BA} - \delta_{BA})}{2}.$$

If the two delay calibrations satisfy

$$|\delta_{AB} - \hat{\delta}_{AB}| \le \varepsilon_{AB}, \qquad |\delta_{BA} - \hat{\delta}_{BA}| \le \varepsilon_{BA},$$

then

$$|\hat{\tau} - \tau| \le \frac{\varepsilon_{AB} + \varepsilon_{BA}}{2}.$$

Proof. Fixing δ_{AB} and δ_{BA} imposes $\Delta\delta_{AB} = 0$ and $\Delta\delta_{BA} = 0$. In the null-space parametrisation, these constraints are $-a - b = 0$ and $-a + b = 0$, hence $a = b = 0$. The formulas for τ and κ follow by adding and subtracting the calibrated equations. The error bound follows from the triangle inequality.

□

3.2.3. GNSS Common-Clock Gauge Freedom

A simplified GNSS residual system illustrates common-clock gauge freedom after geometric and propagation modelling.

Corollary 6 (GNSS common-clock gauge freedom). *After subtracting a specified geometric and propagation model from a raw GNSS pseudorange, a simplified timing residual may be written as*

$$e_i(t) = C_R(t) + q_i(t),$$

where $C_R(t)$ *is the receiver-clock component and* $q_i(t)$ *collects satellite-clock, hardware, and satellite-specific residual timing terms. Both* C_R *and* q_i *are timing nuisance parameters in the notation of Section 2.1; they are not propagation-latency components. The transformation*

$$C'_R(t) = C_R(t) + \alpha(t), \qquad q'_i(t) = q_i(t) - \alpha(t)$$

preserves every residual $e_i(t)$. Thus the raw residual system has a common-clock gauge freedom unless an external clock product, receiver-clock convention, zero-mean gauge, or other closure condition is supplied.

Proof. Substitution gives

$$C'_R(t) + q'_i(t) = C_R(t) + \alpha(t) + q_i(t) - \alpha(t) = C_R(t) + q_i(t) = e_i(t)$$

for every satellite i. The result is a record-preserving null direction in the clock/calibration nuisance subspace. It illustrates the same closure-rank mechanism as the latency theorem. It makes no claim that receiver clock offset is itself a propagation latency. □

3.2.4. Single-Differencing as Parameter Elimination by Projection

Differencing two satellite residuals eliminates their shared receiver-clock term from the reduced equations.

Corollary 7 (Single-differencing as projection). *Form satellite-difference residuals $\Delta e_{ij}(t) = e_i(t) - e_j(t)$. The receiver clock term $C_R(t)$ is eliminated from the reduced equations:*

$$\Delta e_{ij}(t) = q_i(t) - q_j(t).$$

Thus, in the single-differenced observable system, $C_R(t)$ is no longer a parameter of the reduced equations. This should be distinguished from identifying $C_R(t)$ in the unreduced system. Differencing removes the receiver-clock parameter by projecting the record space onto combinations that are insensitive to $C_R(t)$. The operation leaves $C_R(t)$ unestimated and adds no calibration or clock-closure equation to the original system. A closure condition augments the unreduced model with information such as an external clock product, a known receiver-clock constraint, or a stated gauge convention.

Proof. The cancellation of $C_R(t)$ follows by subtraction. Since the operation maps the record vector to a lower-dimensional set of differences, it eliminates the receiver-clock parameter from the projected observation equations. It imposes no new equation on $C_R(t)$ in the original parameterisation. A gauge or clock constraint adds a row to the augmented system and can remove the common-clock null direction when it has nonzero action on that direction.

□

Remark 5. *Double-differencing similarly eliminates selected clock terms by projection. This is useful for forming reduced observables. It does not identify the eliminated clock parameter in the original unreduced system. Parameter elimination by projection and identifiability by closure answer different metrological questions.*

3.3. Finite-Record Demonstrations

3.3.1. Setup and Purpose

The finite-record demonstrations make the abstract identifiability statements operational and reproducible. They are not presented as empirical validation of a new propagation model. They show how the record-preserving null space, admissible decomposition family, and effect of calibration closure appear in a physically standard two-way time-transfer measurement equation.

