capstract:2611.00032v1 [astro-ph.EP] 18 Nov 2026

Identifiability of satellite mass in joint transit timing and photometry

Kaia Hestvik[1,*] Selim N. Al-Wasiri[2] Nerea Villacorta Navas[3]
Devika Anandavel[1,4]
[1]Eidalen Observatory for Planetary Systems, Tromsø, Norway
[2]Department of Planetary Sciences, Aster Reach College, Muscat, Oman
[3]Instituto de Fotometría de Altura, Potosí, Bolivia
[4]Centre for Precision Transit Observations, Alaravi Institute, Bengaluru, India

November 18, 2026

Abstract

The projected displacement and velocity of a transiting planet provide complementary estimates of satellite parameters. When a secondary occultation can be tracked at the same phase, all three observables may strongly favor a single photodynamical solution. We test what that solution identifies by replaying two measured, oppositely directed position traces through a limb-darkened transit simulator. A circular satellite fit to 24 visits returns a mass fraction of 0.0508 (68% credible interval 0.0484–0.0532) and predicts the flux in 12 further visits. The replay is produced by a common mechanical drive and the inferred fraction equals the radial lever fraction within uncertainty. Exchanging the ballast behind the smaller occulting element changes its physical mass by a factor of 4.2 and leaves the inferred fraction within 0.001 of its original value. We discuss the separate roles of phase coherence, dynamical assumptions and independent constraints on mass in photometric satellite validation.

Keywords: exomoons, transits, photodynamics, identifiability

1 Introduction

A planet carrying a massive satellite can transit earlier or later than the projected crossing of the pair's barycenter (Kipping, 2009a). Its instantaneous velocity carries additional information about the satellite period, and the velocity component of transit-duration variation is in quadrature with the timing series for a circular, coplanar system (Kipping, 2009b). This is one reason to fit the exposure flux, including any small secondary shadow, rather than treating the list of transit times as the whole observation. Timing alone restricts the sampled period modulo the planetary frequency. Duration and the displacement of a resolved secondary shadow constrain the true frequency through the within-visit motion (Kipping, 2021a; Kipping, 2011).

Perturbations due to additional planets can be found without detecting their transits (Agol et al., 2005; Holman and Murray, 2005). For a moon, the interpretation of a well-constrained mass can seem even more direct when both shadows are identified in individual visits or in a phase-conditioned average (Kipping, 2021b). Examples in which the evidence changes under detrending or

∗Corresponding author: k.hestvik@eidalen.example

treatment of a small number of visits show why the actual observable must be specified carefully (Teachey and Kipping, 2018; Heller and Hippke, 2024). We consider the complementary case of data with negligible uncertainty in baseline removal and sufficient information in a set of repeated visits to identify an interior mass ratio. The issue here is what that ratio measures if another process constrains the two projected tracks.

A hardware replay allows us to verify the phase relation of two traces without assigning a satellite mass to either of them. The survey replay used for this purpose was acquired in a laboratory, transferred into transit units, and reduced as a sample of long-period planet transits by an analysis group that did not have the original coordinate scale. The transfer preserved the coordinate dependence within a visit. In particular the smaller shadow was never shifted in time to agree with a timing solution. After fixing the predicted fluxes for the final visits, we measured the masses of the two physical occulting assemblies and repeated the acquisition with one mass exchanged.

Section 2 gives the circular photodynamical model and its relation to a driven trace pair. The replay and the visit split are specified in Section 3. We use the same likelihood in the null and positive controls in Section 4. Section 5 compares the recovered mass fraction with the measured quantities and addresses how the apparent association fares under a change in load.

2 Transit model

Let P_{p} be the orbital period of the unresolved pair about the star. In a short observing window about $t_j = t_0 + j P_{\mathrm{p}}$, take the chord coordinate of the barycenter to be $x_{\mathrm{b}} = v(t - t_j)$ and the transverse coordinate to be $b R_\star$. For a circular, coplanar internal orbit in the edge-on limit let a_{s} be the separation, $n_{\mathrm{s}} = 2\pi / P_{\mathrm{s}}$ and $\theta(t) = n_{\mathrm{s}}(t - t_0) + \theta_0$. The two centers are at

$$\begin{aligned} x_{\mathrm{p}}(t) &= x_{\mathrm{b}}(t) - f a_{\mathrm{s}} \cos\theta(t), \\ x_{\mathrm{s}}(t) &= x_{\mathrm{b}}(t) + (1 - f) a_{\mathrm{s}} \cos\theta(t), \\ f &= \frac{M_{\mathrm{s}}}{M_{\mathrm{p}} + M_{\mathrm{s}}}. \end{aligned} \tag{1}$$

