

Characterization of GNSS Observables from LuGRE Lunar Surface Operations

Original

Characterization of GNSS Observables from LuGRE Lunar Surface Operations / Cossu, I., Vouch, O., Minetto, A., Nardin, A., Zocca, S., Dovis, F.. - (2026), pp. 1925-1939. (International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2026) Orlando (USA) 14 - 18 September 2026) [10.33012/2026.20685].

Availability:

This version is available at: 11583/3016636 since: 2026-10-14T09:10:38Z

Publisher:

Institute of Navigation (ION)

Published

DOI:10.33012/2026.20685

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Characterization of GNSS Observables from LuGRE Lunar Surface Operations

Ilaria Cossu, Oliviero Vouch, Alex Minetto, Andrea Nardin, Simone Zocca, Fabio Dovis,
*Department of Electronics and Telecommunications (DET),
Politecnico di Torino, Turin, Italy*

BIOGRAPHY

Ilaria Cossu received the M.Sc. degree (2025) in aerospace engineering from Politecnico di Torino. She is currently a Ph.D. student within the Navigation Signal Analysis and Simulation (NavSAS) group. Her research focuses on the development and application of Bayesian methods for GNSS-based orbit determination and time synchronization in cislunar space.

Oliviero Vouch received the B.Sc. (2018), M.Sc. (2020), and Ph.D. with honors (2025) in electrical, electronics, and communications engineering from Politecnico di Torino. He is currently a post-doctoral researcher focusing on GNSS-based multisensor navigation systems, as well as on-board orbit determination for missions in the Earth–Moon space.

Alex Minetto received the B.Sc. (2013), M.Sc. (2015), and Ph.D. with honors (2020) in telecommunications and electrical engineering from Politecnico di Torino, where he is now Assistant Professor in the Department of Electronics and Telecommunications and a member of the NavSAS research group. His work focuses on advanced GNSS signal analysis, receiver design, and Bayesian state estimation. He has contributed to ESA and EUSPA projects and since 2021 has been involved in the Lunar GNSS Receiver Experiment (LuGRE) as a mission's Science Team member.

Andrea Nardin received the M.Sc. (2018) and Ph.D. with honors (2023) in telecommunications and electrical engineering from Politecnico di Torino, where he is now Assistant Professor in the Department of Electronics and Telecommunications and a member of the NavSAS research group. In 2021, he was a Visiting Doctoral Researcher at Northeastern University, Boston. His research interests include signal processing architectures and signal design for GNSS and next-generation PNT systems.

Simone Zocca received the B.Sc. (2018), M.Sc. (2020) and Ph.D. with honors (2024) in electrical, electronics, and communications engineering from Politecnico di Torino. He is currently a researcher with the NavSAS research group, Department of Electronics and Telecommunications, working on innovative signal processing techniques and Bayesian estimation methods for GNSS.

Fabio Dovis received the M.Sc. (1996) and Ph.D. (2000) from Politecnico di Torino, where he is now Full Professor in the Department of Electronics and Telecommunications and coordinator of the NavSAS Research Group. His research focuses on GNSS receiver design, advanced signal processing for interference and multipath mitigation, and ionospheric monitoring. He has extensive experience in international projects, collaborations with industry and research institutions, and serves on the IEEE AESS Navigation Systems Panel.

ABSTRACT

This paper presents a characterization of the precision of raw GNSS code-pseudorange, carrier-phase, and Doppler-shift observables collected by the Lunar GNSS Receiver Experiment (LuGRE) during lunar-surface operations. Detrended observable time series are used to derive empirical standard-deviation curves as a function of carrier-to-noise density ratio (C/N_0), separately for each of the four tracked signals (GPS L1 C/A, GPS L5, Galileo E1 OS, Galileo E5a). These curves are modeled according to a compact exponential model, yielding an analytical, C/N_0 -conditioned representation of observable precision over the signal-strength range experienced during LuGRE surface operations. The adequacy of this description is assessed by standardizing the de-trended samples using the fitted models and re-binning them against two geometry-related descriptors: the transmit off-boresight angle and the grazing altitude of the line-of-sight path. Results show that the C/N_0 -conditioned models capture the dominant structure of observable dispersion over the well-sampled portion of the dataset. Local departures from unit normalized dispersion are observed for some signal-observable pairs, particularly at high off-boresight angles and within specific grazing-altitude intervals; however, no consistent geometry-dependent trend emerges across signals. The available data therefore do not support the introduction of a general additional geometry-dependent variance term. The resulting flight-derived precision models provide empirical inputs for measurement weighting in cislunar navigation filters and for realistic simulations supporting future lunar missions.

I. INTRODUCTION

Lunar GNSS Receiver Experiment (LuGRE) has provided the first experimental spaceborne Global Navigation Satellite System (GNSS) dataset across the Earth–Moon space and on the lunar surface, offering a unique opportunity to characterize GNSS observables well beyond the operational space service volume (SSV) (Parker et al., 2026). The LuGRE Science Definition Team (SDT) established a structured set of science investigations to be executed as part of the mission science programme (Konitzer et al., 2024). To enable a rapid implementation of these investigations on the flight datasets, a dedicated tool was already developed in support of the mission, automating telemetry parsing, observable processing, and the execution of condensed versions of the experiments defined in the science plan (Minetto et al., 2025). Among the considered analyses, the characterisation of the statistical properties of raw payload telemetry was identified as a primary scientific output, noting that such properties directly impact the performance of on-board navigation estimators relying on GNSS observables, and that flight-proven models are needed to lend credibility to the simulations used in planning future lunar missions exploiting GNSS in cislunar space.

An initial measurement-quality assessment of the LuGRE flight data was already reported by Parker et al. (2026); this work estimated the standard deviation of de-trended code-pseudorange and carrier-phase observables across operations, distinguished according to the signals tracked by the LuGRE receiver, namely GPS L1 C/A, GPS L5, Galileo E1 OS, and Galileo E5a. Such assessment has offered a first quantitative picture of the noise standard deviation affecting the observables constructed under a weak-signal regime given the Earth–Moon baseline. Restricting the perspective on lunar-surface portion of the whole dataset (Parker et al., 2025), the present work extends that initial assessment with the aim of producing and assessing the validity of an empirical stochastic model of observable precision as a function of the carrier-to-noise density ratio (C/N_0) actually experienced by the receiver at lunar altitudes.

The need for such a model follows from a practical issue common to most GNSS estimation problems: observations with different expected precision should not be assigned the same weight. Over the past two decades, a large body of literature has therefore addressed the derivation of stochastic models for code-pseudorange, Doppler and carrier-phase observables. Early efforts in the context of civil-aviation integrity adopted elevation-dependent models in which the variance scaled with the inverse sine of the satellite elevation, with the elevation angle used as a simple proxy for the combined effects of geometry, atmospheric path length (Salós et al., 2010). As the use of GNSS extended to signal-degraded scenarios such as urban environments, where the geometry alone fails to predict measurement quality, models based on the carrier-to-noise density ratio, on combinations of elevation and C/N_0 , were shown to be more representative of the actual observation noise under degraded link conditions (Kuusniemi, 2005; Medina et al., 2018). More recent approaches estimate the residual error variance directly from the raw observations, rather than assembling a measurement-error budget from separately modeled contributions that rely on parametric assumptions (Wang et al., 2013; Dovis et al., 2015; Zhang et al., 2022). This latter approach is particularly relevant to the LuGRE case study, as existing error models cannot simply be carried over to the statistical analysis of observables constructed from tracking GNSS at lunar altitudes or at least their validity may be questionable in absence of other datasets.

