

STRONG LENS TIME DELAY CHALLENGE. II. RESULTS OF TDC1

KAI LIAO^{1,2,20}, TOMMASO TREU^{2,20}, PHIL MARSHALL³, CHRISTOPHER D. FASSNACHT⁴, NICK RUMBAUGH⁴, GREGORY DOBLER^{5,21}, AMIR AGHAMOUSA⁶, VIVIEN BONVIN⁷, FREDERIC COURBIN⁷, ALIREZA HOJJATI^{8,9}, NEAL JACKSON¹⁰, VINAY KASHYAP¹¹, S. RATHNA KUMAR¹², ERIC LINDER^{13,14}, KAISEY MANDEL¹¹, XIAO-LI MENG¹⁵, GEORGES MEYLAN⁷, LEONIDAS A. MOUSTAKAS¹⁶, TUSHAR P. PRABHU¹², ANDREW ROMERO-WOLF¹⁶, ARMAN SHAFIELOO^{6,17}, ANETA SIEMIGINOWSKA¹¹, CHELLIAH S. STALIN¹², HYUNGSUK TAK¹⁵, MALTE TEWES¹⁸, AND DAVID VAN DYK¹⁹

¹ Department of Astronomy, Beijing Normal University, Beijing 100875, China

² Department of Physics, University of California, Santa Barbara, CA 93106, USA

³ Kavli Institute for Particle Astrophysics and Cosmology, P.O. Box 20450, MS29, Stanford, CA 94309, USA

⁴ Department of Physics, University of California, 1 Shields Avenue, Davis, CA 95616, USA

⁵ Kavli Institute for Theoretical Physics, University of California Santa Barbara, Santa Barbara, CA 93106, USA

⁶ Asia Pacific Center for Theoretical Physics, Pohang, Gyeongbuk 790-784, Korea

⁷ EPFL, Lausanne, Switzerland

⁸ Department of Physics and Astronomy, University of British Columbia,

6224 Agricultural Road, Vancouver, BC V6T 1Z1, Canada

⁹ Department of Physics, Simon Fraser University, 8888 University Drive, Burnaby BC V5A1S6, Canada

¹⁰ University of Manchester, School of Physics & Astronomy, Jodrell Bank Centre for Astrophysics, Manchester M13 9PL, UK

¹¹ Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA

¹² Indian Institute of Astrophysics, II Block, Koramangala, Bangalore 560034, India

¹³ Lawrence Berkeley National Laboratory and University of California, Berkeley, CA 94720, USA

¹⁴ Korea Astronomy and Space Science Institute, Daejeon 305-248, Korea

¹⁵ Department of Statistics, Harvard University, 1 Oxford Street, Cambridge, MA 02138, USA

¹⁶ Jet Propulsion Laboratory, California Institute of Technology, M/S 169-506,

4800 Oak Grove Drive, Pasadena, CA 91109, USA

¹⁷ Department of Physics, POSTECH, Pohang, Gyeongbuk 790-784, Korea

¹⁸ Argelander-Institut für Astronomie, Auf dem Hügel 71, D-53121 Bonn, Germany

¹⁹ Department of Mathematics, Imperial College London, London SW7 2AZ, UK

Received 2014 August 27; accepted 2014 December 8; published 2015 February 3

ABSTRACT

We present the results of the first strong lens time delay challenge. The motivation, experimental design, and entry level challenge are described in a companion paper. This paper presents the main challenge, TDC1, which consisted of analyzing thousands of simulated light curves blindly. The observational properties of the light curves cover the range in quality obtained for current targeted efforts (e.g., COSMOGRAIL) and expected from future synoptic surveys (e.g., LSST), and include simulated systematic errors. Seven teams participated in TDC1, submitting results from 78 different method variants. After describing each method, we compute and analyze basic statistics measuring accuracy (or bias) A , goodness of fit χ^2 , precision P , and success rate f . For some methods we identify outliers as an important issue. Other methods show that outliers can be controlled via visual inspection or conservative quality control. Several methods are competitive, i.e., give $|A| < 0.03$, $P < 0.03$, and $\chi^2 < 1.5$, with some of the methods already reaching sub-percent accuracy. The fraction of light curves yielding a time delay measurement is typically in the range $f = 20\%–40\%$. It depends strongly on the quality of the data: COSMOGRAIL-quality cadence and light curve lengths yield significantly higher f than does sparser sampling. Taking the results of TDC1 at face value, we estimate that LSST should provide around 400 robust time-delay measurements, each with $P < 0.03$ and $|A| < 0.01$, comparable to current lens modeling uncertainties. In terms of observing strategies, we find that A and f depend mostly on season length, while P depends mostly on cadence and campaign duration.

Key words: gravitational lensing: strong – methods: data analysis

1. INTRODUCTION

The past decade has seen the emergence of a concordance cosmology, Λ CDM, in which the contents of the universe are dominated by dark matter and dark energy. Even though the basic parameters appear to be robustly measured, more stringent measurements are sought as a way to improve our understanding of the nature of these mysterious components, as well as a way to test the model against signatures of new physics (Suyu et al. 2012; Weinberg et al. 2013).

Achieving better cosmography means two things. On the one hand, increasingly higher quality data are being obtained (e.g., Planck Collaboration et al. 2014) in order to improve the precision of each method. On the other hand, independent observational methods are being exploited to break the degeneracies inherent to each method and to uncover unknown systematic uncertainties, thus improving accuracy. With precision and accuracy rigorously under control, potential inconsistencies might reveal new physics, such as the presence of additional families of neutrinos or deviations from general relativity.

In the past few years, strong lens time delays (Refsdal 1964; Kochanek 2002) have made something of a comeback, becoming an increasingly popular probe of cosmography (Oguri 2007; Coe & Moustakas 2009; Dobke et al. 2009; Paraficz & Hjorth 2010; Treu et al. 2013; Sereno & Paraficz 2014).

²⁰ Also at Department of Physics and Astronomy, University of California, Los Angeles CA 90095, USA.

²¹ Also at Center for Urban Science + Progress, New York University, Brooklyn, NY 11201, USA.

The configuration most suitable for this work consists of a quasar with variable luminosity, being lensed by a foreground elliptical galaxy that creates multiple images of the quasar (e.g., Treu 2010, for a recent review). Differences in optical paths and gravitational potentials give rise to time delays between the images. In turn, the observable time delays, combined with a model of the mass distribution in the main deflector and along the line of sight, provide information on the so-called time-delay distance, which is a combination of angular diameter distances. The time delay distance is primarily sensitive to the Hubble constant (Suyu et al. 2013), but can also constrain other cosmological parameters, especially with large numbers of time delay systems and in combination with other methods (Paraficz & Hjorth 2009; Linder 2011).

At the time of writing, only a fraction of the hundred or so known gravitationally lensed quasars has well-measured time delays, owing to the considerable observational challenge associated with this measurement. Accurate time delays in the optical require long and well-sampled light curves as well as sophisticated algorithms that account for data irregularities and astrophysical effects such as microlensing (e.g., Tewes et al. 2013a). Radio wavelength light curves have been used to determine time delays with great accuracy (e.g., Fassnacht et al. 2002), but unfortunately are restricted to the radio-loud subset of systems. In all cases, the success rate is limited by the intrinsic variability of the sources.

The number of systems with known time delays is about to increase dramatically. In the immediate future, as more lensed quasars are discovered (e.g., via the STRIDES program²²), there will be more opportunities to identify highly variable systems in cosmologically favorable configurations for targeted follow-up. The state-of-the-art project COSMOGRAIL²³ with its newly developed methods (Tewes et al. 2013a) has shown the potential power of extracting time delay data from sparsely sampled photometric data (Tewes et al. 2013b). In the near future, the upcoming cadenced optical imaging surveys will provide light curves for large samples of lensed quasars. For example, the Large Synoptic Survey Telescope (LSST; LSST Science Collaboration et al. 2009; Ivezić et al. 2008) will repeatedly observe approximately 18,000 deg² of sky for 10 years, and is predicted to find and monitor several thousand time delay lens systems (Oguri & Marshall 2010; LSST Dark Energy Science Collaboration 2012).

In preparation for this wealth of light curves, it is crucial to carry out a systematic study of the current algorithms for time delay determination. Such an investigation has two main goals. The first is to determine whether current methods have sufficient precision and accuracy to exploit the kind of data anticipated in the next decade. Identifying limitations and failure modes of current methods is a necessary step to develop the next generation of measurement algorithms. In parallel, the second goal is to test the impact of different observational strategies. For example, what kind of cadence, duration, and sensitivity is required to obtain precise and accurate time delays? Is the LSST baseline strategy sufficient to meet the goals of time delay cosmography or can we identify changes that would improve the outcome?

With these two goals in mind, a time delay challenge (TDC) was initiated in 2013 October. The challenge “Evil” Team (G.D., C.D.F., K.L., P.J.M., N.R., T.T.) simulated large numbers of

time delay light curves, including all anticipated physical and experimental effects. The wider community was then invited to extract time delay signals from these mock light curves, blindly, using their own algorithms as “Good Teams.”²⁴ This invitation was made by the posting of an initial version of Paper I of this series (Dobler et al. 2014) on the arxiv.org preprint server, and on the TDC Web site (<http://timedelaychallenge.org/>).

The two first ladders of this challenge are TDC0 and TDC1. TDC0 consisted of a small set of simulated data, which was used mostly as a debugging and validation tool. TDC0 is discussed in detail in Paper I. Four statistics were used to evaluate the performance of every method’s submitted time delays $\tilde{\Delta}t_i$ and uncertainties δ_i , in light of the true time delay value (defined as positive in the input), Δt_i . These four metrics are as follow: the success fraction

$$f \equiv \frac{N_{\text{submitted}}}{N}, \quad (1)$$

where N is the total number of light curves available for analysis in the ladder; the χ^2 value:

$$\chi^2 = \frac{1}{fN} \sum_i \left(\frac{\tilde{\Delta}t_i - \Delta t_i}{\delta_i} \right)^2; \quad (2)$$

the “precision”

$$P = \frac{1}{fN} \sum_i \left(\frac{\delta_i}{\Delta t_i} \right); \quad (3)$$

and the “accuracy” or “bias”

$$A = \frac{1}{fN} \sum_i \frac{\tilde{\Delta}t_i - \Delta t_i}{\Delta t_i}. \quad (4)$$

In addition to the sample metrics we also define the analogous metrics for each individual point A_i , P_i , and χ_i^2 . Thus, A , P , and χ^2 defined above are the averages of the individual point values.

Target thresholds in each of these sample metrics were set for the teams entering TDC0. The seven “Good” Teams whose methods passed these thresholds were given access to the TDC1 data set, which consisted of several thousand light curves. This large number was motivated by the goals of revealing the potential biases of each algorithm at the sub-percent level and testing the ability of current pipelines to handle large volumes of data.

To put this challenge in a cosmological context, absolute distance measurements with 1% precision and accuracy are highly desirable for the study of dark energy (Suyu et al. 2012; Weinberg et al. 2013) and other cosmological parameters. Therefore, in order for the time delay method to be competitive it has to be demonstrated that the delays can be measured with sub-percent accuracy *and* that the combination of precision for each system and the available sample size is sufficient to bring the statistical uncertainties to sub-percent level in the near future. The total uncertainty on the time delay distance, and therefore on the derived cosmology, depends on both the time delay and on the residual uncertainties from modeling the lens potential and the structure along the line of sight. Thus, controlling the precision and accuracy of the time delay measurement is a

²² strides.physics.ucsb.edu

²³ <http://www.cosmograil.org>

²⁴ We note here that the tongue-in-cheek names “evil” and “good” teams do not denote any despicable intention or moral judgment, but were chosen to capture the desire of the challenge designers to produce significantly realistic (and difficult) light curves as well as an incentive for the outside teams to participate.

Table 1
The Observing Parameters for the Five Rungs of TDC1

Rung	Mean Cadence (days)	Cadence Dispersion (days)	Season (months)	Campaign (yr)	Length (epochs)
0	3.0	1.0	8.0	5	400
1	3.0	1.0	4.0	10	400
2	3.0	0.0	4.0	5	200
3	3.0	1.0	4.0	5	200
4	6.0	1.0	4.0	10	200

necessary, but not sufficient, condition. In this first challenge we focus on just the time delay aspect of the measurement. The assessment of residual systematic uncertainties in the other components of time delay lens cosmography, and the distillation of the time delay measurement biases and uncertainties into a single cosmology metric is left for future work.

This paper focuses on TDC1, the analysis period of which closed on 2014 July 1, and it is structured as follows. Section 2 contains a brief recap of the light curve generation process, and describes the design of TDC1. In Section 3 we describe the response of the community to the challenge and give a brief summary of each method that was applied, and then in Section 4 we analyze the submissions. We look at some of the apparent implications of the TDC1 results for future survey strategies in Section 5, and briefly discuss our findings in Section 6. In Section 7 we summarize our conclusions.

2. DESCRIPTION OF TIME DELAY CHALLENGE TDC1

In TDC1, the “Evil” Team simulated several thousand realistic mock light curve pairs, using the methods outlined in Paper I. In this section, we first describe the general five rung design of TDC1, and then describe the process of generating these light curves step by step, revealing quantitative details of all the elements considered. We emphasize that TDC1 was purely a light curve analysis challenge; no additional information regarding the gravitational lensing configuration, such as positions of the multiple images, or redshifts of the source and deflector, was given. This choice was motivated by the goal of performing the simplest possible test of time delay algorithms. As discussed at the end of this paper, the inclusion of additional lensing information could provide means to further improve the performance of the methods.

2.1. The Rungs of the Challenge

Each rung of TDC1 represents a possible wide-field survey that has monitored sufficient sky area that we are in possession of light curves for 1000 gravitationally lensed active galactic nucleus (AGN) image pairs. The number of lens systems in this sample is somewhat less than 1000: quad systems are presented as two pairs, flagged as coming from the same system but enabling two independent time delay measurements. The five rungs of TDC1 span a selection of possible observing strategies, ranging from a high cadence, long season dedicated survey (such as COSMOGRAIL might evolve into), to the kind of “universal cadence” strategy that might be adopted for an “all-sky” synoptic imaging survey (such as is being designed for LSST). The challenge allows four control variables to be investigated (within small plausible ranges): cadence, sampling regularity, observing season length, and campaign duration. Table 1 gives the values of these control variables for each rung.

To make the mock data generation more efficient, and to better enable comparison of results between the different rungs, we re-used the same catalog of lenses for all the rungs. This trick was disguised from the “Good” Teams by randomly re-allocating the light curve identification labels in each rung. In addition, the random noise was independently generated in each rung. As a consequence, the submissions for different rungs may be deemed independent, as if they had addressed 5000 lensed image pairs.

2.2. Lens Sample

The time delays between the light curves of gravitationally lensed images are determined primarily by the macro structure of the lens galaxy. For the TDC1 sources and lenses we use the mock LSST catalog of lensed quasar systems prepared by Oguri & Marshall (2010, hereafter OM10).²⁵ This sample was drawn from plausible physical distributions for the various key properties of lensed quasar systems and very approximate observing conditions expected with LSST, namely, a characteristic angular resolution of 0.75 arcsec and a 10σ limiting magnitude per monitoring epoch of 23.3 in the i band. Assuming a survey area of 18,000 deg², these numbers correspond to an OM10-predicted mock sample of some 2813 lenses. Given these constraints, we randomly drew 720 doubly imaged and 152 quadruply imaged quasars from this catalog, to give a total of 1024 independent time delayed image pairs. As Figure 1 shows, the mean time delay in TDC1 is several tens of days. We rejected all time delays outside the range 5 to 120 days as we drew the mock sample, since the typical observing cadence and season length are expected to be a few days and a few months, respectively. The same time delay range constraint reduced the parent OM10 mock lens sample by 76%, to 2124 lenses. When analyzing the submissions, we found that very few accurate measurements of time delays less than 10 days were possible, and so in the rest of this paper we focus on the range $10 < \Delta t < 120$ days. Imposing this narrower range on the OM10 mock LSST lens sample results in 1990 systems. While the image pairs with $5 \text{ days} < \Delta t < 10 \text{ days}$ were not used in the analysis, they are still there in the TDC1 data set for potential future use.

To give an overview of this sample, we show the distributions of time delays Δt between images in our 1024 image pairs (in Figure 1), and detection magnitudes i_3 in the 872 lens systems (in Figure 2). The i_3 quantity is the i -band magnitude of the third brightest image in a quad system or the magnitude of the fainter image in a double-image system. (It is an important parameter because it helped OM10 characterize the detectability of lensed quasars: lenses are assumed to be measurable if i_3 is above the 10σ limiting magnitude of a survey.) The lens abundance rises fairly steeply with i_3 , so in order to probe the relationship

²⁵ The OM10 catalog is available from <https://github.com/dphilmarshall/OM10>.

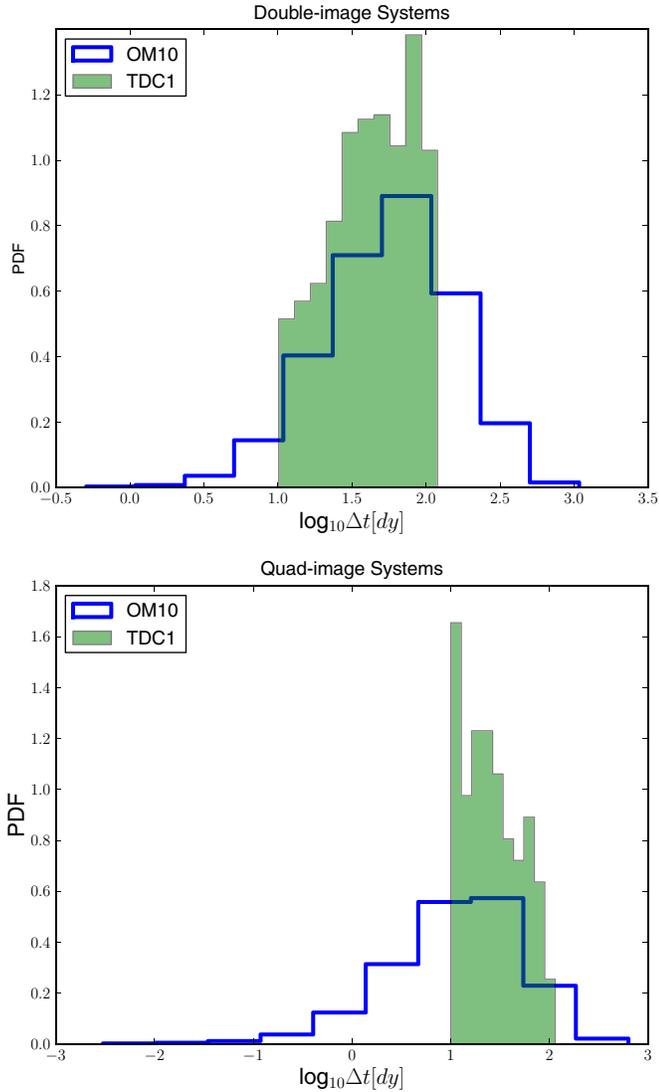


Figure 1. Time delay distributions, from both the parent OM10 catalog and the sample used in the TDC1 analysis, for the double-image (top) and quad-image (bottom) systems.

between it and the time delay measurement accuracy, we split the magnitude range 20–24 into four sub-ranges, and selected approximately equal numbers of systems in each sub-range.