3.3.2. Rank, Nullity, and Gauge Invariance

For the Jacobian J of Corollary 4,

$$J = \begin{pmatrix} 1 & 1 & 1 & 0 \\ 1 & -1 & 0 & 1 \end{pmatrix},$$

the script verifies $\mathrm{rank}(J) = 2$ and $\mathrm{nullity}(J) = 2$, with

$$n_\tau = (1, 0, -1, -1)^T, \qquad n_\kappa = (0, 1, -1, 1)^T$$

satisfying $Jn_\tau = Jn_\kappa = 0$. For the numerical parameters

$$\tau = 1.25, \quad \kappa = 0.47, \quad \delta_{AB} = 0.18, \quad \delta_{BA} = 0.31,$$

the records are

$$y_{AB} = 1.900, \qquad y_{BA} = 1.090.$$

Sampling admissible (a, b) values in the null-space family yielded maximum record distance $4.965e-16$, i.e. numerical zero to floating-point precision.

3.3.3. Admissible Decompositions and Calibration Closure

Under the stated positivity and boundedness constraints, the finite grid bookkeeping count yielded $N_R = 81758$ admissible decompositions ($S_R = \log N_R = 11.312$) and an admissible τ-range

$$[0.550,\ 1.493].$$

This count is reported only to document the grid used in the demonstration. It is not an intrinsic metrological descriptor, entropy, invariant volume, or property of the measurement operator. The intrinsic structural quantities are the rank of J, $\dim \ker J$, $\dim \Pi_\tau(\ker J)$, and the conditioning of the augmented operator after closure.

Independent calibration bounds on δ_{AB} and δ_{BA} collapse the admissible τ-range as the delay-calibration half-width is reduced, consistent with Corollary 5. Figure 2 shows the admissible equivalence class and the effect of calibration closure.

3.4. Synthetic Closure-Identifiability Verification

The results above are structural. To make the closure and null-space claims reproducible, we include a self-contained synthetic verification script that constructs finite measurement operators with known ambiguity and closure structure. The script numerically checks the projected-nullspace, closure-rank, stability, and admissible-range claims proved above. Empirical validation of a physical timing system lies outside this calculation.

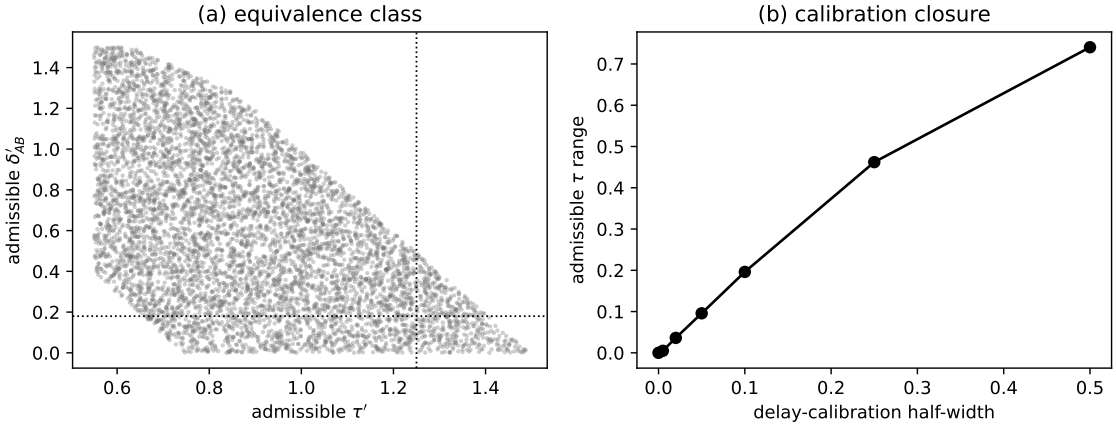


Figure 2. Standard two-way time-transfer closure demonstration. (**a**) Admissible decompositions in the (τ', δ'_{AB}) plane preserving both records under positivity and boundedness constraints; dotted lines mark true parameters. (**b**) Independent delay calibration reduces the admissible τ-range as calibration half-width decreases, consistent with Corollary 5.

The synthetic test uses two finite transfer architectures. The first is a common-mode timing architecture with

$$J = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}, \qquad \theta = (\tau_1, \tau_2, \kappa),$$

so that a common-mode perturbation can exchange latency and calibration structure and preserve the finite record. The second is a two-link architecture with separate calibration components, used to test the effect of independent closure constraints.