The contribution of this work is threefold. First, empirical estimates of the residual error variance are derived for code-pseudorange, carrier-phase, and Doppler observables as functions of C/N_0 , separately for each tracked signal, by aggregating the de-trended measurements collected during all lunar-surface operations. Second, these empirical estimates are approximated by an exponential model, yielding a C/N_0 -conditioned analytical representation of the observable precision for each observable–signal pair, whose empirical validity is established over the range of C/N_0 values estimated by the LuGRE payload during the lunar-surface operations. The functional form is the one conventionally adopted in terrestrial stochastic modeling (Kuusniemi, 2005), and the smoothing applied to the C/N_0 series is fine tuned so as to bring the LuGRE observations into the closest possible agreement with this form. Third, the statistical adequacy of this analytical description is assessed by examining the dispersion of the standardized de-trended samples against the transmit off-boresight angle and the grazing altitude of the line-of-sight (LoS) path. These quantities, which are independent of the particular signal-strength regime experienced by LuGRE, can be considered proxies of the transmit-side conditions of the received signal and the geometry of the atmospheric crossing, respectively. Therefore, they provide a means to assess whether the observed dispersion is sufficiently characterized by C/N_0 alone or whether additional geometry-dependent effects should be considered responsible for an inflation of the error variance.

The significance of this empirical characterization, and the extent to which it carries beyond LuGRE, deserves a clear statement. The empirical curves and fitted models are conditioned on the C/N_0 values reported by LuGRE during surface operations. As discussed in Section 5.1.2 of Parker et al. (2026), these values were found to be 7–12 dB-Hz lower than predicted by pre-flight link-budget simulations for satellites of the same generation. The discrepancy was traced to the integrated lander/payload flight system rather than to the propagation environment and is therefore regarded as a mission-specific feature rather than an intrinsic property of GNSS reception at the Moon. As a consequence, the absolute numerical values of the fitted models inherit this offset, and their direct use in ground-based simulators or for observable weighting in navigation estimators would, in principle, call for a re-calibration accounting for the known C/N_0 gap. However, this same offset also gives the present characterization a practical value: because LuGRE operated at weaker signal levels than expected, the reported precision values represent a conservative, worst-case scenario. A receiver operating under nominal link budget conditions would therefore be expected to achieve equal or better measurement precision. What transfers most directly beyond the mission, irrespective of the C/N_0

offset, is the functional form of the model together with the procedure by which the sufficiency of a C/N_0 -only conditioning is assessed against other descriptors.

II. BACKGROUND & MOTIVATION

This section defines the quantity targeted by the proposed empirical precision assessment, derived from raw observables, and relates its components to the terms of the standard GNSS observation equations. It also motivates deriving residual error statistics directly from raw data, rather than from a classical system-level error budget.

1. Decomposition of raw GNSS observables

Let i denote the satellite index and $j \in \{L1/E1, L5/E5a\}$ the frequency band. For the LuGRE payload, a single signal per band is tracked on each satellite, namely GPS L1 C/A and Galileo E1 OS on the L1/E1 band, and GPS L5 and Galileo E5a on the L5/E5a band; the pair (i, j) therefore uniquely identifies each tracked channel. It is worth remarking that individual signal components (i.e., pilot and data) are not discriminated at the observable level. A generic raw observable on the satellite channel (i, j) , can be described according to the following decomposition:

$$y^{(i,j)} = g^{(i,j)} + \underbrace{w^{(i,j)} + r^{(i,j)}}_{\varepsilon^{(i,j)}}, \quad (1)$$

where the dependence on time is left implicit, and the terms are not only distinguished by their physical origin but also by their expected temporal variation. In particular:

- $g^{(i,j)}$ collects slow enough contributions to be well approximated, at least over short observation spans (i.e., few hundreds of seconds), by low-order polynomials. The dominant term comes from the geometry of the relative motion of the satellite and of the receiver, the latter being tied to the Moon's orbit for a static surface scenario. The remaining contributions are as follows.

Clock and relativistic terms. The satellite clock offset, driven by highly stable on-board atomic standards, exhibits sub-millimeter-per-second rates of change over intervals of a few hundred seconds, resulting in nearly-linear drift profiles (Kaplan and Hegarty, 2017, Ch.10). The net relativistic effects from special and general relativity on the satellite clock, compensated through an in-factory frequency offset calibrated at Earth's mean sea level, no longer cancel for a receiver on the Moon and combine with the periodic eccentricity contribution; both are smooth in time (Ashby, 2003). Similarly, spaceborne miniature atomic clocks for lunar timekeeping can achieve extreme long-term accuracy (Ely et al., 2026); however, when lower-class crystal oscillators are used, their stability can become highly susceptible to the operating conditions, where temperature plays a critical role (Parker et al., 2026). Nonetheless, within the observation spans on which the present analysis operates, the receiver clock term can be considered sufficiently smooth.

Hardware biases. Frequency-dependent transmit and receive hardware biases enter g as well, and are tightly coupled to the clock terms—a hardware bias acts as a frequency-specific offset on the clock. These biases display sub-nanosecond daily variations and are therefore highly stable on the time scales considered here (Su et al., 2026); similar remarks are applicable to the receive-side biases, which are calibrated.

First-order ionospheric delay. At first order, the ionospheric refractive delay on a single frequency is proportional to the integrated electron content along the slant path; in the conventional GNSS observation model, this contribution is evaluated along the geometric satellite–receiver line of sight. For GNSS receptions on the lunar-surface, a subset of the tracked signals exhibits low Earth-limb tangent altitudes, consequently traversing a longer and highly oblique segment of the ionosphere. Over the short observation spans considered here, the dominant first-order contribution is expected to evolve sufficiently smoothly. This is, however, a local modeling assumption rather than a general property, as the electron content integrated by the link depends also on the spatial and temporal evolution of the plasma distribution.

Other contributions include the carrier-phase ambiguity, the phase wind-up, and, where present, any quasi-static reflection from the local terrain around the static receiver. The specific make-up of g differs between observables: the ambiguity and the wind-up enter the carrier-phase only, the Doppler-shift inherits a clock-drift rather than a clock-offset contribution, and the first-order ionospheric delay enters with opposite sign on code and phase, while its rate contributes to the Doppler observable. What all the contributions mentioned above share besides smoothness is that they are not random. Any residual that remains after their correction is attributable to the imperfections of the corrective models.

- $w^{(i,j)}$ isolates random errors originating in the receiver chain, from antenna and front-end thermal noise through correlator processing to tracking-loop dynamics. All these contributions directly affect the available link budget, and therefore reflect on the C/N_0 estimated by the receiver, which averages over the correlators' output. The corresponding error on the observable is mostly driven by tracking loop errors; for delay-lock loop (DLL), jitter induced by thermal noise is the

primary contributor to receiver noise on code-pseudorange observables. Beyond thermal jitter, receiver noise on phase observables is tied to the oscillator’s short-term stability due to the integrative nature of carrier-phase tracking. For high-sensitivity tracking in weak-signal environments, narrowing the loop bandwidth reduces the thermal noise contribution, but simultaneously allows more oscillator phase noise to pass through (Kaplan and Hegarty, 2017, Ch.8). The latter effect contributes less to Doppler noise, which is in turn more affected by thermal noise.

- $r^{(i,j)}$ collects residual, high-frequency contributions that arise from unpredictable events that fit a statistical description but depart from the receiver noise. For the addressed lunar surface scenario, these contributions are mainly expected from main lobe receptions with signals that skim the Earth at low grazing angles, traversing a longer distance across the Earth’s atmosphere. On the large spatial scale, plasma irregularities act as a refractive medium with a non-uniform integrated electron content, and rapid variations of this content along the long crossings can leave a residual dispersion on the observables whose magnitude is expected to depend on the tangent altitude of the path. On the small-scale, signal-diffractive scintillations can affect both the amplitude and the phase of the tracked signal. The amplitude fluctuations perturb the estimated C/N_0 at a level that depends on their persistence (Pi et al., 2017), and the underlying response of the tracking noise cannot be anticipated by the description of w . The phase fluctuations, in turn, introduce additional jitter at the phase-lock loop (PLL) discriminator output, thereby increasing the variance of the phase error. When spiky events occur, these are hardly distinguished from cycle slips.