In summary, our sample is similar to OM10’s, except that the brighter lenses and intermediate time delays are somewhat over-represented. As we will discuss later in this paper, this allows us to sample the range of magnitudes more evenly, while introducing negligible bias in the inferred performance of the methods.

2.3. Generation of Intrinsic Light Curves

The mechanism for generating intrinsic light curves is described in Paper I. In TDC1, we needed to simulate many more data sets; the most time-consuming part was generating the damped random walk (DRW) stochastic process with which we modeled the intrinsic AGN light curves. The interval between discrete epochs had to be 0.01 days in order to enable the counter-image light curve to be simulated with a time delay precision sufficient to not affect the ensemble metrics. Each of these intrinsic light curves took approximately 1–2 CPU hours to make, so for efficiency we created just 500 intrinsic light curves,

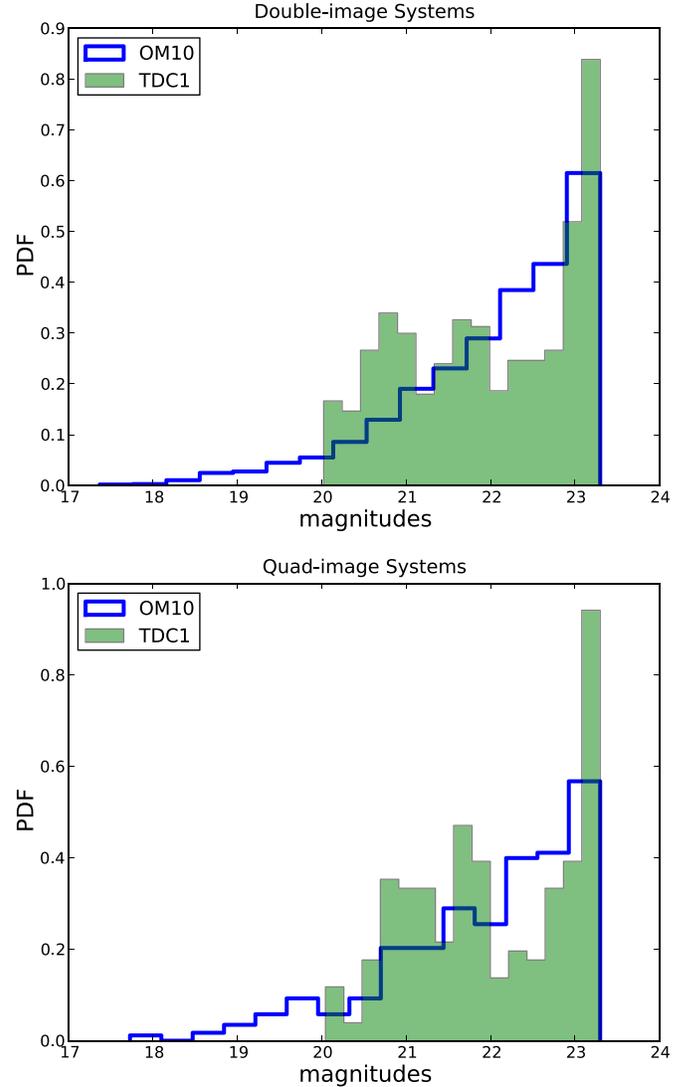


Figure 2. Detection magnitude “ i_3 ” distributions for the double (top) and quad (bottom) systems. For doubles, i_3 is the magnitude of the fainter image, while for quad systems it is the magnitude of the third-brightest image. Distributions are shown both for the parent OM10 sample, and the sample used for TDC1.

each of 10 yr length, and re-cycled them between several mock data sets, with different starting epochs chosen relative to the season gaps, so that all the release data could be considered to be independent.

The DRW light curves represent light curve fluctuations, and have zero mean magnitude. They are determined by only two parameters: the characteristic timescale τ and the characteristic amplitude of the fluctuations σ . These were drawn from distributions designed to match that observed for the spectroscopically confirmed ($i < 19.1$ mag) quasars in MacLeod et al. (2010). Their $\log \tau$ and $\log \text{SF}_\infty$ (asymptotic rms variability on long timescales) parameters were drawn uniformly from the ranges $[1.5 : 3.0]$ and $[-1.1 : -0.3]$, respectively. The endpoints of these ranges correspond to 30 and 1000 days, and 0.08 and 0.5 mag. The rms fluctuation level was derived for each light curve via $\sigma = \text{SF}_\infty / \sqrt{\tau}$.

2.4. Modeling Microlensing

Microlensing is an important source of systematic error because it makes the multiply imaged light curves differ by more

than the time delay and the macrolens magnification ratio. In galaxy-scale lenses, the variability of the microlensing typically has a timescale significantly larger than that of the quasar intrinsic variability (although occasional caustic crossing events can provide some transient rapid variability). We expect the most successful light curve measurement algorithms to model an additional microlensing light curve component individually at each image.

Given an **OM10** catalog convergence κ , shear γ and surface density in stars F_* at each image position, we generated a static stellar field with a mean mass per star of $0.3 M_\odot$ (Schechter et al. 2004). We then calculated its source plane magnification map and convolved this with a Gaussian kernel to represent the extended accretion disk of the source quasar; we drew source sizes s (Gaussian radii) uniformly from the range $[10^{14}\text{--}10^{16}]$ cm. When calculating the microlensing light curves, we assumed Gaussian distributions for the components of the relative velocity v between the source and the stars in the lens, with standard deviation of 500 km s^{-1} in each direction.²⁶ In the [Appendix](#) we show how the scatter in microlensing variability amplitude depends on F_* , κ , and source size. Finally, we note that there are several characteristic timescales in microlensing light curves, ranging from the crossing time of the mean stellar mass Einstein Radius (Paraficz et al. 2006) to the source caustic crossing time, to the density of caustics in the network, and those can give rise occasionally to quasi-periodic features.

2.5. Photometric and Systematic Errors

Following Tewes et al. (2013a) we considered several sources of observational error when generating the light curve fluxes. The main source of statistical uncertainty is the sky brightness, which we assume dominates the photometry. We used the approximate distribution of 5σ limiting point source magnitudes from one of the LSST project operations simulator outputs (L. Jones 2014, private communication), and converted these to flux uncertainties. The mean and standard deviation of the 5σ i -band limiting flux was found to be 0.263 and 0.081 AB nanomaggies²⁷, respectively; to add photometric noise to a light curve flux we first drew an rms photometric uncertainty from a Gaussian of mean 0.053 and width 0.016 nanomaggies (dividing the above numbers by 5), and then drew a noise value from a Gaussian of width equal to this rms. The minimum noise value was set to be 0.001 nanomaggies.

Beyond this basic (though possibly epoch-dependent) Gaussian noise, we might expect additional flux errors to be present as the observing set-up changes over a long monitoring campaign. To mimic such fluctuations, we added the following three types of “evilness” to the light curves.

1. Flux uncertainty under-estimation: for each pair of light curves and for approximately 1 in every 10 epochs, we added noise that was three times larger than standard, but reported it as the normal one.
2. Calibration error: for each pair of light curves and for approximately 1 every 10 epochs, we added correlated noise, i.e., both points were higher or lower than in the normal case.

3. Episodic transparency loss: we took a subset of the data (a few weeks every year), and offset the fluxes by 1% or 3%.

There could be more than one type of “evilness” present in any given light curve: the combinations applied to the TDC1 light curves were as follows. 3% of the light curves, selected randomly, were contaminated with a single type of “evilness.” Another 1% were contaminated with two types, and 3% were contaminated with all three. In total then, 15% of the light curves were contaminated with these simulated bad observational conditions.

2.6. Example TDC1 Light Curves

Figures 3 and 4 illustrate the process of generating TDC1 data in each of the five rungs, using light curves selected randomly from those data sets. The top panels show the AGN intrinsic light curves in magnitudes. The panels beneath them show the microlensing magnifications (also in magnitudes). The third panels show the AGN light curves with microlensing effects, and the effect of sampling is shown in the fourth panels. Finally, the sparsely sampled noisy mock light curves are shown on the bottom panels, in flux units.

Comparing panels 3 and 5, we can easily see how two similar curves become difficult to associate by eye once the sparse sampling and the addition of noise have been applied. Table 2 shows the values of the input parameters τ , σ , v , s , F_* , enabling some intuition to be developed by comparing plots shown for the different rungs.

3. RESPONSE TO THE CHALLENGE

As described in Section 1, the TDC was presented to the community as two “ladders,” TDC0 and TDC1. The TDC0 data were used as a gateway to TDC1; in order to gain access to the TDC1 data, each “Good” Team had to submit a set of time delays inferred from TDC0 that met the targets described in Section 1, and in more detail in Paper I. In total, 13 “Good” Teams participated in TDC0, many of which submitted multiple sets of solutions. Seven teams passed TDC0 and, went on to participate in TDC1. One of the teams submitted results based on three different algorithms: those were considered independent submissions. In addition, the “Evil” Team did an in-house analysis of the TDC1 data, using a relatively simple procedure, to serve as a baseline comparison for the “Good” Team submissions. All 10 of these algorithms are described below and some of their properties are summarized in Table 3. It is worth noting that the teams continued to develop their methods between TDC0 and TDC1 and beyond, and the description given here is for the versions of the methods that were applied to TDC1.

3.1. Benchmark Technique by Rumbaugh (“Evil” Team)

The baseline method used by the “Evil” Team was a χ^2 -based Markov Chain Monte Carlo (MCMC) approach. While the member of the team that wrote and executed this baseline method (N.R.) did not work directly on simulating the light curves, this method should not be considered blind in the same way as the “Good” Teams’.

In practice the method consists of comparing a shifted copy of one of the light curves to the other light curve, and using a χ^2 function to compute the posterior probability distribution function for the time delay. Matching the light curves requires some interpolation, which was carried out using a boxcar kernel

²⁶ The microlensing code used in this work, MULES, is freely available at <https://github.com/gdobler/mules>.

²⁷ One “AB maggy” is the flux corresponding to an AB magnitude of 0.0 (Stoughton et al. 2002). Thus, 0.263 nanomaggies is the flux corresponding to an AB magnitude of 24.

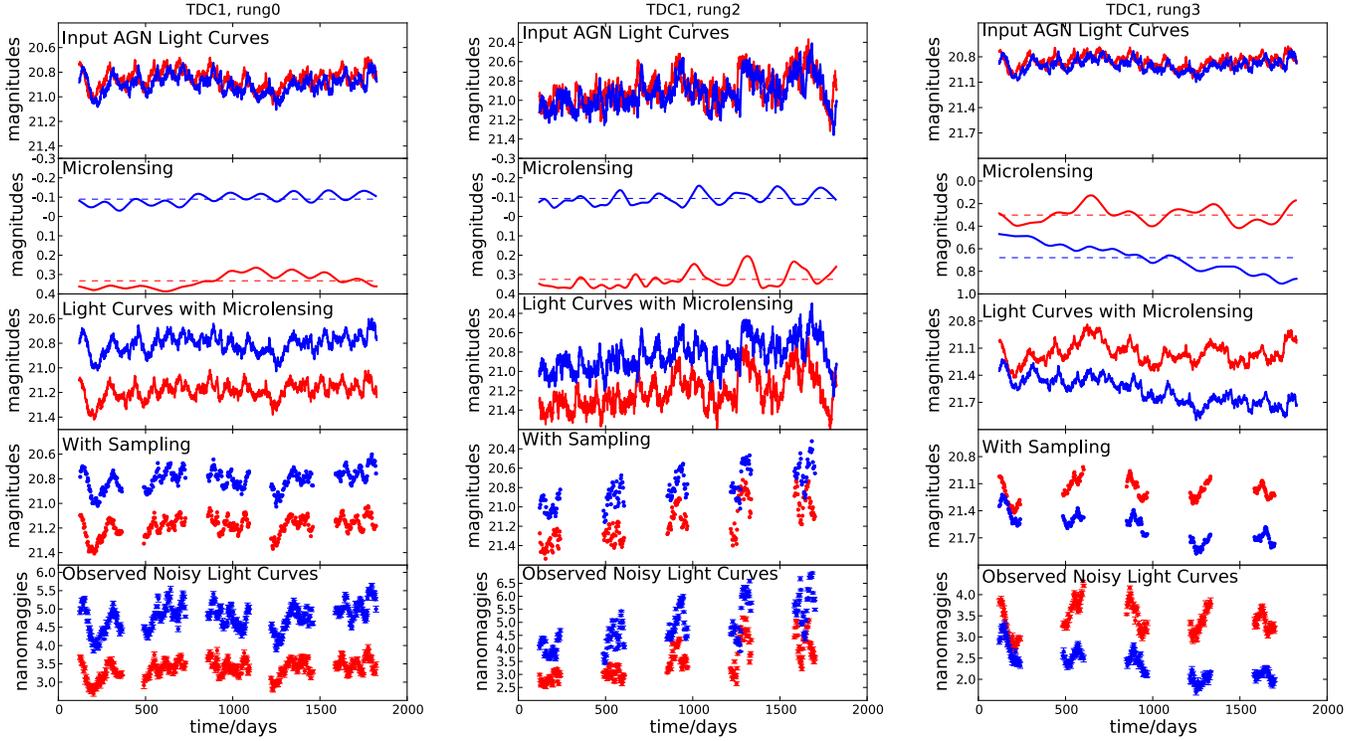


Figure 3. Illustration of the process of generating time delay light curves, with examples taken from the Rung 0 (left), Rung 2 (middle), and Rung 3 (right) samples. The panels in each figure show, going from the top to the bottom, (1) the input AGN light curves, (2) the microlensing contributions in magnitudes, (3) the AGN light curves including the microlensing contributions, (4) the result of down-sampling to the required cadence and season length, and (5) the final sparsely sampled noisy light curves.

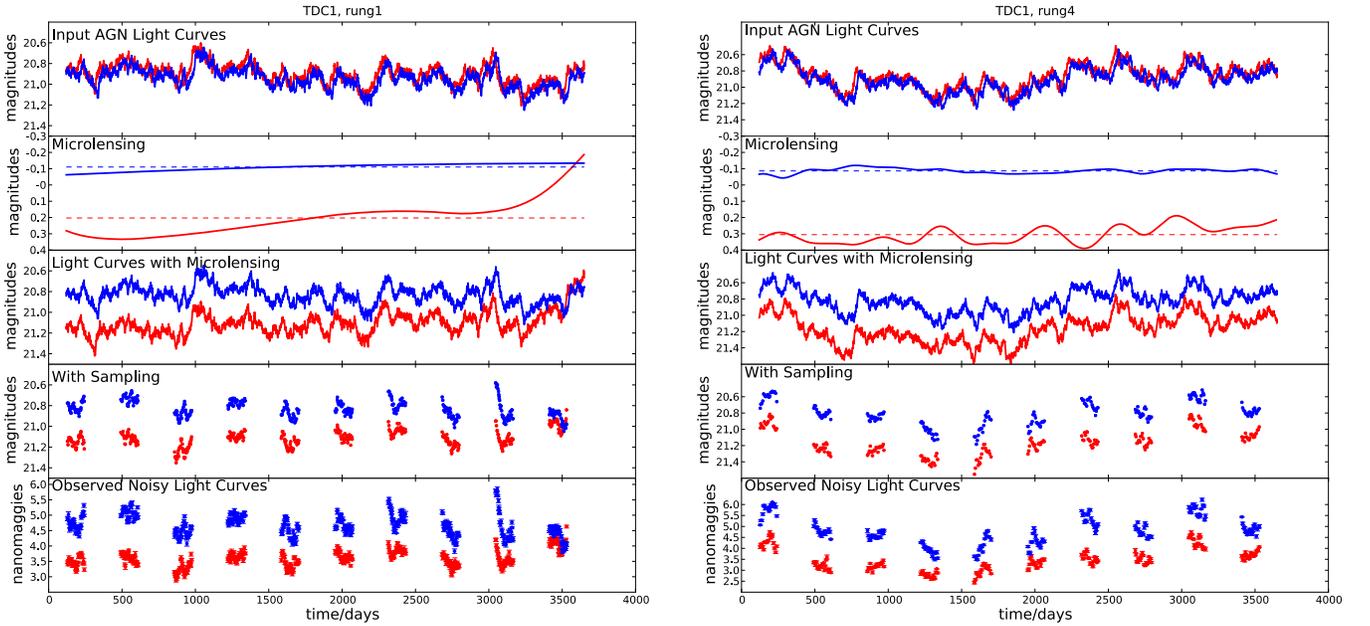


Figure 4. Same as Figure 3, but for the longer campaign-duration light curves of Rungs 1 and 4.

Table 2
The Parameters Used to Make the Simulated Data Shown in Figures 3 and 4, to Enable Study of Their Effects on the Light Curves

Rung	τ (days)	σ (mag/day ^{-1/2})	v (km s ⁻¹)	s (10 ¹⁴ cm)	F_{*A}	F_{*B}
0	37.8	0.017	731	3.87	0.037	0.062
1	83.0	0.017	731	38.7	0.037	0.062
2	40.6	0.039	1462	3.87	0.037	0.062
3	37.8	0.017	731	3.87	0.019	0.031
4	178.0	0.017	365	3.87	0.037	0.062

Table 3
Summary of Methods Explicitly Accounting for Microlensing

Method	Microlensing
Rumbaugh	No
Shafieloo	Yes
PyCS-d3cs	Yes
PyCS-sdi	Yes
PyCS-spl	Yes
Jackson-manchester	Yes
Kumar	Yes
JPL	No
Hojjati	Yes
DeltaTBayes	No

with a full width of 10 days. This particular kernel was chosen to save computational time; however, the choice of the kernel did not have a significant effect on the accuracy or precision of the method. In order to gain additional computational speed, the correlation between temporally close data points introduced by the smoothing kernel was neglected. This approximation reduced the computation time by about an order of magnitude, while providing only marginally worse accuracy. The posterior was sampled using the *emcee* (Foreman-Mackey et al. 2013) software package. For each trial value of the time delay, only the overlapping parts of the time-shifted light curves were used in the computation of the change in χ^2 . To avoid calculations using small overlap regions, a maximum time delay was imposed equal to 75% of the shortest season length of the data set currently being analyzed. Time delay point estimates were chosen to be the median of the output sample values, with the uncertainties chosen to be half the width of the region containing 68.3% of the chain surrounding the median.

