



Microlensing Discovery and Characterization Efficiency in the Vera C. Rubin Legacy Survey of Space and Time

Natasha S. Abrams¹ , Markus P. G. Hundertmark² , Somayeh Khakpash³ , Rachel A. Street⁴ , R. Lynne Jones⁵ , Jessica R. Lu¹ , Etienne Bachelet⁶ , Yiannis Tsapras² , Marc Moniez⁷ , Tristan Blaineau⁷ , Rosanne Di Stefano⁸ , Martin Makler^{9,10} , Anibal Varela⁹ , and Markus Rabus¹¹

¹ University of California, Berkeley, Astronomy Department, Berkeley, CA 94720, USA

² Zentrum für Astronomie der Universität Heidelberg, Astronomisches Rechen-Institut, Mönchhofstr. 12-14, 69120 Heidelberg, Germany

³ Rutgers University, Department of Physics & Astronomy, 136 Frelinghuysen Road, Piscataway, NJ 08854, USA

⁴ Las Cumbres Observatory (LCOGT), 6740 Cortona Drive, Suite 102, Goleta, CA 93117, USA

⁵ Aerotek and Rubin Observatory, Tucson, AZ 85711, USA

⁶ IPAC, Caltech, Pasadena, CA 91125, USA

⁷ Université Paris-Saclay, CNRS/IN2P3, IJCLab, France

⁸ Harvard-Smithsonian Center for Astrophysics, Cambridge, MA 02138, USA

⁹ International Center for Advanced Studies & Instituto de Ciencias Físicas, ECyT-UNSAM & CONICET, San Martín, Buenos Aires, 1650, Argentina

¹⁰ Centro Brasileiro de Pesquisas Físicas, Rio de Janeiro, RJ, 22290-180, Brazil

¹¹ Departamento de Matemática y Física Aplicadas, Facultad de Ingeniería, Universidad Católica de la Santísima Concepción, Alonso de Rivera 2850, Concepción, Chile

Received 2023 September 26; revised 2024 October 24; accepted 2024 October 30; published 2024 December 24

Abstract

The Vera C. Rubin Legacy Survey of Space and Time will discover thousands of microlensing events across the Milky Way, allowing for the study of populations of exoplanets, stars, and compact objects. We evaluate numerous survey strategies simulated in the Rubin Operation Simulations to assess the discovery and characterization efficiencies of microlensing events. We have implemented three metrics in the Rubin Metric Analysis Framework: a discovery metric and two characterization metrics, where one estimates how well the light curve is covered and the other quantifies how precisely event parameters can be determined. We also assess the characterizability of microlensing parallax, critical for detection of free-floating black hole lenses. We find that, given Rubin’s baseline cadence, the discovery and characterization efficiency will be higher for longer-duration and larger-parallax events. Microlensing discovery efficiency is dominated by the observing footprint, where more time spent looking at regions of high stellar density, including the Galactic bulge, Galactic plane, and Magellanic Clouds, leads to higher discovery and characterization rates. However, if the observations are stretched over too wide an area, including low-priority areas of the Galactic plane with fewer stars and higher extinction, event characterization suffers by >10%. This could impact exoplanet, binary star, and compact object events alike. We find that some rolling strategies (where Rubin focuses on a fraction of the sky in alternating years) in the Galactic bulge can lead to a 15%–20% decrease in microlensing parallax characterization, so rolling strategies should be chosen carefully to minimize losses.

Unified Astronomy Thesaurus concepts: [Gravitational microlensing \(672\)](#); [Galactic bulge \(2041\)](#); [the Milky Way \(1054\)](#); [Sky surveys \(1464\)](#); [Optical astronomy \(1776\)](#); [Optical observation \(1169\)](#); [Time domain astronomy \(2109\)](#)

Materials only available in the [online version of record](#): animation, machine-readable tables

1. Introduction

Microlensing occurs when light coming from a distant star (source) is deflected by a foreground object (lens) located along the observer–source line of sight. As a result, multiple images of the source are formed, and as the images are usually unable to be resolved, the images appear blended. The net observational effect is that the source then appears to be photometrically magnified (B. Paczynski 1986). Since the effect depends on the gravitational influence of the lens and not its luminosity, microlensing is a powerful tool to find and weigh dim objects like cool low-mass stars, planets (e.g., B. S. Gaudi 2012; Y. Tsapras 2018), neutron star candidates, and stellar-mass black hole candidates (e.g., J. R. Lu et al. 2016; C. Y. Lam et al. 2022; P. Mróz et al. 2022;

K. C. Sahu et al. 2022; C. Y. Lam & J. R. Lu 2023) that are otherwise hard to observe.

The microlensing discovery rate increases with the stellar density; therefore, it is highest when observing crowded parts of the sky like the Galactic bulge, Galactic plane, and Large and Small Magellanic Clouds (LMC and SMC). Previous and ongoing dedicated microlensing surveys (e.g., the Optical Gravitational Lensing Experiment, OGLE, A. Udalski et al. 2015; Microlensing Observations in Astrophysics, T. Sumi et al. 2003; KMTNet, D. J. Kim et al. 2018; MACHO, C. Alcock et al. 2000; EROS, M. Moniez et al. 2017) have focused on these areas of high stellar density. All-sky surveys offer the opportunity to explore microlensing throughout the Galaxy. Observing throughout the Galaxy gives us the opportunity to probe Galactic structure (e.g., M. Moniez 2010; M. Moniez et al. 2017) and constrain how the mass function of the lenses, such as black holes, changes throughout the galaxy. Observations of the Magellanic Clouds also offer an opportunity to explore compact halo objects and



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

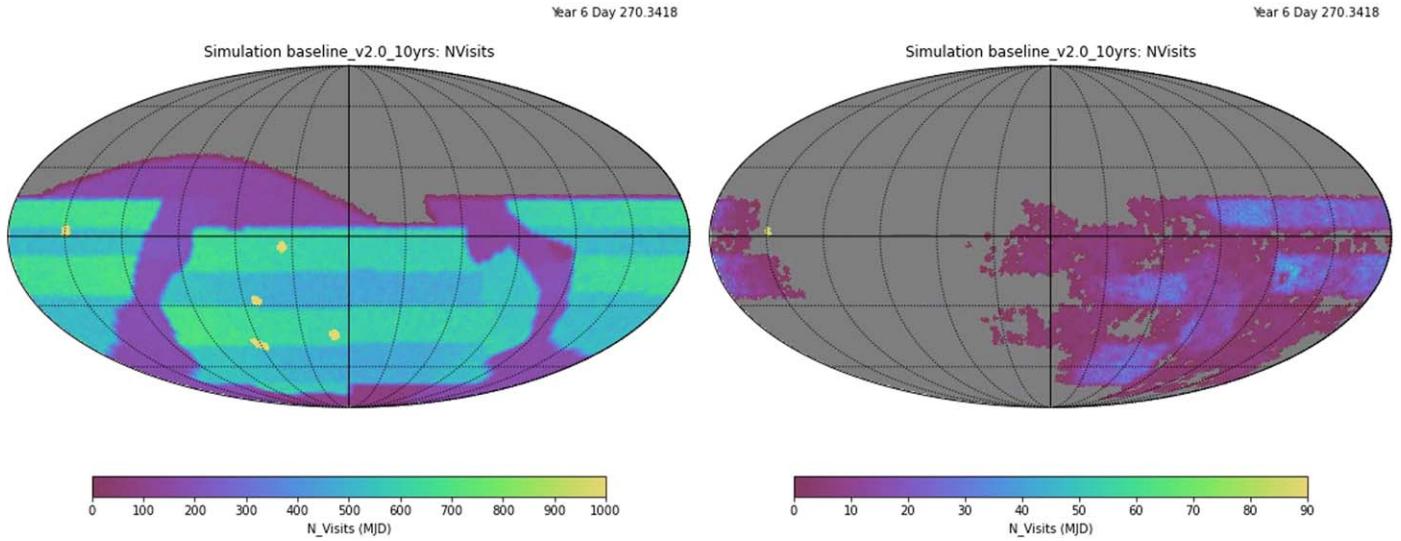


Figure 1. An animation of the `baseline_v2.0_10yrs` OpSim. The left panel shows the observations building up over time, and the right panel shows the observations at a given time as indicated by the date in the upper right corner. The color indicates the number of observations, where purple is fewer observations and yellow is up to 1000 observations. `baseline_v2.0_10yrs` has a rolling cadence in years 2–9. In these years, the sky can be seen split into strips, and alternating strips are focused on by the survey in alternating years.

(An animation of this figure is available in the [online article](#).)

extragalactic stellar remnants, which can then be compared with those of the Milky Way (see T. Blaineau et al. 2021 for a more detailed explanation).

The Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST) will survey 18,000 deg², including parts of the Galactic plane along with the LMC and SMC, as part of its Wide–Fast–Deep (WFD) survey. Surveying at least every 2–3 days is imperative to achieving a high discovery rate of microlensing events (A. Gould 2013; R. Street et al. 2018; S. Sajadian & R. Poleski 2019). The Vera C. Rubin Observatory (Rubin) is going through a community-driven cadence optimization process described in detail in F. B. Bianco et al. (2022). In a community-driven cadence optimization process, it is critical to understand how cadence affects one’s science case. When optimizing, the Survey Cadence Optimization Committee could make many decisions that could render a particular science case impossible, including choosing to exclude a region where most of a particular science case can be done (i.e., excluding the Galactic bulge and plane for microlensing) or choosing to revisit fields on a cadence longer than the timescale of a science case. As this process must balance many science cases with conflicting operational requirements, the community-driven optimization is used to ensure that the survey strategy does not make science cases impossible. In order to evaluate these tensions and to strike a balance between them, statistical metrics have been developed to represent the observational requirements of each science case. The evaluation of these metrics for a range of science goals is described in the collection of papers lead by F. B. Bianco et al. (2022). Here we present an analysis of how potential cadences will affect microlensing.

We determine how cadence will affect microlensing by using a series of hundreds of cadence simulations called Operation Simulations (OpSims), which were created to mock scheduled observations, using the Rubin scheduler (E. Naghib et al. 2019) with the LSST simulation framework (A. J. Connolly et al. 2014). They simulate observations, including their times,

durations, air masses, Moon position, and seeing from mock weather, among other statistics. There are families of simulations that focus on optimizing particular qualities such as the region of sky covered in the WFD (or “footprint”), the frequency of observations (or “cadence”), filter balance (fraction of observations in each filter), and rolling cadence. A rolling cadence is when we divide the sky into multiple sections; in some years Rubin will observe some sections with an increased number of observations, and in other years it will focus its observations on the other sections. An animation of the `baseline_v2.0_10yrs` simulation (which has a rolling cadence in years 2–9) in Figure 1 illustrates how the observations build up over time with a rolling cadence. These are evaluated using the Metric Analysis Framework (MAF; R. L. Jones et al. 2014), which contains metrics both from the Rubin project development team and contributed by the community for particular science cases.

In this work, we have written and tested a multifaceted `MicrolensingMetric` (see Section 2.2) on the set of OpSims from v2.0 to v3.0 that investigates the detection and characterization of microlensing events. We have also tested the effect of changes to footprint and a rolling cadence on the characterization of microlensing events with a parallax signal outside the context of the MAF. The rest of the paper is organized as follows. In Section 1.1 we introduce standard microlensing terminology used throughout. In Section 2 we introduce the metric used to assess the microlensing yields, the sample of microlensing events assessed, and our methodology for determining microlensing parallax characterization. In Section 3 we explore the results of the `MicrolensingMetric` for relevant OpSims and parallax characterization for select OpSims. Finally, in Section 4 we discuss implications and summarize conclusions.

1.1. Microlensing Parameters

We will introduce the standard microlensing parameters that are used to model microlensing events (B. Paczynski 1986).

The characteristic length scale of a microlensing event is known as the angular Einstein radius, which is given by

$$\theta_E = \sqrt{\frac{4GM}{c^2} \left(\frac{1}{D_L} - \frac{1}{D_S} \right)}, \quad (1)$$

where M is the lens mass, D_L is the distance to the lens from the observer, and D_S is the distance to the source from the observer. θ_E and the relative proper motion of the source and lens (μ_{rel}) can be used to define the characteristic timescale, the Einstein crossing time:

$$t_E = \frac{\theta_E}{\mu_{\text{rel}}}. \quad (2)$$

Hence, events with more massive lenses tend to have longer t_E . Neglecting the effects of parallax, the projected separation between the lens and source in units of Einstein radii as a function of time is the impact parameter:

$$u(t) = \sqrt{u_0^2 + \left(\frac{t - t_0}{t_E} \right)^2}, \quad (3)$$

where u_0 is the closest projected separation and t_0 is the time of closest approach. Microlensing events are detected as a temporary photometric magnification. The amplification of a point source–point lens microlensing event is given by

$$A(t) = \frac{u^2 + 2}{u\sqrt{u^2 + 4}}. \quad (4)$$

The amplification is maximized when $u = u_0$. We can use this signal in inferring the lens mass, but a measurement of the lens mass requires either a galactic model or additional observable constraints on D_L , D_S , and μ_{rel} (and/or the lens flux). When the lens is luminous or there are neighboring stars, the light from the source may be blended with them. The total light is characterized by the source flux (F_S) and the blend flux,

$$F_B = F_L + F_N, \quad (5)$$

where F_L is the lens flux and F_N is the neighbor flux. Hence, the flux as a function of time is then

$$F(t) = F_S A(t) + F_B. \quad (6)$$

The blend source flux fraction or blend fraction is $b_{\text{sf}} = F_S / (F_B + F_S)$.