In light of the foregoing, $\varepsilon^{(i,j)}$ is the zero-mean composite error whose standard deviation is the target of the proposed modeling framework. In particular, the dependence of $w^{(i,j)}$ on C/N_0 is understood at the level of C/N_0 experienced on the channel under nominal link conditions, and this is what the empirical stochastic model formulated in Section III.2 aims to characterize. Whether $r^{(i,j)}$ leaves a statistically measurable signature on the dispersion of $\varepsilon^{(i,j)}$ with respect to other geometric descriptors, should any such effect be present in the analyzed dataset, is addressed by the residual variance analysis of Section III.4.

2. Relation to classical error budget in terrestrial GNSS

The system-level characterization of GNSS measurement quality, originally developed within the integrity practice of civil aviation (Parkinson and Spilker, 1996; Axelrad and Brown, 1996), expresses the user equivalent range error (UERE)—namely, the variance of a zero-mean Gaussian distribution overbounding the actual residual error distribution—as a root-sum-square of independent contributions:

$$\sigma_{\text{UERE}}^2 = \sigma_{\text{orb/clk}}^2 + \sigma_{\text{prop}}^2 + \sigma_{\text{mp}}^2 + \sigma_{\text{rec}}^2. \quad (2)$$

This budget is meant *after* the relevant corrections to most of the contributions summarized by g in (1) have been applied, and each term in (2) represents the residual that those corrections leave behind. $\sigma_{\text{orb/clk}}^2$ isolates the residual range error attributable to the space and control segments, commonly referred to as signal-in-space user range error (SISRE). Its magnitude depends on the corrections used, ranging from broadcast clock and ephemeris data in the legacy navigation message to high-precision SSR-type products such as those delivered by the Galileo high-accuracy service. In the Global Positioning System (GPS) legacy framework, the User Range Accuracy (URA) index is the conservative overbound of this residual that integrity-driven users assume in their budget (Renfro et al., 2024). The propagation term σ_{prop}^2 overbound second-or-higher-order ionospheric effects as well as the tropospheric residual. For dual-frequency terrestrial users, the ionospheric residual can be as large as few centimeters. For single-frequency users, standard models such as those discussed in Salós et al. (2010) compute meter-level residual ionospheric budget and turn out to be very conservative against empirical observations. The multipath term σ_{mp}^2 describes contributions that are intrinsic to the user local environment and cannot be mitigated by any correction broadcast by the system; it is the responsibility of the user to adopt a model that bounds the multipath error consistently with the operational scenario (Medina et al., 2018). Finally, the receiver-noise term σ_{rec}^2 is the contribution corresponding to $w_k^{i,j}$ in (1), and is parametrized through models that depend on C/N_0 , the tracked signal modulation, and the parameters of the receiver tracking loops (Kaplan and Hegarty, 2017, Ch.8).

Although very simple to implement, classical residual error models are not directly applicable to spaceborne GNSS observables constructed from signals tracked at lunar altitude. The reasons are twofold, and deducible to a certain extent.

First, several terms in (2) either lack a physical counterpart or lack a validated model for the lunar-surface geometry. The tropospheric contribution is likely to be absent; to the best of the authors’ knowledge, no empirical record of its impact on GNSS signals tracked at lunar altitudes has been provided so far. On the other hand, the overbounding models of the ionospheric residual are tied to correction algorithms which have been developed and calibrated for terrestrial users; the validity of these correction models does not extend up to the nearly-horizontal ray paths that skim the Earth at low tangent altitudes with ionospheric and plasmaspheric crossings of up to few thousands of kilometers. Early results from LuGRE dual-frequency observations confirm this issue: Cesaroni et al. (2026) show that even a global core plasma model struggles to predict the electron density structure in transition regions from high-altitude plasmasphere to top-side ionosphere or fine-scale density variations, yielding integrated electron content values that disagree with the flight measurements. If the underlying model itself is inaccurate for

these geometries, the correction residuals derived from it cannot be expected to provide a meaningful input to a system-level budget.

Second, the tracking-noise contribution is the only term of (2) whose underlying physics remains valid at the Moon, as it depends on receiver parameters rather than the physics of the signal propagation. However, while the models remain applicable, the obstacle lies in configuring their parametrization, as the loop design is proprietary information.

The arguments above motivate the choice, adopted in this work, to characterize $\varepsilon^{(i,j)}$ directly from the raw telemetry, rather than assessing the observables' quality through a model-based error budget. The empirical stochastic model derived from this approach, discussed in Section III, thus targets the dispersion that remains once $g^{(i,j)}$ has been removed. In the language of (2), this corresponds to characterising the receiver-noise term σ_{rec}^2 directly from the observable, complemented by whatever residual signature the contributions in $r^{(i,j)}$ leave on the dataset. The remaining terms of (2) are not captured by this approach: they describe residuals of corrective models or local environmental effects that de-trending removes.

III. METHODOLOGY

This section discusses the procedure which has been followed to derive the stochastic models characterizing the uncertainty present in the GNSS signals collected by LuGRE on the lunar surface as a function of C/N_0 . Different models are derived for the distinct observables making up the raw payload telemetry. Moreover, the analysis is conducted separately for each signal tracked by LuGRE. Given that the empirical stochastic model should be representative of the peculiar environment where the receiver operates rather than of individual surface operations (Parker et al., 2026), all surface telemetry spanning a lunar day has been aggregated into a unique dataset, encompassing different geometries of tracked GNSS satellites and visibility conditions. The logic underpinning the definition of the stochastic model is based on identifying an analytical model that provides the most faithful C/N_0 -conditioned representation of the precision of LuGRE observations. From this perspective, the degree of freedom through which the observations are brought as close as possible to the analytical model is the smoothing applied to the C/N_0 series, introduced in Section III.1 and further developed in Section III.3.

1. De-trending of raw observable time series

In line with (1), the time-series of raw observables can be modeled as discrete-time random processes, where mean stationarity is violated by the presence of slowly-varying contributions summarized by $g^{(i,j)}$. As mentioned in Section II, these contributions should be removed such that the stochastic model can profile the residual dispersion. To this end, raw observable time series are first de-trended within each satellite tracking pass by applying a centered finite-difference operator:

$$\tilde{y}_k^{(i,j)} = \sum_{\ell=-L}^{+L} d_\ell y_{k+\ell}^{(i,j)}, \quad \text{defined only if } k + \ell \in a(k) \quad \forall \ell, \quad (3)$$

where d_ℓ are the coefficients of the stencil, L is its half-width, and $a(k)$ indicates the tracking pass the k -th epoch belongs to. Because satellite tracking during each operation is generally discontinuous, isolating the de-trending operation to individual passes ensures that the operator never crosses tracking discontinuities. This work implements a third-order differencing with fourth-order accuracy, corresponding to a seven-point symmetric stencil $\ell \in \{-3, \dots, +3\}$ with coefficients $d_\ell \in \{+\frac{1}{8}, -1, +\frac{13}{8}, 0, -\frac{13}{8}, +1, -\frac{1}{8}\}$ (Fornberg, 1988). This configuration was selected after testing several combinations of derivative order m and accuracy order p , aiming to identify the lowest-complexity configuration that could reduce the mean of the de-trended series close to zero while preserving a stable estimate of the residual noise standard deviation. Applied to the decomposition in (1), the operator acts linearly on the components:

$$\tilde{y}_k^{(i,j)} = (\mathcal{D}g^{(i,j)})_k + \underbrace{(\mathcal{D}\varepsilon^{(i,j)})_k}_{\tilde{\varepsilon}_k^{(i,j)}}, \quad (4)$$

where any contribution to $y^{(i,j)}$ that can be represented, over the stencil support, by a polynomial of degree at most two is annihilated; this ensures that the low-frequency components are wiped off, i.e., $(\mathcal{D}g^{(i,j)})_k \approx 0$. Therefore, $\tilde{y}^{(i,j)}$ reduces to a zero-mean time series. It should be noted that the centered form discards L samples at the two boundaries of each pass, where the stencil cannot be symmetrically applied.