Before applying the benchmark technique to TDC1 data, it was tested on the TDC0 data, as well as on an additional set of simulated data designed to be similar to TDC0. In this testing, the smoothing kernel was varied, as well as several other aspects of the method as indicated above (including whether or not the full covariance matrix was used). The accuracy and precision of the inference were found to not depend significantly on these choices.

Time delay estimates from three implementations of this method were submitted, with the aim of producing answers of different degrees of reliability. The three implementations were obtained by restricting the submissions to those systems with estimated time delay uncertainty below 6, 10, and 20 days. The submissions resulting from these cuts are named Gold, Silver, and Bronze, respectively.

3.2. Gaussian Processes by Hojjati & Linder

This ‘‘Good’’ Team implemented Gaussian process (GP) regression to estimate the time delays (see Hojjati et al. 2013, for the basic approach). GPs are widely used as a model-independent technique for reconstructing an underlying function from noisy measurements. The GP is specified by a mean function, and a covariance (kernel) function characterized by a set of hyperparameters, describing the time delay, relative magnitude shift, QSO variability and coherence length, microlensing variability and coherence length, and measurement noise. This approach is very flexible, not assuming a physical model for the quasar or microlensing input, but allowing the data to decide how best to describe the signal in terms of a GP. The hyperparameters were fitted to data using the GP likelihood through a

Bayesian analysis. The parallel and highly efficient fitting code employed two covariance kernels, two optimization methods, and variation of priors to cross-check the results for robustness. The team passed or rejected a system, based on the consistency of fits and their likelihood weights, and then assigned a final best fit, uncertainty, and confidence class to the passed systems.

The overall philosophy emphasized complete automation and accuracy of estimation, rather than precision (e.g., fitting down to five day delays and placing no cut on precision) or numbers of fits. Within this, the team fine-tuned samples based on their confidence in the fit, and to a lesser extent the error estimation. Six samples were submitted, with the basic three representing progressively more inclusive fit confidence along the lines of, e.g., gold, silver, bronze estimation. These correspond to the samples nicknamed Lannister, Targaryen, and Baratheon, respectively. In addition, a more conservative sample (nicknamed Tully) and one with tighter error assignment (nicknamed Stark) were submitted. Catastrophic outliers were identified by running selected samples (e.g., especially short or long time delays) with controlled priors, and also an analysis of the best-fit parameters for the selected systems. The sample nicknamed ‘‘Freefolk’’ was the result of such analysis.

A correction to the mean function treatment in the code significantly increased the consistency of the fits. However, since this modification was made after the TDC1 submission deadline, this is not reflected in the results presented in this paper; see the updates and discussion by Hojjati & Linder (2014). Furthermore, the method has benefited from, and was improved after, a reanalysis of the fits and the investigation of the hyperparameter behavior using the unblinded TDC1 data.

3.3. FOT by Romero-Wolf & Moustakas

The Full of Time (FOT) team’s GP inference algorithm took a Bayesian approach to solve for the delay between a pair of light curves. The probability of the light curve parameters \bar{M} (mean magnitude), σ (characteristic amplitude of the fluctuations), and τ (characteristic timescale) given the data is proportional to the product of the likelihood function for a CAR process (Kelly et al. 2009; MacLeod et al. 2010) and uniform priors. Details about the CAR process can be also found in Paper I. The *emcee* (Foreman-Mackey et al. 2013) MCMC ensemble sampler provides an estimate of the posterior probability distribution for the light curve parameters. To reconstruct the delay, the pair of light curves were combined into a single time series assuming a delay and magnitude offset. The probability of the delay and magnitude offset, along with light curve parameters, is given by the CAR process likelihood function of the combined light curve and uniform priors. The light curve delay and its uncertainty were then inferred from the marginalized posterior distribution for the time delay given the light curves. The algorithm did not characterize or fit for microlensing, although it identifies the data sets that are most likely to have microlensing variations. A more thorough description of this method and internal tests are being written up by L. A. Moustakas & A. Romero-Wolf (2014, in preparation).

The procedure was tested by generating tens of thousands of ‘‘blind’’ time-delayed light curves through the CAR process, with varying (irregular) observational patterns and campaigns, photometric uncertainties, magnitude offsets, and time delays. These were then processed with the inference technique described above. Both the successful recovery rate and the precision of the (marginalized) time delay and magnitude offset were then studied as a function of each ‘‘observational’’

parameter (i.e., the observational campaign factors and the assumed photometric precision).

To avoid outliers, a set of consistency requirements between the posterior distributions for the individual and combined light curve parameters were required. A solution was rejected if the mean of the posterior σ distributions from each light curve and their combinations differed by more than 2.6 root-sum-squared standard deviations. The means of the posterior $\log_{10} \tau$ distributions for each light curve must also agree to within one standard deviation, forcing a consistency in the physical behavior of the reconstructed “stitched” data set compared to the input data. Additional quality cuts were included from inspection of the reconstructed time delay and time delay uncertainty scatter relation. These required that delays less than 100 days have uncertainties smaller than 10 days. The ratio of the delay uncertainty to the delay was also required to be smaller than 2.

3.4. Smoothing and Cross-correlation by Aghamousa & Shafieloo

This “Good” Team combined various statistical methods of data analysis in order to estimate the time delay between different light curves. At different stages of their analysis they used iterative smoothing, cross-correlation, simulations and error estimation, bias control, and significance testing to prepare their results. Given the limited timeframe (they started the project in early 2014 May), they had to make some approximations in their error analysis.

In their approach to estimate the time delay between a pair of light curves A_1 and A_2 , they first smoothed over both light curves using an iterative smoothing method (Shafieloo et al. 2006; Shafieloo 2007, 2012; Shafieloo & Clarkson 2010), producing the smoothed light curves A_1^{smooth} and A_2^{smooth} . During smoothing, they recorded the ranges with no data available (which would have resulted in unreliable smoothing). The algorithm was set to automatically detect such ranges. Then, they calculated the cross-correlation between A_1 and A_2^{smooth} and also between A_2 and A_1^{smooth} for different time delays, and found the maximum correlations. These two maximum correlations should be for the same time delays (that is, the absolute values of the time delays should be consistent with each other). The difference between these two estimated time delays (with maximum correlations) was part of the total uncertainty considered for each pair (in the estimated time delay). To estimate the error on each derived time delay, the team also simulated many realizations of the data for each rung, and for various time delays. Knowing the fiducial values, they derived the expected uncertainties in the estimated values of the time delays.

This team also performed bias control, since long time delays have a limited data overlap between the two light curves. In the case of the quad sample, they used different combinations of the smoothed and raw light curves to test the internal consistency of the results and relative errors. These internal consistency relations can be used to adjust the estimated error-bars for each pair (considering the consistency of all light curves as a prior). The team selected for cross-correlations between the two light curves with more than 50% or 60% correlation coefficients. Pairs with potentially high bias were cut as well. In this methodology the light curves are compared in multi-segments. The effect of micro-lensing can be considered as a linear distortion in these segments. While the correlation coefficient is unchanged under linear transformation, there is no concern for micro-lensing in

this algorithm and the method is unaffected. Additional details of this method will be described in a separate paper Aghamousa & Shafieloo (2014).

3.5. Supervised Pelt by Jackson

All pairs of joint light curves were inspected by eye by this team, using a Python tool developed for the purpose. An initial Pelt et al. (1994) statistic was calculated for a large range of time delays, and its minimum found, but this resulted in catastrophic errors in many cases and was frequently over-ridden by visual inspection. Time delays were regarded as believable if (1) at least three coincident points of inflection were detected in the light curves, (2) if no discordant features were seen (i.e., differences between the light curves which could not be plausibly attributed to microlensing) and (3) if the plot of the Pelt statistic against time delay showed a smooth and well-defined minimum.

In the process of assessing the light curves by eye, the following operations were available to find a time delay fitting the above criteria: (1) smoothing of either light curve to match the scatter of the other; (2) adjustment of the zero-point of each segment of the light curve to match the zero-point of the segment of the other light curve that it overlapped using the current time delay; (3) manual adjustment of the current time delay; (4) deletion of one or more segments of the light curve if they were judged to be severely affected by microlensing. In practice, this was the case if a simple rescaling of a whole segment of data between the two light curves produced residuals much larger than those of other rescaled segments. This will happen if the microlensing produces a large change in flux over the period of one data segment; the method therefore roughly corresponds to assuming that microlensing produces variations on a timescale larger than those of the intrinsic brightness variations of the quasar, and deleting regions of data where this is not the case. In most cases, the delay and its error bar were calculated after this process using 100 instances of resampling of the data set using the observed flux errors and a small Gaussian error in each time stamp. This allowed the calculation of a set of delays, in each case using the delay from the Pelt statistic minimum, from which the mean and scatter was used for the delay and its error. In a few cases, mostly those in which the Pelt statistic versus time delay plot had a local minimum around the optimum, an additional error, or in some cases a minor adjustment to the value, was estimated by eye. The error bar was also adjusted in cases where the optimization using smoothing and adjustment of the zero point resulted in a significant reduction of the error estimated by the resampling process. With practice, about 100 pairs of light curves per hour could be processed, so that thousands or tens of thousands of light curve pairs could in principle be analyzed using this method.

The same basic algorithm was used for all submissions, but different submissions were made by separating the objects into three categories, again by eye, according to confidence that the time delay was correct within the stated error. Evaluations with less confidence corresponded to violation of one or more of the believability conditions, and the least certain category usually involved light-curves with only two clearly detected points of inflection. (For each of the three categories, subsidiary submissions were also made with a smaller number of rungs). Three catastrophic errors in rung 0 of the original blind submission were due to incorrect entry of a minus sign during the manual adjustment process in three objects; these were corrected in a non-blind submission which consisted of the original blind submission for all rungs, and all three confidence levels with the

three signs corrected. The program was accordingly modified to question the user in the case of large changes imposed by hand.

3.6. PYCS by Bonvin, Tewes, Courbin, & Meylan

The PyCS team made submissions using three time delay measurement methods: `d3cs`, `sp1`, and `sdi`. The latter two build on initial estimations provided by the former. The following subsections summarize each of these three methods (see V. Bonvin et al., in preparation, for more details).

3.6.1. `d3cs`: D3 Curve Shifting

This first method is based on human inspection of the light curves, in the spirit of citizen science projects. The PyCS team has developed a dedicated browser-based visualization interface, using the D3.js JavaScript library²⁸ by Bostock et al. (2011). The tool is now publicly available online.²⁹

The main motivation behind this time-consuming yet simple approach were to obtain, for each light curve pair, (1) a rough initial estimate for the time delay and its associated uncertainty, and (2) a robust characterization of the confidence that this estimate is not a catastrophic error. The interface asks each user to pick a confidence category for the proposed solution, among four choices:

1. “doubtless” if a catastrophic error can be virtually excluded,
2. “plausible” if the solution yields a good fit and no other solutions are seen,
3. “multimodal” if the proposed solution is only one among two or more possible solutions,
4. “uninformative” if the data does not reveal any delay.

At least two human estimates were obtained for each pair of curves. The database of `d3cs` estimates was then carefully reduced to a single estimate per pair, resolving any conflicts between estimates in a conservative way. A key result of this step was a sample of 1628 “doubtless” time-delay estimates, which the team hoped to be free from any catastrophic outliers. Through this exercise, the team demonstrated that such an approach remains tractable for about 5000 light curves, with typical human inspection times of a minute per light curve pair and user.

3.6.2. `sp1`: Free-knot Spline Fit

The `sp1` method is a simplified version of the “free-knot spline technique” described by Tewes et al. (2013a) and implemented in the PyCS software package. Using the `d3cs` estimate as the starting point, the method simultaneously fits a single spline representing the intrinsic QSO variability, and a smoother “extrinsic” spline representing the differential microlensing variability, to the light curves. During this iterative process, the curves were shifted in time so as to optimize the fit. The fit was repeated 20 times, starting from different initial conditions, to test and improve the robustness of the resulting delay against local minima of the χ^2 hyper surface. Such a model fit was then used to generate 40 simulated noisy light curves with a range of true time delays around the best-fit solution. By re-running the spline fit on these simulated curves, and comparing the resulting delays with the true input time delays, the delay measurement uncertainty was estimated.

The `sp1` method for TDC1 is simpler, faster, and significantly less conservative in the uncertainty estimation than the free-knot spline technique that was applied to the COSMOGRAIL data³⁰ by Tewes et al. (2013b) and Rathna Kumar et al. (2013). In particular, the temporal density of spline knots was automatically determined from signal-to-noise ratios measured on the two light curves, and only white noise was used in the generative model. With these simplifications, the team expects the resulting TDC1 error estimates to be rather optimistic. The entire `sp1` analysis took about 5 CPU-minutes for an average TDC1 pair.

3.6.3. `sdi`

The third method, `sdi` (for spline difference) was inspired by the “regression difference technique” of Tewes et al. (2013a), replacing the GP regressions by spline fits to speed up the analysis. The method involves fitting a different spline to each of the two light curves, and then minimizing the variability of the difference between these two splines by shifting them in time with respect to each other. The advantage of this approach is that it does not require an explicit microlensing model. To estimate the uncertainty, this method uses the simulated light curves provided by the `sp1` technique. As in the `sp1` technique, the estimates from `d3cs` were used as the starting point to define the time delay intervals in which `sdi` optimizes its cost function.

3.6.4. Identification of Catastrophic Failures

To prevent catastrophic failures, this team relied solely on the `d3cs` “doubtless” sample. The `sp1` and `sdi` methods do not alter this confidence classification. Furthermore, a small number of `sp1` and `sdi` measurements that did not lie within 1.5σ of the corresponding `d3cs` estimates were rejected.

3.6.5. Differences between Submissions

For all three methods, the submissions were named following the scheme A–B–C–D. dt, where

- A gives the method, `d3cs`, `sp1`, or `sdi`;
- B gives the method parameters, with `vanilla` denoting the a priori best or simplest;
- C gives the confidence category, with `dou` for doubtless and `doupl` for both doubtless and plausible light curve pairs. The `doupl` submissions are expected to be contaminated by some catastrophic outliers, but feature more than twice the number of time delays than the `dou` sample; and
- D gives the filter that selects systems according to different criteria across all rungs, mostly based on the blind relative precision $\delta_i/|\Delta t_i|$. The code `full` corresponds to no filter. `XXXbestP` selects the XXX “best” systems in terms of blind relative precision, `P3percent` selects the largest number of systems so that the average blind relative precision is approximately 3%, and `100largestabstd` is the selection of the 100 largest delays.

Submissions that share the same method and method parameters (A and B) differ only in the selection of systems, and not in the numerical values of the estimates. They can thus be seen as subsamples of the A–B–`dou`/`doupl`–`full` submissions.

²⁸ Data-Driven Documents, <http://www.d3js.org/>.

²⁹ <https://www.astro.uni-bonn.de/~mtewes/d3cs/tdc1/>—see “Read me first” for help.

³⁰ <http://www.cosmograil.org>

3.7. Difference-smoothing by Rathna Kumar, Stalin, & Prabhu

The difference-smoothing technique, originally introduced by Rathna Kumar et al. (2013), is based on the principle of minimizing the residuals of a high-pass filtered difference light curve between the lensed quasar images. The method is a point estimator that determines an optimal time delay between two given light curves, and an optimal shift in flux to one of the light curves, besides allowing for smooth extrinsic variability. To estimate the uncertainty of the measured time delay in Rathna Kumar et al. (2013), this team made use of simulations produced and adjusted according to Tewes et al. (2013a). However, for participation in the TDC, they made use of a modified version of the difference-smoothing technique as presented by Rathna Kumar et al. (2014). In that paper, they describe an optimal way to adjust the two free parameters in the technique according to the peculiarities of the light curves under analysis and also introduce a recipe for simulating light curves having true delays at discrete intervals in a plausible range around the optimal time delay found. These simulations were used to estimate the uncertainty of the measured value of the time delay. Outliers were identified by noting when the team’s technique was found to return random time delays which were uncorrelated with the true delays in their simulated light curves.

The free parameters in the technique are decorrelation length and smoothing timescale. For participation in the TDC, the value of decorrelation length was set equal to the mean temporal sampling of the light curves and the value of smoothing timescale was set equal to the largest integer multiple of decorrelation length for which the amplitude of residual extrinsic variability was less than the 3σ level of noise for each of the light curves. In the absence of significant extrinsic variability between the light curves, the value of smoothing timescale was set equal to ∞ .

3.8. Δt -Bayes by Tak, Meng, van Dyk, Siemiginowska, Kashyap, & Mandel

A fully Bayesian approach was developed by this team, based on the key assumption that one of the unobserved underlying light curves is a shifted version of the other. The horizontal shift is the time delay (Δt), and the vertical shift is the magnitude offset (c). Both shifts are treated as unknown parameters. Specifically, from the state-space modeling perspective, it was observed that $\mathbf{x}(\mathbf{t}) \equiv \{x(t_1), x(t_2), \dots, x(t_n)\}$ and $\mathbf{y}(\mathbf{t})$, transformed into the logarithm of flux, around the irregularly sampled underlying light curves, $\mathbf{X}(\mathbf{t})$ and $\mathbf{Y}(\mathbf{t}) \equiv \mathbf{X}(\mathbf{t} - \Delta t) + c$ each, with measurement errors in log scale. The posterior distribution for Δt is of primary interest. Also, it was assumed that the unobserved true process $\mathbf{X}(\mathbf{t})$ follows an Ornstein-Uhlenbeck process (also known as CAR) as described by Kelly et al. (2009), although a different parameterization was used for more efficient model fitting. Harva & Raychaudhury (2006) proposed a similar idea, but they assumed a different model for the underlying process.