The script verifies five statements. First, it computes $\ker J$, projects it onto latency coordinates, and constructs the identifiable latency-functional basis as the set of latency functionals that are invariant under $\Pi_\tau(\ker J)$. Figure 3 shows this projected-nullspace relation. Second, it verifies the codimension relation

$$\dim(\mathcal{I}) = \dim(\Theta_\tau) - \dim \Pi_\tau(\ker J).$$

Third, it checks the closure-rank condition by comparing a closure row that removes the common-clock null direction with a differential constraint that does not. Fourth, it tests the approximate-closure stability bound using

the Moore–Penrose pseudoinverse of the augmented operator. Fifth, it compares the analytic admissible range along a null direction with a finite grid bookkeeping count.

In the common-mode case, the test gives

$$\operatorname{rank}(J) = 2, \qquad \dim \ker J = 1, \qquad \dim \Pi_\tau(\ker J) = 1.$$

Thus, before closure, the identifiable latency-functional dimension is

$$\dim(\mathcal{I}) = 1,$$

corresponding to differential latency. After imposing the clock-closure constraint, the identifiable latency dimension becomes

$$2.$$

The maximum record distance along the admissible null direction is

$$3.553e - 15,$$

confirming numerical record invariance up to floating-point precision. The admissible null-direction range is

$$17.320508.$$

For the chosen finite grid, the bookkeeping log-count is

$$S_R = \log(1793) = 7.491645.$$

This value only records the chosen grid resolution and admissibility box. It is not used as an intrinsic metrological descriptor of the operator.

The stability test perturbs the closure equation with the finite record held fixed (Figure 4). The bound

$$\|\delta\tau_{\mathrm{id}}\| \le \|\Pi_\tau M^\dagger\| \, \|\delta b\|$$

is satisfied for all tested perturbation levels, with

$$\|\Pi_\tau M^\dagger\| = 1.732051.$$

All checks in the synthetic verification script passed. The complete reproducibility package, including the script, tables, figures, README file, manifest, and JSON summary, is provided as Supplementary Material.

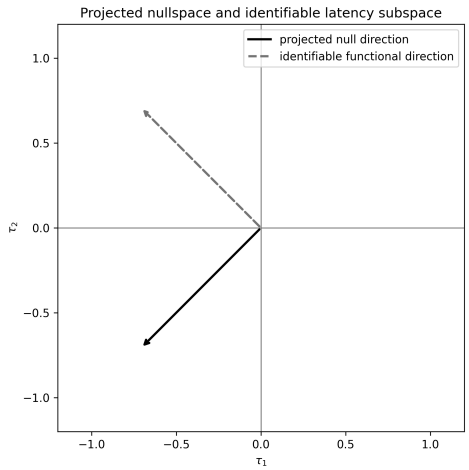


Figure 3. Synthetic projected-nullspace verification. The solid black vector shows a latency change that the record cannot detect because a nuisance-parameter change preserves the record. The dashed grey vector is orthogonal to that projected null direction. Its inner product remains fixed along the record-equivalence class and defines the identifiable latency functional.

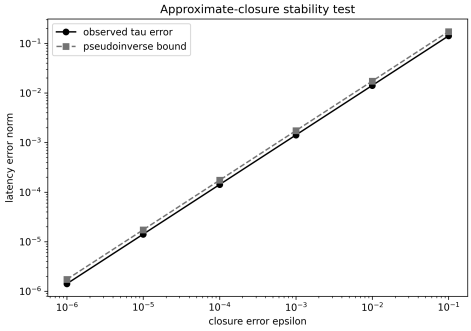


Figure 4. Synthetic approximate-closure stability check. The observed latency error remains below the Moore–Penrose pseudoinverse bound for all tested closure perturbations.

3.5. GNSS/RINEX Time-Transfer Application

The GNSS/RINEX material serves as a public-record diagnostic for common-clock gauge freedom. Standard GNSS processing supplies additional modelling, products, and conventions beyond the simplified residual system considered here.

Status of GNSS material.