The combined choice of the differencing order and stencil accuracy is driven by a second aspect. The relation $\text{Var}(\tilde{\varepsilon}^{(i,j)}) = \gamma^2 \text{Var}(\varepsilon_k^{(i,j)})$ with $\gamma^2 = \sum_\ell d_\ell^2$ holds only when the original samples are uncorrelated. At the typical GNSS sampling rate of one sample per second, consecutive samples of raw observables may carry correlated error contributions inherited from the tracking loops and from the slow evolution of the propagation conditions, in which case the simple gain correction is no longer exact; the residual correlation inflates or deflates the de-trended variance (Dovis et al., 2015). To preserve the validity of the

normalization to the largest extent possible, the differencing order and the accuracy order have been jointly selected so as to minimize the residual correlation in the de-trended series. Combinations with derivative order up to the fifth and accuracy order up to the eighth have been compared among the admissible combinations, and the third-order derivative with fourth-order accuracy yielded the lowest residual correlation, justifying its adoption.

However, differently from Dosis et al. (2015), sub-sampling of the raw time series has been deliberately avoided. Many of the passes for signals tracked at the Moon are often short, and a sub-sampling factor of the order of tens of seconds would compromise the statistical significance of the stochastic model, especially at the extremes of the evaluated signal-strength range. It follows that the empirical variance reported here may therefore be interpreted as a slight overestimate of the variance that would be obtained on a strictly decorrelated subsample, with the practical consequence that the fitted model represents a conservative indicator of observable precision.

2. Construction of the analytical model

The next step pairs each $\tilde{y}_k^{(i,j)}$ with the corresponding C/N_0 estimate reported by the receiver on that channel at that epoch (Parker et al., 2025). However, these C/N_0 time series $c^{(i,j)}$ are estimates produced from the outputs of the tracking loops' correlators. Given the weak-signal conditions of LuGRE surface operations, these estimates can exhibit significant fluctuations from epoch to epoch that do not fully represent the underlying link conditions. Evidence of this can be seen in Figures 14 and 31 of Parker et al. (2026), where pronounced fluctuations are often observed as distinct features of ionospheric occultation passes.

To obtain a stable indicator of C/N_0 against which the de-trended samples can be binned, the C/N_0 time series are smoothed within each tracking pass by a moving-median filter with window length W , yielding $\bar{c}_k^{(i,j)}$. As for the de-trending, smoothing is separately operated on distinct tracking passes of the same satellite channel, and the smoothing window is adapted at the edges of each pass. Moreover, samples whose smoothed C/N_0 falls below the receiver tracking sensitivity (see Table II in Parker et al. (2026)) are excluded to avoid skewing the precision statistics in the lower C/N_0 regime.

The retained de-trended samples are then grouped by C/N_0 into bins indexed by b , with bin center c_b and fixed-width Δb . Let $\mathcal{S}_b^j$ denote the set of de-trended samples on satellite channel j that contribute to the b -th bin, defined as

$$\mathcal{S}_b^j = \left\{ \tilde{y}_k^{(i,j)} : |\bar{c}_k^{(i,j)} - c_b| \leq \frac{\Delta b}{2}, \forall i \right\}, \quad (5)$$

where the membership condition pools the samples across all satellites i tracked on that channel. To reduce the sensitivity to outliers and occasional residual transients, the empirical dispersion within each bin is estimated using the robust Median Absolute Deviation (MAD) estimator (Hampel et al., 2005):

$$\hat{\sigma}_{y,b}^j = 1.4826 \text{ MAD}(\mathcal{S}_b^j), \quad (6)$$

where $\hat{\sigma}_{y,b}^j$ is taken as a robust empirical estimate of σ_{ε^j} , namely the standard deviation of the composite residual error on band j for the generic observable y , conditional on the C/N_0 range represented by $\mathcal{S}_b^j$.

An analytical functional form is then fitted to the empirical curve to obtain a compact representation of the dependence of σ_{ε^j} on the generic C/N_0 value c . In agreement with the fundamental limits for the maximum likelihood (ML) estimation of signal synchronization parameters (Medina et al., 2020; Kuusniemi et al., 2004), the following exponential form is selected:

$$\sigma_y^j(c; \pi_y^j) = \sqrt{\beta + \alpha 10^{-\eta c/10}}, \quad \alpha > 0, \beta \geq 0, \eta > 0, \quad (7)$$

where $\pi_y^j = (\alpha, \beta, \eta)$ collects the model parameters of the observable-channel pair (y, j) . Starting from the empirical sequence $\{(c_b, \hat{\sigma}_{y,b}^j)\}_b$, the fit is solved as a nonlinear least squares problem weighted by the bin-dependent sample count N_b^j :

$$\hat{\pi}_y^j = \arg \min_{\pi \in \Pi} \sum_b N_b^j \left(\hat{\sigma}_{y,b}^j - \sigma_y^j(c_b; \pi) \right)^2, \quad (8)$$

where Π denotes the admissible parameters' set defined by the constraints in (7). The weighting is intended to give larger influence to bins with larger samples size, whose corresponding empirical estimates $\hat{\sigma}_{y,b}^j$ have higher statistical significance. Problem (8) is solved via a trust-region method, initialized from values derived from the empirical curve itself: β_0 from the variance at the highest C/N_0 bin, α_0 from the residual variance at the lowest C/N_0 bin once β_0 has been subtracted, and η_0 estimated from the decay rate of standard deviation across adjacent bins.

Table 1: Goodness-of-fit metrics for the selection of the smoothing window. For each observable $y \in \{P, \Phi, D\}$, both RMSE_y and $R^2_{\text{adj},y}$ are reported across the candidate window lengths W , split in two consecutive blocks ($W = 2\text{--}26$ s above, $W = 28\text{--}50$ s below). Entries marked with ‘–’ correspond to window lengths excluded by the enforced positivity constraint on the adjusted coefficient of determination. The selected W_y^* for each observable is highlighted in bold.

Observable	W [s]	2	4	6	8	10	12	14	16	18	20	22	24	26
P	RMSE_P	1.29	1.26	1.53	1.42	1.48	1.58	1.59	1.49	1.21	1.28	1.16	1.15	1.25
	$R^2_{\text{adj},P}$	0.98	0.98	0.93	0.94	0.94	0.94	0.94	0.96	0.99	0.99	0.99	0.99	0.99
Φ	RMSE_Φ	–	1.76	–	1.75	–	–	–	–	–	1.74	–	–	1.27
	$R^2_{\text{adj},\Phi}$	–	0.05	–	0.01	–	–	–	–	–	0.01	–	–	0.01
D	RMSE_D	1.40	1.23	1.06	1.20	1.32	1.27	1.23	1.18	1.18	1.12	1.20	1.21	1.27
	$R^2_{\text{adj},D}$	0.97	0.98	0.98	0.97	0.96	0.97	0.97	0.98	0.97	0.98	0.97	0.96	0.95
	W [s]	28	30	32	34	36	38	40	42	44	46	48	50	
P	RMSE_P	1.18	1.09	1.20	1.12	1.24	1.32	1.19	1.44	1.28	1.72	1.90	1.83	
	$R^2_{\text{adj},P}$	0.99	0.99	0.99	0.99	0.99	0.98	0.98	0.95	0.98	0.96	0.93	0.94	
Φ	RMSE_Φ	1.70	1.43	1.89	2.01	2.40	1.67	2.14	2.36	2.39	–	–	–	
	$R^2_{\text{adj},\Phi}$	0.08	0.01	0.11	0.09	0.01	0.10	0.02	0.01	0.02	–	–	–	
D	RMSE_D	1.25	1.29	1.24	1.27	1.28	1.44	1.39	1.31	1.34	1.44	1.38	1.32	
	$R^2_{\text{adj},D}$	0.96	0.96	0.96	0.96	0.96	0.95	0.96	0.95	0.94	0.95	0.93	0.94	