The motion of Earth around the Sun introduces a higher-order effect in microlensing events known as the microlensing parallax (A. Gould 1992), π_E , the magnitude of which is defined as

$$\pi_E = \frac{\pi_{\text{rel}}}{\theta_E}, \quad (7)$$

where π_{rel} is the relative lens–source parallax ($\pi_{\text{rel}} = (1 \text{ au}) \left(\frac{1}{D_L} - \frac{1}{D_S} \right)$). The signal strength depends on the difference between the lens parallax and the source parallax, the direction of the lens–source proper motion, and the time of the year. The changing position of an observer following Earth’s orbit changes the optical axis of the lensing configuration. Microlensing events much shorter than a year tend to exhibit a negligible parallax signal. π_E is a vector that can be broken

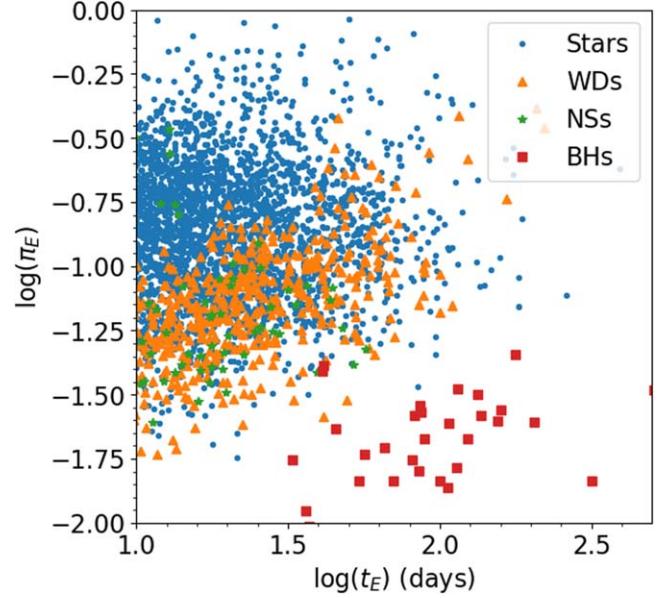


Figure 2. Microlensing parallax (Equation (7)) vs. Einstein crossing time (Equation (2)) for microlensing events simulated using PopSyCLE (C. Y. Lam & J. R. Lu 2020). Blue circles are stars, orange triangles are white dwarfs, green stars are neutron stars, and red squares are black holes. The ranges of t_E and π_E are cropped to match Figures 12–15. As shown in Equation (8), more massive lenses tend to be toward the lower right corner of this plot.

down into east and north components, $\pi_E = [\pi_{E,E}, \pi_{E,N}]$. It is pointed in the direction of μ_{rel} , the relative lens–source proper motion. Along with μ_{rel} , which can have a prior put on it by a galactic model (e.g., J. P. Beaulieu et al. 2006; D. P. Bennett et al. 2014), and t_E , π_E can be used to determine the mass of the lens:

$$M = \frac{\mu_{\text{rel}} t_E}{\kappa \pi_E}, \quad (8)$$

where $\kappa \equiv \frac{4G}{(1 \text{ au})^2} \approx 8.144 \text{ mas } M_{\odot}^{-1}$.

π_E and t_E thus give you a statistical constraint on the mass of the lens. For a physical understanding of where microlensing events lie in the $\pi_E - t_E$ plane, see Figure 2. Here we plot simulated events from Population Synthesis for Compact-object Lensing Events (PopSyCLE; C. Y. Lam & J. R. Lu 2020). Events with less massive, often stellar, lenses tend to have shorter t_E and larger π_E , whereas events with more massive, often compact-object, lenses tend to have longer t_E and smaller π_E .

2. Methodology

In this section we describe OpSims, the MAF framework, and a suite of metrics that are used to assess how the relative number of detected and characterized microlensing events is affected by the survey strategy. This is not expected to produce a realistic yield for Rubin LSST. Instead, it is expected to inform which strategies are beneficial for microlensing, which make it impossible to do microlensing science, and which may negatively affect the microlensing yield without making the science case nonviable. There are three metrics within the MAF framework that are simulated without microlensing parallax and one metric outside the MAF framework simulated with microlensing parallax (see Table 1). The results for the three metrics in the `MicrolensingMetric` are described in

Table 1
Description, Name, and Input Sample of Metrics Used to Analyze Microlensing Efficiency

Metric	Simulated Sample	Simulated Mags	Metric Description
MicrolensingMetric: Discovery Metric	t_E -only (Section 2.3)	1 mean star	2 points on rise with $\geq 3\sigma$ difference
MicrolensingMetric: Npts Metric	t_E -only (Section 2.3)	1 mean star	10 points (with $\text{SNR} > 3\sigma$) within $t_0 \pm t_E$
MicrolensingMetric: Fisher Metric	t_E -only (Section 2.3)	1 mean star	Analytic Fisher with $\frac{\sigma_E}{t_E} < 0.1$
Parallax Characterization Metric	$t_E + \pi_E$ (Section 2.4)	Rand TRILEGAL	Numerical Fisher with $\pi_E > 2\sigma_{\pi_E}$ and $t_E > 2\sigma_{t_E}$

Note. The Simulated Sample column refers to the sample of microlensing events that were evaluated with the metric. The Simulated Mags column refers to the source magnitudes of the simulated microlensing events. For those in the `MicrolensingMetric`, there was a single average star (with different magnitudes in each filter) from the TRILEGAL map used (see Section 2.3). For the parallax characterization, random stars from TRILEGAL were drawn and their magnitudes were used for source magnitudes of the microlensing event (see Section 2.4). The three metrics in the `MicrolensingMetric` are not simulated with parallax and are all-sky. The parallax characterization metric used a sample with parallax for two small patches of the sky.

Sections 3.1–3.6, and the parallax characterization metric results are in Section 3.7.

2.1. OpSims and MAF Framework

The Rubin OpSim team ran the `MicrolensingMetric` on 360+ OpSims. The OpSims evaluated in this paper are summarized in Table 2. See R. L. Jones et al. (2020) and The Rubin Observatory Survey Cadence Optimization Committee (2023) for more detailed descriptions. We also include summary statistics of OpSims discussed in this paper for reference in Table 4.

These OpSims have “families” that vary a particular aspect of the cadence so that they can be tested in isolation and the many aspects of the cadence can be optimized. For example, OpSims in the `vary_gp_` family vary the fraction of the Galactic plane included in the WFD footprint. The first-order OpSims are known as the baseline OpSims. They are marked with version numbers, and the most recent baseline OpSim is considered the nominal current suggested plan for LSST. The baseline OpSims are also the jumping-off point for the other OpSims. So OpSims marked with a v2.1 are changing an aspect relative to the `baseline_v2.1_10yrs` OpSim.

There are a few main quantities that OpSims tend to vary:

1. *Footprint*. This mostly refers to which areas will be included in the WFD and which will be included but at a lower cadence. This also refers to the location of the small deep drilling fields (DDFs), which will be surveyed at a higher cadence.
2. *Cadence*. This is the frequency at which observations are taken. This can be both internight cadence, for example, if triplet observations should be taken in a single night, and intranight cadence, for example, “rolling.”
3. *Filter Balance*. This refers to what fraction of exposures are taken in which filter.
4. *Exposure Length*. This refers to the durations of the exposures that may affect both the science in each individual exposure and the overall mission efficiency.
5. *Additional Surveys*. Rubin may have time to do small additional surveys that are not part of the DDFs or WFD, such as observing near-Earth objects in a twilight survey. These OpSims explore those possibilities.

There are also a number of OpSims that explore technical aspects such as how the performance of the telescope or how different simulated weather will change the survey strategy,

which we do not evaluate here. In this paper we also do not evaluate filter balance or exposure length, due to using a single average magnitude in `MicrolensingMetric` (see Section 2.3).

There are hundreds of metrics that are evaluated for each OpSim in the MAF (R. L. Jones et al. 2014), including general metrics, for example, one that measures the 5σ limiting magnitude in each filter, and specific ones for each science case. These metrics must be lightweight and are generally designed to test one result so that they can be run on all the OpSims and interpreted collectively (LSST Science Collaboration et al. 2017). In this paper we analyze the effect on microlensing, so we use the metrics described in Section 2.2. We also analyze the ability to characterize events with microlensing parallax outside of the context of the MAF due to computation speed (Section 2.4) on a few key OpSims (Section 3.7).

2.2. MicrolensingMetric

The `MicrolensingMetric` is integrated with the `rubin_sim`¹² package and provides a set of metrics for evaluating the efficacy of cadences in detecting, alerting, and characterizing microlensing events. The metric relies on a simulated population of microlensing events described in detail in Section 2.3. We can calculate the discovery efficiency using the Discovery Metric in default mode by computing the fraction of events with at least two points on the rise of the light curve with at least a 3σ difference between the highest and lowest magnitude points. Events that meet this criterion are considered “discovered”/“detected” (the terms are used interchangeably). We can also specify a number of days before the peak time that the event must be “triggered” by, since ensuring the observation of the peak in follow-up is important for light-curve characterization. (Note that while this functionality is in the metric and is important for event follow-up, including exoplanet and black hole candidates, an exploration of alerting efficiency is beyond the scope of this paper). The Discovery Metric can also compute the fraction of events with at least two points on the rise and fall of the light curve with at least a 3σ difference in magnitude. This serves a similar function to the Npts and Fisher Metrics described below, so we do not explore the results of that metric in this paper. The default mode of the Discovery Metric is based on a common first step in detecting microlensing events (e.g., P. Mróz et al. 2017). These typically

¹² https://github.com/lsst/rubin_sim

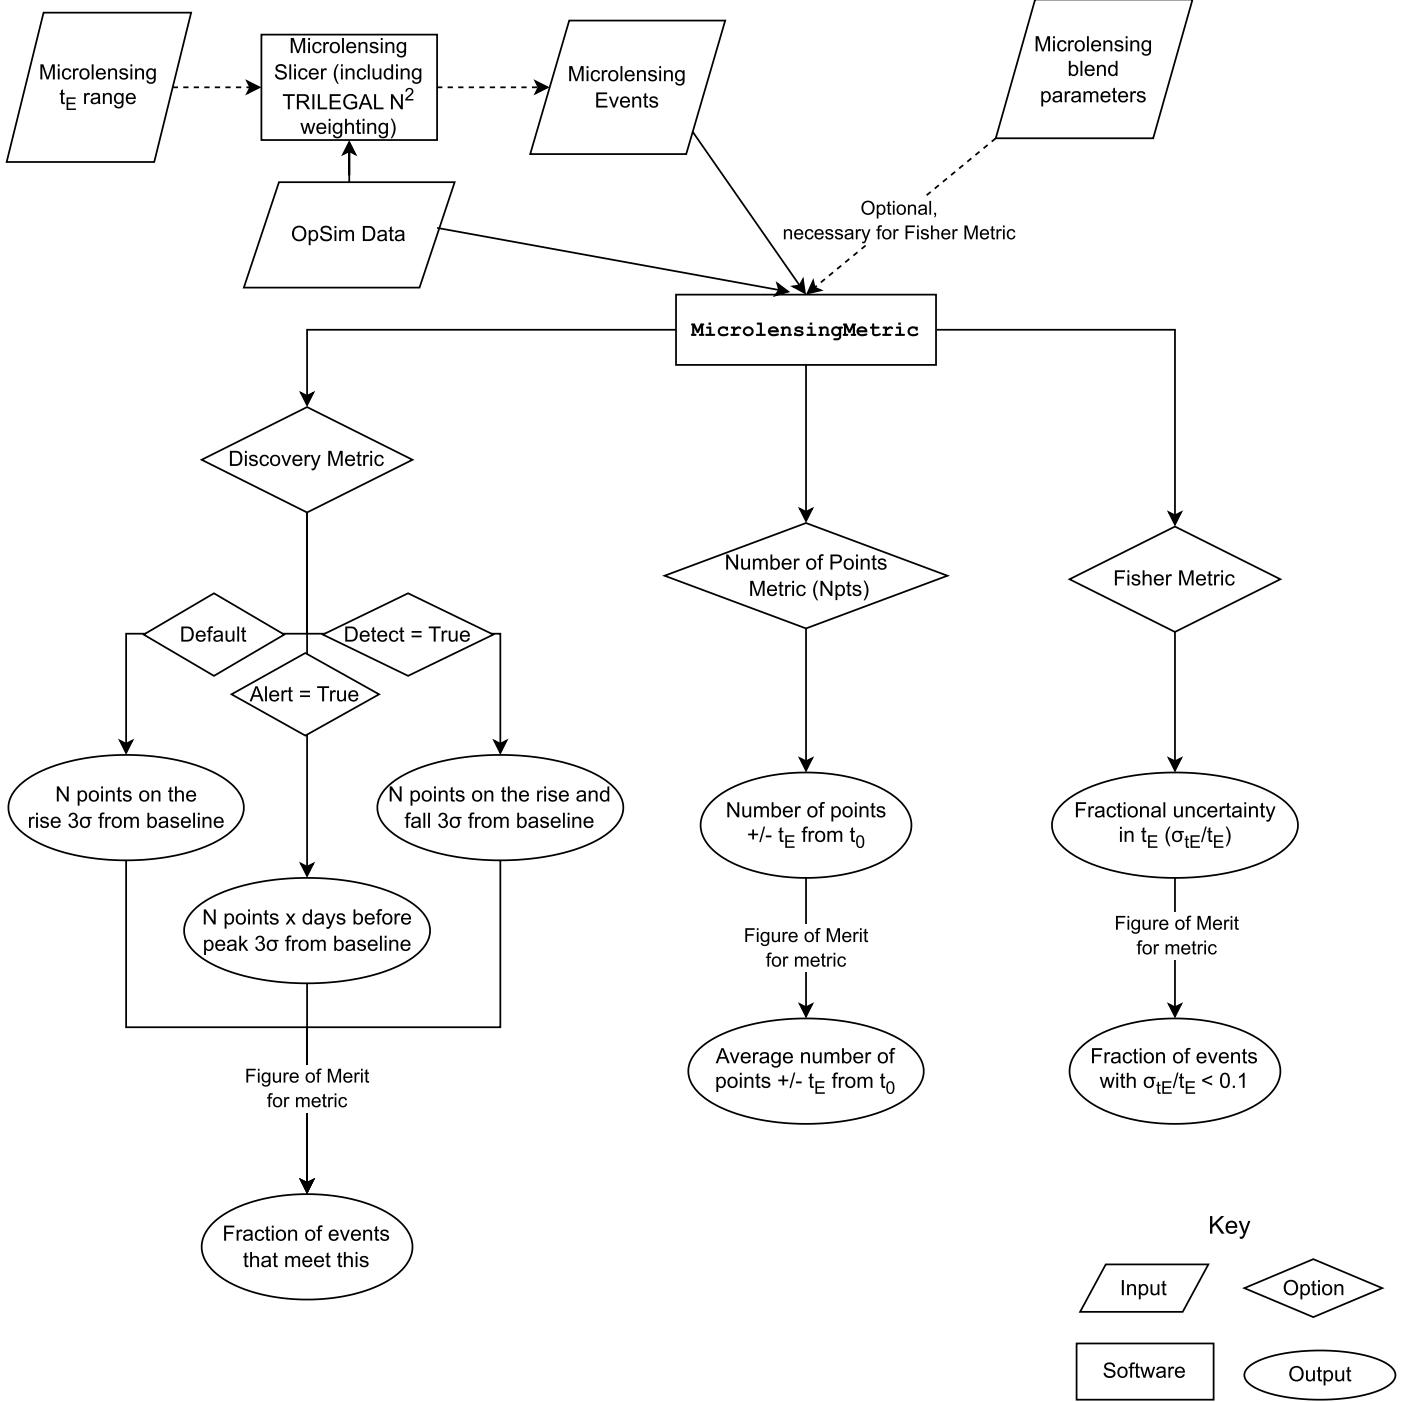


Figure 3. Flowchart of the `MicrolensingMetric` functionality. Either the user can input a t_E range and generate microlensing events with a function known as a “Slicer,” or the user can input the events. The user can then choose to use the Discovery Metric, Npts Metric, or Fisher Metric. The Discovery Metric, by default, finds the fraction of events with N points on the rising side with at least a 3σ difference from baseline. The user can select `detect = True` to find at least N points on both the rising and falling side or `alert = True` to require the N points on the rising side to be a certain number of days before the peak. In this paper, we only explore the default mode of the Discovery Metric. The Npts Metric returns the number of points within $t_0 \pm t_E$ with $\text{SNR} > 3\sigma$, and the Fisher Metric returns the fraction of events with a fractional t_E uncertainty < 0.1 .

use three points above baseline, but we found minimal difference in relative OpSim performance between requiring two and three points.