The GNSS analyses in this section apply the closure formalism to the clock/calibration subspace of timing residuals. The residual equations provide a diagnostic example of common-clock gauge freedom in finite timing records.

3.5.1. Raw Pseudorange Records

The anchor analysis used the public EPN RINEX observation file `BRUX0010.10D.Z` [25,29,30], parsed with GeoRinex [31]. The file contained 2880 epochs, 30 satellites, and C1/P1/P2 observables. This stage operated directly on the raw pseudorange matrix. For the algebraic decomposition $\rho_i(t) = q_i(t) + C_R(t)$, the common-clock shift of Corollary 6 preserved every pseudorange: $d(R, R') = 0$ with maximum absolute difference 0. A fixed grid of 1586 tested decompositions all preserved the records exactly ($S_R = \log(1586) = 7.369$). This S_R value reflects the fixed grid size, not any intrinsic property of this station or observable; it is a grid bookkeeping count as defined in Section 2.2.3.

3.5.2. Cross-Station and Cross-Observable Replication

The same test was applied to 15 public EPN RINEX files across 7 stations (BRUX, WTZR, POTS, ZIMM, MATE, GOPE, BOR1), 6 receiver model strings, and 4 observables (C1, P1, P2, C2), yielding 39 successful file-observable analyses. In all 39 cases, the common-mode transformation preserved raw pseudorange records exactly ($d(R, R') = 0$). The repeated $S_R = 7.369$ across all analyses reflects that the same fixed grid was applied in each case. It confirms consistency with the null direction identified in Corollary 6. The result makes no claim of a shared intrinsic equivalence-class volume across distinct architectures.

3.5.3. Broadcast-Ephemeris Residualisation Test

To test whether the common-clock gauge freedom survives a broadcast-ephemeris geometry and satellite-clock layer, the observation file was paired with the broadcast navigation file `brdc0010.10n.Z` [32–34]. Post-geometry residuals were formed as $e_i(t) = \rho_i(t) - r_i(t) + c\,\delta t_i(t)$, where $r_i(t)$ is broadcast-ephemeris geometric range and $\delta t_i(t)$ is the broadcast satellite clock correction [24, 35–38]. This test is a gauge diagnostic. Operational GNSS processing also uses precise orbit and clock products, atmospheric models, hardware-delay calibration, and receiver-specific corrections [23,27,35,38,39].

After residualisation, the common-clock transformation preserved residual records exactly ($d(R, R') = 0$); 1066 tested decompositions were invariant ($S_R = 6.972$). A zero-mean closure constraint was then imposed. The baseline satisfies that constraint. The script applies a chosen nonzero common shift of approximately 50 m, which produces a closure violation of the same magnitude. This test checks whether the constraint acts on the common-clock null direction, in direct agreement with Theorem 2(ii). The chosen shift is a diagnostic input and carries no empirical weight.

Figure 5 shows the residual records and the closure diagnostic.

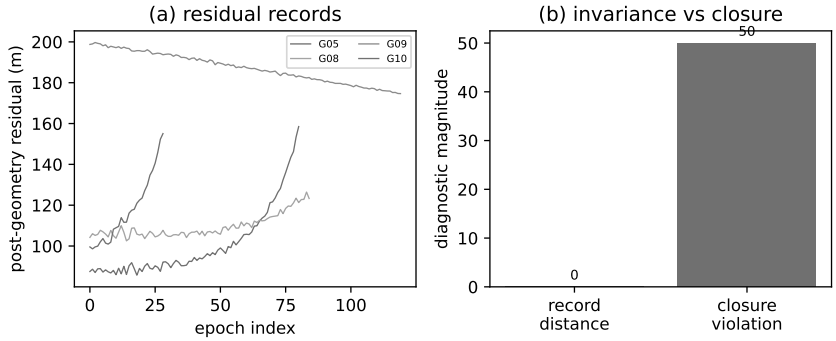


Figure 5. Broadcast-ephemeris residualisation diagnostic. (**a**) Post-geometry residuals after broadcast satellite-geometry and clock correction. (**b**) The common-mode transformation gives record distance $d(R, R') = 0$ and a zero-mean closure violation of approximately 50 m. The shifted decomposition is invisible to the record. The closure constraint detects it.