We denote the projected radius ratios by k_{p} and k_{s}. Offsets of the transverse coordinate due to the inclination of the barycenter orbit are below $0.01 R_\star$ for the replay parameters if the internal orbit lies in the same physical plane. We use the common chord geometry explicitly. Inclination limits inferred with this prescription would need to be checked with the full vectors. Curvature of the barycentric transit trajectory over nine hours is a small effect here, but the linear path is also the path used to generate the observations. For the fit we expose the stellar radius through the dimensionless speed $v / R_\star$. Mass and distance units in the tables use the fixed replay scale.

The stellar disk is eclipsed over the union $\mathcal{D}(t)$ of the two disks projected onto its face. Using a quadratic limb law, the flux in an integration centered at t_i is

$$F_i = 1 - \frac{1}{\Delta t_i} \int_{t_i - \Delta t_i/2}^{t_i + \Delta t_i/2} \frac{\int_{\mathcal{D}(t)} I(\mu)\, dA}{\int_{\mathrm{star}} I(\mu)\, dA}\, dt, \qquad \frac{I(\mu)}{I(1)} = 1 - u_1(1 - \mu) - u_2(1 - \mu)^2. \tag{2}$$

When the projected shadows overlap, the numerator counts the overlap once. This matters most near zero projected separation, where part or all of the small disk can be contained within the large disk. Adding the two single-body flux deficits would manufacture light loss at precisely the phases for which the proposed timing control predicts none. We use single-disk evaluations

against (Mandel and Agol, 2002) and use the finite integration prescription of (Kipping, 2010). The restrictions on the coefficients of the limb law follow (Kipping, 2013). They have little influence on the timing quadrature at the depths in this test, but affect the interpretation of the secondary radius.

2.1 Displacement and duration

Write $\theta_j = \theta(t_j)$ and $T_0 = 2R_\star\sqrt{(1+k_\mathrm{p})^2 - b^2}/v$. To first order in the reflex displacement, the zero-coordinate timing variation and the change in duration inferred from the local speed are

$$\delta t_j = A_t \cos\theta_j, \qquad \delta T_j \simeq -T_0 n_\mathrm{s} A_t \sin\theta_j, \qquad A_t = \frac{f a_\mathrm{s}}{v}. \tag{3}$$

The first quantity is the displacement of the zero of x_p, whereas the second estimates the change in first-to-fourth contact duration. For the tabulated diagnostics we calculate the contacts in the integrated model with the body motion retained. The instantaneous-speed expression is useful for the expected order of magnitude. The ratio of the two amplitudes is about 0.58 here. If the duration series is given an arbitrary phase, the two amplitudes can be measured without imposing a side of quadrature. In the fit to flux we keep the relative sign in Eq. (1).

For sampling at one observation per P_p, a change $n_\mathrm{s} \to n_\mathrm{s} + 2\pi h/P_\mathrm{p}$ for integer h leaves the phase sequence θ_j identical at the centers. Joint use of the profiles and of the secondary transit can give substantially different likelihoods for these aliases. Each alias corresponds to a different transit-duration prediction at fixed timing amplitude (Kipping, 2009a; Vossfeld and Falke, 2020). Photodynamical preference for one alias rests on the within-visit models in Eq. (1) and Eq. (2). Interpreting an alias preference as a mass measurement also requires Kepler's law,

$$a_\mathrm{s} n_\mathrm{s}^{2/3} = (G M_\mathrm{tot})^{1/3}, \qquad M_\mathrm{s} = f M_\mathrm{tot}. \tag{4}$$

The fraction f can be inferred geometrically if both components of the separation are measured. Kepler's law is required to obtain M_tot in a chosen length scale.