There are also two metrics for the characterization of the light curves purely from Rubin observations. There is a basic metric that quantifies the number of points (Npts) within $t_0 \pm t_E$ that estimates the coverage of the light curve. We require that

the signal-to-noise ratios (SNRs) of the points are above a 3σ threshold. The SNR is calculated using the `rubin_sim` function `m52snr()`, which requires the magnitude and 5σ depth and takes into account a number of effects, including weather and instrumental effects. The Npts Metric was used as a proxy for characterization, and figure of merit is the fraction of events with at least 10 points (with $\text{SNR} > 3\sigma$) within $t_0 \pm t_E$.

Requiring 10 points as a proxy is rooted in inspection of microlensing light curves (e.g., from inspecting Zwicky Transient Facility (ZTF) light curves, it was determined that 10 points were necessary to establish a detection; M. S. Medford et al. 2023). There is also a metric that calculates the Fisher matrix for each event and returns the fractional uncertainty in t_E (see Section 2.2.1). See Figure 3 for a summary of the `MicrolensingMetric` functionality.

When discussing the results of these metrics, we may roughly refer to regions in the sky that are affected differently by OpSims. For reference, the Galaxy is $|b| < 10^\circ$, the Galactic bulge is $|l| < 20^\circ$, and the Galactic plane are the l values outside that (though some of this is not observable by Rubin).

2.2.1. Fisher Matrix

To characterize events, we want to ensure an adequate photometric cadence during the information-bearing part of the light curve. This is essential for determining microlensing parameters, such as t_E , that allow us to infer physical properties of the lens and source populations. We adapt the Fisher matrix approach, which is widely used in many fields, such as cosmology (e.g., G. Jungman et al. 1996; A. Albrecht et al. 2006). To quantify how precisely parameters can be recovered, we use a fiducial model, namely assuming that our simulated parameters correspond to the actual parameter estimates. According to the Cramer–Rao inequality, the Fisher matrix allows us to calculate a lower bound on the uncertainty.

In the MAF, since we have simulated the microlensing event population, we can use their known parameters to evaluate the Fisher matrix

$$I_{i,j} = \sum_{k=1}^{N_{\text{data}}} \frac{1}{\sigma_F^2(t_k)} \left(\frac{\partial F_{\text{model}}(t_k)}{\partial p_i} \right) \left(\frac{\partial F_{\text{model}}(t_k)}{\partial p_j} \right), \quad (9)$$

where p_i and p_j denote the event parameters, t_E , t_0 , u_0 , and the blend and baseline flux parameters for each passband; t_k is each time of observation; N_{data} is the length of the entire data set (including all passbands); F is the flux; and σ_F is the error in the flux. Assuming Gaussian errors on each observable, the Fisher information matrix is approximately the inverse of the covariance matrix. One element of this matrix will be the uncertainty in t_E , so we can calculate the fractional uncertainty σ_{t_E}/t_E . We treat an event as well characterized if $\sigma_{t_E}/t_E < 10\%$. This threshold indicates whether we can constrain the lens mass or its error budget is dominated by the unknown t_E . An advantage of using a Fisher information matrix is that it can account for the contribution of both the light-curve coverage and the uncertainty of the observations, $\sigma_F^2(t_k)$, to the uncertainties of the model parameters.

We can evaluate the Fisher matrix by taking analytic derivatives with respect to each of the parameters (t_E , t_0 , u_0 , F_S , and F_B) using `Sympy` (A. Meurer et al. 2017). Speeding up the calculation of the `MicrolensingMetric` is key for evaluating many OpSims, and Equation (9) is the best-suited approximation since it only relies on first derivatives. To optimize the evaluation of the analytic Fisher Metric, we use the common subexpression elimination part of `Sympy` jointly finding suitable substitutions for the parameters to reduce the

number of coefficients. We then use the analytic equations that `Sympy` had optimized to carry out the calculation.

2.3. Sample of Microlensing Events

In order to cover the phase space of microlensing events in a heuristic way, we simulate across the whole sky in `HEALPIX`¹³ and weight the probability of having a microlensing event at a particular R.A. and decl. with the number of stars squared in the TRILEGAL stellar map simulated for LSST (P. Dal Tio et al. 2022). The number of events should scale with the square of the visible density, which traces the square of all compact objects. While the purpose of this paper is to compare between OpSims, if one were to determine the number of detectable or characterizable microlensing events, one could use a compact-object tracer and source-star tracer (e.g., R. Poleski 2016) or carry out a full mock-microlensing survey using a Galactic model (e.g., C. Y. Lam & J. R. Lu 2020). A full simulation is beyond the scope of this paper but will be the subject of future work (see Section 4.1).

For each metric, we split the population into t_E bins and generated 10,000 events for each t_E bin. The events were simulated with uniform distributions of t_E from minimum to maximum t_E in the bin and t_0 from the minimum to maximum observation date in the given OpSim. For the discovery metric we make eight t_E bins: 1–5, 5–10, 10–20, 20–30, 30–60, 100–200, 200–500, and 500–1000 days. For the Npts and Fisher Metrics we analyze a subset of these bins for computational efficiency (10–20, 20–30, 30–60, and 200–500 days). We break the events up into t_E bins so that we can analyze different populations of objects separately. Since $t_E \propto \frac{\sqrt{M_{\text{lens}}}}{\mu_{\text{rel}} \sqrt{\pi_{\text{rel}}}}$, various t_E values are related to different populations of objects, from low-mass stars and free floating planets with short t_E to black holes with long t_E . We also break the events into t_E bins because longer- t_E events are more likely to be detected and characterized than shorter- t_E events (see Figure 5) and a rolling cadence (which is where Rubin focuses on a fraction of the sky in alternating years) could leave large parts of long- t_E events unobserved (see Section 3.6). Hence, if we evaluated events with 1 days $< t_E < 1000$ days together, then we would not see the strength of the effect of changing the cadence. The events were also simulated with uniform distribution of u_0 from 0 to 1. While Rubin will likely detect events with lower magnifications than this, as OGLE and other surveys do, this is an approximation to compare between OpSims.

To simulate a typical star as the source of our microlensing events, we found the mean stellar magnitude in each filter of the TRILEGAL map and used that as the source magnitude for all the events. These values are as follows: $u = 25.2$, $g = 25.0$, $r = 24.5$, $i = 23.4$, $z = 22.8$, $y = 22.5$. Note these are close to the detection limit of Rubin. Using a brighter star does not change the overall effects of cadence in most cases. However, in the case of exposure time, using a faint star can lead to an effect where it is detected in longer exposures but not in shorter ones, leading to longer exposures being preferred. On the other hand, if it had been brighter, shorter exposures mean that there would be more time for more data points, so shorter exposures are preferred. Given this limitation in this paper, we do not include an assessment of exposure time. In general, brighter

¹³ A division of the spherical sky into surface area pixels (see <https://healpix.sourceforge.io/>).

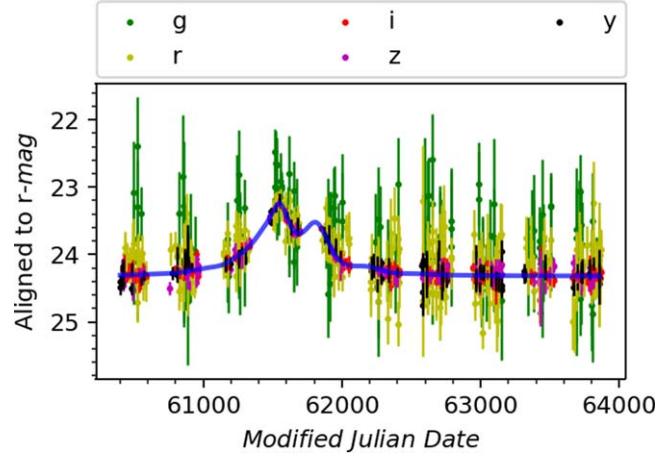


Figure 4. Example plot of a characterizable event with $t_E = 161$ days, $\pi_E = 0.21$, $u_0 = 1.5$, and $\phi = 0.79$. It was simulated in a Galactic disk patch (R.A. = 288.34° , decl. = 9.66°) with `baseline_v3.0_10yrs`. Details of the simulation are described in Section 2.4. The parallax is characterized with a relative uncertainty of $\sigma_{\pi_E}/\pi_E = 0.06$, and the Einstein time is characterized with $\sigma_{t_E}/t_E = 0.06$. As is standard for achromatic microlensing events, all light curves are aligned and rescaled to the best data set. In this case, the r band serves as the reference baseline.

stars are less strongly affected by the cadence, as they are over the detection threshold a higher fraction of the time. In addition, using a single set of assumed magnitudes leads to the limitation that we cannot effectively assess filter balance since it would only be affected by the single magnitudes used.

For the Fisher matrix calculation (see Section 2.2.1), it is important to take blending into account. A high blend fraction can increase the uncertainty of parameters, since for blended events the apparent baseline becomes brighter and the blend fraction, t_E , and u_0 are degenerate (e.g., J. C. Yee et al. 2012). We estimated that in the locations of high stellar density where most of these events occur the blend fraction is $\sim 50\%$, i.e., the flux from neighboring stars + the lens (F_B) is approximately equal to the source flux (F_S) (see Figure 3 of Y. Tsapras et al. 2016). This blend fraction of $\sim 50\%$ is used in our Fisher matrix calculation.

2.4. Parallax Characterization

The methods for evaluating the effect of the cadence on microlensing in Section 2.2 did not include the parallax effect. Characterizing the microlensing parallax is important for inferring the mass and nature of the lenses (see Equation (8); see Figure 13 of C. Y. Lam & J. R. Lu 2020). Hence, on a subset of OpSims we simulated 100,000 events in a representative bulge and disk field including parallax to determine how the characterization of light curves including parallax is affected by cadence. We are particularly interested in how a rolling cadence affects the characterization of microlensing parallax, as this is a periodic effect in long-enough events.

We determined how well we could characterize each event by taking numerical derivatives of each of the parameters (t_0 , t_E , u_0 , $\pi_{E,E}$, $\pi_{E,N}$, blend parameter, and source magnitude) and applied Equation (9) to determine the Fisher information matrix. Numerical derivatives were computed by simulating models where the parameters differed by a tolerance of $(0.01 \times$ value of the parameter) and the slope was calculated. The errors on the magnitude of each observation were determined using the `calc_mag_error_m5()` in `rubin_sim`. The light curves were modeled using Bayesian Analysis of Gravitational

Lensing Events (BAGLE),¹⁴ an open-source microlensing event modeling and fitting code.

We simulated two patches of 100,000 events, one at R.A. = 263.89° , decl. = -27.16° ($l = 0.33^\circ$, $b = 2.82^\circ$) and the other at R.A. = 288.34° , decl. = 9.66° ($l = 44.02^\circ$, $b = -0.42^\circ$). The first is a representative bulge field and the second is a representative disk field in one of the pencil beam fields described in R. A. Street et al. (2023b). We used the observations in a square field of view 3.5° across to mimic a Rubin field of view. The events were simulated with uniform distributions of u_0 from -1 to 1 , $\log(t_E)$ from 0.70 to 2.78 (i.e., 5–600 days), t_0 from the minimum to maximum observation date in the given OpSim, $\log(\pi_E)$ from -2 to 0 (i.e., 0.01 – 1), and ϕ from 0 to 2π . Here ϕ determines the direction of the parallax by

$$\pi_{E,N} = \pi_E \sin(\phi) \quad (10)$$

$$\pi_{E,E} = \pi_E \cos(\phi). \quad (11)$$

Note that, in reality, ϕ is not uniform and is dependent on Galactic location owing to the spatial and velocity distributions of the Galaxy. A variation of characterization depending on ϕ changes on the $\sim 15\%$ level, but this is constant across all evaluated OpSims, so it does not affect our relative comparisons that are done on individual Galactic patches.

Instead of a single average source magnitude, as was used for the metrics in the `MicrolensingMetric`, here we used a distribution of magnitudes. This is because brighter stars enable events to be more easily characterized, so we are better able to explore the entire parameter space where t_E and π_E may be characterized. We selected 10^6 random stars in each representative patch with an r -mag brighter than 24.03 (the single image 5σ r -mag depth¹⁵) in the TRILEGAL LSST simulation (P. Dal Tio et al. 2022). We then randomly matched a star to a microlensing event and kept the star’s color.