This Bayesian approach treats the unknown parameters as random variables and this team uses specific prior distributions for the time delay and magnitude offset: $p(\Delta t, c) \propto \delta_{\{|\Delta t| \in [0, (t_n - t_1)]\}}$. A uniform prior on c is a typical choice because this y -shift is related to the mean of observed data or the underlying process. The uniform prior on Δt constrains its values to ensure that the shifted light curves overlap in time. This naively informative hyperprior distribution on the parameters governing the underlying process is $p(\bar{M}, \sigma, \tau) \propto \tau^{-2} e^{-1/\tau}$, where \bar{M} , σ ,

τ are CAR parameters as defined above and in Paper I. This puts a uniform prior on \bar{M} and σ , and an inverse- $\Gamma(1, 1)$ prior on τ .

The full posterior distribution was obtained by multiplying together (1) the likelihood for the state-space representation, (2) the prior for the underlying process, Δt , and c , and (3) the hyperpriors for \bar{M} , σ , and τ . The team proposed a Gibbs sampler for this full posterior distribution (algorithm 2) and its approximation (algorithm 1) in TDC1. Details of the two samplers were submitted to the “Evil” Team and will appear in a separate paper, in preparation. In order to obtain the time delay from its posterior distribution, three Markov chains were combined with starting values chosen randomly around the most likely values. Rigorous convergence checks of the Markov chains were conducted using trace plots, autocorrelation plots, and the Gelman–Rubin diagnostic statistic, applied to all of the model parameters.

The model did not account for the microlensing. However, when it was suspected it after a visual inspection, this team accounted for its polynomial long-term effect (linear or quadratic) by the regression and ran the model on the residuals. This worked well because the intrinsic variability of quasar data did not disappear even after the long-term trend was removed.

4. ANALYSIS OF THE SUBMISSIONS

4.1. Lessons from TDC0 Applied to TDC1

During the analysis of the TDC0 submissions, the “Evil” Team noticed that several teams were affected by outliers: most of their submitted time delay estimates were good, but a few differed from the truth by more than would be expected, given their uncertainties. To characterize this, an additional metric X was introduced: X is the fraction of pairs with $\chi_i^2 < 10$, i.e., the fraction without outliers. $X = 1$ means that none of the submitted delays is an outlier. Outliers in this category could stem from underestimated error bars, or for example by convergence on the wrong solution in the presence of light curve features (due to, e.g., microlensing) that are not taken into account by the method’s model.

We will return to the issue of outliers, and how they can be identified based on lensing geometry or cosmological analysis, after we present the main results of TDC1. In this section, we give the unfiltered statistics as well as the metrics calculated after points with $\chi_i^2 > 10$ have been removed, in order to give an idea of how well a method *could* do if outliers could be identified and rejected.

We also consider an additional cut, based only on the accuracy parameter $|A_i| < 0.1$, and the related quantity X_A , which counts the fraction of systems satisfying this alternative criterion, i.e., we take $|A_i| > 0.1$ as outliers rather than $\chi_i^2 > 10$ in this case. This cut was chosen to quantify the number of systems for which the time-delay would be much more uncertain than the 3%–5% modeling error that can be obtained in the reconstruction of the difference in gravitational potential between two images in the best cases (Suyu et al. 2013, 2014). In some sense this cut filters out the systems that are not cosmologically consistent and thus could be rejected by a joint cosmological analysis.

Finally, as a third way to illustrate the potential of each method once outliers have been removed, we also consider the median, 16 and 84 percentile of the statistics A_i , P_i , and χ_i^2 for each method, as opposed to the means defined in Section 1.

4.2. Blind and Non-blind Submissions

One of the main goals of this TDC is to achieve a true blind testing of the algorithms. To achieve this, TDC0 truth files were not revealed until after the deadline of TDC1, lest they give too much away about the data generation process. In addition, upon requests from each “Good” Team we provided only minimal feedback after each submission, in the form of the metrics listed above rounded to two significant digits. This was deemed to be a reasonable compromise between preserving the blindness of the challenge, and helping teams to identify coding errors that had nothing to do with their actual chosen algorithms. Only submissions made prior to any feedback were considered truly blind, even though resubmissions by the teams who decided to take advantage of this opportunity were accepted. Resubmissions were considered not fully blind for the purpose of this analysis. Note that all of the “representative” submissions referred in later sections were made fully blind.

4.3. Basic Statistics

The metrics for each submission are shown in Tables 4 and 5, separated by challenge rung. In order to visually compare the different algorithms in a relatively clear manner, we have chosen to show only one submission for each team. This “representative” algorithm was chosen by each team after the true time delays were unblinded, and therefore it is somewhat indicative of the best performance of each method. Results for all the other submissions are available at the TDC Web site. Importantly, it should be kept in mind that this is a multi-dimensional problem, and there is not necessarily a “best” submission, not even within each method. Rather, each submission is a tradeoff between competing needs of achieving low P and A , while keeping χ^2 reasonable and f and X as high as possible. Some of the statistics are mathematically interdependent. For example, χ^2 and P both contain the submitted uncertainty estimates: teams could decide to reduce their χ^2 at the price of increasing their P , and vice versa.

The metrics obtained by these submissions are plotted in Figures 5–9. The plots show the metrics that have been computed directly from the submitted values, together with the recomputed metrics after rejecting the outliers using the $\chi_i^2 < 10$ cut. The corner plots in Figures 5–9 also show a shaded region that represents the TDC1 soft targets that were estimated in Paper I as the metric values needed for methods to be competitive, namely:

1. $f > 0.5$
2. $\chi^2 < 1.5$
3. $|A| < 0.03$ [goal 0.002]
4. $P < 0.03$.

As discussed in Paper I, in this exploratory challenge, these targets were deemed sufficient given the current lensed quasar sample of a few tens of systems. In the long run, for samples of thousands of lenses, a desirable goal is to improve the accuracy or bias to sub-percent level ($|A| < 0.2\%$; see Hojjati & Linder 2014 for the cosmological requirement derivation), such that the contribution of time delay measurement to the error budget of cosmological parameters would be smaller than the projected statistical uncertainties. We emphasize that these targets are approximate and only with a fully cosmological challenge would they be translated into a single indicator of performance, as we outline in the final section of this paper.

As is shown in the figures, most of the algorithms achieved the $|A|$ and χ^2 criteria, especially after the rejection of outliers in the submissions. The “Evil” Team’s baseline method had a large fraction of outliers, but once those were rejected, it did not perform significantly worse than many of the “Good” Teams submissions. The criterion that proved more difficult to meet was the one on the success fraction f , where teams were typically closer to the threshold for TDC0 (shown also in the cornerplot as a lighter shaded region) than for TDC1. As we discuss below, this is due to the strategy that most teams followed, i.e., to have high standards of acceptance in order to reduce outliers. Notably, for many of the methods $|A|$ is at the sub-percent level—well below the target of 0.03—which is very promising in view of future cosmological studies.

Interestingly, the “Evil” light curves did not yield significantly poorer statistics than the regular ones. From this comparison we infer that the methods used are generally robust to small and realistic unknown light curve systematics like the ones introduced by the “Evil” Team. This is encouraging and bodes well for the application of the methods to real data.

4.4. Trends with Intrinsic Properties of the Light Curves and Implications for Future Work

We now investigate how the quality of the inferred time delays depends on the intrinsic properties of the light curves. We wish to discover general trends that are not inherent to the peculiarities of each method. In order to carry out this investigation, in Figure 10 we plot the individual accuracy, precision and goodness of fit of each submission (A_i , P_i , and χ_i^2) as a function of true time delay, the variability parameters of the intrinsic quasar light curves (τ and σ), and the magnitude of the fainter image of each pair (i_2). In this illustration we show the results for Rung 1; the other rungs give similar results. Figure 11 shows summary statistics of the same data, represented by the average statistics in bins of the variable on the abscissa—the color scheme is the same as described in the legend to Figure 5.

We can see in these figures a few global trends. The most prominent appears to be between P and the true time delay. P decreases with time delay consistent with the time delay uncertainty being approximately constant in days, as expected if the absolute precision is driven by the sampling of the light curves. Qualitatively, P_i and A_i also appear to decrease (i.e., improve) as σ increases, also as expected: the light curves with the highest variability amplitudes should be easier to interpret and therefore should yield higher precision and fewer outliers.

Remarkably, we see very little dependency on i_2 , as if the signal-to-noise ratio of the fainter image is not as important, once it is passes some minimum threshold. This suggests that the simulated data are of sufficient quality and that the photometric uncertainty is subdominant with respect to the uncertainties introduced by microlensing and sampling. The weak dependency on the magnitude of the fainter image i_2 implies that the statistics we derive from the TDC1 data set are very similar to what we would have derived from a random subset of OM10. In fact, by recomputing weighted averages of the statistics to match the OM10 i_2 magnitude distribution, we verified that the changes of the statistics would have been comparable to their uncertainty.

Finally, we investigated the level of agreement between the algorithms to see whether success was due solely to the properties of the light curves or whether it depended on the specifics of each algorithm. The results are shown in Figure 12 for three

Table 4
Mean and Median Statistics for the “Representative” Submissions

Method	Rung	f	χ^2	P	A	χ^2_{median}	P_{median}	A_{median}
0	0	0.36	195000 ± 76000	0.078 ± 0.004	-0.181 ± 0.065	0.085 ^{0.189} _{0.078}	0.055 ^{0.083} _{0.036}	-0.00 ^{0.025} _{0.86}
0	1	0.36	390000 ± 150000	0.08 ± 0.005	-0.281 ± 0.061	0.47 ^{0.2046} _{0.46}	0.052 ^{0.088} _{0.039}	-0.02 ^{0.04} _{0.98}
0	2	0.32	3996 ± 1052	0.082 ± 0.005	-0.28 ± 0.042	0.42 ^{0.199} _{0.4}	0.059 ^{0.088} _{0.041}	-0.02 ^{0.05} _{0.97}
0	3	0.33	920000 ± 500000	0.08 ± 0.005	-0.247 ± 0.053	0.37 ^{0.2527} _{0.36}	0.05 ^{0.098} _{0.036}	-0.013 ^{0.034} _{0.97}
0	4	0.35	950000 ± 240000	0.042 ± 0.004	-0.712 ± 0.03	16136 ^{0.71657} ₁₆₁₃₆	0.008 ^{0.087} _{0.007}	-1.0 ^{0.99} _{0.007}
1	0	0.53	0.579 ± 0.047	0.038 ± 0.001	-0.018 ± 0.001	0.26 ^{0.77} _{0.22}	0.034 ^{0.028} _{0.016}	-0.015 ^{0.016} _{0.024}
1	1	0.37	0.543 ± 0.049	0.045 ± 0.001	-0.022 ± 0.001	0.24 ^{0.69} _{0.22}	0.04 ^{0.025} _{0.015}	-0.02 ^{0.017} _{0.022}
1	2	0.35	0.89 ± 0.19	0.053 ± 0.001	-0.025 ± 0.002	0.23 ^{0.92} _{0.21}	0.047 ^{0.034} _{0.021}	-0.02 ^{0.024} _{0.038}
1	3	0.34	0.524 ± 0.077	0.059 ± 0.002	-0.021 ± 0.002	0.17 ^{0.67} _{0.15}	0.051 ^{0.037} _{0.02}	-0.018 ^{0.025} _{0.029}
1	4	0.35	0.608 ± 0.072	0.056 ± 0.002	-0.024 ± 0.002	0.2 ^{0.84} _{0.18}	0.051 ^{0.036} _{0.024}	-0.019 ^{0.024} _{0.035}
2	0	0.53	0.125 ± 0.011	0.205 ± 0.007	-0.017 ± 0.004	0.043 ^{0.178} _{0.039}	0.151 ^{0.198} _{0.078}	-0.008 ^{0.046} _{0.062}
2	1	0.27	0.138 ± 0.016	0.233 ± 0.01	-0.025 ± 0.006	0.054 ^{0.216} _{0.05}	0.19 ^{0.17} _{0.1}	-0.008 ^{0.05} _{0.086}
2	2	0.21	0.043 ± 0.004	0.242 ± 0.01	-0.015 ± 0.004	0.021 ^{0.058} _{0.019}	0.201 ^{0.207} _{0.092}	-0.009 ^{0.04} _{0.056}
2	3	0.3	0.099 ± 0.013	0.247 ± 0.011	-0.03 ± 0.006	0.039 ^{0.121} _{0.035}	0.17 ^{0.266} _{0.085}	-0.013 ^{0.046} _{0.08}
2	4	0.21	0.178 ± 0.018	0.363 ± 0.015	-0.059 ± 0.011	0.097 ^{0.252} _{0.084}	0.32 ^{0.27} _{0.15}	-0.04 ^{0.12} _{0.15}
3	0	0.53	1.068 ± 0.069	0.043 ± 0.003	-0.0 ± 0.003	0.46 ^{1.67} _{0.4}	0.022 ^{0.041} _{0.012}	0.0 ^{0.025} _{0.025}
3	1	0.26	1.031 ± 0.097	0.04 ± 0.003	0.008 ± 0.003	0.49 ^{1.47} _{0.46}	0.027 ^{0.034} _{0.014}	0.004 ^{0.033} _{0.026}
3	2	0.21	1.02 ± 0.13	0.043 ± 0.004	-0.002 ± 0.004	0.38 ^{1.43} _{0.34}	0.026 ^{0.037} _{0.013}	0.003 ^{0.02} _{0.033}
3	3	0.3	0.813 ± 0.074	0.068 ± 0.006	-0.004 ± 0.006	0.39 ^{1.04} _{0.37}	0.034 ^{0.066} _{0.019}	-0.002 ^{0.032} _{0.032}
3	4	0.21	1.07 ± 0.23	0.098 ± 0.014	0.0 ± 0.008	0.24 ^{1.41} _{0.22}	0.064 ^{0.06} _{0.034}	0.003 ^{0.054} _{0.04}
4	0	0.53	0.497 ± 0.047	0.033 ± 0.002	-0.0 ± 0.001	0.15 ^{0.75} _{0.14}	0.018 ^{0.038} _{0.011}	0.0 ^{0.012} _{0.012}
4	1	0.27	0.528 ± 0.066	0.028 ± 0.002	0.0 ± 0.002	0.16 ^{0.78} _{0.15}	0.02 ^{0.021} _{0.01}	-0.001 ^{0.015} _{0.012}
4	2	0.21	0.464 ± 0.069	0.028 ± 0.002	-0.001 ± 0.002	0.15 ^{0.54} _{0.13}	0.02 ^{0.023} _{0.011}	0.0 ^{0.013} _{0.01}
4	3	0.3	0.542 ± 0.074	0.042 ± 0.003	-0.003 ± 0.003	0.16 ^{0.76} _{0.14}	0.023 ^{0.038} _{0.013}	-0.001 ^{0.017} _{0.015}
4	4	0.21	0.665 ± 0.065	0.045 ± 0.003	0.001 ± 0.003	0.31 ^{0.94} _{0.29}	0.032 ^{0.035} _{0.015}	-0.001 ^{0.035} _{0.028}
5	0	0.68	0.91 ± 0.092	0.032 ± 0.001	0.003 ± 0.002	0.24 ^{1.19} _{0.23}	0.024 ^{0.034} _{0.014}	0.001 ^{0.022} _{0.015}
5	1	0.27	1.76 ± 0.42	0.037 ± 0.002	-0.002 ± 0.003	0.39 ^{1.86} _{0.36}	0.03 ^{0.029} _{0.015}	-0.001 ^{0.026} _{0.026}
5	2	0.32	1.57 ± 0.21	0.043 ± 0.001	-0.003 ± 0.004	0.44 ^{1.93} _{0.41}	0.036 ^{0.036} _{0.017}	-0.001 ^{0.036} _{0.043}
5	3	0.35	1.89 ± 0.31	0.036 ± 0.001	0.002 ± 0.003	0.42 ^{2.3} _{0.4}	0.029 ^{0.03} _{0.015}	0.001 ^{0.029} _{0.031}
5	4	0.18	7.2 ± 2.7	0.05 ± 0.003	-0.021 ± 0.007	1.5 ^{4.3} _{1.4}	0.043 ^{0.04} _{0.021}	-0.016 ^{0.072} _{0.068}
6	0	0.04	0.32 ± 0.071	0.077 ± 0.017	0.005 ± 0.011	0.11 ^{0.66} _{0.1}	0.044 ^{0.06} _{0.027}	0.0 ^{0.027} _{0.037}
6	1	0.02	66 ± 64	0.175 ± 0.055	2.3 ± 2.2	0.36 ^{0.25} _{0.27}	0.093 ^{0.13} _{0.037}	0.042 ^{0.056} _{0.047}
6	2	0.03	0.71 ± 0.21	0.142 ± 0.021	0.027 ± 0.032	0.37 ^{0.78} _{0.36}	0.117 ^{0.098} _{0.064}	0.029 ^{0.077} _{0.089}
6	3	0.02	1.7 ± 1.2	0.168 ± 0.031	0.14 ± 0.1	0.33 ^{0.9} _{0.3}	0.118 ^{0.079} _{0.068}	0.02 ^{0.119} _{0.056}
6	4	0.01	0.19 ± 0.1	0.55 ± 0.12	0.169 ± 0.058	0.066 ^{0.169} _{0.051}	0.48 ^{0.2} _{0.25}	0.16 ^{0.24} _{0.19}
7	0	0.33	65 ± 51	0.04 ± 0.003	-0.011 ± 0.009	0.6 ^{3.44} _{0.55}	0.021 ^{0.057} _{0.015}	-0.0 ^{0.029} _{0.034}
7	1	0.24	2.71 ± 0.5	0.036 ± 0.003	0.002 ± 0.006	0.67 ^{3.14} _{0.62}	0.021 ^{0.045} _{0.015}	0.001 ^{0.034} _{0.029}
7	2	0.37	3.21 ± 0.55	0.04 ± 0.003	0.008 ± 0.006	0.74 ^{3.39} _{0.69}	0.023 ^{0.051} _{0.015}	-0.0 ^{0.036} _{0.029}
7	3	0.3	2.39 ± 0.39	0.051 ± 0.004	0.02 ± 0.012	0.65 ^{2.49} _{0.6}	0.025 ^{0.067} _{0.018}	-0.0 ^{0.035} _{0.035}
7	4	0.22	185 ± 119	0.062 ± 0.005	-0.03 ± 0.02	0.6 ^{3.47} _{0.53}	0.035 ^{0.104} _{0.026}	-0.001 ^{0.061} _{0.064}
8	0	0.44	109 ± 58	0.047 ± 0.004	-0.025 ± 0.032	0.16 ^{1.21} _{0.15}	0.025 ^{0.047} _{0.016}	0.0 ^{0.019} _{0.021}
8	1	0.22	88 ± 38	0.101 ± 0.05	-0.02 ± 0.019	0.17 ^{2.4} _{0.16}	0.029 ^{0.066} _{0.016}	0.0 ^{0.026} _{0.04}
8	2	0.18	91 ± 72	0.07 ± 0.006	-0.006 ± 0.019	0.14 ^{0.81} _{0.14}	0.046 ^{0.076} _{0.028}	0.0 ^{0.032} _{0.032}
8	3	0.19	27 ± 21	0.059 ± 0.004	-0.008 ± 0.013	0.24 ^{1.34} _{0.23}	0.041 ^{0.064} _{0.025}	0.001 ^{0.033} _{0.041}
8	4	0.16	2.6 ± 1.1	0.068 ± 0.004	-0.0 ± 0.006	0.3 ^{1.29} _{0.28}	0.055 ^{0.07} _{0.032}	0.001 ^{0.045} _{0.049}
9	4	0.27	8.7 ± 3.5	0.035 ± 0.002	0.003 ± 0.006	0.55 ^{2.79} _{0.47}	0.024 ^{0.037} _{0.014}	0.0 ^{0.031} _{0.042}

Notes. Method: 0, Rumbaugh-Gold; 1, Shafieloo-Arman7; 2, PyCS-d3cs-vanilla-dou-full; 3, PyCS-sdi-vanilla-dou-full; 4, PyCS-spl-vanilla-dou-full; 5, Jackson-manchester2_0_3_4; 6, Kumar; 7, JPL; 8, Hojjati-Stark; 9, DeltaTBayes-DeltaTBayes1.

representative rungs. Clearly, some light curves do not contain enough information for any method to be successful (hence the peak at zero). In Rung 0, there is a bump around 6 indicating that for very good light curve a majority of the methods are

successful. However, as the quality of data degrades in the next rungs it appears that there is a continuum distribution. Therefore we conclude that different methods pick up different features of the light curves and accuracy varies widely between methods.