3. Selection of the smoothing window.

The empirical curve on which (8) is solved is not invariant with respect to the length of the smoothing window applied to the C/N_0 time series $c^{(i,j)}$: the smoothed sequence $\bar{c}^{(i,j)}$ governs the assignment of each $\tilde{y}_k^{(i,j)}$ to its bin, so a different W yields a different partition of the samples across the C/N_0 range. Different smoothing intensities therefore produce different empirical estimates of the intra-bin dispersion at any given c_b , which in turn produce different fitted parameter sets $\hat{\pi}_y^j$. It is therefore necessary to identify the window length W^* that yields the parameter set providing the best fit of the analytical model (7) to the empirically observed residual dispersion.

The adopted approach selects the window length by aggregating the goodness of fit (GoF) over the tracked channels, yielding an optimal window length W_y^* for each observable type which is shared across the channels. Operationally, a set of candidate lengths $\mathcal{W}$ is defined; for each $W \in \mathcal{W}$, $c^{(i,j)}$ are smoothed with W , the samples are re-binned, and the problem (8) is solved. Two GoF metrics are then evaluated: the root-mean-square error $\text{RMSE}_y^{(j)}(W)$ and the adjusted coefficient of determination $R^2_{\text{adj},y}(W)$.

First, not all windows are admissible, and the search is restricted to window lengths at which the analytical model produces a non-degenerate fit on each of the tracked channels,

$$\mathcal{W}_y^+ = \left\{ W \in \mathcal{W} : R^2_{\text{adj},y}(W) > 0 \text{ for all } j \right\}. \quad (9)$$

Only the window lengths in $\mathcal{W}_y^+$ are considered in the two subsequent stages. The selection of W_y^* then proceeds in two stages.

In the first stage, the RMSE estimates are made comparable across satellite channels by normalizing each value with respect to $\min_{W \in \mathcal{W}_y^+} \text{RMSE}_y^{(j)}(W)$, i.e., the smallest RMSE attained over $\mathcal{W}$ within the same channel. The normalization removes the dependence on the absolute scale of the RMSE, and the normalized values are then averaged across the channels, producing a single indicator per window length, denoted as $\text{RMSE}_y(W)$. A subset $\mathcal{W}_y^\dagger \subseteq \mathcal{W}_y^+$ is then retained, consisting of those window lengths whose (averaged normalized) RMSE does not exceed the minimum over the tested set by more than τ ,

$$\mathcal{W}_y^\dagger = \left\{ W \in \mathcal{W}_y^+ : \text{RMSE}_y(W) \leq (1 + \tau) \min_{W \in \mathcal{W}} \text{RMSE}_y(W) \right\}, \quad (10)$$

with $\tau = 0.02$ being chosen in this work. The tolerance τ allows to define a sub-class of window lengths that perform comparably well in terms of the metric.

In the second stage, R^2_{adj} is examined on the subset $\mathcal{W}_y^\dagger$. For each $W \in \mathcal{W}_y^\dagger$ the worst-channel adjusted R^2 ,

$$R^2_{\text{adj},y}(W) = \min_j R^2_{\text{adj},y}^{(j)}(W), \quad (11)$$

provides a single GoF indicator at the considered window length. Eventually, the optimal length is obtained as

$$W_y^* = \arg \max_{W \in \mathcal{W}_y^\dagger} R_{\text{adj},y}^2(W), \quad (12)$$

implementing a minmax criterion that selects the window offering the most robust fit across the tracked satellite channels for that observable. The full procedure is repeated independently for each observable $y \in \{P, \Phi, D\}$, yielding three distinct optimal window lengths W_P^*, W_Φ^*, W_D^* , each applied in the subsequent processing of the corresponding empirical curves. The numerical outcome of the procedure is reported in Table 1. Observable-specific smoothing windows are adopted because the window length determines the C/N_0 value used to assign each de-trended sample to a bin and therefore affects the resulting empirical precision curve. Since pseudorange, carrier phase, and Doppler may exhibit different sensitivities to short-term C/N_0 variations, a single smoothing window may not be equally suitable for all observables. The window is therefore selected separately for each observable type, while being kept common across its tracked channels to preserve consistency and avoid channel-specific overfitting.

4. Residual variance analysis against geometric descriptors

The analytical model $\sigma_y^j(c; \hat{\pi}_y^j)$ is mainly intended to capture the receiver-noise component $w^{(i,j)}$, whose variance reflects the average signal strength conditions existing on the channel after correlators' processing. The next step is to assess whether this C/N_0 -based description is sufficient, or whether the residual term $r^{(i,j)}$ in (1) leaves a measurable effect on $\tilde{y}^{(i,j)}$. This term was introduced to account for additional high-frequency contributions that are physically plausible in the lunar-surface geometry, even though their presence must be verified from the data.

To this end, the de-trended samples are standardized by the model prediction,

$$\bar{y}_k^{(i,j)} = \frac{\tilde{y}_k^{(i,j)}}{\gamma \sigma_y^j(c_k^{(i,j)}; \hat{\pi}_y^j)}, \quad (13)$$

where γ brings the model prediction—expressed in the units of the original observable—to the units of the de-trended sample, yielding dimensionless samples. If (7) fully describes the dispersion of $y^{(i,j)}$, the standardized samples $\bar{y}_k^{(i,j)}$ should have unit variance across the different conditions represented in the dataset. It should be noticed that (13) uses the receiver-reported estimates $c_k^{(i,j)}$ of C/N_0 , not the smoothed version used for binning in Section III.2. This is coherent with the run-time process of an estimator or a simulator, which would query the analytical model on unsmoothed C/N_0 estimates. In this perspective, this step is inherently underpinned by an assessment of model conformity to the data.

The standardized samples are then re-binned with respect to two geometric descriptors that characterize the transmit and propagation geometry rather than the receiver-side processing itself: the transmit off-boresight angle $\theta_k^{(i,j)}$, and the grazing altitude $h_k^{(i,j)}$. The latter, which quantifies the proximity of signal propagation to the terrestrial ionosphere, has been geometrically evaluated considering the point on the satellite–receiver LoS at minimum radial distance above the Earth's ellipsoid. For each descriptor, the variance of $\bar{y}_k^{(i,j)}$ is evaluated within each bin.