The criteria for characterization of a light curve are both $t_E > 2\sigma_{t_E}$ and $\pi_E > 2\sigma_{\pi_E}$. The criteria on the characterization of each parameter here are less strict than the set in MAF

¹⁴ https://github.com/MovingUniverseLab/BAGLE_Microlensing/tree/main

¹⁵ <https://www.lsst.org/scientists/keynumbers>

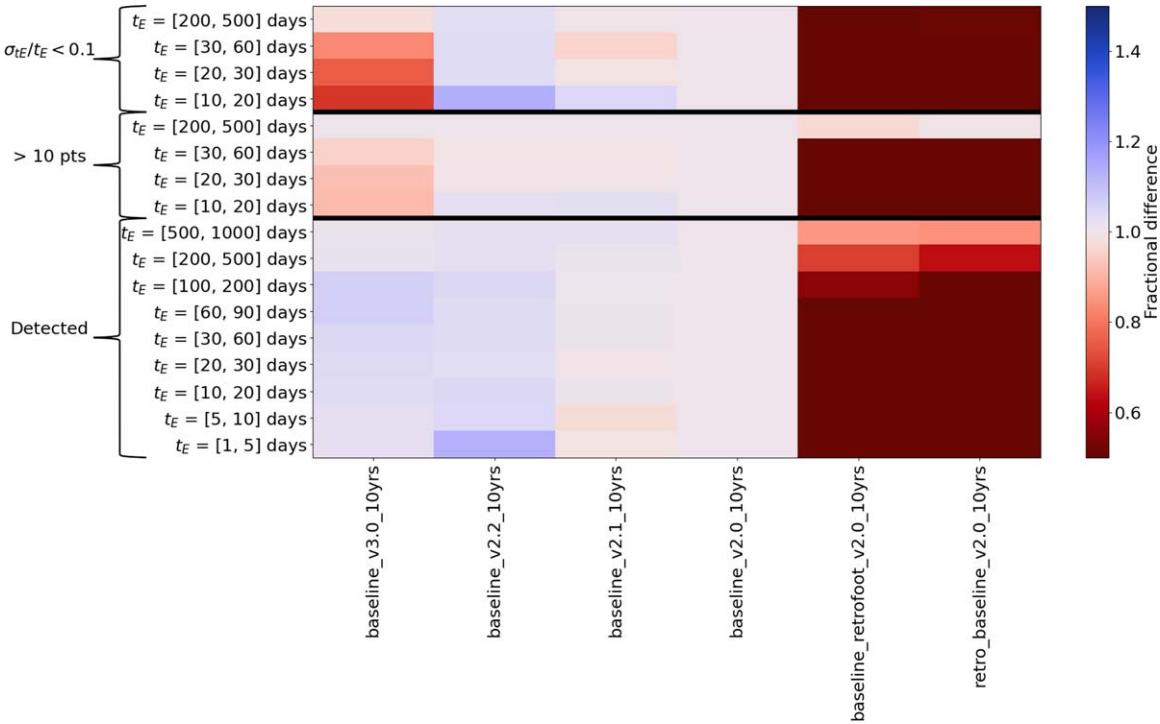


Figure 5. Comparison between `MicrolensingMetric` results for several distinct iterations of the Rubin baseline strategy. The colors show the fractional improvement relative to `baseline_v2.0_10yrs`, where blue means that the metric has performed better relative to the baseline metric and red means that it has performed worse. On the x-axis are each of the OpSims, where `10yrs` means that it was simulated for a 10 yr LSST survey. On the y-axis are each of the metrics for a t_E range compared to their values for `baseline_v2.0_10yrs`. Shorter events (smaller t_E) correspond to less massive objects such as brown dwarfs and low-mass stars, whereas longer events (larger t_E) can correspond to black holes. The y-axis groups each set of metrics, and each line is a different t_E range. The black horizontal lines indicate the separate metrics. “ $\sigma_{t_E}/t_E < 0.1$ ” refers to the Fisher Metric, “ > 10 pts” refers to the Npts Metric, and “Detected” refers to the Discovery Metric. For example, “ $\sigma_{t_E}/t_E < 0.1$ $t_E = [30, 60]$ days” refers to the fraction of the 10,000 events with 30 days $\leq t_E \leq 60$ days simulated as described in Section 2.3 with $\frac{\sigma_{t_E}}{t_E} < 0.1$ as calculated by the Fisher information matrix. All of the current baselines show a $\gtrsim 50\%$ improvement over the retro baseline.

$(\sigma_{t_E}/t_E < 0.1 \rightarrow t_E > 10\sigma_{t_E})$ since we are characterizing two parameters instead of just one. An example characterized event can be seen in Figure 4. These results are outside the context of the `MicrolensingMetric` and explored in Section 3.7.

3. Results

We will describe the results of the `MicrolensingMetric` for each of the families of OpSims. See Table 2 for descriptions of the OpSims discussed here and their relevance to microlensing and Galactic science. In this simulation, the Discovery Metric, Npts Metric, and Fisher Metric were all run. The Discovery Metric was configured such that two points 3σ above baseline were required on the rising side of the light curve. When we refer to the discovery efficiency, this refers to the fraction of simulated events that meet the Discovery Metric criteria, and when we refer to characterization efficiency, this refers to the fraction of simulated events that meet the Npts or Fisher Metric criteria. The retro baseline is described in detail in R. L. Jones et al. (2020), and the v2.0–v3.0 are described in detail in The Rubin Observatory Survey Cadence Optimization Committee (2023); where `baseline_v3.0_10yrs` is the same as `draft2_rw0.9_uz_v2.99_10yrs`.

In general, the larger the footprint dedicated to the Galactic bulge and plane, the more microlensing events we can see and characterize. There is also a trend where longer-duration microlensing events are less affected by observing strategy since they last long enough that most strategies eventually accumulate enough observations; however, the exact cadence of observations is still important, especially for characterization

of t_E and π_E . We quantify how well the metrics perform by comparing their performance to the `baseline_v2.0_10yrs` OpSim. The results of the sample without parallax analyzed with the `MicrolensingMetric` are in Sections 3.1–3.6, and the parallax characterization metric results are in Section 3.7. A few OpSim families that include additional surveys unrelated to microlensing or do not have a strong effect on microlensing were analyzed with the `MicrolensingMetric` and are included in Appendix A. We do not include OpSims that explicitly vary the filter balance or exposure length, due to using a single set of average stellar magnitudes (see Section 2.3). All results from the `MicrolensingMetric` are summarized in Table 5 (Discovery Metric) and Table 6 (Npts and Fisher Metrics).

3.1. Baseline Family

The baseline strategies show the general evolution over time of the survey, and the rest of the OpSims are variations on these baselines. In most of this paper, we compare OpSims to `baseline_v2.0_10yrs`, which includes the Galactic bulge and parts of the Galactic plane at a WFD cadence. It does not perform rolling (see Section 3.6 for a detailed definition) in the Galaxy and does perform rolling in the rest of the sky in years 2–9 of the survey. Most estimates of the number of microlensing events Rubin is expected to discover are based on the `baseline_2018a` OpSim (i.e., S. Sajadian & R. Poleski 2019), which does not include the Galactic bulge or denser parts of the Galactic plane. Along with updating the survey strategy, the software to simulate survey strategies has changed, so instead of

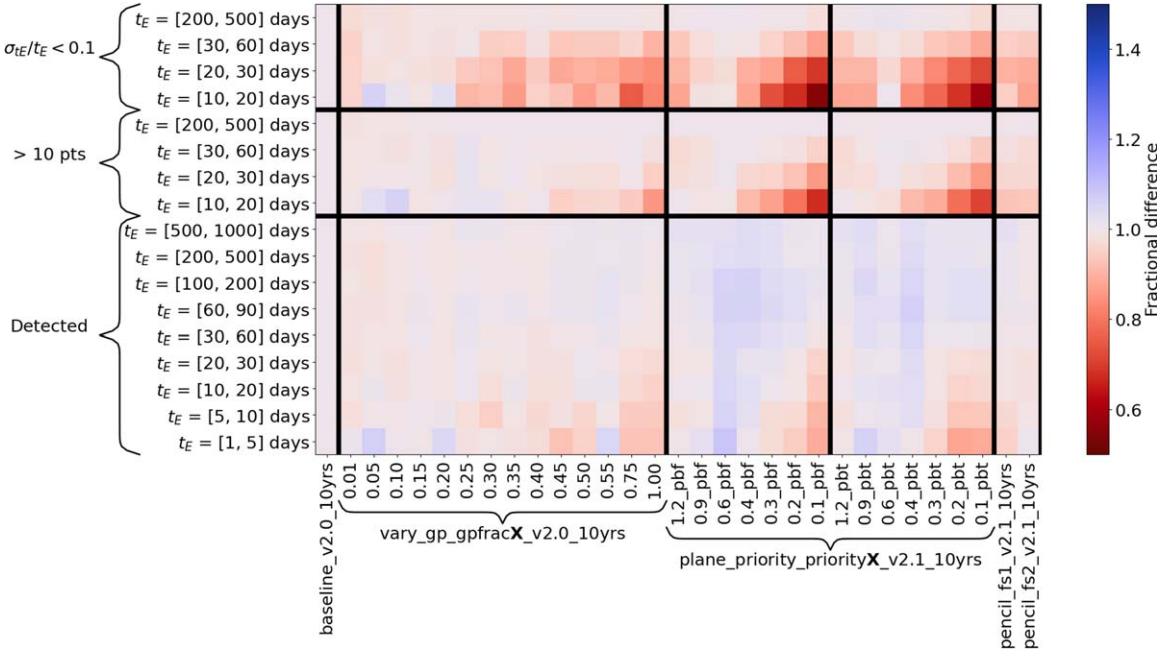


Figure 6. Same as Figure 5, but for the `vary_gp`, `plane_priority`, `pencil`, and `galplane_priority` families. The colors show the fractional improvement relative to `baseline_v2.0_10yrs`. Black vertical lines indicate separate OpSim families (summarized in Table 2), and black horizontal lines indicate separate metrics. Covering regions of lower stellar density and/or high extinction (areas of priority less than 0.4 as defined in R. A. Street et al. 2023b) leads to significant decrease in microlensing characterization efficiency. The `pbf` strategies do not include pencil beams selected in R. A. Street et al. (2023b) of scientific interest, whereas the `pbt` strategies include them. While the number of detected and characterized microlensing events does not significantly differ between them, since the pencil beams were specifically chosen to optimize our ability to probe the Galaxy, the strategy is preferable. The size of the pencil beams does not appear to significantly affect microlensing efficiency.

using `baseline_2018a` we use `retro_baseline_v2.0_10yrs`, which uses a similar observation strategy to `baseline_2018a`, but implemented with updated software and updated throughput and weather inputs.

More recent versions of the baseline (v2.0 and higher) of the OpSims lead to a $\gtrsim 50\%$ improvement in both discovery and characterization over the `retro` baseline due to the inclusion of the Galactic bulge and parts of the Galactic plane. `baseline_retrofoot_v2.0_10yrs` adopts the old footprint but uses the v2.0 baseline strategy, whereas `retro_baseline_v2.0_10yrs` is a version of the retro footprint and strategy. See Figure 5 for a comparison of the metrics for the baseline OpSims. v2.0 performs slightly better than v2.1, since v2.1 includes the Virgo Cluster, which is not a traditional microlensing target and takes time away from other areas. v2.2 included optimizations to the code and a change in DDF strategy that should not significantly affect microlensing. `baseline_v3.0_10yrs` spends less time on the Galactic bulge and spreads out observations across the plane (see R. A. Street et al. 2023b for detailed discussion). Covering this larger area leads to $\sim 10\% - 20\%$ fewer events being characterized (Fisher and Npts Metrics) but $\sim 5\% - 10\%$ more being discovered (Discovery Metric), since there are fewer events in the Galactic plane than in the Galactic bulge, due to the decrease in stellar density. We also note that the t_E distribution changes as a function of position. Due to changes in relative proper motion, t_E is on average longer in the Galactic plane than the bulge (P. Mróz et al. 2020; M. S. Medford et al. 2023). We include Table 3 with results broken down by general location for `baseline_v2.0_10yrs` for reference. A strategy similar to this could allow Rubin to better probe

galactic structure but may require increased follow-up to characterize the discovered microlensing events.

3.2. Galactic Bulge and Plane Coverage and Footprint

There are a number of families of OpSims that explore the Galactic plane coverage and footprint. The `vary_gp` family varies the visits to fields in the Galactic plane as a percentage of the WFD survey from `gpfrac` = 1% to 100%. We see a significant decrease of microlensing characterization in strategies with `gpfrac` ≥ 0.55 . We see that technically if we cover the Galactic plane more we characterize fewer microlensing events overall, since many of the microlensing events are concentrated toward the Galactic bulge (see Figures 6 and 7). However, it is scientifically interesting to be able to probe microlensing events throughout the Galactic plane.

Besides time spent on the Galactic plane, we can also probe the optimal Galactic bulge and plane footprint. We used a map of the Galactic plane with scientifically motivated priorities assigned to each region (R. A. Street et al. 2023b) ranging between 0.1 and 1.2. Generally regions of lower stellar density and/or high extinction correspond to lower-priority regions. The `plane_priority_priority` OpSim family adds regions of progressively less priority to the WFD footprint, so `plane_priority_priority0.4` includes regions assigned priority ≥ 0.4 . We find that there is a drop in characterization efficiency for the long-duration events in the plane priority map when it covers regions of priority of 0.4 or lower (see Figure 6), as a finite number of visits are distributed over too large a spacing in time for characterization (see Figures 6, 8, and 9). This matches what is found by the general Galactic plane metrics (R. A. Street et al. 2023b).

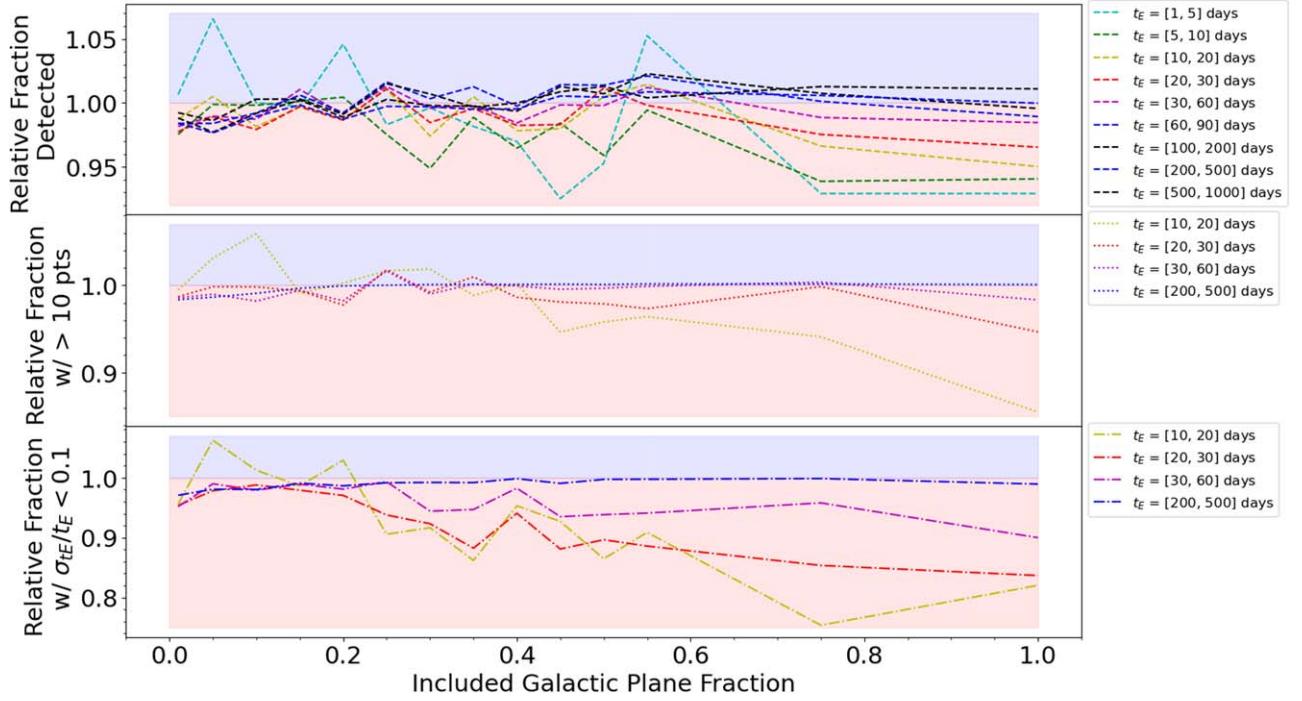


Figure 7. Plot of fraction of events detected (top), fraction of events with at least 10 points (middle), and fraction of events with $\sigma_{tE}/tE < 0.1$ (bottom), relative to baseline_v2.0_10yrs as a function of the fraction included of the Galactic plane. These correspond to the results from the vary_gp family. The colors of the lines correspond to the timescale of the events, and the line style corresponds to the type of metric (dashed for Discovery, dotted for Npts, and dotted-dashed for Fisher). If the line is in the blue shaded region, this indicates that the OpSim performed better than baseline_v2.0_10yrs, and if it is in the red shaded region, it performed worse. In general, shorter- t_E events are more affected by the change in cadence including random fluctuations. There is a significant decrease of microlensing characterization above gpfrac ≥ 0.55 since most microlensing events are in the Galactic bulge.