Table 5
Filtered Statistics for the “Representative” Submissions

Method	Rung	$f_{3.3\sigma}$	$\chi_{3.3\sigma}^2$	$P_{3.3\sigma}$	$A_{3.3\sigma}$	X	f_A	χ_A^2	P_A	A_A	X_A
0	0	0.29	0.379 ± 0.072	0.087 ± 0.005	−0.003 ± 0.004	0.8	0.28	0.299 ± 0.056	0.08 ± 0.004	−0.0 ± 0.002	0.77
0	1	0.23	0.577 ± 0.095	0.096 ± 0.006	−0.01 ± 0.007	0.65	0.22	3.9 ± 2.3	0.082 ± 0.005	−0.004 ± 0.002	0.62
0	2	0.23	0.8 ± 0.11	0.098 ± 0.005	−0.007 ± 0.006	0.73	0.21	0.74 ± 0.23	0.09 ± 0.005	−0.002 ± 0.003	0.66
0	3	0.22	0.59 ± 0.1	0.097 ± 0.006	0.0 ± 0.007	0.66	0.21	1.26 ± 0.4	0.087 ± 0.006	−0.002 ± 0.002	0.64
0	4	0.11	0.37 ± 0.11	0.119 ± 0.009	−0.009 ± 0.006	0.3	0.1	0.26 ± 0.058	0.112 ± 0.008	−0.003 ± 0.004	0.28
1	0	0.53	0.552 ± 0.04	0.038 ± 0.001	−0.017 ± 0.001	1.0	0.52	0.53 ± 0.038	0.038 ± 0.001	−0.017 ± 0.001	0.99
1	1	0.37	0.543 ± 0.049	0.045 ± 0.001	−0.022 ± 0.001	1.0	0.36	0.497 ± 0.041	0.044 ± 0.001	−0.021 ± 0.001	0.99
1	2	0.35	0.673 ± 0.068	0.053 ± 0.001	−0.025 ± 0.002	0.99	0.33	0.73 ± 0.19	0.052 ± 0.001	−0.021 ± 0.002	0.95
1	3	0.34	0.458 ± 0.039	0.059 ± 0.002	−0.02 ± 0.002	1.0	0.33	0.419 ± 0.036	0.058 ± 0.002	−0.018 ± 0.002	0.97
1	4	0.35	0.559 ± 0.052	0.056 ± 0.002	−0.024 ± 0.002	1.0	0.33	0.535 ± 0.069	0.055 ± 0.002	−0.021 ± 0.002	0.97
2	0	0.53	0.125 ± 0.011	0.205 ± 0.007	−0.017 ± 0.004	1.0	0.45	0.081 ± 0.008	0.17 ± 0.006	−0.005 ± 0.002	0.83
2	1	0.27	0.138 ± 0.016	0.233 ± 0.01	−0.025 ± 0.006	1.0	0.21	0.078 ± 0.01	0.191 ± 0.008	−0.006 ± 0.003	0.79
2	2	0.21	0.043 ± 0.004	0.242 ± 0.01	−0.015 ± 0.004	1.0	0.19	0.033 ± 0.004	0.217 ± 0.009	−0.007 ± 0.003	0.9
2	3	0.3	0.099 ± 0.013	0.247 ± 0.011	−0.03 ± 0.006	1.0	0.25	0.056 ± 0.005	0.201 ± 0.01	−0.007 ± 0.003	0.83
2	4	0.21	0.178 ± 0.018	0.363 ± 0.015	−0.059 ± 0.011	1.0	0.12	0.063 ± 0.008	0.287 ± 0.018	−0.007 ± 0.005	0.55
3	0	0.53	1.048 ± 0.066	0.043 ± 0.003	−0.0 ± 0.003	1.0	0.5	0.956 ± 0.068	0.037 ± 0.003	0.001 ± 0.001	0.94
3	1	0.26	0.977 ± 0.081	0.04 ± 0.003	0.006 ± 0.003	1.0	0.25	0.858 ± 0.069	0.037 ± 0.003	0.004 ± 0.002	0.95
3	2	0.21	0.94 ± 0.1	0.043 ± 0.004	−0.002 ± 0.004	0.99	0.2	0.92 ± 0.13	0.035 ± 0.002	−0.003 ± 0.002	0.93
3	3	0.3	0.813 ± 0.074	0.068 ± 0.006	−0.004 ± 0.006	1.0	0.27	0.747 ± 0.073	0.05 ± 0.004	−0.003 ± 0.002	0.92
3	4	0.21	0.804 ± 0.096	0.098 ± 0.015	0.005 ± 0.006	0.99	0.19	0.64 ± 0.11	0.069 ± 0.004	0.005 ± 0.003	0.86
4	0	0.53	0.472 ± 0.04	0.033 ± 0.002	−0.0 ± 0.001	1.0	0.52	0.483 ± 0.048	0.029 ± 0.001	0.0 ± 0.001	0.98
4	1	0.27	0.528 ± 0.066	0.028 ± 0.002	0.0 ± 0.002	1.0	0.27	0.467 ± 0.051	0.027 ± 0.002	−0.0 ± 0.001	0.99
4	2	0.21	0.464 ± 0.069	0.028 ± 0.002	−0.001 ± 0.002	1.0	0.21	0.431 ± 0.064	0.028 ± 0.002	−0.001 ± 0.001	0.99
4	3	0.3	0.494 ± 0.057	0.042 ± 0.003	−0.001 ± 0.003	1.0	0.29	0.455 ± 0.052	0.037 ± 0.003	−0.001 ± 0.001	0.97
4	4	0.21	0.665 ± 0.065	0.045 ± 0.003	0.001 ± 0.003	1.0	0.2	0.571 ± 0.056	0.041 ± 0.002	0.0 ± 0.002	0.95
5	0	0.68	0.741 ± 0.053	0.032 ± 0.001	0.004 ± 0.002	0.99	0.65	0.659 ± 0.054	0.03 ± 0.001	0.002 ± 0.001	0.95
5	1	0.27	0.926 ± 0.098	0.037 ± 0.002	−0.003 ± 0.003	0.97	0.26	1.42 ± 0.42	0.034 ± 0.002	−0.001 ± 0.002	0.93
5	2	0.31	1.083 ± 0.096	0.043 ± 0.001	−0.002 ± 0.003	0.97	0.29	1.08 ± 0.13	0.04 ± 0.001	−0.001 ± 0.002	0.92
5	3	0.34	1.165 ± 0.099	0.036 ± 0.001	0.002 ± 0.003	0.98	0.32	1.23 ± 0.17	0.032 ± 0.001	0.0 ± 0.002	0.91
5	4	0.16	2.12 ± 0.2	0.052 ± 0.003	−0.015 ± 0.007	0.92	0.15	5.4 ± 3.1	0.044 ± 0.002	−0.011 ± 0.004	0.82
6	0	0.04	0.32 ± 0.071	0.077 ± 0.017	0.005 ± 0.011	1.0	0.04	0.32 ± 0.073	0.063 ± 0.01	−0.004 ± 0.006	0.97
6	1	0.02	0.334 ± 0.051	0.121 ± 0.016	0.04 ± 0.014	0.95	0.02	0.31 ± 0.053	0.111 ± 0.016	0.027 ± 0.012	0.86
6	2	0.03	0.71 ± 0.21	0.142 ± 0.021	0.027 ± 0.032	1.0	0.02	0.333 ± 0.087	0.111 ± 0.012	0.019 ± 0.011	0.75
6	3	0.02	0.51 ± 0.15	0.155 ± 0.03	0.037 ± 0.02	0.95	0.02	0.278 ± 0.095	0.13 ± 0.034	−0.003 ± 0.011	0.64
6	4	0.01	0.19 ± 0.1	0.55 ± 0.12	0.169 ± 0.058	1.0	0.0	0.024 ± 0.011	0.358 ± 0.075	−0.005 ± 0.026	0.33
7	0	0.31	1.42 ± 0.12	0.041 ± 0.003	−0.001 ± 0.004	0.95	0.3	1.82 ± 0.28	0.033 ± 0.003	−0.001 ± 0.002	0.89
7	1	0.23	1.39 ± 0.13	0.037 ± 0.003	−0.0 ± 0.006	0.95	0.22	2.25 ± 0.47	0.028 ± 0.002	0.002 ± 0.002	0.91
7	2	0.35	1.41 ± 0.1	0.04 ± 0.003	0.006 ± 0.004	0.94	0.33	2.06 ± 0.34	0.032 ± 0.002	−0.001 ± 0.002	0.89
7	3	0.28	1.28 ± 0.11	0.051 ± 0.004	0.007 ± 0.007	0.95	0.26	1.82 ± 0.32	0.033 ± 0.002	−0.003 ± 0.002	0.87
7	4	0.21	1.33 ± 0.14	0.063 ± 0.005	0.003 ± 0.007	0.93	0.18	1.93 ± 0.44	0.043 ± 0.004	0.002 ± 0.003	0.79
8	0	0.42	0.531 ± 0.054	0.047 ± 0.004	−0.0 ± 0.002	0.95	0.41	0.81 ± 0.14	0.041 ± 0.003	−0.001 ± 0.001	0.93
8	1	0.2	0.596 ± 0.087	0.105 ± 0.056	−0.004 ± 0.004	0.9	0.2	0.76 ± 0.14	0.101 ± 0.057	−0.001 ± 0.002	0.89
8	2	0.17	0.62 ± 0.11	0.07 ± 0.006	0.003 ± 0.004	0.96	0.16	0.354 ± 0.064	0.064 ± 0.005	−0.001 ± 0.003	0.88
8	3	0.18	0.78 ± 0.12	0.06 ± 0.004	−0.003 ± 0.005	0.96	0.17	1.03 ± 0.34	0.053 ± 0.004	0.0 ± 0.003	0.89
8	4	0.16	0.89 ± 0.14	0.07 ± 0.004	0.002 ± 0.005	0.98	0.15	1.59 ± 0.69	0.063 ± 0.004	0.002 ± 0.003	0.9
9	4	0.25	1.2 ± 0.1	0.036 ± 0.003	−0.006 ± 0.004	0.94	0.25	3.7 ± 1.4	0.03 ± 0.002	−0.002 ± 0.002	0.91

Notes. Method 0:Rumbaugh-Gold, 1:Shafieloo-Arman7, 2:PyCS-d3cs-vanilla-dou-full, 3:PyCS-sdi-vanilla-dou-full, 4:PyCS-spl-vanilla-dou-full, 5:Jackson-manchester2_0_3_4, 6:Kumar, 7:JPL, 8:Hojjati-Stark, 9:DeltaTBayes-DeltaTBayes1.

5. IMPLICATIONS FOR OBSERVING STRATEGY

By comparing the results from the different TDC1 rungs, we can now answer the following question: How does time delay measurement accuracy depend on observing cadence, season length and campaign length?

Figure 13 shows the variation in the $|A|$, P , and f metrics with cadence and season length, assuming outliers to have been rejected by $|A_i| > 0.1$. Each pair of connected points plotted in the panels of this figure represents a simple test where the control variable (cadence or season length) is varied, while keeping the others constant. Campaign length and cadence regularity were also investigated in a similar manner, but the results—which are less striking—are not shown here. The six tests we carried out

in total are summarized in Table 6. The top two rows in the table correspond to the plots shown in the left and right columns of the figure, respectively.

Figure 13 shows some interesting diversity between methods. Despite this, some approximate general trends can be seen. Greater accuracy and success fractions seem to be associated primarily with longer seasons, but there is considerable scatter between submissions, perhaps due to residual outliers in some cases. In most methods, little dependence of accuracy on cadence, campaign lengths beyond five years, or the regularity of the sampling was seen. The success fraction seems to be somewhat dependent on cadence but less so on campaign length. In general, the trends in precision with cadence and season length seem to be less marked, and show less scatter, than those

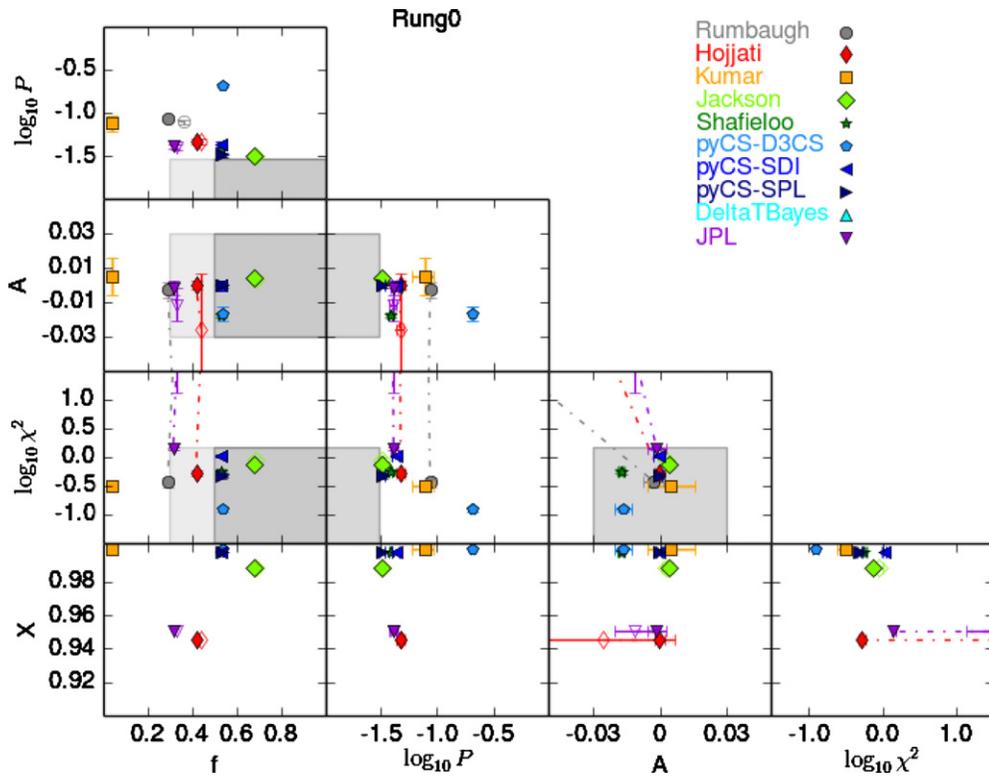


Figure 5. Results for TDC1 Rung 0, showing metrics for the “representative” submission for each of the 10 algorithms. This includes the baseline submission by the “Evil” Team (“Rumbaugh”). The f , P , A , and χ^2 metrics are defined in Section 1, while X is defined in Section 4.1. The shaded regions of each plot represent the soft targets for TDC1, as presented in the TDC0 paper. Both unfiltered results (open symbols) and results filtered by $\chi_i^2 < 10$ (solid symbols) are presented, and they are connected by dashed lines to show the improvements. Rung 0 simulates 3 day cadence and 8 month seasons over a 5 yr campaign with 400 observations in total (Table 1).

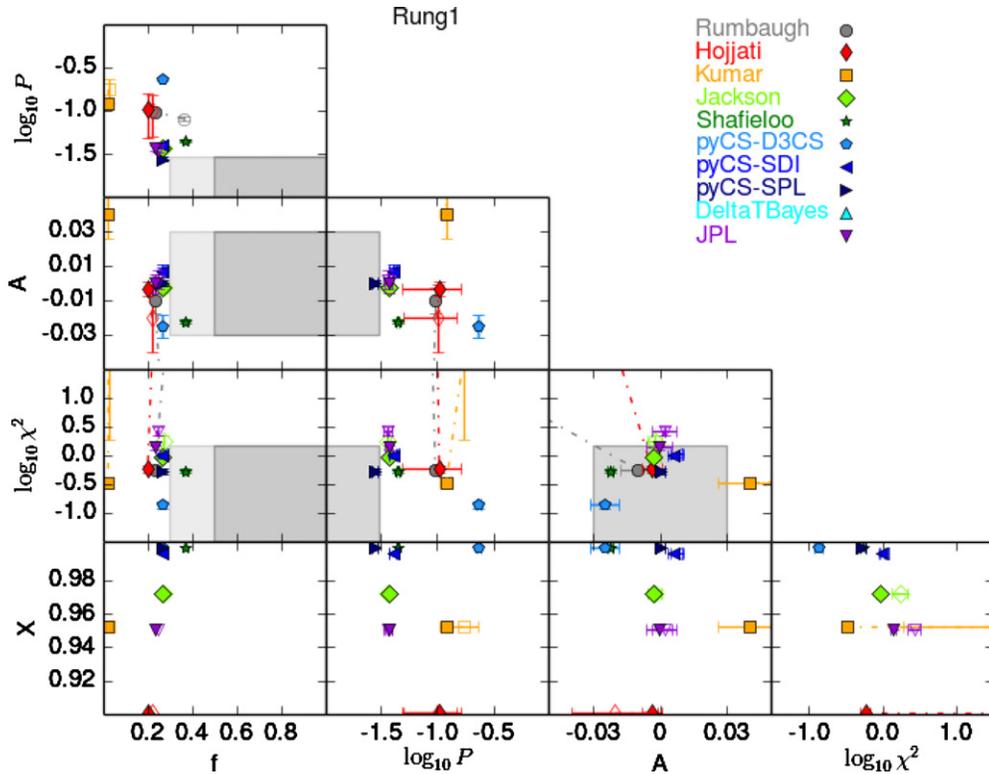


Figure 6. Same as Figure 5, but showing the results of TDC1 Rung 1, which simulates 3 day cadence and 4 month seasons over a 10 yr campaign with 400 observations in total (Table 1).