5. C/N_0 -conditioned precision models

IV. RESULTS

Figure 1 shows the empirical code-pseudorange, carrier-phase, and Doppler measurement precision as a function of tracking-stage C/N_0 . Markers show the empirical bin-wise standard deviation estimates, while solid lines represent the fitted analytical models. As described in Section III, the empirical curves were derived from $\tilde{y}_k^{(i,j)}$ on the different tracked satellite channels. To increase the statistical representativeness of the analysis, the samples from all the considered surface operations were pooled together on a channel-by-channel basis. Using the observable-specific smoothing windows identified in Section III.3 (i.e., 30 s for code-pseudorange, 26 s for carrier-phase, and 6 s for Doppler-shift according to Table 1), the C/N_0 time series were binned with a bin width of 0.5 dB–Hz. For each bin and observable type, the empirical standard deviation was computed from the corresponding de-trended samples, thus providing an estimate of the intra-bin standard deviation. The considered signal-strength range extends approximately from 20 to 40 dB–Hz with non-uniform level of confidence due to limited intra-bin sample statistics. The lower bound is determined by the receiver tracking sensitivity, whereas the upper bound corresponds approximately to the maximum smoothed C/N_0 value. Overall, the measurement dispersion generally decreases as the received signal strength increases, consistently with the behavior expected from conventional GNSS tracking models (Medina et al., 2020). The reduction is more pronounced in the low- C/N_0 region, whereas several curves tend to flatten at higher signal strengths, indicating that the contribution directly associated with signal strength becomes progressively less dominant. For pseudorange observables, the separation among curves is consistent with the different modulation characteristics and effective




Figure 1: Empirical observables' precision curves and corresponding fitted analytical models for code-pseudorange, carrier-phase, and Doppler observables as functions of C/N_0 for the different tracked signals. Markers represent empirical bin-wise standard deviation estimate using the smoothed C/N_0 series, whereas solid lines show the fitted model over the C/N_0 range.

bandwidths of the tracked signals. In particular, GPS L1 C/A exhibits the largest code-pseudorange dispersion at a given C/N_0 , while Galileo E1 OS shows better performance. The GPS L5 and Galileo E5a measurements are characterized by substantially lower noise variance on code-pseudorange over most of the investigated range, with empirical precision values decreasing from approximately 2 m to 30 cm. This behavior is consistent with the higher chip rate and wider effective bandwidth of the L5/E5a signals, which produce a sharper correlation function and therefore support more precise delay estimation. The observed ordering of the empirical curves is thus broadly consistent with the signal-dependent performance trends commonly expected in terrestrial GNSS applications, despite the markedly different reception conditions encountered on the lunar surface.

Carrier-phase and Doppler dispersion exhibit the same general improvement with increasing C/N_0 , although the differences among signal components are less uniform than for code-pseudorange measurements. The curves for carrier-phase in Figure 1b show a peculiar trend in the lowest C/N_0 region which is markedly signal-dependent. GPS L1 C/A has the largest dispersion, reaching approximately 6.6 cm near 20 dB – Hz, followed by a rapid decrease as the signal strength increases. Galileo E1 OS also exhibits a higher dispersion than the L5/E5a signals at low C/N_0 , although its profile rapidly converges to lower values. In contrast, GPS L5 and Galileo E5a display a considerably weaker dependence on C/N_0 , remaining close to 2.8 cm. Above approximately 25 dB – Hz, the four carrier-phase curves approach a common dispersion level. The Doppler curves show a clearer reduction in dispersion over the entire observed C/N_0 interval. Galileo E1 OS generally exhibits the highest Doppler dispersion, followed by GPS L1 C/A, whereas GPS L5 and Galileo E5a remain closer and provide the lowest dispersion levels over most of the range. Some local irregularities and deviations from a strictly monotonic trend can also be observed. These irregularities should be interpreted in light of the non-uniform distribution of the available samples across the C/N_0 bins. Consequently, the empirical standard deviations associated with different bins do not have the same statistical significance: densely-populated bins provide estimates which are more statistically reliable, whereas bins containing fewer samples should be interpreted with some caution, particularly close to the boundaries of the observed C/N_0 range. The number of samples contributing to each bin is therefore retained in the subsequent model-fitting procedure (see Section III.2) to account for the different statistical support of the empirical points. The different shapes and asymptotic behaviors observed in Figure 1 are reflected in the parameters and model formulations reported in Table 2.

The nested formulation was retained for all code-pseudorange measurements, indicating that estimating an additional exponent did not provide a sufficient improvement in the fit, as the relative reduction in the fitting error remained below the adopted threshold of 5%. For carrier phase, the selection of the full formulation for L1 C/A and E1 OS reflects the sharper transition

Table 2: Parameters of the selected exponential precision models for each observable type and signal. For the nested model, $\eta = 1$.

Observable	Band	Selected model	α	β	η
P	L1 C/A	Nested	6.397e4	1.275e-12	1
	L5	Nested	3.933e2	3.748e-3	1
	E1 OS	Nested	9.544e3	5.885e-1	1
	E5a	Nested	4.105e2	3.306e-3	1
Φ	L1 C/A	Full	1.427e4	8.016e-4	3.385
	L5	Nested	1.937e-2	7.847e-4	1
	E1 OS	Full	4.567e2	8.200e-4	3.075
	E5a	Nested	1.846e-2	7.960e-4	1
D	L1 C/A	Nested	6.888e2	1.134e-1	1
	L5	Nested	5.637e2	7.432e-2	1
	E1 OS	Full	7.498e2	1.021e-1	0.923
	E5a	Full	4.162e2	6.552e-2	0.943




Figure 2: Standard deviation of the standardized de-trended samples, evaluated per bin of raw C/N_0 , for the different tracked signals (top), and relative 95% confidence interval of the corresponding bin-wise dispersion estimates (bottom). A normalized dispersion close to unity indicates agreement between the C/N_0 -conditioned model and the observed dispersion, while the lower panel quantifies the statistical uncertainty of each estimate.

observed in the low- C/N_0 region, whereas the nested model is sufficient to describe the comparatively flatter L5 and E5a curves. At the same time, the similar β values obtained for all carrier-phase signals are consistent with their convergence towards a common dispersion level at higher C/N_0 . For Doppler, the full formulation is selected only for E1 OS and E5a. However, the corresponding exponents remain close to unity, indicating that the additional degree of freedom introduces only a moderate adjustment to the decay rate rather than a substantially different functional behavior.

Using the derived analytical functions, the de-trended samples were standardized (see Section III.4). Figure 2 shows the resulting dispersion of standardized de-trended samples for code-pseudorange, carrier-phase and Doppler observables as functions of C/N_0 . To quantify the finite-sample uncertainty associated with each bin-wise dispersion estimate, the figure also reports an approximate 95% confidence interval. Under the Gaussian approximation, specialization of the asymptotic variance of the MAD gives the relative standard error (Arachchige and Prendergast, 2026)

$$\text{SE}_{\text{rel},b}^j \simeq \frac{1.166}{\sqrt{N_b^j}}. \quad (14)$$

A two-sided interval is constructed on the logarithmic scale using the delta-method approximation (Oehlert, 1992). Defining $h_b^j = z_{1-\alpha/2} \text{SE}_{\text{rel},b}^j$, with $\alpha = 0.05$, the confidence bounds are obtained by back-transformation as

$$\text{CI}_{95\%,b}^j = \left[\hat{\sigma}_{y,b}^j e^{-h_b^j}, \hat{\sigma}_{y,b}^j e^{+h_b^j} \right] \quad (15)$$

where $\hat{\sigma}_{y,b}^j$ denotes the bin-wise MAD-based dispersion estimate of the standardized samples. The logarithmic construction guarantees positive confidence bounds. The lower panels of Figure 2 report the relative extension of the corresponding 95% confidence interval, expressed as a percentage of the estimated bin-wise dispersion. Larger values indicate greater sampling uncertainty associated with less populated bins. Over most of the observed C/N_0 range, the normalized dispersion remains close to unity, indicating that the fitted models reproduce the dominant signal-strength-dependent behavior of the observables. Larger uncertainty is observed toward the edges of the sampled C/N_0 domain, where the number of available measurements decreases. Above approximately 35 dB-Hz, a positive shift becomes apparent, particularly for code pseudorange, for which $\sigma_{P,\text{norm}}$ reaches values up to approximately 1.7. This positive shift indicates that the fitted model tends to underestimate the observed code dispersion in this region.