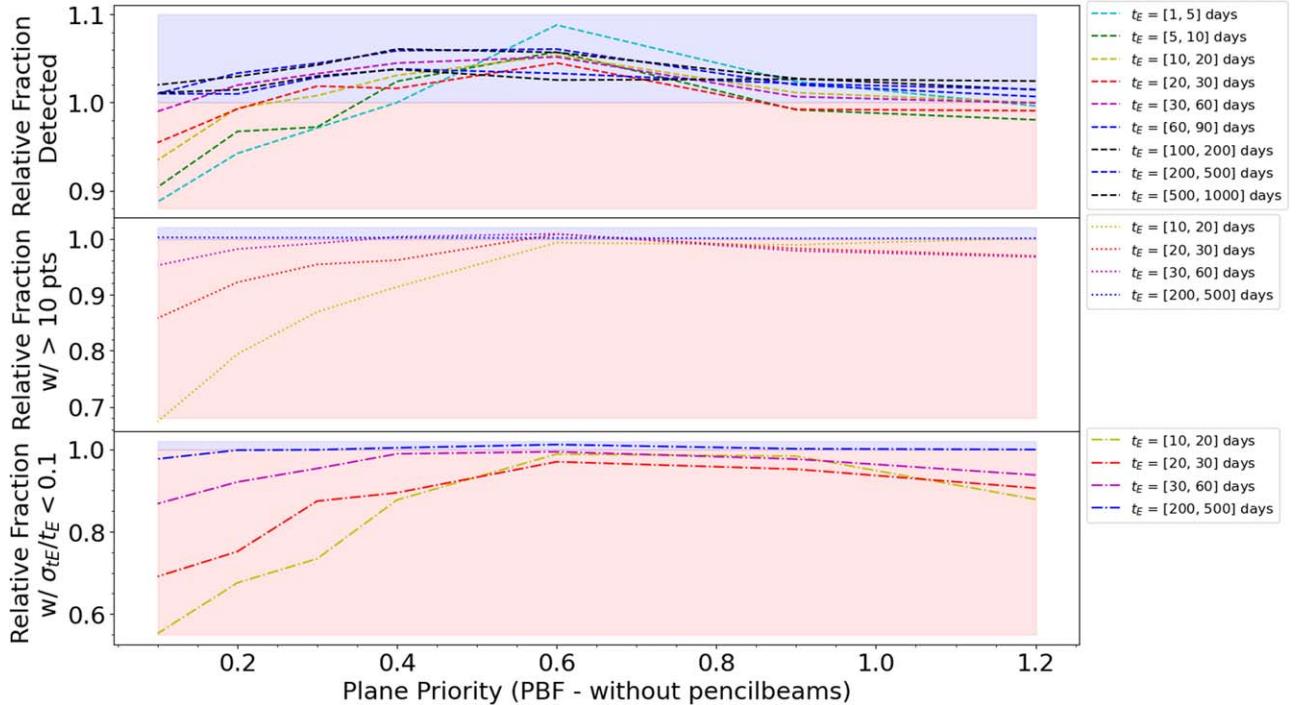


Figure 8. Same as Figure 7, but for each metric as a function of plane priority without pencil beams (pbf = pencil beams false). If the line is in the blue shaded region, this indicates that the OpSim performed better than baseline_v2.0_10yrs, and if it is in the red shaded region, it performed worse. As plane priority decreases, this means that areas of the Galaxy of lower priority are covered at a WFD-level cadence. The priorities are scientifically motivated regions of interest as defined in R. A. Street et al. (2023b). There is a peak around 0.6, but there is a significant drop when areas of priority 0.4 and lower are covered. This is because a finite number of visits are distributed over too large an area.

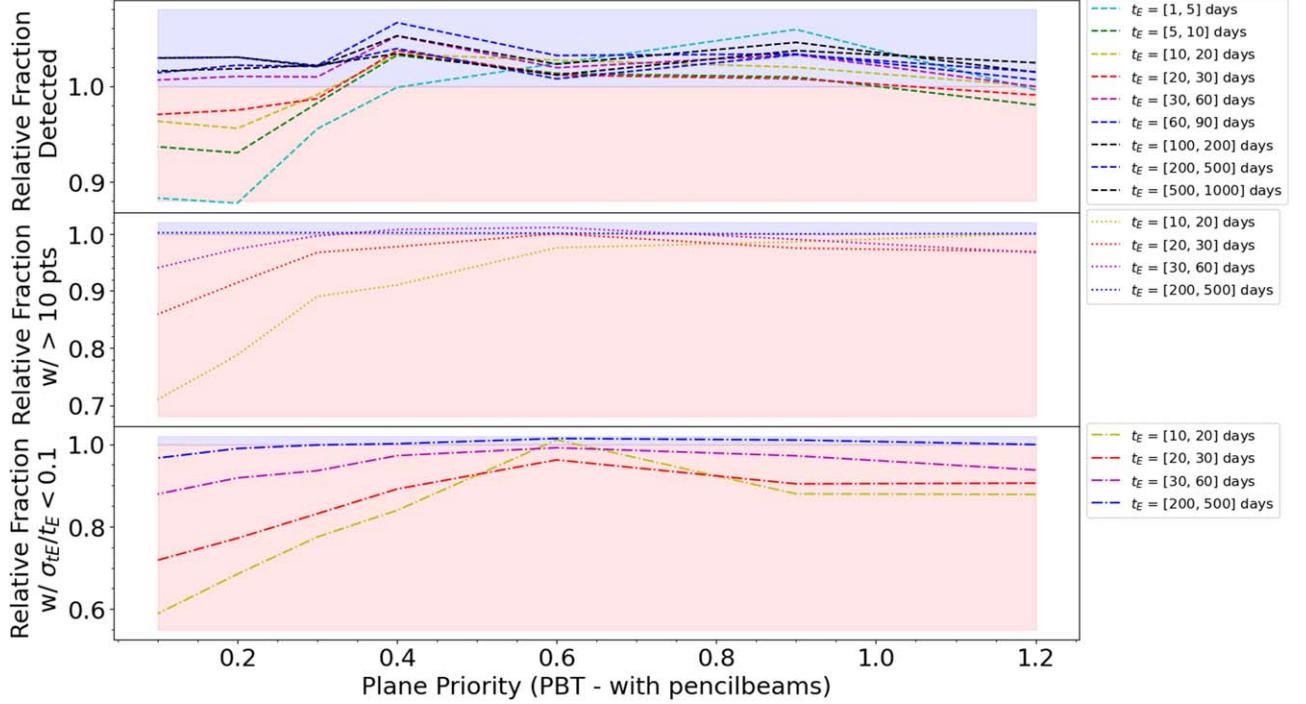


Figure 9. Same as Figure 8, but for each metric as a function of plane priority with pencil beams (pbt = pencil beams true). If the line is in the blue shaded region, this indicates that the OpSim performed better than `baseline_v2.0_10yrs`, and if it is in the red shaded region, it performed worse. Comparatively, the results are similar to those without pencil beams, with a significant drop when areas with priority 0.4 and below are included. Though a similar number of events are detected and characterized, the pencil beam fields were specifically chosen in R. A. Street et al. (2023b) to probe Galactic structure (and other scientific goals) throughout the Galactic plane.

In addition to areas ranked by their priority, R. A. Street et al. (2023b) define a series of designated “pencil beam” fields selected for their scientific interest. The `plane_priority_priorityX.X_pbff`_OpSims do not include those pencil beams (pencil beams false), and `plane_priority_priorityX.X_pbt`_OpSims include them (pencil beams true). We technically find similar results to the OpSims without pencil beams since we detect and characterize a similar number of events (see Figures 6, 8, and 9). However, the pencil beam fields were picked specifically to optimize our ability to probe Galactic structure (along with other goals) throughout the Galactic plane. Decades of microlensing surveys have looked at the Galactic bulge, with some surveys delving into the Galactic plane (e.g., Gaia, Ł. Wyrzykowski et al. 2023; ZTF, A. C. Rodriguez et al. 2022; M. S. Medford et al. 2023; R. Zhai et al. 2023; and OGLE, P. Mróz et al. 2020), but Rubin will enable us to look much deeper across a larger area of the Galactic plane, so looking in strategic spots is helpful. The `pencil_` OpSim family varies the size of the pencil beam fields. As can be seen in Figure 6, the size of the pencil beams does not appear to affect the microlensing results.

3.3. Image Quality

The `good_seeing_` OpSim family adds the requirement of at least three good seeing images per year per pointing. By adding a requirement, here we mean that a requirement is added in the scheduler. So when the scheduler decides what part of the sky to look at, it will attempt to ensure that at least three good seeing images per year per pointing are obtained. This is balanced against other requirements such as maintaining a particular cadence. As the good seeing metric is prioritized,

detection and characterization metrics worsen on the 10% level relative to `baseline_v2.1_10yrs` for characterization since it appears as though the footprint decreases and we end up with fewer events (see Figure 17). However, better template images for difference image analysis could improve alerts and photometric accuracy, but there are insufficient data to assess a suitable trade-off.

3.4. Triplet Observations

The `presto_gapX`, `presto_half_gapX`, and `long_gaps_nightsoffX` sets of families explore “triplets” of observations described in detail in F. B. Bianco et al. (2019). This means that there will be a third visit on the same night (in the case of the `presto` family; see Figure 10), a third visit on the same night after X hours (in the case of the `presto_gapX` family), every other night (in the case of the `presto_half_gapX` family), or a third visit every X nights (in the case of the `long_gaps_nightsoffX` family). Microlensing events decrease in discovery and characterization by 10%–30% in the `presto` family. In general, microlensing events do not change sufficiently in a single night to warrant a third visit that night, and taking time away from looking at more varied points in time greatly decreases the efficacy of microlensing detection and characterization. In some strategies, there is an improvement in discovery of events with $t_E = 1$ –10 days, which do change at this timescale but at the large expense of the majority of events ($t_E > 10$ days), and of their characterization. There could also be an improvement to events with short-duration features such as microlensing events with a binary lens, though events with such features are a small fraction of events.

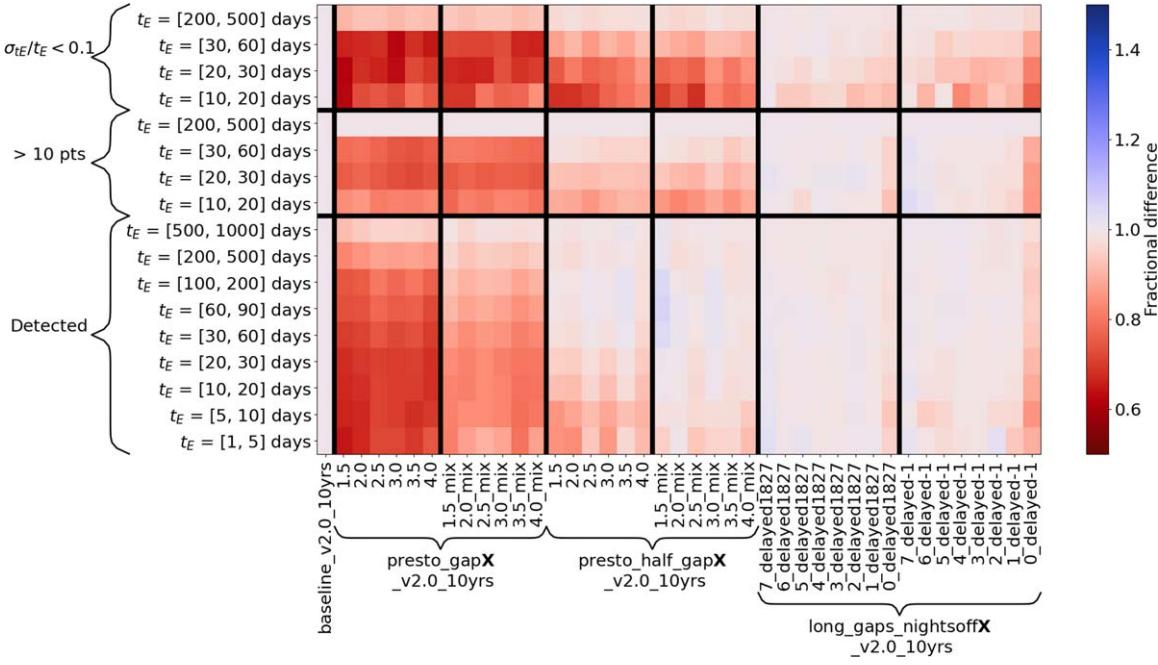


Figure 10. Same as Figure 5, but for the presto family, which explores triplets of observations within one night. The colors show the fractional improvement relative to baseline_v2.0_10yrs. Black vertical lines indicate separate OpSim families (summarized in Table 2), and black horizontal lines indicate separate metrics. Since most microlensing events do not vary significantly over the course of a night, when observations are taken over a less varied time period, this decreases detection and characterization efficiency.

3.5. Microsurveys

Microsurveys are “micro” observing surveys that take up to a few percent of the LSST observing time (explored in detail in Z. Ivezic 2022). The two microsurveys of relevance for microlensing are roman_v2.0_10yrs and smc_movie_v2.0_10yrs. Since the rest of the surveys do not focus on microlensing targets, they only negatively impact microlensing on the 5%–10% level since it takes time away from microlensing targets. See Figure 16 for a summary.