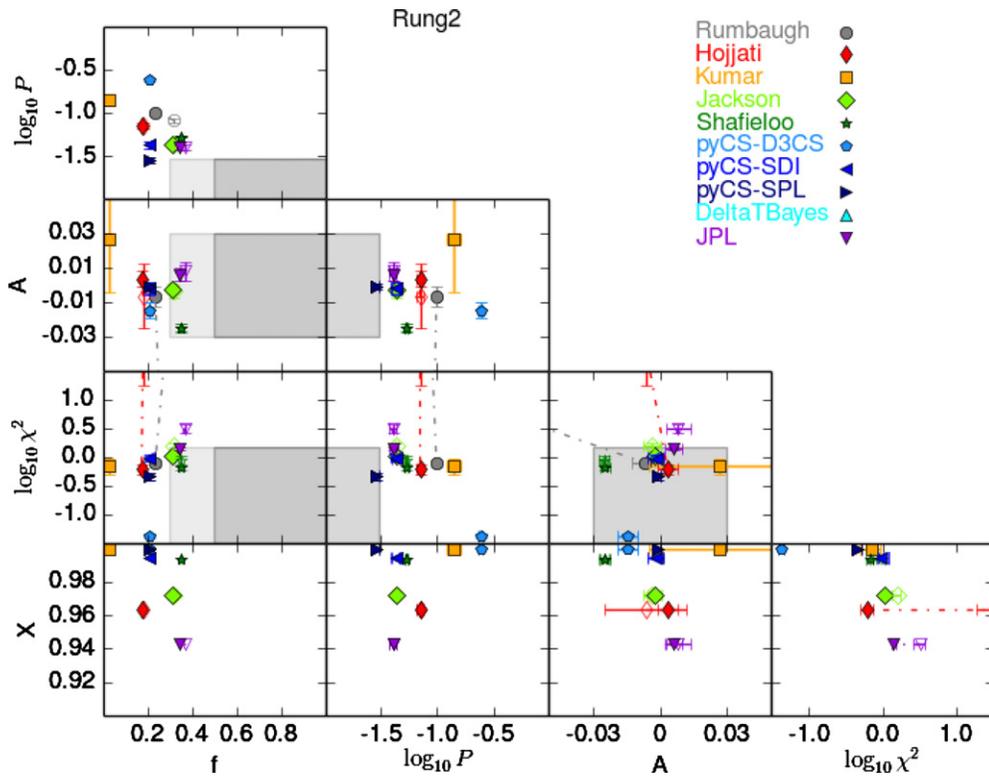


Figure 7. Same as Figure 5, but showing the results of TDC1 Rung 2, which simulates 3 day cadence and 4 month seasons over a 5 yr campaign with 200 observations in total (Table 1), and exactly regular time sampling.

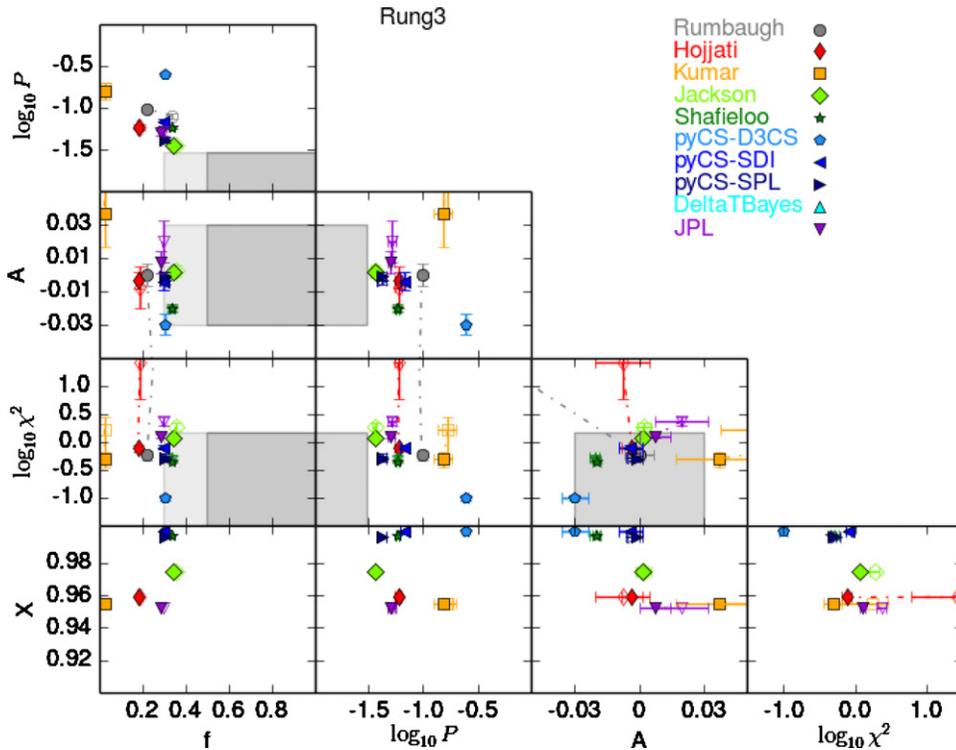


Figure 8. Same as Figure 5, but showing the results of TDC1, Rung 3, which simulates 3 day cadence and 4 month seasons over a 5 yr campaign with 200 observations in total (Table 1), and with 1 day scatter in the separations between observations.

in accuracy and success fraction. In general, cadence seems to be the most important factor for precision.

While the variation of time delay measurement with observing strategy seems to be somewhat algorithm-dependent, we can nevertheless hope to capture the general trends just

described. Focusing on the PyCS-SPL results, we derived a very approximate power-law model for how the A , P , and f metrics varied with the main three quantities that describe the observing strategies in the rungs, mean cadence (cad), mean season length (sea), and campaign length (camp).

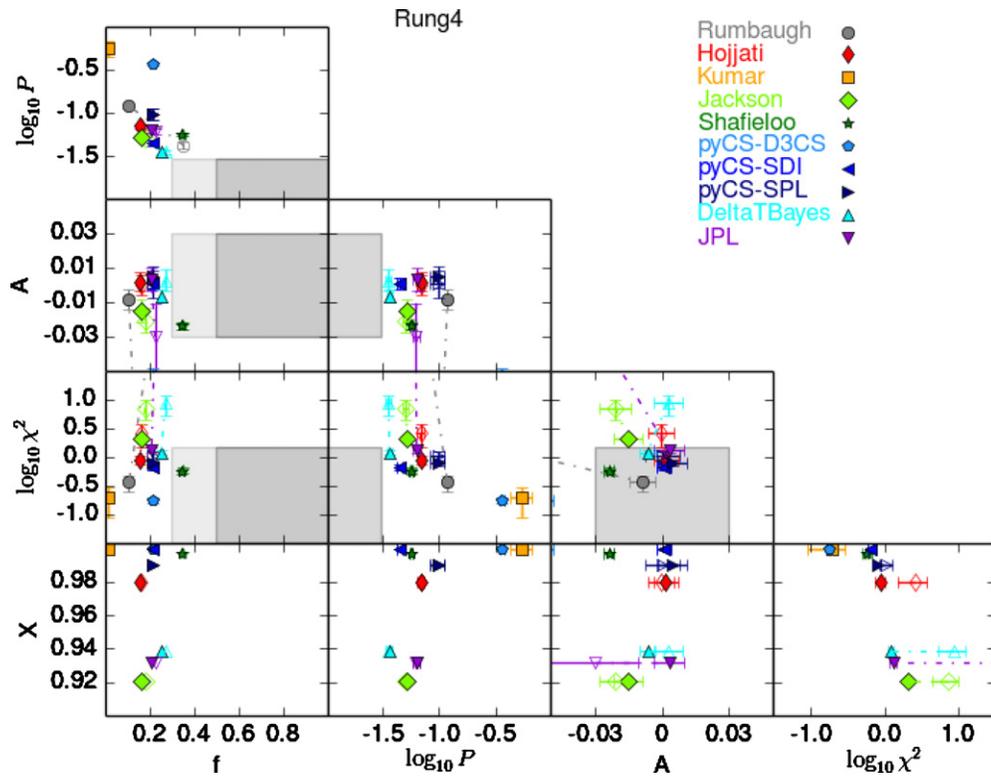


Figure 9. Same as Figure 5, but showing the results of TDC1, Rung 4, which simulates 6 day cadence and 4 month seasons over a 10 yr campaign with 200 observations in total (Table 1).

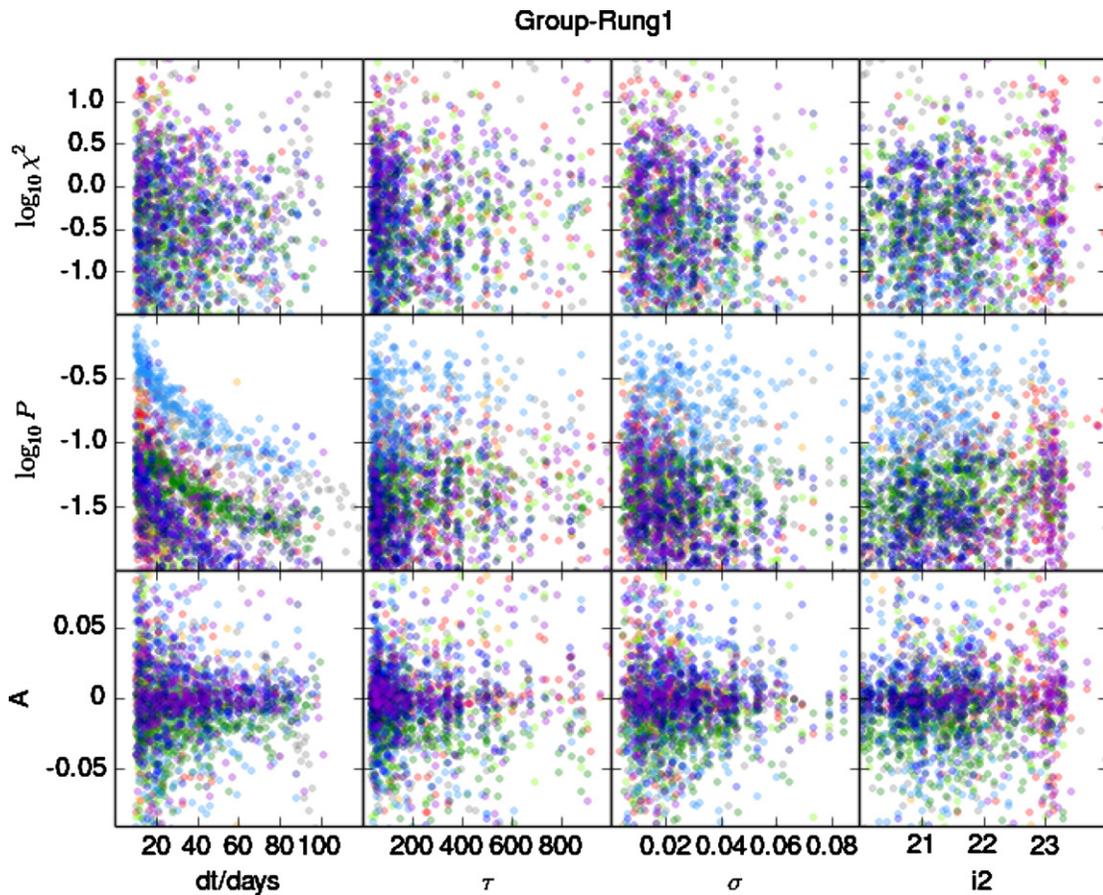


Figure 10. Unfiltered results of Rung 1: individual metrics of each “representative” submission (A_i, P_i, χ_i^2) as a function of true time delay dt , the variability parameters of the intrinsic quasar light curves (τ, σ), and the magnitude of the fainter image of each pair (i_2). The color scheme is the same as that described in the legend of Figure 5.

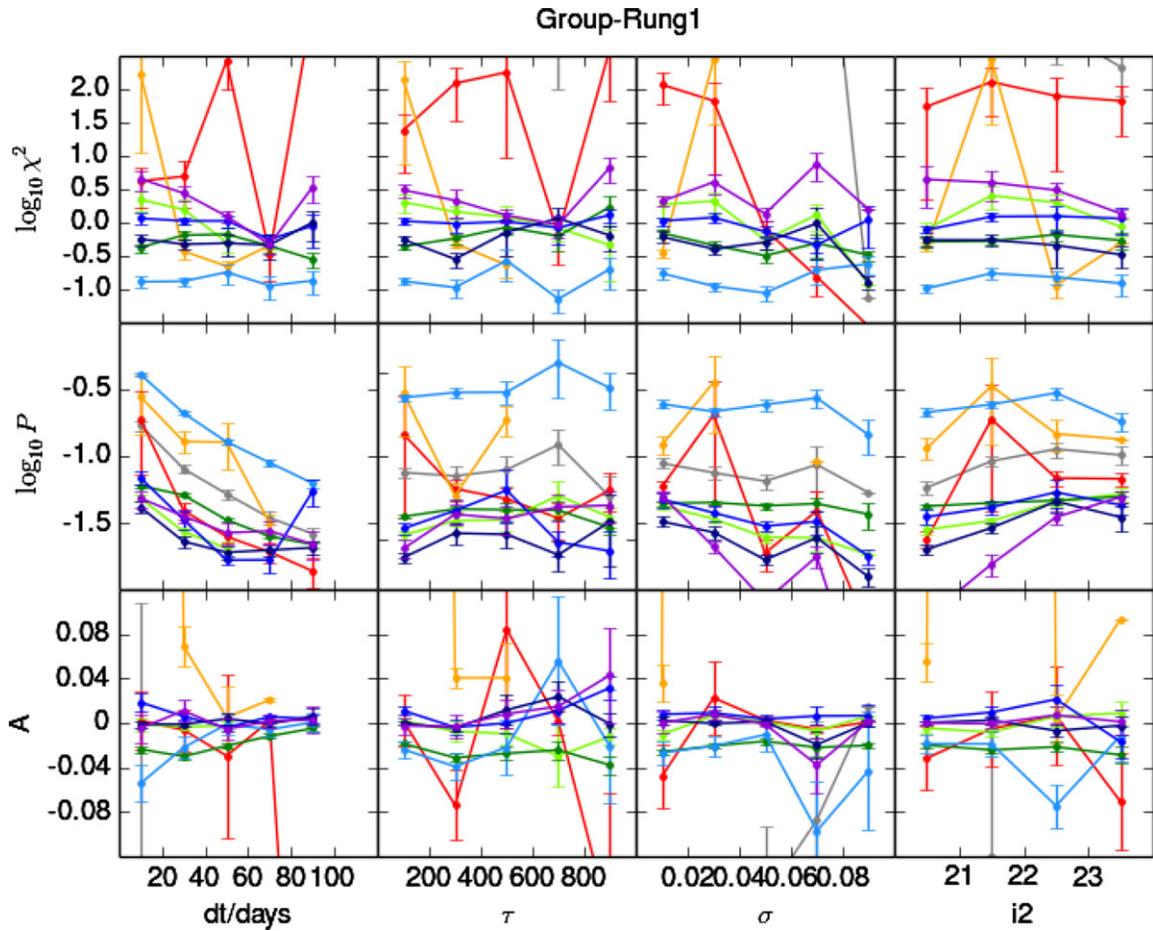


Figure 11. Summary unfiltered statistics of the same data in Figure 10, represented by the average statistics in bins of the variable on the abscissa. The color scheme is the same as described in the legend to Figure 5.

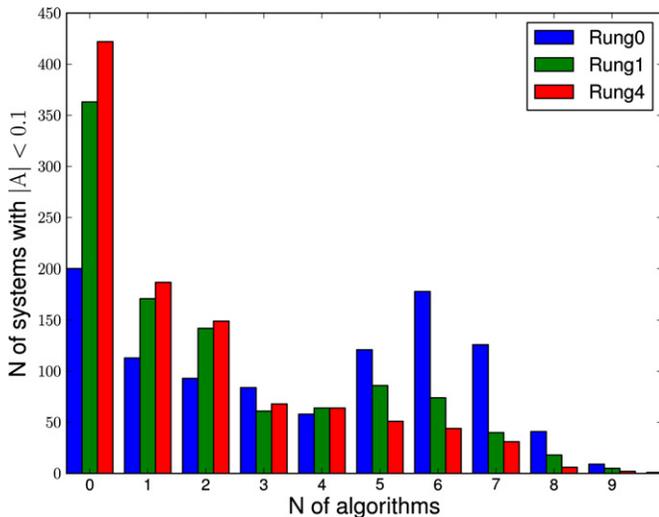


Figure 12. Distribution of the number of systems for which the time delay was successfully measured to a level of $|A_i| < 0.1$, plotted as a function of the number of algorithms (out of 10) that measure the time delays to this level. The plot shows Rungs 0, 1, and 4, which represent COSMOGRAIL-like, “optimistic” LSST, and “realistic” LSST programs, respectively. For Rung 0, there were more than ~ 200 systems for which none of the algorithms achieved the desired A , but also a large number of systems for which five, six, or seven of the algorithms successfully recovered this level. For Rungs 1 and 4, fewer of the systems were successfully recovered at the $|A_i| < 0.1$ level.