The mismatch should, however, be interpreted together with the associated confidence intervals. The model fit is primarily constrained by the more densely populated portion of the C/N_0 domain, whereas the highest C/N_0 bins provide weaker statistical support. Moreover, the models are derived from empirical curves constructed using smoothed C/N_0 , while the sample-wise standardization is performed using the raw receiver-reported C/N_0 . Overall, the models reproduce the dominant C/N_0 -dependent dispersion over the well-sampled range, while localized model mismatches remain at its boundaries.




Figure 3: Standard deviation of the standardized de-trended code-pseudorange samples for the different signals, evaluated as a function of the transmit off-boresight angle (left) and of the grazing altitude (right). In each panel, the normalized dispersion is shown on top and the relative 95% confidence interval of the corresponding bin-wise estimates is reported below. Departures from the expected unit value would signal a residual dependence of the code precision on the geometric descriptor not captured by the C/N_0 -conditioned model.

1. Model validity against geometric descriptors

Having assessed the ability of C/N_0 -conditioned models to reproduce the observed measurement standard deviation, the dispersion of standardized de-trended samples was further examined as a function of the transmit off-boresight angle and grazing altitude to identify possible residual dependencies not captured by the analytical model (7).

Figure 3 shows the residual code-pseudorange dispersion for the two geometric descriptors. As in Figure 2a, up to $20^\circ - 25^\circ$, the dispersion of standardized samples remains close to unity for all signal bands and is supported by large sample populations, generally ranging from several thousand to several tens of thousands of samples per bin. Beyond this region, the behavior becomes more signal-dependent. Between approximately 25° and 35° , a gradual increase in normalized dispersion is observed for some signals, together with increasing sampling uncertainty.

At larger off-boresight angles, where only GPS observations are available, a more pronounced positive shift is observed. In particular, GPS L5 reaches normalized-dispersion values of approximately 1.35–1.45 between about 58° and 66° . This behavior is compatible with a reception-regime effect associated with side-lobe tracking, although the available data do not allow a unique causal attribution. Moreover, these high-angle bins contain only a few hundred samples, approximately 98% fewer than those in the central angular range, and the corresponding deviations should therefore be interpreted with greater caution. For grazing altitudes below approximately 4000 km, the observed residual dispersion generally remains close to unity for all signal bands, with several thousand samples contributing to each bin. Larger discrepancies emerge at higher altitudes, although their interpretation depends strongly on the available statistical support. The pronounced fluctuations for Galileo E1 OS near 4000 km coincide with a rapid reduction of sample count, while the GPS estimates between 8000 km and 10 000 km are based on only a few tens to a few hundred samples per bin. At higher grazing altitudes, the behavior becomes more structured but does not follow a common monotonic trend across signals. Some signal-dependent curves return toward unit dispersion, while positive shifts reappear at the largest sampled altitudes. Therefore, grazing altitude alone does not provide a unique explanation for the residual dispersion. Rather, the data show localized geometry-related signatures whose magnitude and statistical support vary across signals.




Figure 4: Standard deviation of the standardized de-trended carrier-phase samples for the different signals, evaluated as a function of the transmit off-boresight angle (left) and of the grazing altitude (right). In each panel, the normalized dispersion is shown on top and the relative 95% confidence interval of the corresponding bin-wise estimates is reported below. Departures from the expected unit value would signal a residual dependence of the carrier-phase precision on the geometric descriptor not captured by the C/N_0 -conditioned model.

In Figure 4, the carrier-phase results show an overall more stable behavior. In the densely populated angular and grazing-altitude regions, the residual dispersion remains very close to unity for all signal bands. Larger fluctuations appear mainly in bins with reduced statistical support, particularly at intermediate and high off-boresight angles and in the more sparsely sampled grazing altitude intervals. Even in these regions, however, the deviations are generally less systematic than for code-pseudorange, and no clear monotonic dependence on either geometric descriptor emerges.

For Doppler, the standard deviation of standardized samples exhibits a broadly similar trend. Figure 5a shows the standard deviation close to unity for signals coming from the main-lobe of the transmit antenna pattern, indicating a good agreement with the C/N_0 -based model. At the largest off-boresight angles, a systematic positive shift is visible for the GPS signals: GPS L1 C/A reaches approximately 1.16–1.39 between 52.5° and 61.5° , while GPS L5 reaches values of approximately 1.22–1.42




Figure 5: Standard deviation of the standardized de-trended Doppler samples for the different signals, evaluated as a function of the transmit off-boresight angle (left) and of the grazing altitude (right). In each panel, the normalized dispersion is shown on top and the relative 95% confidence interval of the corresponding bin-wise estimates is reported below. Departures from the expected unit value would signal a residual dependence of the Doppler precision on the geometric descriptor not captured by the C/N_0 -conditioned model.

between 58.5° and 66.5° . Although these bins have lower statistical support than the central angular region, the recurrent positive shift suggests a localized reception-regime dependence that deserves further investigation. The grazing-altitude dependence of Doppler is less regular. Local departures from unity are observed across several altitude intervals, including positive shifts at intermediate and high grazing altitudes. However, these features are not reproduced consistently across the tracked signals and do not define a unique monotonic dependence on grazing altitude. Taken together, the code-pseudorange, carrier-phase, and Doppler results indicate that the C/N_0 -conditioned models capture the dominant structure of the short-term observable dispersion over the conditions most strongly represented in the LuGRE surface dataset. Local geometry-related departures from unity are nevertheless observed for individual signal–observable pairs. Their statistical support varies considerably across the sampled domain, and no common monotonic dependence on either off-boresight angle or grazing altitude emerges across signals.

Consequently, the present dataset does not provide sufficient evidence to introduce a general geometry-dependent variance term into the analytical model. This does not exclude the presence of localized geometry-related effects; rather, it indicates that the available observations are insufficient to characterize them through a unique correction applicable across signals and observables.

V. CONCLUSIONS

This work characterized the precision of LuGRE raw pseudorange, carrier-phase, and Doppler observables collected on lunar-surface operations. Empirical precision curves, derived as functions of C/N_0 for GPS L1 C/A, GPS L5, Galileo E1 OS, and Galileo E5a, were fit by compact, exponential analytical models obtained directly from flight data. The results support the use of a C/N_0 -conditioned description as the primary model for the short-term dispersion of LuGRE observables.

The adequacy of the C/N_0 -conditioned description was then probed by examining the dispersion of the standardized observables against two geometric descriptors—transmit off-boresight angle and grazing altitude—chosen because they reflect the transmit and propagation geometry rather than the receiver-side processing, and would therefore expose a contribution attributable to the residual term of the observable decomposition. Normalized dispersion remains close to unity over the most well-sampled portion of the C/N_0 domain, confirming that signal strength captures the dominant structure of the observable dispersion. Local departures are nevertheless observed, including a positive code-dispersion shift at high C/N_0 , increasing dispersion at large off-boresight angles, and structured variations across the grazing-altitude domain. These effects are signal dependent and do not define a consistent cross-signal monotonic trend. Within the statistical support of the analyzed LuGRE surface dataset, the results therefore do not justify introducing a general additional geometry-dependent variance term. The resulting flight-derived precision models can support measurement weighting in cislunar navigation filters and improve the realism of simulations for future lunar missions, subject to appropriate calibration of the receiver-reported C/N_0 level.

Future work should extend the analysis to other LuGRE mission phases and investigate the operation-wise transferability of the derived models. A correlation-aware assessment of the de-trending strategy is also required to quantify the effect of temporal correlation in the original observable time series on the inferred residual dispersion. Additional geometry-related and operational descriptors should be investigated to better characterize the localized residual signatures observed in the present analysis. Finally, future lunar GNSS experiments could target the regions of the observation domain where the largest deviations or weakest statistical support are currently found, providing denser measurements at high off-boresight angles and selected grazing-altitude intervals and enabling a more exhaustive characterization of GNSS observables in cislunar space.