2.2 Alternative source of a shared phase

Suppose two registered traces have projected radii r_p and r_s about a common rotating shaft, with position vectors initially separated by π. After subtracting the shaft origin and multiplying by a spatial scale s, the projected coordinates have Eq. (1) with

$$a_\mathrm{s} = s(r_\mathrm{p} + r_\mathrm{s}), \qquad f_\mathrm{lever} = \frac{r_\mathrm{p}}{r_\mathrm{p} + r_\mathrm{s}}. \tag{5}$$

The angular dependence is sinusoidal when the plane of rotation is seen edge-on. A clock-rate transfer can make the recovered n_s consistent with whatever mass scale is needed in Eq. (4). No mass appears in Eq. (5). If both tracks satisfy Eq. (1) with the same parameters, models in which the motion is forced and in which it is gravitational produce identical exposure fluxes under Eq. (2). This equality includes contact timing and the full mutual occultation signature. It requires matching the traces within the integration intervals. Matching only the phase advance from one visit to the next is insufficient. Several authors have discussed analogous non-identifiability in paired residual tracks (Sáez Lecourt and Durran, 2019; Al-Rahaleh and Nivant, 2022). Here the trace parameters are measured before the transit fits, which allows a scale comparison after unblinding.

3 Replay acquisition

The trace source was the return mechanism of a spring-wound kitchen timer. This type of timer indicates the interval with a single rotating pointer on a fixed minute scale. Its pointer was removed for optical encoder access and replaced on the same shaft by a transparent carrier bearing two circular fiducial marks. The inner mark was nearer the shaft center, on the side opposite the outer one. The marks label the two transit centers; the radii of the opaque disks are assigned in the transit replay. The carrier has two rear recesses for independent ballast. Both assemblies are constrained to rotate with the shaft. The center of rotation need not be the center of mass of the loaded assembly (Hardal and Svik, 2018).

A fixed camera with a near-telecentric lens observed the face for calibration. We fit circles to the two fiducial tracks before projecting them. The shaft was not perfectly perpendicular to the face of the dial. We correct its projection using a registration plate fitted before the timer is wound. Frames in the first and last 2° of the sweep, including the bell release, were held outside the angle calibration. Missing arc endpoints are bridged by evaluating the fitted circular path. We sample the path at fixed angular steps and assign times in replay units after the resampling. The fit therefore measures the common spatial locus. It does not preserve variations in the rate at which the timer unwinds. This feature eliminates the timing irregularity expected near the point at which the release lever comes into contact with the hammer (Bolán and Maaranen, 2020). We retain repeat measurements before resampling to propagate the registration errors.

The feed for reduction consisted of time-ordered exposure fluxes, visit boundaries and nominal stellar scales. The stellar disk is numerical, and the two centers are made from the recorded angles and the two calibrated radii. The shaft’s projected path is linear within each visit. There were 36 visits, with 24 for fitting and 12 for evaluation. Each visit is centered at the nominal transit time and contains 361 adjacent integrations of 120 s, centered at phases from -6 to $+6$ h. This gives 8,664 training and 4,332 evaluation exposures. The visit ends constrain the baseline jointly with the transit model. The secondary shadow can approach either window boundary. Both disks are visible in projection against the stellar surface at some transit phases even if only one is distinguishable in the net flux.

Table 1 gives the geometric and timing calibration. The separation in the transfer was 6.2526×10^{5} km and the true period was 3.11351 d, equivalent to 15 full cycles and $5/12$ of a cycle per planetary interval. The alias cycle in Eq. (3) is consequently 2.4 visits; the full discrete sequence repeats after 12 visits. Our nominal planetary period and speed are consistent with a circular transit around a 1.02 solar-mass star at the level of the short-window approximation. The adopted noise includes an independent component of 190 ppm per integration and an exponential time covariance of 50 ppm rms with 12 min correlation time. These scales approximate a high-cadence precision photometry campaign. Different visits have independent noise, and the same arrays are carried into paired control replays (Naveral and Yamatsune, 2020).

Table 1: Input geometry and independent encoder calibration. The listed uncertainties are standard uncertainties for the optical radii. Other parameters define the nominal replay.