We find

$$|A|_{\text{model}} \approx 0.06\% \left(\frac{\text{cad}}{3 \text{ days}} \right)^{0.0} \left(\frac{\text{sea}}{4 \text{ months}} \right)^{-1.0} \left(\frac{\text{camp}}{5 \text{ yr}} \right)^{-1.1}$$

$$P_{\text{model}} \approx 4.0\% \left(\frac{\text{cad}}{3 \text{ days}} \right)^{0.7} \left(\frac{\text{sea}}{4 \text{ months}} \right)^{-0.3} \left(\frac{\text{camp}}{5 \text{ yr}} \right)^{-0.6}$$

$$f_{\text{model}} \approx 30\% \left(\frac{\text{cad}}{3 \text{ days}} \right)^{-0.4} \left(\frac{\text{sea}}{4 \text{ months}} \right)^{0.8} \left(\frac{\text{camp}}{5 \text{ yr}} \right)^{-0.2}$$

We can see that in this model, the accuracy metric A is the most sensitive to the observing strategy. It is also the case that it is the metric most sensitive to how the outliers are rejected. Rejecting outliers that have $\chi_i^2 > 10$ gives similar conclusions to those drawn here, but slightly different model parameters, in the sense that there is even stronger dependence of A on the observing strategy. In both cases the dependence of A on cadence is relatively weak. The season length and campaign length seem to be more important parameters: doubling either of these results in approximately a factor of two improvement in A . We note that constraining the total number of observations weakens these dependencies somewhat: for example, at fixed cadence, lengthening the season means shortening the campaign, and in our model, $|A|$ then decreases only as the ratio of the season length to the campaign length to the power of 0.1. The results of the fixed epoch number tests in Table 6 bore this out.

The precision and success fraction metrics’ dependence on observing strategy is weaker, but it is interesting to note that

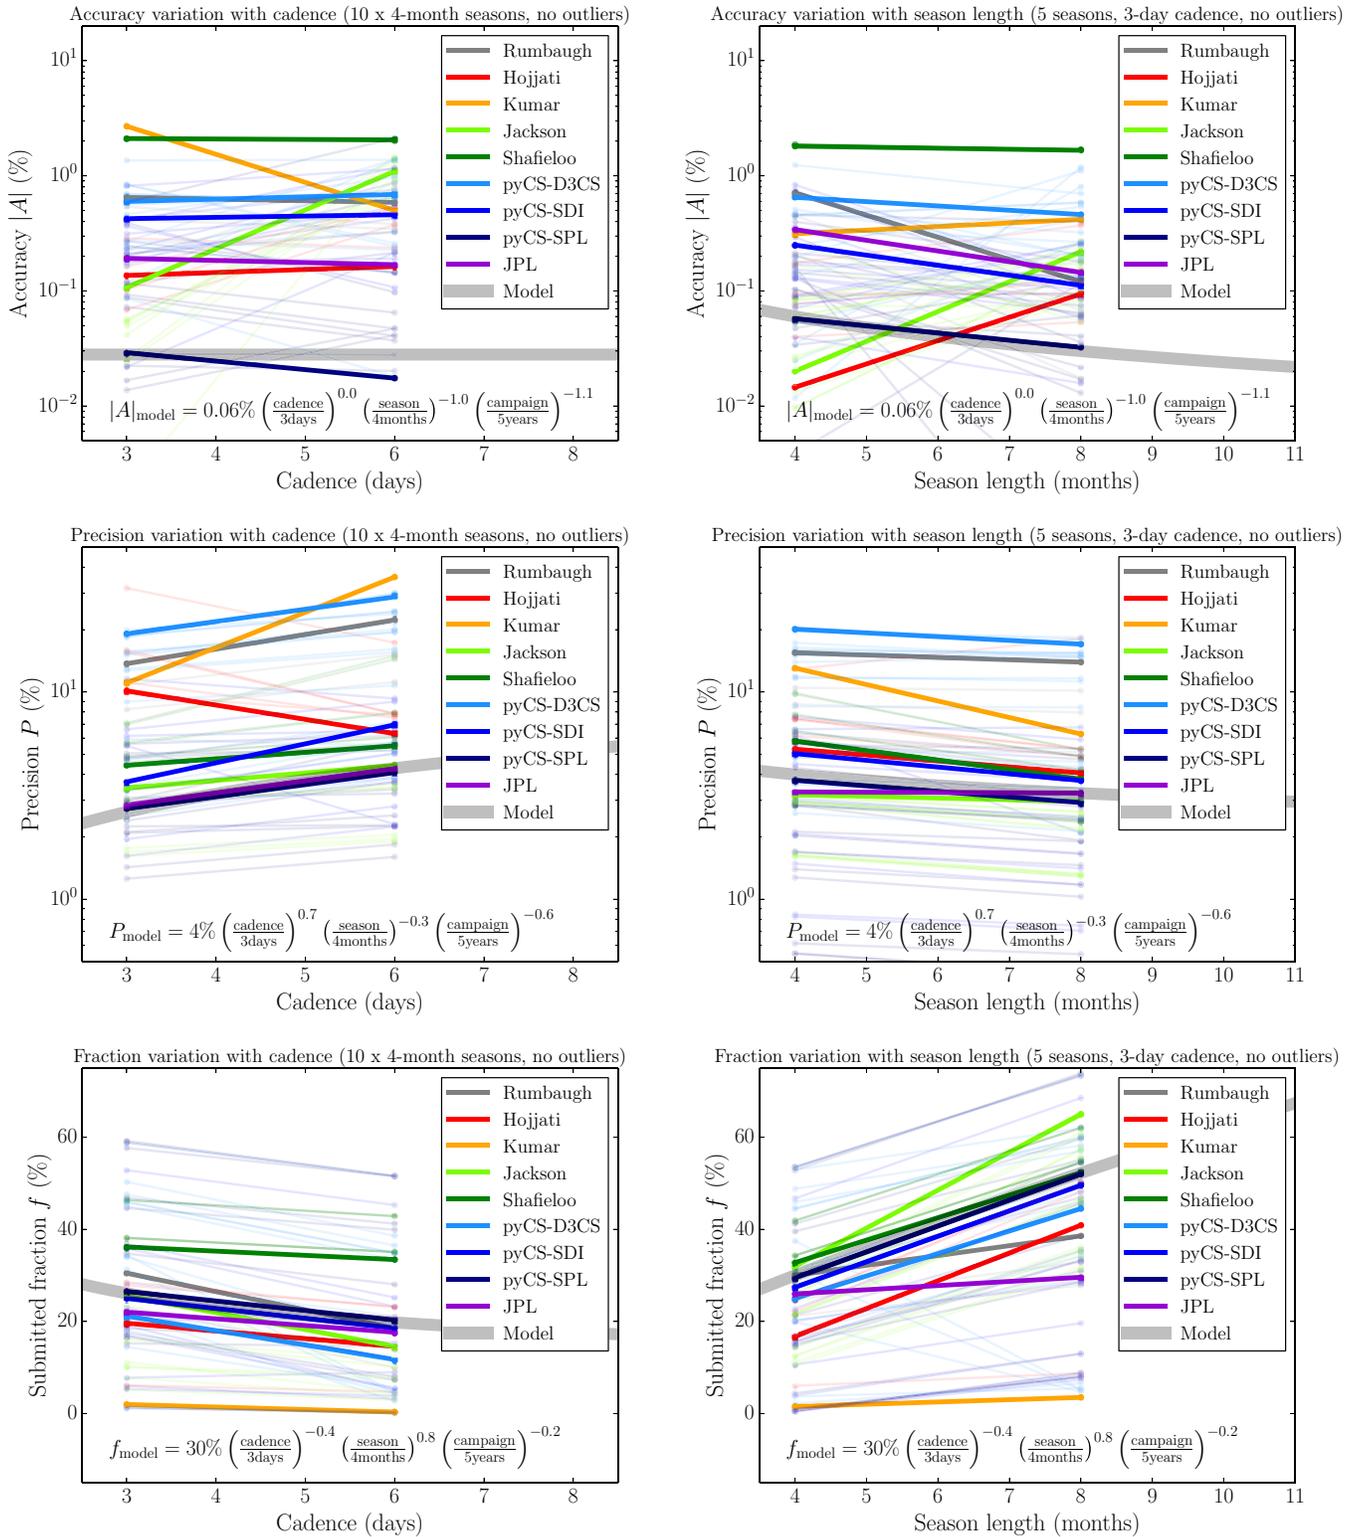


Figure 13. Changes in accuracy A (top row), precision P (middle row) and success fraction f (bottom row) with cadence (left) and season length (right), seen in the different TDC1 submissions. The gray approximate power-law model was derived by visual inspection of the pyCS-SPL results; the signs of the indices were pre-determined according to our expectations.

the precision depends more strongly on cadence than the season length, while the opposite is true for the success fraction. This can be understood qualitatively as the presence of large gaps reducing the overlap between light curves, making it more difficult to reliably and uniquely identify common features

between them. Conversely, if the signal is properly identified, then the precision is driven by the total number of observation points, i.e., a combination of cadence and campaign duration. As a rough rule of thumb, we might have in mind that season length largely determines bias, while cadence controls precision.

Table 6
Exploring Time Delay Estimation Performance Against Observing Strategy

Rungs	Variable Parameter	Fixed Parameters
1, 4	Cadence (3, 6 days)	4 month seasons, 10 yr campaign
0, 3	Season (4, 8 months)	3 day cadence, 5 yr campaign
3, 4	Cadence (3,6 days)	4 month seasons, 200 epochs length
0, 1	Season (4, 8 months)	3 day cadence, 400 epochs length
1,3	Campaign (5, 10 yr)	3 day cadence, 4 month seasons
2, 3	Cadence dispersion (0, 1 days)	3 day cadence, 4 month season, 5 yr campaign

Note. The tests defined in the top two rows (above the line) are illustrated in Figure 13.

The precision of an ensemble average parameter, such as the cosmological parameters, may yet depend primarily on season length, however, through the success fraction.

These simple model conclusions represent small extrapolations—we did not, for example, test doubling the season length and cadence simultaneously—but they represent a first approximation to the response of the more accurate time delay estimation routines to variations in observing strategy.

Finally, we note briefly the implications of this model for the sample of lensed quasars that was forecast for LSST by Oguri & Marshall (2010). Rung 4 represents something like the “universal cadence” planned for LSST (Ivezic et al. 2008), albeit with slightly shorter seasons. A cadence of 6 days would be well within the reach of such a strategy, but would require using observations from most of the filters in the set. While in this work we have only simulated and analyzed single filter light curves, AGN variability has been observed to be significantly correlated across the optical and near infrared bands (see, e.g., Schmidt et al. 2012), and microlensing variability is expected, and observed, to vary smoothly with wavelength due to source size effects (e.g., Poindexter et al. 2008). With sufficiently sophisticated algorithms we might expect to be able to measure time delays from multi-filter light curves with fidelity not dissimilar to that shown by the TDC1 methods tested here. The three day cadence of Rung 1 could be achieved by LSST without changing the total number of visits; the impact of such a strategy on the various different LSST science cases would need to be investigated. We take Rungs 1 and 4 to span the range of possibilities for the LSST time sampling.

Our model suggests that, if outliers with $|A_i| > 0.1$ can be rejected (perhaps during a joint analysis of the ensemble), the cadence is effectively unimportant for time delay measurement bias, and with LSST we might expect to achieve an accuracy metric of $|A| = 0.03\%–0.06\%$. Such a small time delay measurement bias is well below the systematic errors expected from lens modeling. Meanwhile, the expected precision achievable per lens in the Rung 1 and 4 cadences would be 2.6%–4.3%, and the success fractions would be 20%–26%. The mock lenses used in this data challenge were not quite randomly drawn from the OM10 catalog, but instead had approximately uniformly distributed i_3 image magnitudes within four broad magnitude bins (Section 2.2). Correcting for this, we find that we might, with the present-day algorithms (tested here and represented by our simple model), expect to be able to make time delay measurements with the above accuracy in at least 20% of an LSST sample of 1990 lenses selected to have $i_3 < 23.3$ and $10 \text{ days} < \Delta t < 120 \text{ days}$. This would correspond to a well-measured sample of around 400 lensed quasars. We must expect these numbers to be refined as the LSST observing strategy is defined, and further time delay measurement tests are carried out.

6. DISCUSSION

In this section, we give a brief analysis of each method’s performance, discussing how they performed and what can be improved in the future. We note that the performance of each method must be evaluated in multi-dimensional metric space. Each “Good” Team had to make choices with respect to which metric to optimize. Some teams decided to favor inclusiveness (high f) at the cost of a higher fraction of outliers (lower X) or lower precision P , and vice-versa. In fact, some of the teams submitted multiple entries spanning the range of parameter space, and illustrating these competitive needs. Therefore, at this stage it is not possible, nor useful, to identify a “best” submission, not even within each method. It is more fruitful to understand the tradeoffs and explore the range covered by each method, and then identify areas for improvement.

6.1. Gaussian Processes, by Hojjati & Linder

The GP method attained its twin goals of an automated fitting pipeline and very good fit accuracy. The main issue to address is one of outliers, which can be handled in two ways: global clipping and image information. This team found that the outliers were not due to multi-modal fit distributions—indeed the fits often have better likelihood for the data than the truth. However, the cosmology derived from the outliers would be discrepant from the cosmology from the global fit ensemble, and in this way, outliers could be recognized and clipped. Another approach would be to use information such as image separation (not provided in TDC1) to recognize and discard discrepant fits. While these considerations would lower the accepted fraction of fits, the correction of the mean function discussed in Section 3.2 raises the fraction over those given here. This, and a set of new but simple selection criteria (no limits on precision were imposed by this team for TDC1 submissions), discussed in a follow-up paper by Hojjati & Linder (2014), give considerable improvement in the precision and fraction, and further improvement in accuracy.

6.2. FOT, by Romero-Wolf & Moustakas

The unblinding of the TDC1 simulated data provided valuable information on the behavior of this team’s Bayesian inference algorithm. For the most part, the technique identified catastrophic outliers. However, some light curve pairs still resulted in large contributions to the χ^2 estimator. Identifying this subset of outliers that pass the quality cuts has provided valuable insight into the behavior of this technique, and will allow for future refinement and development to reduce the probability of mis-reconstructions.

6.3. Smoothing and Cross-correlation, by Aghamousa & Shafieloo

Throughout the challenge this team’s main concern was to achieve a high f value without having any outliers. This was achieved with $f > 0.3$ for all five rungs. This conservative approach yielded average χ^2 values of around 0.5–0.9 for different Rungs with P of about 0.03–0.06. As noted before, since χ^2 and P are correlated, by simply dividing all estimated errors by a factor of $\sqrt{2}$, χ^2 of ~ 1 , and P of ~ 0.02 – 0.04 could be achieved trivially. After the true time delays were revealed, a calibration bias of 0.5 days for all the submissions was discovered, resulting in $A \sim 1.8\%$ – 2.5% (the method had been calibrated only on TDC0 Rung 0, owing to lack of time). By adding a calibration correction of 0.5 days to all this team’s submissions’ delay estimates, the bias was removed, improving A to 0.1%–0.6%. To summarize, this method seems promising in both reliability and precision, and is automated in all steps. There is also the potential to improve the error estimation by doing appropriate simulations for each set of light curves separately.

6.4. Supervised Pelt, by Jackson

After the release of the true time delays, this submission was re-examined to try to understand the reasons for the most severe errors, especially those in which the true time delay differed from the inference by $>3\sigma$ (between 9 and 18 cases in each rung out of a few hundred submitted). In four of the worst cases, the problem appeared to be unrealistically low errors fitted during the resampling process, possibly due to a small number of anomalous points, and not corrected by eye. This suggests that for a given set of light-curves, a minimum error based on the fits to the ensemble should be adopted. A significant fraction of the remaining severe errors were characterized by a Pelt statistic versus time delay plot with a relatively bumpy and irregular minimum, even when the eye detected a good fit in terms of the number of coincident points of inflection. This is more difficult to quantify, but suggests that an addition to the resampling-derived error based on the shape of the Pelt statistic may be useful.

6.5. PyCS d3cs, spl and sdi

The d3cs classification of the light-curve pairs into different confidence categories proved valuable. All the resulting “doubtless” (dou) submissions ($f = 0.31$, averaging across all rungs) are free from any catastrophic outliers. As an example, none of the point estimates from the vanilla spl method is farther than $3.7\sigma_i$ or 12.0 days from the truth. For this same method, the less pure doup1a submission ($f = 0.65$) is contaminated by 1.0% of delays that are off by more than 20 days, or, alternatively, $5\sigma_i$. Interestingly, the d3cs estimates for time delays shorter than 50 days are systematically biased low, leading to a significant A of approximately -0.03 for d3cs. We speculate that this bias is perceptual and due to users involuntarily trying to maximize the overlap in the light curves. The sdi and spl techniques were not influenced by this bias in their initial conditions, and both reached a high accuracy, consistent with being unbiased. For these two numerical techniques, the χ^2 metric values are close to unity, suggesting adequate to slightly over-estimated delay uncertainties. The implemented simplifications to the original techniques from Tewes et al. (2013a) seem therefore acceptable for the level of complexity present in the TDC1 data.

6.6. Difference-smoothing, by Rathna Kumar, Stalin, & Prabhu

From the TDC1 feedback, it was realized that this procedure overestimates the uncertainties in the measured time delays, and hence was more prone to reporting catastrophic failures. This problem can be solved by using a Gaussian filter of width equal to the median rather than the mean temporal sampling of the light curves in the process of simulating light curves having known time delays. With this choice, the intrinsic variability in the simulated light curves does not get smoothed out on short timescales. Also, there were a few cases in the submissions where the measured and true time delays were discrepant at the level of $\chi_i^2 > 10$. This points to a need to increase the plausible range of time delays around the measured delay over which the simulated light curves are generated to at least the 3σ confidence interval implied by the inferred uncertainty, rather than the 2σ confidence interval used in the TDC1 submissions. The time delay measurements can be improved further by exploring a range of reasonable values of free parameters, and selecting those which result in the smallest uncertainty in the measured time delay. These changes are now being rigorously tested on the TDC1 light curves and will be described in the paper by Rathna Kumar et al. (2014) during the revision process.