3.5.4. Interpretation

These results show a concrete RINEX instance consistent with the common-mode null direction of Corollary 6, both before and after a broadcast-ephemeris residualisation layer. Uniqueness is recovered only upon imposing an explicit closure constraint, consistent with Theorem 2(ii) and Proposition 3: a single linear constraint whose inner product with the null direction is nonzero suffices. Operational GNSS processing supplies additional information through ephemerides, receiver- clock estimation, atmospheric models, hardware-delay calibration, and gauge conventions. Differencing changes the observable subspace. Explicit clock, calibration, and gauge constraints augment the parameter system in the sense of M_G.

The common-clock direction is one of several ambiguity sources in operational GNSS time transfer. Atmospheric delays, hardware code or phase biases, orbit errors, and satellite-clock errors enter as additional components of κ or through the forward operator used to construct the residuals. Treating these quantities as unknowns adds columns to the operator and can create record-preserving directions. Supplying values through atmospheric models, orbit and clock products, or hardware calibration adds exact or approximate closure information. Their effect on latency identifiability is determined by the projected null space of the resulting augmented operator, $\Pi_\tau(\ker M_G)$.

4. Discussion

4.1. Calibration, Traceability, and Uncertainty

The inverse-problem result has a direct measurement consequence. Calibration and closure do not always act after a measurand has already been fixed. In finite time-transfer records, closure may be the operation that removes the latency-changing record-equivalence directions and therefore specifies the latency functional that the record can support.

Traceability begins with whether the record identifies the latency before calibration or closure is added. If $\Pi_\tau(\ker J) \neq \{0\}$, then at least one latency change is invisible to the finite record. A reported latency value must then be read as a value selected by an added convention, calibration, regularisation, or closure assumption. The uncertainty budget applies to that completed model and its selected latency component. It does not describe a latency determined by the record alone.

The proposed diagnostic requires the forward operator, the latency-changing null-space projection, and the closure constraints acting on it. A zero projection identifies the intended latency functional from the finite record. A nonzero projection requires a statement of the added information that selects the latency value. The action of calibration on the projection reveals whether it reduces uncertainty for an identified latency or specifies the measurand by removing an admissible ambiguity.

A regularised inverse solver can produce a reproducible latency estimate even when the record admits several latency-changing decompositions. Numerical stability of the returned estimate does not establish structural identifiability of the latency. The closure test asks what changes without changing the record. That question has to be answered before precision claims can be interpreted physically.

4.2. Scope

The theory concerns finite-record identifiability, not the full design of a time-transfer system. It asks whether the record and stated closure information determine a latency functional uniquely. It does not claim that all practical time-transfer analyses are structurally ambiguous, nor that existing calibration practice is deficient. The claim is

conditional: when the latency projection of the record-preserving null space is nonzero, the finite record alone has not identified that latency component.

4.3. Limits

The analysis is exact for the finite-dimensional linear or linearised models stated in the paper. Nonlinear architectures require local linearization or a corresponding nonlinear record-equivalence analysis. Environmental propagation models, relativistic corrections, receiver-specific effects, and laboratory calibration histories are not replaced by the theorem. They enter through the forward operator and closure constraints. The GNSS/RINEX analysis is a public-record gauge diagnostic, not a full operational GNSS time-transfer solution.

For a genuinely nonlinear operator F, the Jacobian result is local at the nominal parameter θ_0. The kernel $\ker J(\theta_0)$ detects infinitesimal record-preserving directions near that point. Separate admissible values $\theta_1 \neq \theta_2$ may still satisfy $F(\theta_1) = F(\theta_2)$. The Jacobian may also change rank elsewhere in the admissible domain. Global latency identifiability requires analysis of the nonlinear level sets $[\tau, \kappa]_F$ over the full domain, or a proof that the target latency functional is invariant on every such level set. The operational check is to evaluate the Jacobian across the admissible region and search for separated record-equivalent parameter pairs.

4.4. Relation to Inverse-Problem Practice

A finite record can be compatible with multiple decompositions that differ in the latency coordinate of interest. Regularisation, priors, conventions, and calibration constraints can select one decomposition. The proposed test distinguishes selection from identification by checking whether the intended latency is invariant on the record-equivalence class. Analytic examples are given in Appendix A.