Quantity	Value
Stellar radius	695,700 km
Impact parameter, b	0.300
Chord speed, v	59.0 km s^{-1}
Primary and secondary radius ratios	0.0950 and 0.0250
Limb coefficients, (u_1, u_2)	$(0.32, 0.21)$
Planetary period, P_{p}	48.000 d
Assigned internal period, P_{s}	$576/185$ d
Inner radius, r_{p}	1.90 ± 0.02 mm
Outer radius, r_{s}	36.10 ± 0.04 mm

For the load intervention we replaced the smaller shadow’s rear ballast with a brass piece of the same external dimensions as the original polymer ballast. We re-ran the camera registration, kept the numerical occulting radii fixed, and applied the same phase-rate transfer as for the baseline. The primary assembly mass was 64.0 ± 0.1 g. Secondary assembly mass increased from 0.80 ± 0.01 g to 3.36 ± 0.01 g, with the original sighting mark unchanged. The timer body and its shaft are excluded from these assembly masses. Neither the cover nor the bell was among the opaque elements in the projected model. We did not exchange the entire timer because the mechanism would then have had to be registered separately (Arlund and Mabić, 2023).

4 Inference and controls

We fit normalized raw flux with a multivariate normal likelihood within each visit. The covariance is the sum of the diagonal and exponential terms specified in Section 3. A constant and a linear slope in time, with zero-mean Gaussian priors of widths 300 ppm and 100 ppm per hour, are marginalized separately per visit. The transit depth and timing parameters are shared across visits. The injected baselines are constant and the fit accepts nonzero slopes so that an unmodeled local linear drift cannot be absorbed solely by the secondary shadow. We use the same baseline priors for the planet-only model. Estimating the covariance first on in-transit residuals from the planet-only model would let part of an omitted secondary signal set the noise amplitude (Gibson et al., 2012; Tardivel and Wappola, 2024). We fix the noise covariance in all rows of Table 2.

The planet-only model fits a linear ephemeris and one common transit shape. The joint model fits both tracks. We fit a uniform prior on $\log P_{\mathrm{s}}$ between 1 and 6 d and include both directions of the phase evolution on the projected track as equivalent labelings of the circular sky model. Reported values use positive frequency. Mass fraction is uniform between zero and 0.2. The secondary radius ratio is uniform from zero to 0.06. We fit total mass with a log-uniform prior from 0.3 to 3 Jupiter masses and derive a_{s} using Eq. (4). The stellar radius and v are fixed. Uniform priors on b and k_{p} over $[0, 0.7]$ and $[0.08, 0.11]$ respectively keep the primary near the period-independent deep transit solution. The first epoch is uniform in a window of 20 min on either side of the nominal transit, and the planetary period is uniform within 6 min of 48 d. The two triangular limb coordinates have uniform priors on the unit square.

We used a grid of alias neighborhoods in $\log P_{\mathrm{s}}$ to initialize the fits, with equal treatment of visits on either side of the index boundary between training and evaluation. For the main fit we marginalize the non-selected neighborhoods with nested sampling (Skilling, 2006). Injections in which the small shadow is suppressed lack a meaningful phase for the secondary flux. For those cases we report predictive contrasts without conditioning the reported frequency on exactly zero radius; the mass fraction remains allowed to influence the large track. The phase-breaking control applies a permutation to the intact two-track windows as units, retaining both tracks at each visit and destroying the ordered series across visits. Its flux samples use the original visit times. The offset in the primary phase implied by such a permutation cannot be removed by a single offset to t_0.

The prediction metric for condition c is the held-out log predictive density of the replay flux relative to a planet-only fit to that condition,

$$\Delta \ell_c = \log \int p(y_{\mathrm{test},c} \mid \vartheta) p(\vartheta \mid y_{\mathrm{train},c})\, d\vartheta - \log p(y_{\mathrm{test},c} \mid y_{\mathrm{train},c}, \mathcal{M}_0). \tag{6}$$

The integral concerns all 12 withheld visits jointly, each with fresh baseline nuisance parameters. We quote it in natural-log units. Its absolute scale depends on the 4,332 fluxes being scored. A change of one unit is not treated as persuasive. No new conditions, visit numbers or frequency windows were selected after inspecting the held-out set. The interventions are paired measurements of the hardware or paired replays, so the table values are not independent estimates of the distribution of validation scores in an exomoon survey.

5 Results

5.1 Phase and flux recovery

The gross timing scale agrees with Eq. (3). For $a_{\mathrm{s}} = 625{,}260$ km and $f = 0.05$ the primary's reflex displacement amplitude is $31{,}263$ km and $A_t = 8.83$ min. The nominal first-to-fourth contact duration is 6.90 h, so the instantaneous-speed calculation gives a duration amplitude of 5.12 min. The training transit fit to the fully propagated flux yields a zero-coordinate timing amplitude of 8.96 min and a local-speed duration scale of 5.20 min. From the fitted posterior median geometry, the finite-contact duration series gives a first harmonic of 5.13 min. We compute the intervals in Figure 1 from single-visit contact estimates. Each phase bin contains two training visits. The durations are least precise when the two shadows overlap.