ACKNOWLEDGMENTS

This study was funded within the contract n. 2021-26-HH.0 ASI/Politecnico di Torino "Attività di R&S inerente alla Navigazione GNSS nello Space volume Terra/Luna nell'ambito del Lunar GNSS Receiver Experiment". The authors gratefully acknowledge NASA for supporting the Lunar GNSS Receiver Experiment (LuGRE) through the Commercial Lunar Payload Services (CLPS) program. We also thank Qascom s.r.l. for the development of the LuGRE GNSS payload and for providing the experimental datasets used in this study. Finally, the authors acknowledge Firefly Aerospace Inc. for hosting the LuGRE payload onboard the Blue Ghost Mission 1 lander and for providing the ground truth data, payload telemetry, and operational support that were essential to this work.

REFERENCES

- Arachchige, C. N. P. G. and Prendergast, L. A. (2026). Confidence intervals for median absolute deviations. *Communications in Statistics - Simulation and Computation*, 55(1):13–22.
- Ashby, N. (2003). Relativity in the Global Positioning System. *Living Reviews in Relativity*, 6(1):1.
- Axelrad, P. and Brown, R. G. (1996). GPS Navigation Algorithms. In Parkinson, B. W. and Spilker, J. J., editors, *Global Positioning System: Theory and Applications, Volume I*, volume 163 of *Progress in Astronautics and Aeronautics*, chapter 9. American Institute of Aeronautics and Astronautics, Washington, DC.
- Cesaroni, C., Spogli, L., Guerra, M., Ghidoni, R., Ventriglia, V., Minetto, A., Nardin, A., Alfonsi, L., and Dovis, F. (2026). Observing the Earth's Plasmasphere and Ionosphere From the Lunar Surface. *Geophysical Research Letters*, 53(12):e2026GL121811. e2026GL121811 2026GL121811.
- Dovis, F., Muhammad, B., Cianca, E., and Ali, K. (2015). A Run-Time Method Based on Observable Data for the Quality Assessment of GNSS Positioning Solutions. *IEEE Journal on Selected Areas in Communications*, 33(11):2357–2365.
- Ely, T., Burt, E., Turyshv, S. G., and Jenson, E. (2026). A Study of Lunar Coordinated Time. Technical presentation, International Technical Meeting of The Institute of Navigation (ION ITM 2026), Anaheim, CA, USA. Presented on January 29, 2026.
- Fornberg, B. (1988). Generation of finite difference formulas on arbitrarily spaced grids. *Mathematics of Computation*, 51(184):699–706.
- Hampel, F. R., Ronchetti, E. M., Rousseeuw, P. J., and Stahel, W. A. (2005). One-dimensional estimators. In *Robust Statistics: The Approach Based on Influence Functions*, pages 78–186. John Wiley & Sons, Inc.
- Kaplan, E. and Hegarty, C. (2017). *Understanding GPS/GNSS: Principles and Applications*. Artech House.
- Konitzer, L., Parker, J. J. K., Ashman, B., Esantsi, N., Facchinetti, C., Dovis, F., Minetto, A., Nardin, A., Bauer, F., Ansalone, L., and Impresario, G. (2024). Science Objectives and Investigations for the Lunar GNSS Receiver Experiment (LuGRE). In *Proceedings of the 37th International Technical Meeting of The Satellite Division of The Institute of Navigation (ION GNSS+ 2024)*, pages 1061–1081.
- Kuusniemi, H. (2005). *User-Level Reliability and Quality Monitoring in Satellite-Based Personal Navigation*. Doctor of technology thesis, Tampere University of Technology, Tampere, Finland.
- Kuusniemi, H., Lachapelle, G., and Takala, J. (2004). Reliability in personal positioning. In *Proceedings of ENC GNSS 2004, The European Navigation Conference, 16-19 May, 2004, Rotterdam, The Netherlands*, page 14. Contribution: organisation=ttkt,FACT1=1.
- Medina, D., Gibson, K., Ziebold, R., and Closas, P. (2018). Determination of pseudorange error models and multipath characterization under signal-degraded scenarios. In *Proceedings of the 31st International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2018)*.

- Medina, D., Ortega, L., Vilà-Valls, J., Closas, P., Vincent, F., and Chaumette, E. (2020). Compact CRB for delay, Doppler, and phase estimation – application to GNSS SPP and RTK performance characterisation. *IET Radar, Sonar & Navigation*, 14(10):1537–1549.
- Minetto, A., Zocca, S., Vouch, O., Nardin, A., Dovis, F., Musmeci, M., and Facchinetti, C. (2025). LuNART-q: The LuGRE Quick Navigation Analysis and Reporting Tool. In *Proceedings of the 38th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2025)*, page 1377–1390, Baltimore, Maryland, USA. Institute of Navigation.
- Oehlert, G. (1992). A note on the delta method. *American Statistician*, 46(1):27–29.
- Parker, J. J. K., Dovis, F., Konitzer, L., Esantsi, N., Ashman, B., Minetto, A., Nardin, A., Vouch, O., Zocca, S., Bernardi, F., et al. (2026). GNSS Reception at the Moon: First Results of the Lunar GNSS Receiver Experiment (LuGRE). *NAVIGATION*, 73.
- Parker, J. J. K., Dovis, F., Pozzobon, O., Facchinetti, C., Anderson, B., Ansalone, L., Ashman, B., Bauer, F. H., Bernardi, F., Boschiero, M., Campagnolo, G., D’Amore, G., Fantinato, S., Fiorina, F., Ghedin, M., Guzzi, S., Heighes, C., Impresario, G., Konitzer, L., Martini, D., McKim, S., Miller, J. J., Minetto, A., Miotti, E., Montini, N., Morichi, L., Musmeci, M., Nardin, A., Pino, T., Pulliero, M., Sanathanamurthy, S., Sciacca, L., Tedesco, S., Valencia, L., Varacalli, G., Vicari, D., Vouch, O., and Zocca, S. (2025). Lunar GNSS Receiver Experiment (LuGRE) Mission Data. Dataset, Zenodo. Version 1.
- Parkinson, B. W. and Spilker, J. J., editors (1996). *Global Positioning System: Theory and Applications, Volume I*, volume 163 of *Progress in Astronautics and Aeronautics*. American Institute of Aeronautics and Astronautics, Washington, DC.
- Pi, X., Iijima, B. A., and Lu, W. (2017). Effects of Ionospheric Scintillation on GNSS-Based Positioning. *NAVIGATION*, 64(1):3–22.
- Renfro, B. A., Stein, M., Reed, E. B., and Finn, A. (2024). An Analysis of Global Positioning System Standard Positioning Service Performance for 2022. Technical Report TR-SGL-23-02, Space and Geophysics Laboratory, Applied Research Laboratories, The University of Texas at Austin, Austin, TX, USA. NAVSEA Contract N00024-17-D-6421, Task Order 5101030.
- Salós, D., Macabiau, C., Martineau, A., Bonhoure, B., and Kubrak, D. (2010). Nominal GNSS pseudorange measurement model for vehicular urban applications. In *IEEE/ION Position, Location and Navigation Symposium*, pages 806–815.
- Su, K., Xu, Y., Jiao, G., Yang, Y., Wu, Z., and Gao, P. (2026). Resilient code bias for assured PNT service. *Journal of Geodesy*, 100(6):39.
- Wang, L., Feng, Y., and Wang, C. (2013). Real-Time Assessment of GNSS Observation Noise with Single Receivers. *Journal of Global Positioning Systems*, 12(1):73–82.
- Zhang, Q., Zhao, L., and Zhou, J. (2022). A resilient adjustment method to weigh pseudorange observation in precise point positioning. *Satellite Navigation*, 3:16.