5. Conclusions

A finite time-transfer record does not automatically identify the propagation latency it is intended to measure. The decisive object is the latency-changing part of the record-preserving null space. If an admissible null direction changes τ and leaves the record unchanged, then the corresponding latency functional is not determined by the finite record alone.

The paper gave a finite-record inverse-problem test for this condition. Compute the null space of the forward operator, project it onto latency coordinates, and ask whether the intended latency functional is invariant on that projection. Closure succeeds structurally only when calibration, loop, gauge, or synchronisation constraints remove the latency-changing directions. Approximate closure gives controlled degradation governed by the projected pseudoinverse of the augmented operator.

The central conclusion is metrological and inverse-theoretic. Precision has a physical interpretation after the finite record and its closure information uniquely specify the latency to be estimated. If uniqueness fails, a stable returned value is selected by additional assumptions and is not identified by the record alone.

Supplementary Materials

Supplementary Material 1 contains the self-contained standard two-way transfer and synthetic closure-identifiability packages, including scripts, generated tables, machine-readable summaries, figures, and a file manifest. The public-data RINEX packages are archived under the DOI given in the Data Availability statement.

Author Contributions

A.H.: conceptualisation, methodology, formal analysis, software, investigation, data curation, visualisation, writing, review, and editing. The author has read and agreed to the published version of the manuscript.

Funding

This research received no external funding. The article processing charge is waived under the journal's invitation.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

The code, public-data retrieval scripts, input manifests, generated tables, and machine-readable summaries used for the numerical checks are archived at Zenodo: doi:10.5281/zenodo.20439799. The files associated with this manuscript cover the standard symmetric two-way model, the synthetic closure-identifiability checks, the raw RINEX cross-station test, and the broadcast-ephemeris residualisation diagnostic. The Zenodo record also preserves historical packages; Supplementary Material 1 identifies the files and numerical outputs used here. All RINEX inputs are public and are identified by source and filename in the scripts. No human-participant, personal, or confidential data were used.

Conflicts of Interest

The author declares no conflict of interest.

Use of AI and AI-Assisted Technologies

OpenAI ChatGPT (OpenAI, 2026) was used for editorial and manuscript-preparation assistance, including language editing, organisation, formatting, mathematical and logical consistency checks, and identification of potentially relevant literature for the author to verify against the original sources.

The theory, causal architecture, hypotheses, examples, predictions, experimental designs, methodological logic, interpretation of the literature, and conclusions were developed solely by the author. The model did not originate any scientific idea, claim, hypothesis, causal relation, prediction, experimental design, interpretation, or conclusion.

Any model-assisted wording was used only to express or edit author-developed content. The author will determine the final wording and remains solely responsible for the submitted manuscript.

Appendix A. Analytic Examples of Record Equivalence

These examples state direct consequences of the kernel criterion. They require no numerical evidence and introduce no additional empirical claims.

Appendix A.1. Additive Propagation–Latency Degeneracy

For the model $R(t) = E(t - (\theta + \phi))$, the record depends on the sum $\theta + \phi$. Every perturbation $(a, -a)$ preserves the record. The sum is identifiable, and the individual components remain unidentifiable without a constraint that acts on this direction.

Appendix A.2. Finite-Bandwidth Measurement-Kernel Degeneracy

Let O be a linear measurement operator and let $h \in \ker O$. Then

$$O(s + h) = O(s).$$

A nonzero out-of-band component in the kernel therefore generates an exact record-equivalence direction. The identifiable functionals are those that vanish on the signal-space kernel relevant to the stated target.

Appendix A.3. Finite-Window Reconstruction Degeneracy

Let F_T denote restriction of a predicted record to $[0, T]$. A parameter perturbation that changes the prediction only outside this interval belongs to $\ker F_T$. Distinct global parameter values can therefore be record-equivalent on the finite observation window. The applicable identifiability statement is tied to the sampled operator used in the analysis.