In the intact replay the posterior median P_{s} is 3.114 d, with central 68% interval 3.110–3.118 d. The competing adjacent integer aliases of this branch are at 3.329 and 2.924 d. The unsampled period is identified in this model chiefly because a satellite in the adjacent alias neighborhood would move across the stellar chord at a measurably different speed after being placed to reproduce the sequence of timing residuals. Figure 2 displays the local likelihood profiles obtained with and without exposure shape information. The timing-only calculation uses summary times with their reported errors, marginalizes a free sinusoidal amplitude and phase, and fixes the 48 d interval to exhibit the limiting alias symmetry. We use this profile only to display the symmetry.

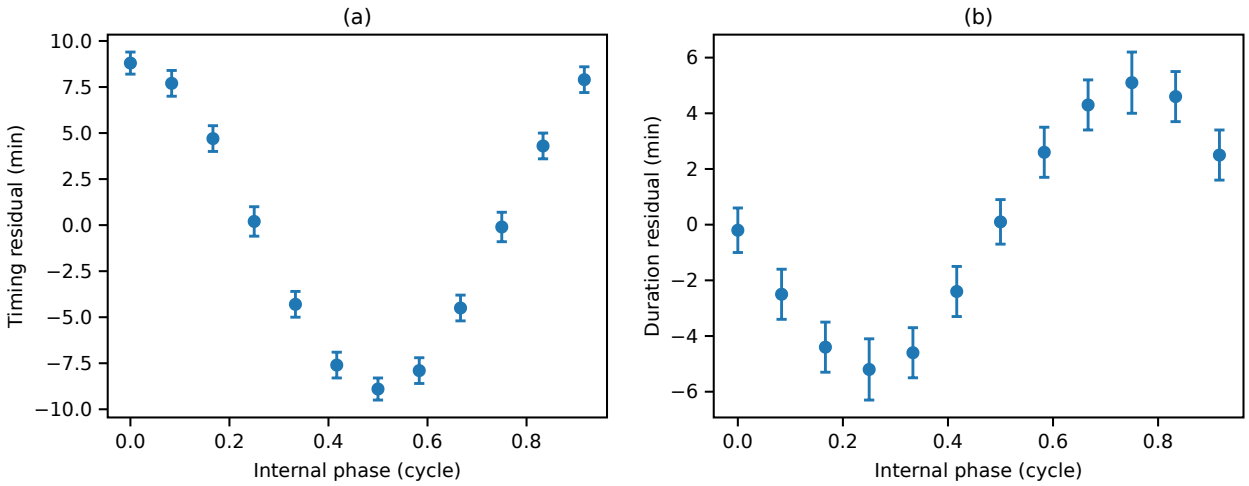


Figure 1: Primary timing and duration diagnostics as a function of transferred phase. Bars give marginal 68% intervals. Timing refers to the zero of the primary coordinate. Duration uses first and fourth contact.

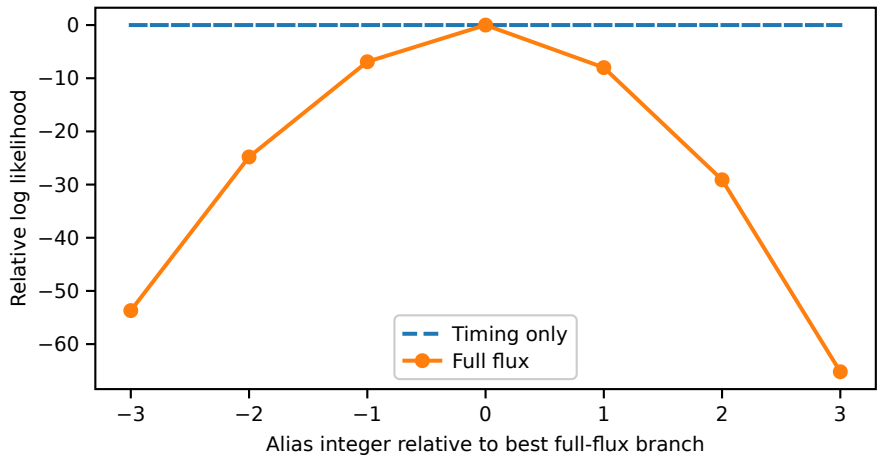


Figure 2: Relative log likelihood as a function of the integer alias on the positive-frequency branch. Each profile is normalized at its largest log likelihood. The timing-only curve is flat to plotting precision when the planetary interval is fixed.