The roman_v2.0_10yrs microsurvey is designed to look at the footprint of the Nancy Grace Roman Galactic Bulge Time Domain Survey (GBTDS; D. Spergel et al. 2015; M. T. Penny et al. 2019). Observing the Roman field during Roman’s survey “seasons” of \sim 60–72 days and also filling in the multimonth gaps between its observations would be impactful. During Roman’s observing windows, concentrating more of the Galactic bulge observations on the Roman field could allow for simultaneous observations that could be used to measure satellite parallax that can be used to constrain the mass of the lens (J. C. Yee et al. 2014). The number of increased visits to the Roman field should not be at a level such that visits to the rest of the Galactic bulge are significantly reduced, but perhaps \sim rw0.5 rolling (i.e., the scheduler allocating 50% more observations to the Roman field than if no rolling were performed) since that did not seem to significantly negatively impact detection and characterization of events >30 days (see Section 3.6). The nominal GBTDS is planned to have seasons of 60–72 days with multiple months-long gaps. While some of these gaps are at times when Rubin cannot observe the Galactic bulge, there are some where Rubin could fill in Roman gaps. Filling in the photometry is particularly beneficial for characterizing long-duration events that span multiple Roman seasons (C. Y. Lam et al. 2023). The impact of a lack of space-based astrometry during those times, though, is still to be determined. Work is in progress to further quantify the synergy

between Roman and Rubin; see R. A. Street et al. (2023a) for more details.

The smc_movie_v2.0_10yrs includes two nights of high-intensity observations of the SMC. Though the smc_movie_v2.0_10yrs survey decreases the characterization fraction of short-duration events by 5%–10%, the SMC is a target of scientific interest for microlensing for compact halo objects and to probe galactic structure in a nearby dwarf galaxy.

3.6. Rolling Cadence

In OpSims with a rolling cadence, the sky is broken up into two to six regions (indicated by the number following ns) and observations are focused on one region each year, alternating between years. These regions are not necessarily continuous. The strength of the rolling indicates the aggressiveness of reshuffling visits into active or inactive rolling seasons, with higher numbers corresponding to pushing more visits into active rolling seasons. For example, rw0.8 or strength0.8 means 80% rolling, so if there were 100 visits in a nonrolling season, the scheduler would try to have 180 visits in an active rolling season and 20 visits in an inactive rolling season. However, since there are other scheduling requirements (i.e., coverage of the WFD footprint with observations with good seeing), it tends to perform rolling less than the programmed strength. There is already a rolling cadence incorporated in the baseline, though the rolling, roll_, and six_rolling OpSim families explore different configurations of strips and the effect of having a rolling cadence in the Galactic bulge. As can be seen in Figures 16 and 17, the rolling cadence generally does increase discovery of short events (<10 days) by \sim 5%, especially OpSims with a rolling cadence in the Galactic bulge. However, events with $t_E > 10$ days have a decrease in detection and characterization.

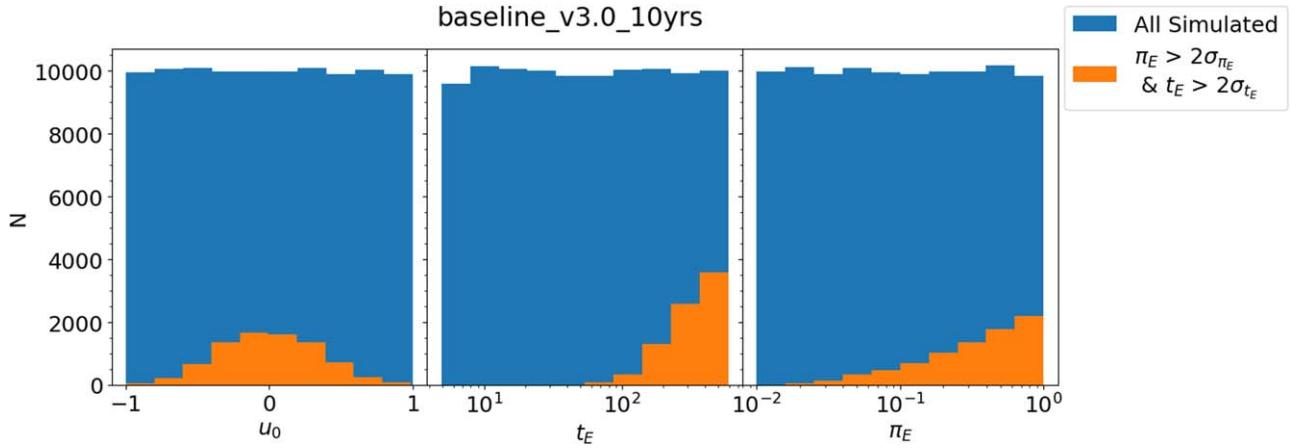


Figure 11. Distribution of u_0 , t_E , and π_E for simulated and characterized events. In blue are all of the simulated events, and in orange are those that are characterized ($\pi_E > 2\sigma_{\pi_E}$ and $t_E > 2\sigma_{t_E}$) by the `baseline_v3.0_10yrs` OpSim in the bulge field R.A. = 263.89° , decl. = -27.16° .

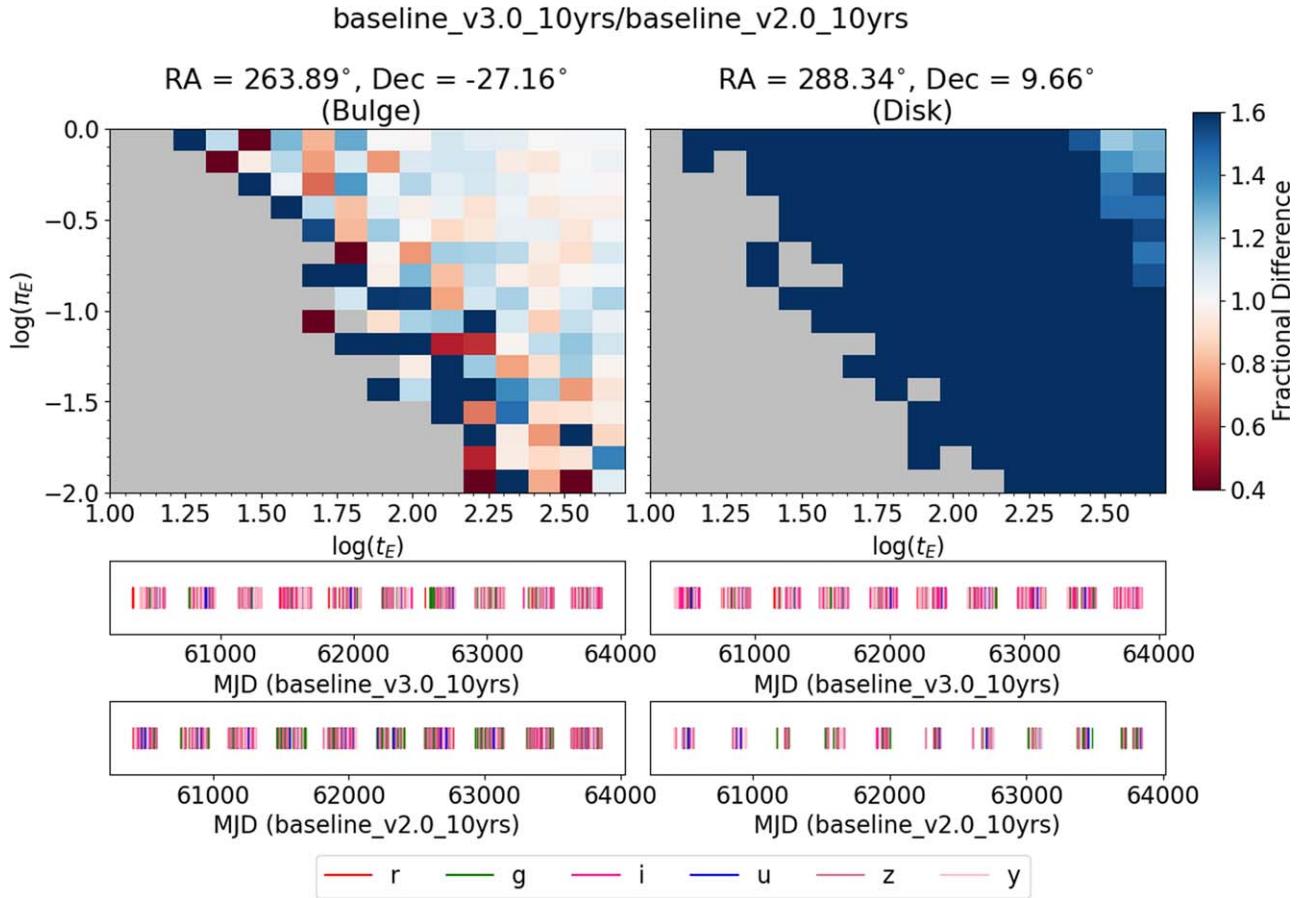


Figure 12. 2D histogram comparing the fraction of events that are characterized ($\pi_E > 2\sigma_{\pi_E}$ and $t_E > 2\sigma_{t_E}$) in `baseline_v3.0_10yrs` and `baseline_v2.0_10yrs` (`baseline_v3.0_10yrs/baseline_v2.0_10yrs`) for any given t_E and π_E . Bluer squares mean that more events were characterized in `baseline_v3.0_10yrs`, and redder squares mean that more were characterized in `baseline_v2.0_10yrs`. Gray squares mean that no events were characterized in either OpSim, in large part due to π_E being intrinsically hard to measure with a single telescope for short-duration events (see Section 3.7). The bottom panel show the cadence of observations, where different colors represent different filters, as indicated in the legend, for the given OpSim in a $3.5^\circ \times 3.5^\circ$ square centered on the R.A. and decl. indicated in the title of the left and right panels. The middle row shows the cadence of the OpSim in the numerator (`baseline_v3.0_10yrs`), and the bottom row shows the cadence of the OpSim in the denominator (`baseline_v2.0_10yrs`). The left panel is a representative bulge field, and the right panel is a representative disk field in a pencil beam. See Figure 2 for a realistic distribution of simulated events in $\pi_E - t_E$ space. Since the `baseline_v2.0_10yrs` footprint included Galactic bulge and neglected the Galactic plane, the fraction of events characterized in the disk field increased by ~40%–50% in `baseline_v3.0_10yrs` and the fraction of events characterized in the Galactic bulge stayed about the same.

In strategies in which the Galactic bulge and plane are explicitly included in areas with a rolling cadence, detection efficiency decreases by 5%–15% and characterization efficiency decreases by 10%–40%. Most microlensing events have

$t_E > 10$ days, and it was expected that Rubin would send an alert on shorter-duration events, not that it would be able to completely characterize them. Beyond cadences that explicitly include the Galactic bulge and plane in a rolling cadence, even

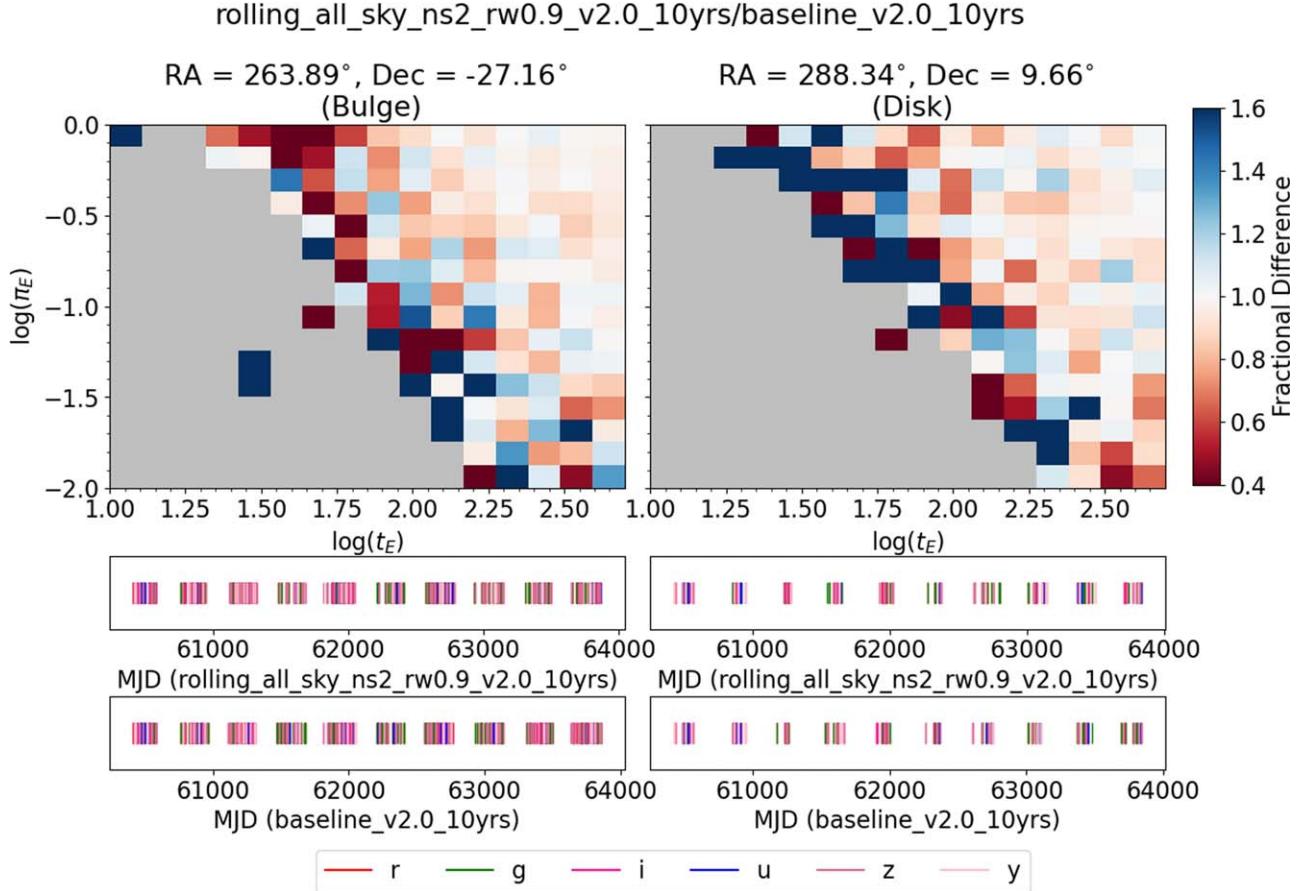


Figure 13. Same as Figure 12, but comparing `rolling_all_sky_ns2_rw0.9_v2.0_10yrs` to `baseline_v2.0_10yrs`. This OpSim has the same footprint as `baseline_v2.0_10yrs` but with a rolling cadence. While this does not explicitly include the Galactic plane as part of the footprint with a rolling cadence, the fraction of events characterized in the bulge is $\sim 10\%$ less than `baseline_v2.0_10yrs`.

if a region is not explicitly part of the rolling cadence, due to slew times and survey efficiency, those regions may also be affected. There is a decrease in characterization efficiency by 10%–20% in `six_rolling_ns2_rw0.9_v2.0_10yrs` (see Figure 16), even though no region with significant microlensing population is part of the footprint with a rolling cadence. Note that we have not included microlensing parallax here; see Section 3.7.