References

1. Joint Committee for Guides in Metrology. *Evaluation of Measurement Data—Guide to the Expression of Uncertainty in Measurement*; JCGM, 2008. Available online: https://www.bipm.org/documents/20126/2071204/JCGM_100_2008_E.pdf (accessed on 26 October 2026). https://doi.org/10.59161/JCGM100-2008E.
2. Joint Committee for Guides in Metrology. *International Vocabulary of Metrology—Basic and General Concepts and Associated Terms (VIM)*; JCGM: Sèvres, France, 2012. Available online: https://www.bipm.org/documents/20126/2071204/JCGM_200_2012.pdf (accessed on 26 October 2026).
3. Bellman, R.; Åström, K.J. On structural identifiability. *Math. Biosci.* **1970**, *7*, 329–339. https://doi.org/10.1016/0025-5564(70)90132-x.

4. Walter, E.; Pronzato, L. *Identification of Parametric Models from Experimental Data*; Springer: Berlin/Heidelberg, Germany, 1997.
5. Ljung, L. *System Identification: Theory for the User*, 2nd ed.; Prentice Hall: Upper Saddle River, NJ, USA, 1999.
6. Miao, H.; Xia, X.; Perelson, A.S.; et al. On Identifiability of Nonlinear ODE Models and Applications in Viral Dynamics. *SIAM Rev.* **2011**, *53*, 3–39. https://doi.org/10.1137/090757009.
7. Raue, A.; Karlsson, J.; Saccomani, M.P.; et al. Comparison of approaches for parameter identifiability analysis of biological systems. *Bioinformatics* **2014**, *30*, 1440–1448. https://doi.org/10.1093/bioinformatics/btu006.
8. Villaverde, A.F. Observability and Structural Identifiability of Nonlinear Biological Systems. *Complexity* **2019**, *2019*, 8497093. https://doi.org/10.1155/2019/8497093.
9. Richmond, J. Identifiability in Linear Models. *Econometrica* **1974**, *42*, 731–736. https://doi.org/10.2307/1913941.
10. Tuncer, Y. Equivalence of parametric identifiability and estimability. *Commun. Stat.-Theory Methods* **1985**, *14*, 393–400. https://doi.org/10.1080/03610928508828920.
11. Penrose, R. A generalized inverse for matrices. *Math. Proc. Camb. Philos. Soc.* **1955**, *51*, 406–413. https://doi.org/10.1017/S0305004100030401.
12. Ben-Israel, A.; Greville, T.N.E. *Generalized Inverses: Theory and Applications*, 2nd ed.; Springer: New York, NY, USA, 2003. https://doi.org/10.1007/b97366.
13. Hadamard, J. *Lectures on Cauchy's Problem in Linear Partial Differential Equations*; Yale University Press: New Haven, CT, USA, 1923.
14. Tikhonov, A.N.; Arsenin, V.Y. *Solutions of Ill-Posed Problems*; V. H. Winston & Sons: Washington, DC, USA, 1977.
15. Engl, H.W.; Hanke, M.; Neubauer, A. *Regularization of Inverse Problems*; Kluwer Academic Publishers: Dordrecht, The Netherlands, 1996. https://doi.org/10.1007/978-94-009-1740-8.
16. Hansen, P.C. *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion*; Society for Industrial and Applied Mathematics: Philadelphia, PA, USA, 1998. https://doi.org/10.1137/1.9780898719697.
17. Aster, R.C.; Borchers, B.; Thurber, C.H. *Parameter Estimation and Inverse Problems*, 3rd ed.; Elsevier: Amsterdam, The Netherlands, 2018. https://doi.org/10.1016/C2015-0-02458-3.
18. Tarantola, A. *Inverse Problem Theory and Methods for Model Parameter Estimation*; SIAM: Philadelphia, PA, USA, 2005. https://doi.org/10.1137/1.9780898717921.
19. Allan, D.W.; Weiss, M.A. Accurate time and frequency transfer during common-view of a GPS satellite. In Proceedings of the 34th Annual Symposium on Frequency Control, Philadelphia, PA, USA, 28–30 May 1980; pp. 334–346. https://doi.org/10.1109/FREQ.1980.200424.
20. Levine, J. A review of time and frequency transfer methods. *Metrologia* **2008**, *45*, S162–S174. https://doi.org/10.1088/0026-1394/45/6/s22.
21. ITU-R. Recommendation ITU-R TF.1153-4: The Operational Use of Two-Way Satellite Time and Frequency Transfer Employing Pseudorandom Noise Codes; ITU-R, August 2015. Available online: https://www.itu.int/dms_pubrec/itu-r/rec/tf/R-REC-TF.1153-4-201508-I!!PDF-E.pdf (accessed on 26 October 2026).
22. Defraigne, P.; Petit, G. CGGTTS-Version 2E : an extended standard for GNSS Time Transfer. *Metrologia* **2015**, *52*, G1. https://doi.org/10.1088/0026-1394/52/6/g1.
23. Kouba, J.; Héroux, P. Precise Point Positioning Using IGS Orbit and Clock Products. *GPS Solut.* **2001**, *5*, 12–28. https://doi.org/10.1007/pl00012883.
24. Ashby, N. Relativity in the Global Positioning System. *Living Rev. Relativ.* **2003**, *6*, 1. https://doi.org/10.12942/lrr-2003-1.
25. Gurtner, W.; Estey, L. RINEX: The Receiver Independent Exchange Format Version 2.11; IGS/RTCM RINEX Working Group, 10 December 2007; Minor Clarifications, 26 June 2012. Available online: https://files.igs.org/pub/data/format/rinex211.txt (accessed on 26 October 2026).
26. International GNSS Service RINEX Working Group; Radio Technical Commission for Maritime Services Special Committee 104. RINEX: The Receiver Independent Exchange Format, Version 3.03, 14 July 2015. Available online: https://files.igs.org/pub/data/format/rinex303.pdf (accessed on 26 October 2026).
27. Defraigne, P.; Achkar, J.; Coleman, M.J.; et al. Achieving traceability to UTC through GNSS measurements. *Metrologia* **2022**, *59*, 064001. https://doi.org/10.1088/1681-7575/ac98cb.
28. Jiang, Z.; Zhang, V.; Parker, T.E.; et al. Improving two-way satellite time and frequency transfer with redundant links for UTC generation. *Metrologia* **2019**, *56*, 025005. https://doi.org/10.1088/1681-7575/aafced.
29. EUREF Permanent GNSS Network Central Bureau. EPN GNSS Data Repository. Available online: https://gnss.be/epndata.php (accessed on 25 May 2026).
30. EUREF Permanent GNSS Network Central Bureau. RINEX Formats. Available online: https://epncb.oma.be/ _documentation/formats/rinex.php (accessed on 25 May 2026).
31. Hirsch, M. GeoRinex: RINEX 2/3/4 Reader and Batch Conversion to NetCDF4/HDF5. Available online: https://pypi.org/project/georinex/ (accessed on 25 May 2026).
32. BKG GNSS Data Center. GNSS Observation and Navigation Data and Products. Available online: https://igs.bkg.bund.de/ (accessed on 25 May 2026).