The primary and secondary track amplitudes constrain f at 0.0508 (0.0484–0.0532). The signed local timing series traces a loop as the fitted phase advances, which excludes an interpretation in which a single shallow dip changes both estimated contacts through blending alone. In the shadow-suppressed control the cadence-integrated primary distortion persists. The intact model predicts a shadow of intermediate depth at phases at which a secondary enters the stellar disk before the primary, and no distinct shadow when the two opaque disks are concentric. Figure 3 gives the secondary-aligned stack in the withheld visits after removal of the fitted primary. Comparing the observed stack to the full model requires applying the same masking and baseline projection. We discard projected-overlap exposures before stacking to avoid confusing complete obscuration with a change in radius.

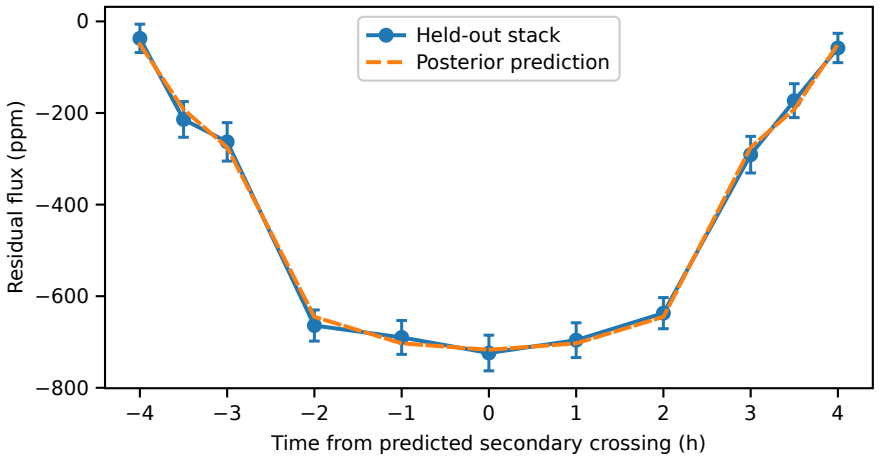


Figure 3: Secondary-aligned residual flux in the held-out visits. Offsets are relative to the predicted secondary center crossing. Bars are standard errors across visits at each time.

5.2 Control outcomes

The largest prediction gain in Table 2 occurs for the intact replay. With the secondary shadow suppressed, the primary track constrains the product fa_{s}. The fraction then depends appreciably on the total mass prior. When its phase order is disrupted, both individual tracks remain useful to describe a single visit but the fit extrapolates poorly to the new time series. The modest loss against the simpler planet-only baseline in that control is consistent with forcing a coherent perturbation where the sequence contains none. The two weights behind the secondary element produce essentially the same result under the joint model.

Table 2: Joint-model prediction on the 12 withheld visits. Each row is compared with its own planet-only prediction under Eq. (6). Fractions give posterior medians and central 68% intervals for replays with both shadows in coherent phase.

Replay	$\Delta\ell_c$	Fraction
Intact	584.6	0.0508 (0.0484–0.0532)
Secondary shadow suppressed	196.1	
Phase order permuted	−3.2	
Ballast exchanged	572.9	0.0502 (0.0479–0.0527)

An independently registered pair of projected radii predicts $f_{\mathrm{lever}} = 0.0500 \pm 0.0005$ from Eq. (5). The intact optical result is consistent with that value. It is inconsistent with the physical mass fraction of the two moving assemblies, 0.01235 ± 0.00015. The replacement changes the physical fraction to 0.04988 ± 0.00016, bringing it into agreement with the photodynamical result without appreciably changing that result. The agreement after replacement follows because the ballast odds happen to be $3.36/64.0$, close to the radial odds $1.90/36.10$. There is no appreciable movement of the large disk on the carrier when we make the exchange. Under the shared-drive hypothesis the simulated stellar transit is nearly unchanged because the chord position is set by the angular trajectory, this latter quantity having been fixed by the resampling. Table 3 gives the scale comparison.