3.7. Parallax Characterization

In this subsection we explore the characterizability of events with a microlensing parallax signal as described in Section 2.4 ($t_E > 2\sigma_{t_E}$ and $\pi_E > 2\sigma_{\pi_E}$). In Figure 11, we show a histogram of simulated (blue) and characterized (orange) event parameters. This is for the `baseline_v3.0_10yrs` in a field around R.A. = $263^\circ.89$, decl = $-27^\circ.16$. We see that those with smaller $|u_0|$ are characterized more easily owing to their higher magnification. Longer t_E are characterized more frequently since events with longer t_E are more likely to be covered eventually. Large π_E are characterized more frequently owing to the larger measurable signal.

In Figures 12–15, we plot 2D histograms of $\log(\pi_E)$ versus $\log(t_E)$ dividing the fraction of events characterized by the OpSim in the numerator by the fraction of events characterized by `baseline_v2.0_10yrs`. This compares each OpSim relative to `baseline_v2.0_10yrs` such that the redder squares indicate that `baseline_v2.0_10yrs` characterized a higher

fraction of events and bluer squares indicate that the compared OpSim characterized a higher fraction of events. In the middle and bottom panels, we plot the cadence of observations for that OpSim and field, where lines of different colors represent observations taken with different filters, as indicated by the legend. The left panel in each figure is for a representative bulge field, and the right panel in each figure is for a representative disk field. Most events fall to the left side of these plots; see Figure 2 for a realistic distribution of simulated events in $\pi_E - t_E$ space. The gray squares indicate that neither OpSim characterized any events there. In the low- t_E parameter space, there are some events with characterizable t_E values, in particular bright events. However, these are not characterizable here because we demand that both t_E and π_E are characterizable and it can be intrinsically difficult to have a measurable π_E signature in short-duration events. When an event is short compared to the duration of a year, it is not usually possible to measure the π_E signal from a single telescope (e.g., S. Poindexter et al. 2005; B. S. Gaudi 2012). Instead, one can use satellite parallax, in which a single event is measured simultaneously from an Earth-based observatory and faraway space-based satellite (S. Refsdal 1966; A. Gould 1994). This may be a possible synergy between Roman and Rubin, although most likely only for short- t_E events owing to the relatively short baseline between the telescopes (e.g., J. C. Yee 2013; R. A. Street et al. 2023a).

In Figure 12, we compare `baseline_v2.0_10yrs` and `baseline_v3.0_10yrs`. We can see in the disk that there is an improvement in characterization by 40%–50%. This is

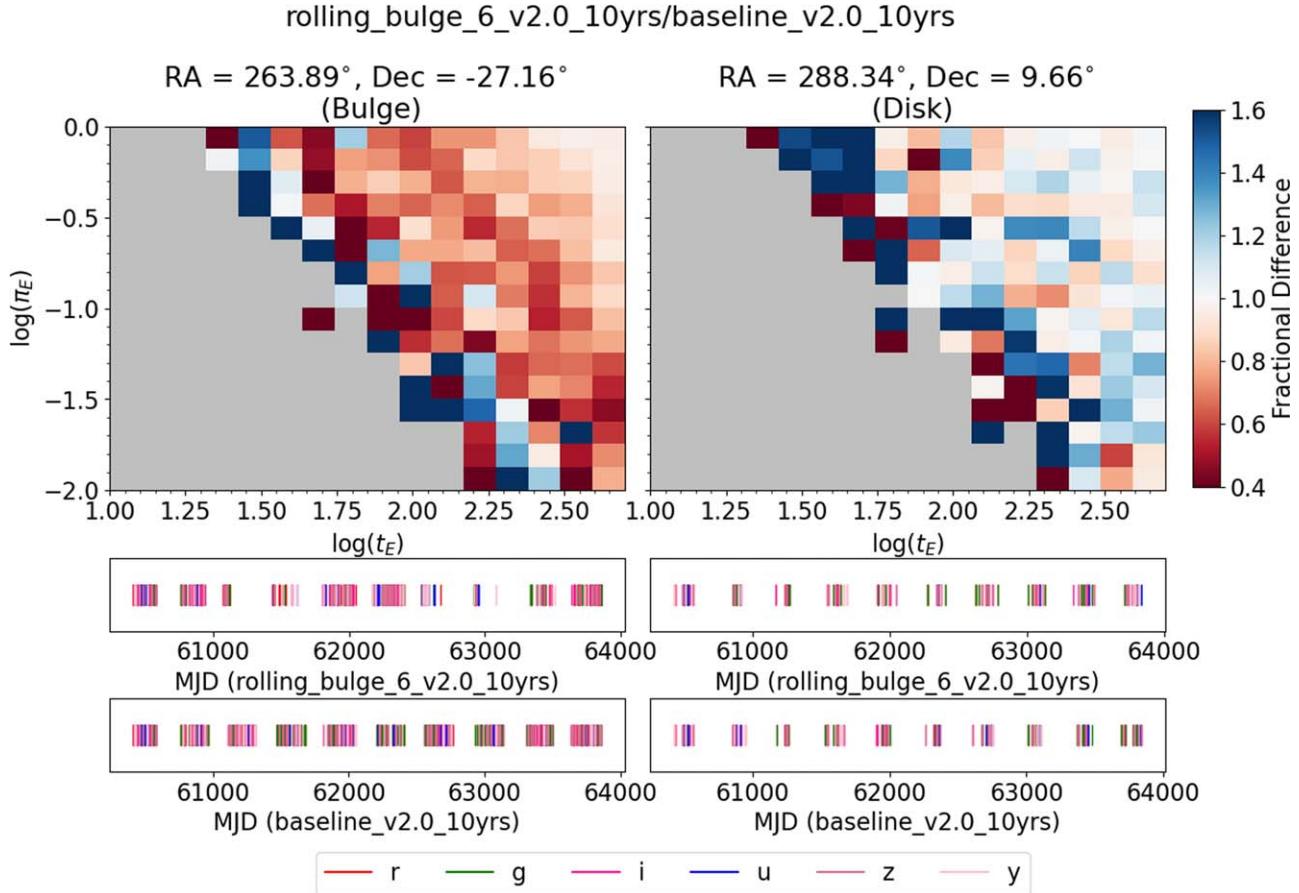


Figure 14. Same as Figure 12, but comparing `rolling_bulge_6_v2.0_10yrs` to `baseline_v2.0_10yrs`. This has the same footprint as `baseline_v2.0_10yrs` but splits the bulge into six sections on which it alternates focus, as can be seen in the middle left panel. This leads to $\sim 20\%$ fewer events being characterized in the Galactic bulge.

because in `baseline_v3.0_10yrs` there are more observations spread out across the plane, maximizing Rubin’s capability to do a Galaxy-wide study.

We can compare `baseline_v2.0_10yrs` to the rolling cadences since they share the same footprint. In `rolling_all_sky_ns2_rw0.9_v2.0_10yrs` (Figure 13) there is a drop in characterization efficiency on the 5%–10% level due to longer periods with long gaps between observations, whereas in `rolling_bulge_6_v2.0_10yrs` (Figure 14) there is a drop by $\sim 15\%-20\%$, particularly for high-parallax events, due to seasons with very sparse coverage and long gaps. In both, there is little change in the disk field, as it has such sparse cadence and has a relatively constant cadence between years.

For reference, we can compare these cadences to `retro_baseline_v2.0_10yrs` (Figure 15). This covers the Galactic bulge and plane very sparsely and causes events to be 60%–80% more difficult to characterize than the current baseline. This is indicative that the strongest determiner of characterization is the footprint. If time is not dedicated to the Galactic plane and bulge, most events will not be characterizable.

4. Discussion and Conclusion

Given the simplifying assumptions made in the paper, the main survey cadence optimizations that have a major effect on

microlensing discovery and characterization can be summarized as follows:

1. Footprint, to the first order, makes the most significant impact on microlensing detection and characterization. This can be seen by comparing the current baselines to the `retro_baseline`, which largely avoids the Galactic bulge and plane (see Figure 5 and Table 2 for summaries of OpSims).
2. Rubin will be able to use its uniquely deep and wide survey area to detect and characterize microlensing events across the Galactic plane. However, if areas of low stellar density and high extinction are included, this can lead to a decrease in the fraction of characterized events (see Figures 6–9).
3. A rolling cadence in the Galactic bulge and plane has the potential to improve synergies with Roman but should be approached with caution. The survey should avoid long gaps, since many current rolling strategies decrease detection and characterization of most microlensing events.

The survey cadences besides the `retro` footprints have also incorporated the LMC and SMC. Their inclusion will allow us to probe microlensing events caused by objects in the halo.

This paper has mostly discussed what Rubin can do alone with cadence optimization. Rubin will also send out nightly alerts that could be used to do follow-up on candidates and is of

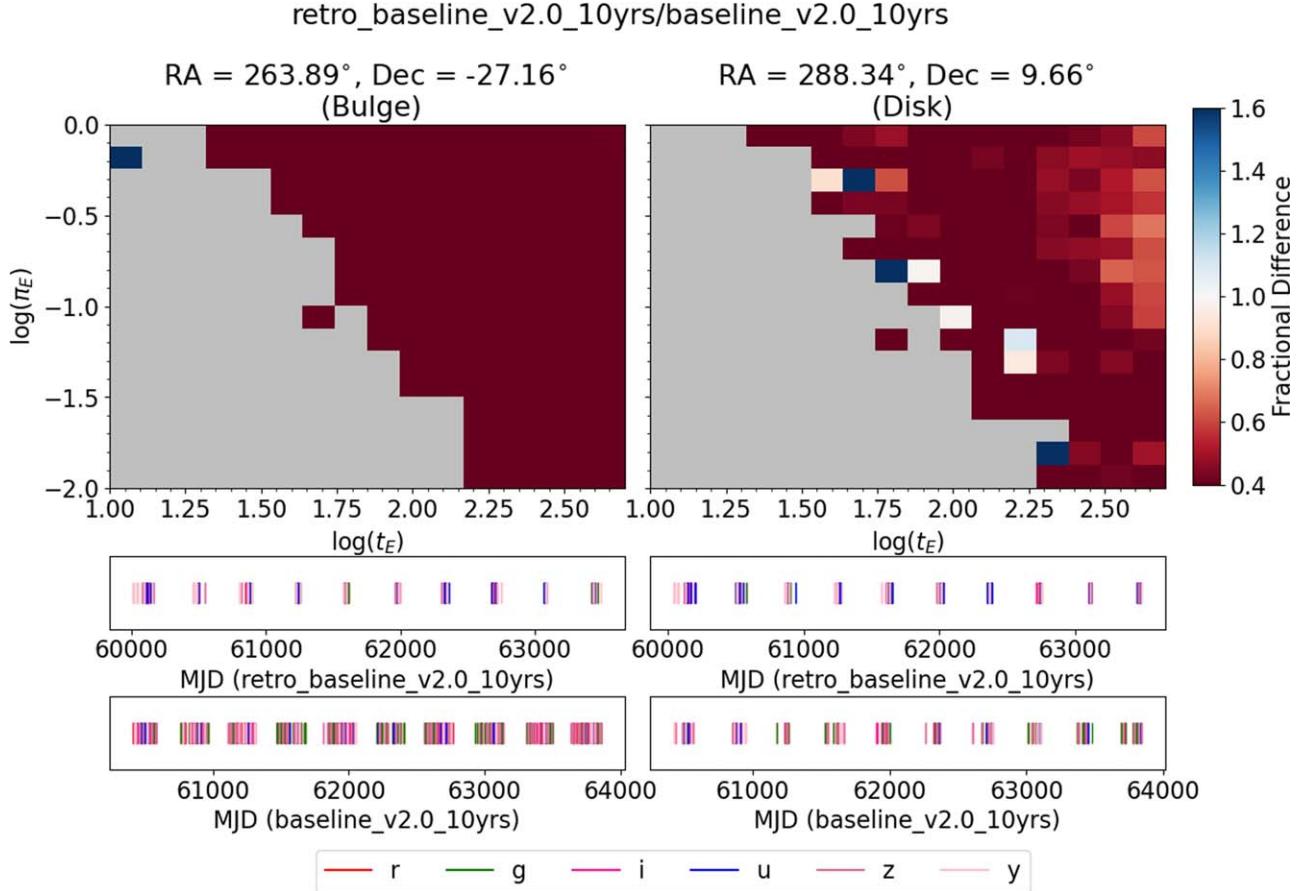


Figure 15. Same as Figure 12, but comparing `retro_baseline_v2.0_10yrs` to `baseline_v2.0_10yrs`. Given the extremely sparse cadence in the Galactic bulge and plane in the retro baseline, we would characterize 60%–80% fewer microlensing events.

particular interest to exoplanet science (see Hundertmark et al. 2021). Follow-up observations could fill in gaps in Rubin’s coverage, but Rubin would still require adequate cadence in the wings of the event to reliably send alerts on events in progress. It may also be difficult to follow up on most faint events from most ground-based facilities.

4.1. End-to-end Pipeline

The above assessment of cadence strategy has all been relative between cadences. We have simulated ranges of microlensing parameters such as t_E and π_E or plotted results as a function of those values. We have not incorporated either a Galactic model with a simulated microlensing survey or a distribution of magnitudes and colors, which are both necessary for predicting a realistic number and distribution of detected and characterized microlensing events and understanding the effect of filter balance and exposure length on the survey.

Beyond the cadence, the reduction and analysis pipelines will also play a large part in detection and characterization of microlensing events. We have seen with the first 3 yr of ZTF data, another all-visible-sky survey, covering one hemisphere, that ~ 100 events were discovered (A. C. Rodriguez et al. 2022; M. S. Medford et al. 2023). In order to find a pure sample, many real microlensing events were likely excluded, due to the $> 10^9$ initial light curves requiring strict cuts to fit all of the events (see M. S. Medford et al. 2023, Section 6). The analysis pipeline will have a significant effect on the microlensing yield. Given that microlensing is more likely to occur in crowded

fields, carefully deblending and extracting photometry will be imperative to maximizing the number of detected and well-characterized events. In a true end-to-end Rubin microlensing simulation, everything from initial physical Galactic parameters to reduction pipeline would be incorporated.

Since the Rubin Survey Optimization Committee plans to assess the cadence multiple times throughout Rubin’s operation, it is important to continue to develop our ability to assess cadences, including folding in real data.

Acknowledgments

We thank Federica Bianco for helpful conversations and guidance. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a US Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under contract No. DE-AC02-05CH11231 using NERSC awards HEP-ERCAP0023758 and HEP-ERCAP0026816. N.S.A. and J.R.L. acknowledge support from the National Science Foundation under grant No. 1909641 and the Heising-Simons Foundation under grant No. 2022-3542. R.A.S. gratefully acknowledges support from NSF grant No. 2206828. Y.T. acknowledges the support of DFG priority program SPP 1992 Exploring the Diversity of Extrasolar Planets (TS 356/3-1). Support for M.R. is provided by the Dirección de Investigación of the Universidad Católica de la Santísima Concepción with the project DIREG 10/2023.