6.7. DeltaTBayes, by Tak, Meng, van Dyk, Siemiginowska, Kashyap, & Mandel

This team considered TDC1 to be a great opportunity to develop and improve their Bayesian approach. Considering the team’s late entry into the challenge, the pragmatic Bayesian perspective was taken (Lee et al. 2011), developing the approximate Gibbs sampling scheme (algorithm1) and applying it only to the most realistic rung (Rung 4). The main advantage of this pragmatic approach was the fast convergence of its Markov chains, saving some computational time, a desirable characteristic for analyzing large number of data sets. The method performs well in terms of precision and accuracy. However it produces error bars that are smaller than those from a fully Bayesian approach, though larger than an empirical Bayesian approach, leading to a relatively high χ^2 . To be balanced, several Gibbs sampling schemes are being tested for the future.

7. SUMMARY AND CONCLUSIONS

In the next decade, dedicated efforts and the LSST survey will deliver thousands of light curves for lensed quasars, ushering in a revolution in time-delay cosmology (Treu et al. 2013). In order to prepare for and make the most of this wealth of data, it is essential to ascertain whether current algorithms are sufficiently accurate, fast, and precise. It is also important to investigate the optimal observing strategies for time delay determination, both in dedicated monitoring campaigns and for LSST.

In order to investigate these two issues, we carried out the first strong lens TDC. After the preliminary time delay challenge TDC0 (Dobler et al. 2014), the challenge “Evil” Team simulated several thousand time delay light curves and made them available to the community on the challenge Web site. Seven “Good” Teams responded to the challenge, and blindly measured the time delays for TDC1 using 9 independent algorithms. A simple method implemented by the “Evil” Team as a baseline was also included. Our main findings from analyzing the blind TDC1 submissions can be summarized as follows.

1. The measurement of time delays from thousands of realistic light curves in manageable amounts of CPU and

investigator time has been demonstrated. This is a considerable achievement given that traditionally this process has been carried out only for very small numbers of light curves (allowing investigators to spend significant amounts of time on each system). Several independent approaches were successful, ranging from cross-correlation, to scatter minimization, to data modeling with GPs and other suitable sets of basis functions. Some methods relied heavily on visual inspection, while others were almost completely automated.

2. In Rung 0—which simulates the typical observing parameters of a dedicated monitoring campaign like COSMOGRAIL—the best current algorithms can recover time delays with negligible bias (often sub-percent) and 3% precision for over 50% of the light curves. The error bars are generally reasonable, resulting in χ^2 of the order of unity, while the fraction of outliers is also just a few percent. These were the requirements for a method to be competitive, as described in Paper I. When enough information was present in the light curves, typically six independent algorithms were able to recover time delays within 10% of the truth.
3. As the data quality was degraded in the subsequent Rungs 1–4 (emulating some observing strategies possible with LSST), the fraction of usable light curves also decreased, hovering between 20% and 30%. Outliers became more common, although they can be contained by suitably conservative algorithms, or by visual inspection. Once outliers are excluded, the algorithms perform as well as in Rung 0, albeit with a smaller fraction of the light curves (10%–30%) yielding robust results with competitive precision and accuracy. A success fraction of 20% translates to an expected sample size of around 400 lensed quasars detected and measured by LSST to very high accuracy—well within the systematic error requirements of time delay cosmography.
4. We have derived approximate scalings for the time delay metrics as a function of observing parameters. Season and campaign length appear to be the dominant terms controlling accuracy (or bias) and success rate, while the precision of the time delay is most strongly related to the cadence and campaign duration.

Much has been learned from this first blind TDC, and the results provide useful guidance and reference for designing future experiments and improving the measurement algorithms. However, it should be emphasized that this challenge was designed to be somewhat simplistic. In particular, TDC1 consisted of a pure time delay estimation challenge from light curves alone: teams were not given the image positions, nor the deflector and source redshifts. It is likely therefore that the results of this challenge might be overly pessimistic. In real life, investigators will have access to the full lensing configurations, and will be able to use this information as a prior for their time delay inference (for example using the lensing geometry for quads).

Furthermore, a fully cosmological challenge should enable outlier rejection based on cosmological self-consistency in a joint analysis of the ensemble of lenses. It should be possible to identify and reject outliers that lead to cosmological parameters (chiefly H_0) that are inconsistent with those inferred from the majority of sample. Another limitation of the simplicity of TDC1 is that the metrics measure how well an algorithm performs on time-delay estimation, not directly on cosmological parameter inference.

Given the encouraging results of TDC1, we plan to overcome these two limitations in the future. In the short term, we plan

to translate the simple metrics adopted here into a full cosmological estimation tool by introducing the available additional information, and justifiable assumptions about the underlying lens models. In the longer term, we plan to organize a second TDC, to further test our ability to handle outliers, and to investigate the measurement of time delays from multi-band data, and in which more information will be provided for each system with the ultimate goal for the “Good” Teams of inferring cosmological parameters, rather than just time delays.

The TDC0 and TDC1 data will remain available at <http://timedelaychallenge.org> for any team who might be interested in using them for developing algorithms for strong lens time delay measurement.

We acknowledge the LSST Dark Energy Science Collaboration for hosting several meetings of the “Evil” Team, and the private code repository used in this work. We thank the referee for constructive criticism which helped improved this paper. T.T., C.D.F., and K.L. acknowledge support from the National Science Foundation collaborative grant “Collaborative Research: Accurate cosmology with strong gravitational lens time delays” (AST-1312329 and AST-1450141). T.T. gratefully acknowledges support by the Packard Foundation through a Packard Research Fellowship. K.L. is supported by China Scholarship Council. The work of P.J.M. was supported by the U.S. Department of Energy under contract number DE-AC02-76SF00515. V.B. and F.C. are supported by the Swiss National Science Foundation (SNSF). M.T. acknowledges support by the DFG grant Hi 1495/2-1. A.A. and A.S. wish to acknowledge support from the Korea Ministry of Education, Science and Technology, Gyeongsangbuk-Do and Pohang City for Independent Junior Research Groups at the Asia Pacific Center for Theoretical Physics. A.S. would like to acknowledge the support of the National Research Foundation of Korea (NRF-2013R1A1A2013795). E.L. is supported by DOE grant DE-SC-0007867 and contract No. DE-AC02-05CH11231. A.H. is supported by an NSERC grant and thanks the Institute for the Early Universe, Korea, for computational resources. A.A., A.S., A.H., and E.L. thank IBS Korea for hospitality. The work of L.A.M. and A.R.W. was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. K.M. is supported at Harvard by NSF grant AST-1211196.

APPENDIX

FACTORS AFFECTING THE RMS MICROLENSING MAGNIFICATION

How sensitive is the distribution of mock light curves to the random realizations of the positions of the stars in the lens? We generated 30 star field realizations, over fields 30 Einstein Radii (R_E) by 30 R_E in area, with different random seeds for each fixed F_* or κ , and calculated the mean of their standard deviations as a characteristic measure of the rms fluctuation in the microlensing magnification. Figure 14 shows how this rms fluctuation varies as a function of F_* . The top panel shows the case where the image arises at the minimum of the time delay surface (where the eigenvalues of the Hessian matrix are both positive and the image has positive parity); the bottom panel shows the case where the image arises at a saddle point of the time delay surface (where the eigenvalues have opposite signs and the parity is flipped compared to the original source). For both figures, significant trends, increasing when F_* is small, and

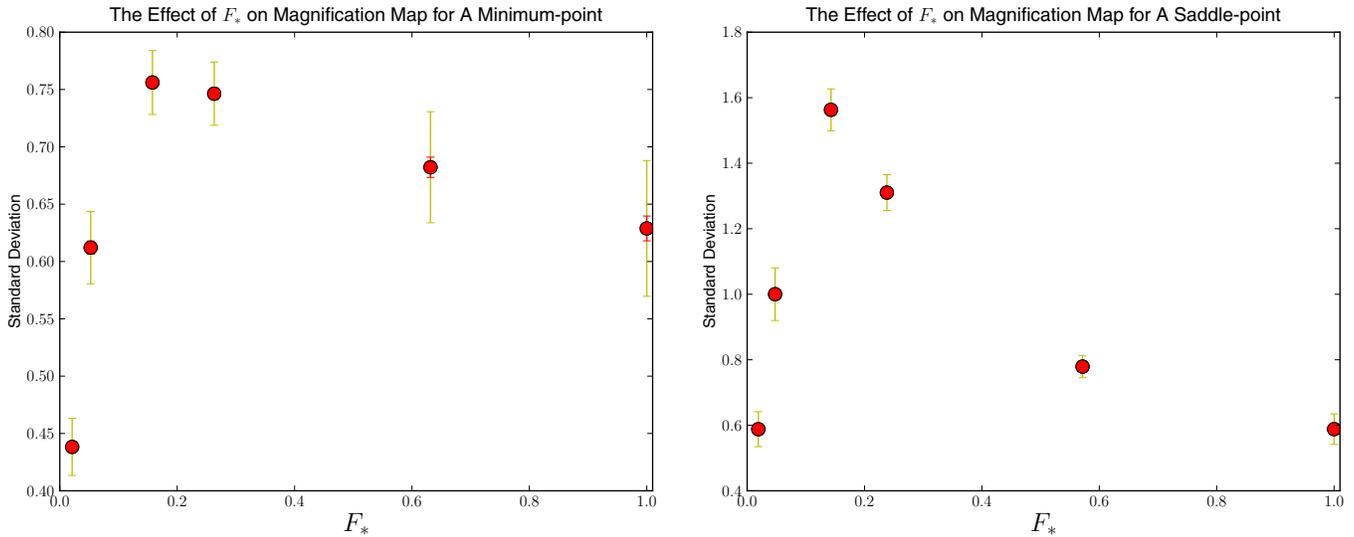


Figure 14. Mean Standard Deviation of the magnification map as a function of F_* . Each point is the result from 30 realizations with different position seeds. We show two errors: Standard Deviation (yellow) and Standard Deviation of the mean value (red). Top figure is for a minimum-image with $\kappa = 0.475$, $\gamma = 0.425$. Bottom figure is for a saddle-image with $\kappa = 0.525$, $\gamma = 0.575$. Both have the same macro magnification $\mu = 10$.

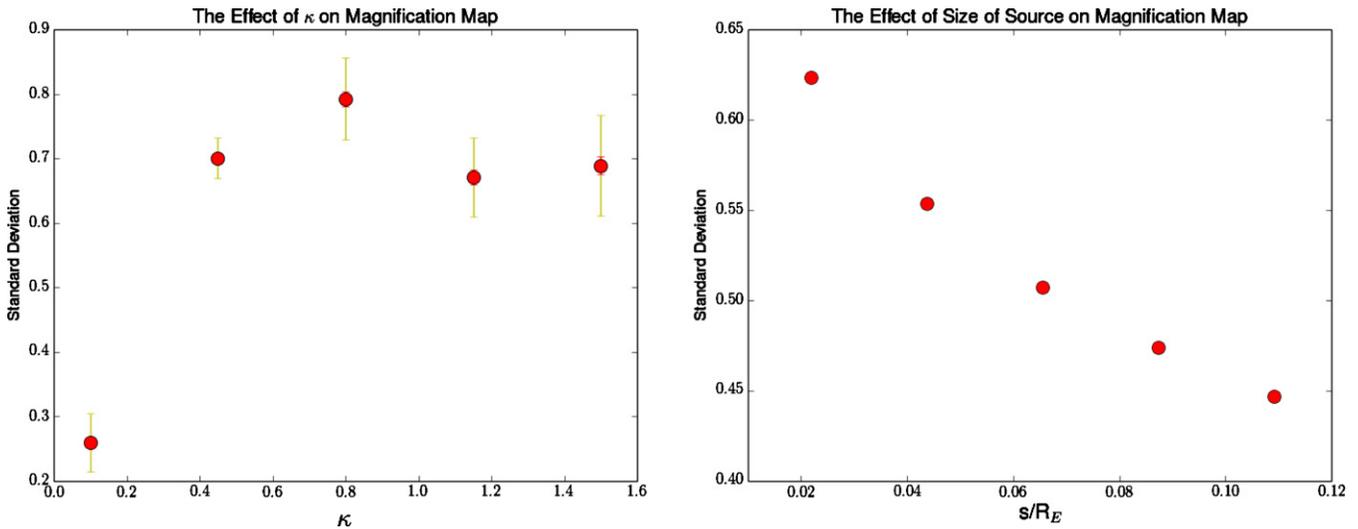


Figure 15. Mean Standard Deviation of the magnification map as a function of local convergence κ (left) and source size s (right). In the left panel, $\kappa = \gamma$ and $F_* = 0.1$ are fixed for each point, while in the right panel $\kappa = \gamma = 0.45$, $F_* = 0.1$ are fixed for each point.

decreasing at larger F_* , are apparent. These can be explained as follows.

At small F_* , when there are few stars, sparsely distributed in the field, the magnification of each position is dominated by the nearest individual star, and the variation of the map increases with more stars that bring more caustics. However, when F_* grows large, the magnification at any position becomes less affected by the addition of more stars, and the magnification and demagnification attributed to different stars will average away. The saddle-point images are more vulnerable to demagnification and hence show larger variations in their magnification maps (see Schechter 2003, for more on the differences of microlensing between minima and saddle-point images).

The left-hand panel of Figure 15 shows the effect of the macrolens convergence κ on the standard deviation of the source plane magnification map. κ affects the variation in two ways, changing the stellar density fraction, and also the macro magnification. These effects appear to approximately balance each other at high κ . At low convergence, the magnification and shear are also low, and the microlensing effects weaker.

Meanwhile, the right-hand panel of Figure 15 shows the rms microlensing magnification fluctuation as a function of source size. As expected, the fluctuations are smoothed out at large source size, reducing the amplitude of the microlensing fluctuations and ensuring that the average microlensing magnification is unity.

REFERENCES

- Aghamousa, A., & Shafieloo, A. 2014, arXiv:1410.8122
 Bostock, M., Ogievetsky, V., & Heer, J. 2011, IEEE Trans. Visualization & Comp. Graphics (Proc. InfoVis), 17, 2301
 Coe, D., & Moustakas, L. A. 2009, *ApJ*, 706, 45
 Dobke, B. M., King, L. J., Fassnacht, C. D., & Auger, M. W. 2009, *MNRAS*, 397, 311
 Dobler, G., Fassnacht, C., Treu, T., et al. 2014, *ApJ*, in press
 Fassnacht, C. D., Xanthopoulos, E., Koopmans, L. V. E., & Rusin, D. 2002, *ApJ*, 581, 823
 Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, 125, 306
 Harva, M., & Raychaudhury, S. 2006, in Proc. 16th IEEE Signal Processing Society Workshop, Machine Learning for Signal Processing, 111

- Hojjati, A., Kim, A. G., & Linder, E. V. 2013, [PhRvD](#), **87**, 123512
- Hojjati, A., & Linder, E. V. 2014, [PhRvD](#), **90**, 123501
- Ivezic, Z., Tyson, J. A., Acosta, E., et al. 2008, arXiv:0805.2366
- Kelly, B. C., Bechtold, J., & Siemiginowska, A. 2009, [ApJ](#), **698**, 895
- Kochanek, C. S. 2002, [ApJ](#), **578**, 25
- Lee, H., Kashyap, V. L., van Dyk, D. A., et al. 2011, [ApJ](#), **731**, 126
- Linder, E. V. 2011, [PhRvD](#), **84**, 123529
- LSST Dark Energy Science Collaboration 2012, arXiv:1211.0310
- LSST Science Collaboration, Abell, P. A., Allison, J., Anderson, S. F., et al. 2009, arXiv:0912.0201
- MacLeod, C. L., Ivezić, Ž., Kochanek, C. S., et al. 2010, [ApJ](#), **721**, 1014
- Oguri, M. 2007, [ApJ](#), **660**, 1
- Oguri, M., & Marshall, P. J. 2010, [MNRAS](#), **405**, 2579
- Paraficz, D., & Hjorth, J. 2009, [A&A](#), **507**, L49
- Paraficz, D., & Hjorth, J. 2010, [ApJ](#), **712**, 1378
- Paraficz, D., Hjorth, J., Burud, I., Jakobsson, P., & Eliasdottir, A. 2006, [A&A](#), **455**, L1
- Pelt, J., Hoff, W., Kayser, R., Refsdal, S., & Schramm, T. 1994, [A&A](#), **286**, 775
- Planck Collaboration, Ade, P. A. R., Aghanim, N., Armitage-Caplan, C., et al. 2014, [A&A](#), **571**, A16
- Poindexter, S., Morgan, N., & Kochanek, C. S. 2008, [ApJ](#), **673**, 34
- Rathna Kumar, S., Stalin, C. S., & Prabhu, T. P. 2014, submitted (arXiv:1404.2920)
- Rathna Kumar, S., Tewes, M., Stalin, C. S., et al. 2013, [A&A](#), **557**, A44
- Refsdal, S. 1964, [MNRAS](#), **128**, 307
- Schechter, P. L. 2003, arXiv:0304480
- Schechter, P. L., Wambsganss, J., & Lewis, G. F. 2004, [ApJ](#), **613**, 77
- Schmidt, K. B., Rix, H.-W., Shields, J. C., et al. 2012, [ApJ](#), **744**, 147
- Sereno, M., & Paraficz, D. 2014, [MNRAS](#), **437**, 600
- Shafieloo, A. 2007, [MNRAS](#), **380**, 1573
- Shafieloo, A. 2012, [JCAP](#), **08**, 002
- Shafieloo, A., Alam, U., Sahni, V., & Starobinsky, A. A. 2006, [MNRAS](#), **366**, 1081
- Shafieloo, A., & Clarkson, C. 2010, [PhRvD](#), **81**, 083537
- Stoughton, C., Lupton, R. H., Bernardi, M., et al. 2002, [AJ](#), **123**, 485
- Suyu, S. H., Auger, M. W., Hilbert, S., et al. 2013, [ApJ](#), **766**, 70
- Suyu, S. H., Treu, T., Blandford, R. D., et al. 2012, arXiv:1202.4459
- Suyu, S. H., Treu, T., Hilbert, S., et al. 2014, [ApJL](#), **788**, L35
- Tewes, M., Courbin, F., & Meylan, G. 2013a, [A&A](#), **553**, A120
- Tewes, M., Courbin, F., Meylan, G., et al. 2013b, [A&A](#), **556**, A22
- Treu, T. 2010, [ARA&A](#), **48**, 87
- Treu, T., Marshall, P. J., Cyr-Racine, F.-Y., et al. 2013, arXiv:1306.1272
- Weinberg, D. H., Mortonson, M. J., Eisenstein, D. J., et al. 2013, [PhR](#), **530**, 87