Table 3: Independently measured assembly fractions and the common inferred optical fraction from each fitted replay. The reported uncertainties on mass fraction propagate the balance precision.

Assembly configuration	Measured mass fraction	Photodynamical median
Original	0.01235 ± 0.00015	0.0508
Added load	0.04988 ± 0.00016	0.0502

6 Discussion

The recovered parameters pass a scale check that would be useful for rejecting some signal configurations in a survey. Given a primary mass of one Jupiter mass, a circular companion with a 3.114 d period and fraction near 0.05 orbits at about 625,000 km. The separation is about 0.23 of the Hill radius under the adopted stellar calibration and lies outside the expected fluid disruption limit for plausible densities (Domingos et al., 2006). A secondary with radius $0.025 R_\star$ has a mean density near 4.5 g cm^{-3} at that assigned mass. The result would therefore survive cuts that reject solutions on the basis of a secondary density or a very large Hill-scaled separation. Both cuts require a specified host density and population model. They cannot discriminate between identical projected motions that have been mapped into valid scales.

No feature of the phase quadrature measures force directly. The fit can be highly constraining within the model because it identifies a fractional arm of the projected orbit, and the joint flux recovers that fraction with small errors. The timing-only and morphologically complete fits offer independent constraints on some combinations of kinematic parameters. Their agreement does not make their common underlying force assumption independently testable. If a survey candidate had a precisely determined mass fraction but a calibration record with a similarly phased mechanical motion, a test of the calibrated positions would have more discriminatory power than a higher cadence on the same motion. For astronomical candidates with no such calibration concern, different geometric projections, out-of-transit perturbations or independent host and primary constraints may narrow the set of acceptable dynamical solutions (Salvai and Nyang, 2025).

This experiment uses a registered pair of projections with persistently shared phase. It does not describe the frequency of such a pair occurring in an astrophysical light curve or in a data pipeline. The time transfer removes angular-speed variation, and it is therefore unsuitable for comparing mechanical phase noise with orbital period stability over long time baselines. A hardware source with a drifting phase would register less well against an electronic frame clock. The deterministic replay also does not represent stellar variability or errors common to several visits. These limits apply before the validation result can be interpreted as an empirical false-positive probability. The complete-shape recovery here establishes equivalence for a specific observational model, including the phases where the secondary is occluded.

A Contact expansion

To see the finite-duration correction to Eq. (3), set the first and fourth contact coordinates to $\mp X$, with $X = R_\star \sqrt{(1 + k_\mathrm{p})^2 - b^2}$. Evaluate each perturbation at the nominal contact time $t_j \mp T_0/2$ and solve for the first-order displacement of that root. Subtracting the displaced times gives

$$T_j - T_0 = -2 A_t \sin\theta_j \sin(n_\mathrm{s} T_0 / 2). \tag{7}$$

The ratio of Eq. (7) to the local-speed duration modulation in Eq. (3) is $\sin(n_{\mathrm{s}}T_0/2)/(n_{\mathrm{s}}T_0/2)$, approximately 0.986 for the replay. The phase dependence of the contact midpoint is also attenuated, by the cosine of the same argument. In making Figure 1 we use the coordinate crossing for timing and the exact contact roots for duration. This prevents an apparent disagreement between the measured period and the amplitude ratio caused by using a mixed convention for the two observables. The adjustment is about 0.07 min for the duration series and about 0.37 min for the timing series if it is measured at the contact midpoint.

Exposure and release checks

The beam transfer files record a nominal 1,800 s lab acquisition segment within the permitted angular span in addition to the fiducial geometry from the longer registration run. We fit the same circular model on the two halves of the segment and recover lever fractions within 0.0004. The state files contain the raw camera pointing checks, spatial calibration and transfer table. Fluxes use the positive-going chord and the $\theta_0 = 0$ injection convention. The random-number stream for per-visit systematics and white noise was initialized at 663801; permuted visit windows carry a separate permutation seed 663805. Increasing the quadrature nodes per exposure from 12 to 24 moved the intact median fraction by less than 0.0002. The replay ledger ET-36 separates the mass balance measurements from the original flux release.