33. International GNSS Service. Data Access. Available online: https://igs.org/data-access/ (accessed on 25 May 2026).
34. NASA Earthdata / CDDIS. Broadcast Ephemeris Data Product. Available online: https://www.earthdata.nasa.gov/data/space-geodesy-techniques/gnss/broadcast-ephemeris-data-product (accessed on 25 May 2026).
35. Sanz Subirana, J.; Juan Zornoza, J.M.; Hernández-Pajares, M. *GNSS Data Processing, Volume I: Fundamentals and Algorithms*; ESA Communications: Noordwijk, The Netherlands, 2013.
36. GPS Directorate. Navstar GPS Space Segment/Navigation User Segment Interfaces, IS-GPS-200N, 1 August 2022. Available online: https://www.gps.gov/sites/default/files/2025-07/IS-GPS-200N.pdf (accessed on 26 October 2026).
37. Misra, P.; Enge, P. *Global Positioning System: Signals, Measurements, and Performance*, 2nd ed.; Ganga-Jamuna Press: Lincoln, MA, USA, 2006.
38. Teunissen, P.J.G.; Montenbruck, O., Eds. *Springer Handbook of Global Navigation Satellite Systems*; Springer: Cham, Switzerland, 2017. https://doi.org/10.1007/978-3-319-42928-1.
39. Håkansson, M.; Jensen, A.B.O.; Horemuz, M.; et al. Review of code and phase biases in multi-GNSS positioning. *GPS Solut.* **2017**, *21*, 849–860. https://doi.org/10.1007/s10291-016-0572-7.