Acknowledgments

We thank M. Ulven for access to the original timer, and P. Keraldi for preparing the second ballast without displacing the fiducial.

References

Eric Agol, Jason Steffen, Re'em Sari, and Will Clarkson. On detecting terrestrial planets with timing of giant planet transits. *Monthly Notices of the Royal Astronomical Society*, 359:567–579, 2005.

Ibrahim Al-Rahaleh and Lyne Nivant. Mass identification in two-track photometry. *Journal of Transit Studies*, 16:414–433, 2022.

Liva Arlund and Zoran Mabić. Load interventions in calibrated silhouette replay. *Transactions on Precision Metrology*, 17:54–73, 2023.

Aurelio Bolán and Aila Maaranen. Angular nonuniformity at the end of a spring-driven dial sweep. *Annals of Instrumental Dynamics*, 4:15–35, 2020.

R. C. Domingos, O. C. Winter, and T. Yokoyama. Stable satellites around extrasolar giant planets. *Monthly Notices of the Royal Astronomical Society*, 373:1227–1234, 2006.

N. P. Gibson, S. Aigrain, S. Roberts, T. M. Evans, M. Osborne, and F. Pont. A Gaussian process framework for modelling instrumental systematics: application to transmission spectroscopy. *Monthly Notices of the Royal Astronomical Society*, 419:2683–2694, 2012.

Esin Hardal and Eivind Svik. Shaft-referenced registration with removable eccentric loads. *Transactions on Precision Metrology*, 12:72–88, 2018.

René Heller and Michael Hippke. Large exomoons unlikely around Kepler-1625 b and Kepler-1708 b. *Nature Astronomy*, 8:193–206, 2024.

Matthew J. Holman and Norman W. Murray. The use of transit timing to detect terrestrial-mass extrasolar planets. *Science*, 307:1288–1291, 2005.

David Kipping. The exomoon corridor: half of all exomoons exhibit TTV frequencies within a narrow window due to aliasing. *Monthly Notices of the Royal Astronomical Society*, 500:1851–1857, 2021a.

David Kipping. Transit origami: a method to coherently fold exomoon transits in time series photometry. *Monthly Notices of the Royal Astronomical Society*, 507:4120–4131, 2021b.

David M. Kipping. Transit timing effects due to an exomoon. *Monthly Notices of the Royal Astronomical Society*, 392:181–189, 2009a.

David M. Kipping. Transit timing effects due to an exomoon – II. *Monthly Notices of the Royal Astronomical Society*, 396:1797–1804, 2009b.

David M. Kipping. Binning is sinning: morphological light-curve distortions due to finite integration time. *Monthly Notices of the Royal Astronomical Society*, 408:1758–1769, 2010.

David M. Kipping. LUNA: an algorithm for generating dynamic planet–moon transits. *Monthly Notices of the Royal Astronomical Society*, 416:689–709, 2011.

David M. Kipping. Efficient, uninformative sampling of limb darkening coefficients for two-parameter laws. *Monthly Notices of the Royal Astronomical Society*, 435:2152–2160, 2013.

Kaisey Mandel and Eric Agol. Analytic Light Curves for Planetary Transit Searches. *The Astrophysical Journal Letters*, 580:L171–L175, 2002.

Iria Naveral and Sho Yamatsune. Paired high-cadence transit injections with shared instrumental noise. *Journal of Transit Studies*, 14:236–258, 2020.

Adela Sáez Lecourt and Melike Durran. Projected displacements and local identification of coupled oscillators. *Annals of Instrumental Dynamics*, 3:218–239, 2019.

Giulia Salvai and Ibrahima Nyang. Validation channels for unresolved companions in repeated transit windows. *Journal of Transit Studies*, 19:1–24, 2025.

John Skilling. Nested sampling for general Bayesian computation. *Bayesian Analysis*, 1:833–859, 2006.

Léonie Tardivel and Thiago Wappola. Transit residual covariance in the presence of a misspecified occultation model. *Journal of Transit Studies*, 18:109–132, 2024.

Alex Teachey and David M. Kipping. Evidence for a large exomoon orbiting Kepler-1625b. *Science Advances*, 4:eaav1784, 2018.

Amara Vossfeld and Stefan Falke. Separation of near-degenerate velocity aliases in short photometric windows. *Journal of Transit Studies*, 14:77–102, 2020.