Software: LSST Metric Analysis Framework (https://github.com/lsst/rubin_sim), Numpy (C. R. Harris et al. 2020), Matplotlib (J. D. Hunter 2007), BAGLE, SymPy (A. Meurer et al. 2017), PopSyCLE (C. Y. Lam & J. R. Lu 2020).

Appendix A Additional OpSims

Figures 16 and 17 summarize other select OpSims discussed in the paper, but without dedicated plots. Figure 16 has v2.0 OpSims, and Figure 17 has v2.1–v3.0 OpSims. In Table 2, the OpSims discussed in this paper are summarized with descriptions relevant to microlensing and other Galactic science. See R. L. Jones et al. (2020) and The Rubin Observatory Survey Cadence Optimization Committee (2023) for more detailed descriptions. Tables 3–6 contain metric values for the OpSims in this paper. Table 4 has summary statistics for each OpSim, Table 5 has Discovery Metric results, and Table 6 has Npts and Fisher Metric results. In addition, Table 3 contains metric values for baseline_v2.0_10yrs broken down by Galactic location.

There are a few OpSim families that include additional surveys unrelated to microlensing to LSST that we will discuss here. The `ddf_OpSim` family evaluates changes to the DDF strategy. Since the DDFs do not cover the Galactic plane, the more they have visits dedicated to them, the fewer visits are available for regions in the Galactic plane. At the current

level of the DDFs in the OpSims, it decreases short-event characterization at the 10%–15% level, but besides that it does not appear to significantly affect the microlensing science case.

The `vary_nes` OpSim family varies the coverage of the North Ecliptic Spur (NES) as a percentage of the WFD survey. The more the survey strategy covers the NES, the less we are able to cover the Galactic bulge and plane, which causes the microlensing metric to suffer. For an NES coverage of `nesfrac` = 60%–75%, there is a significant drop of about 15% in fraction of microlensing events that can be characterized (see Figure 16).

The Twilight NEO OpSim family (`twilight_neo_nightpatternX`) explores adding a twilight observing strategy, primarily looking for near-Earth objects. The SNR of observations is reduced, so the events, especially the short stellar events, suffer in characterizations by 5%–10%. Some of the long events technically have more observations that overcome the SNR downsides, but the quality loss and systematic effects would make an analysis challenging despite the technically better relative assessment. This is a surprising result since we did not think that the twilight observations would lead to a poorer coverage of the nighttime events. A representative 3 of the 84 `twilight_neo` runs are plotted in Figure 17.

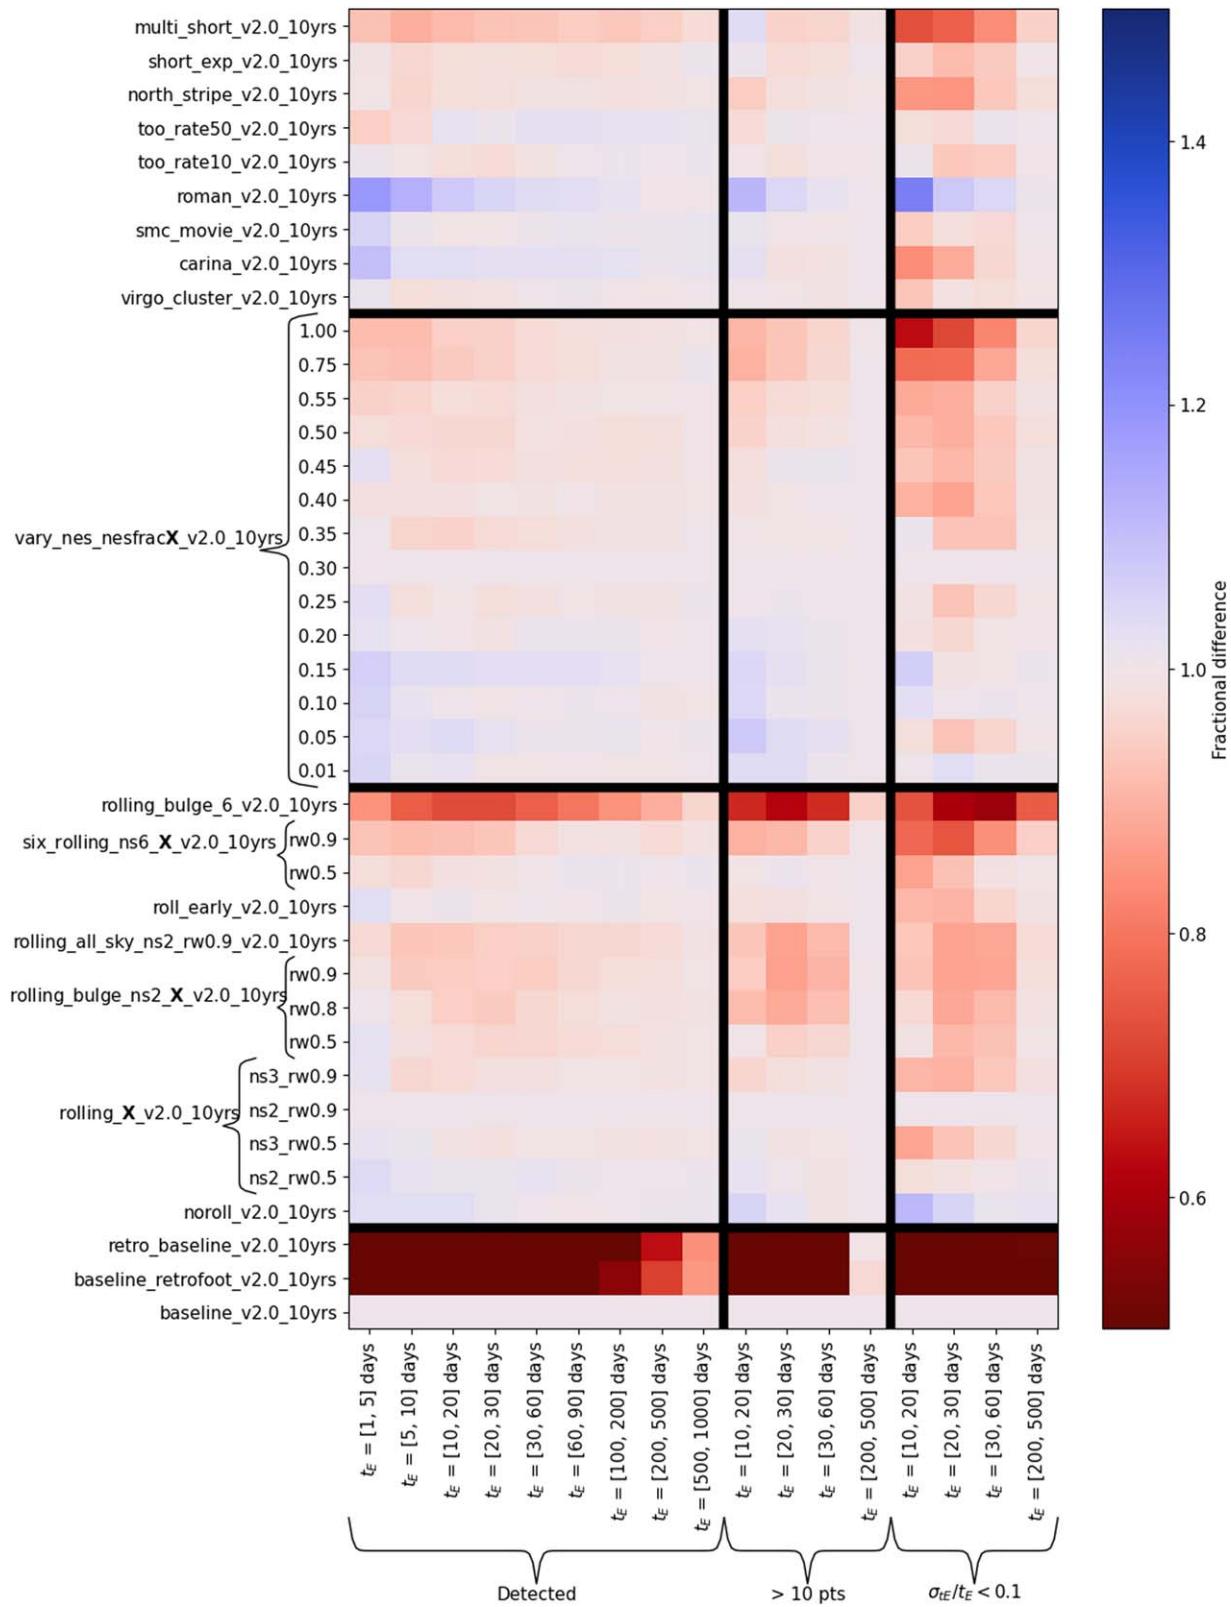


Figure 16. Same as Figure 5, but for select other v2.0 OpSims, with baseline metrics for reference. This has the OpSims on the y-axis and the metrics on the x-axis for ease of plotting.

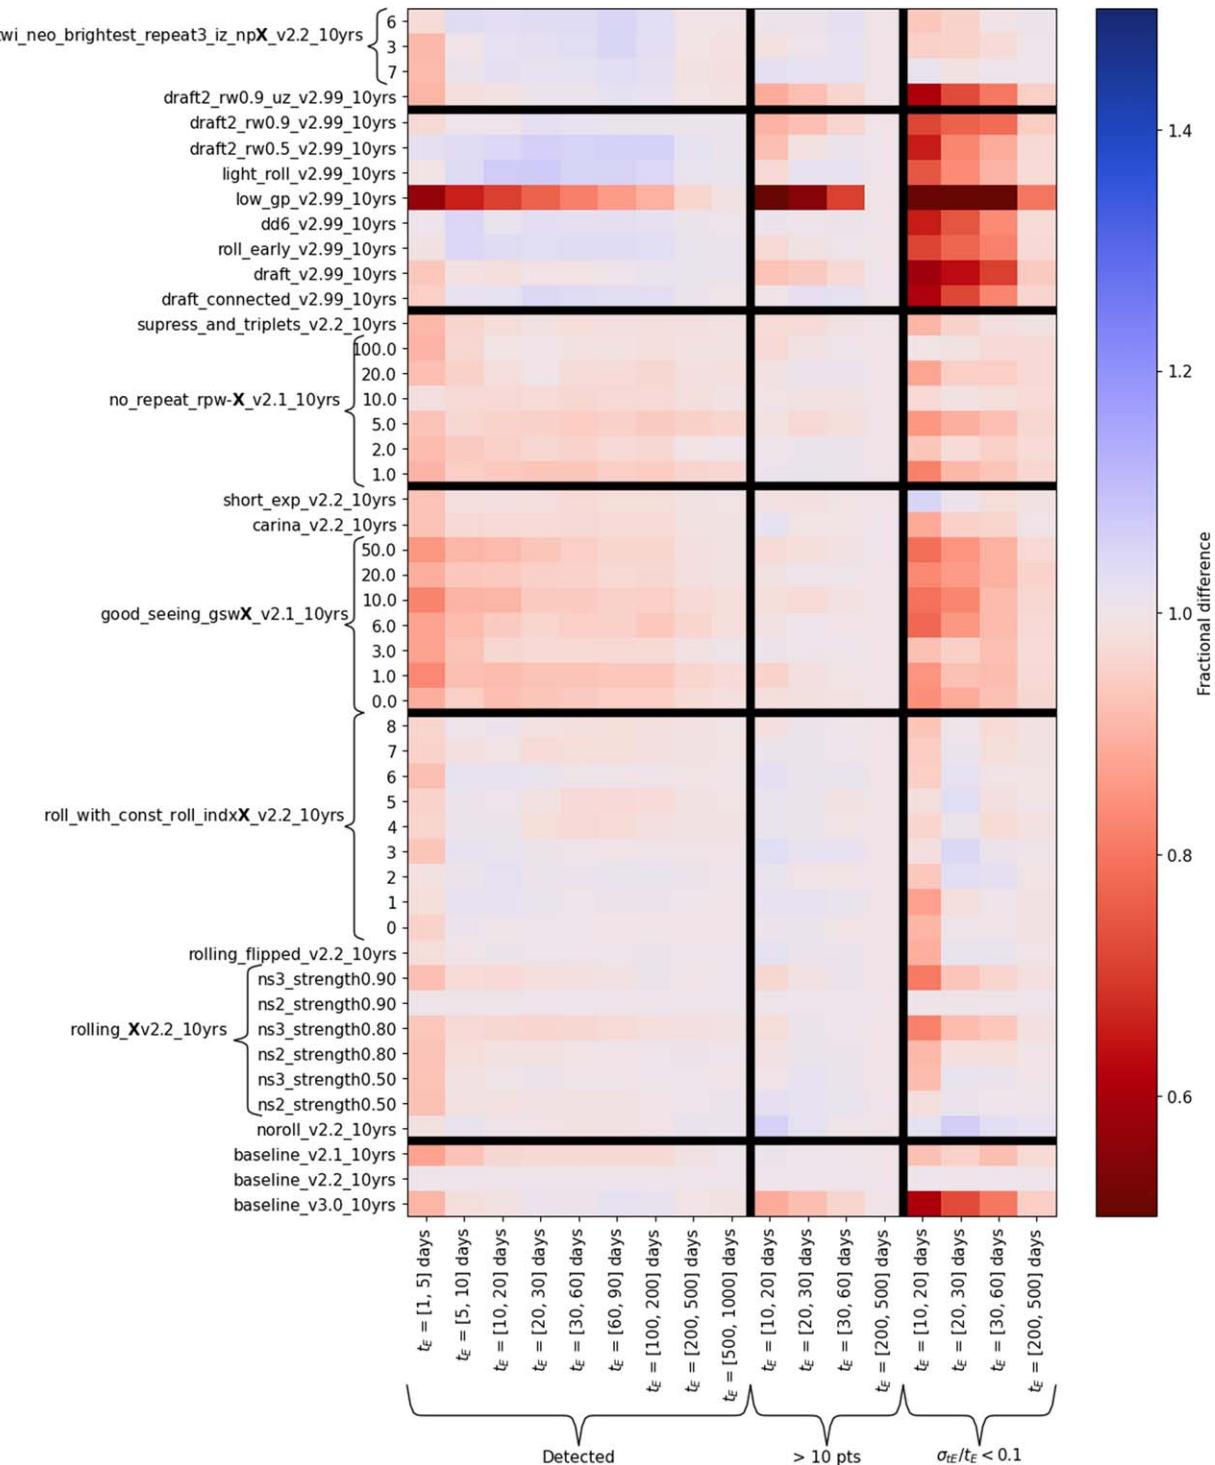


Figure 17. Same as Figure 5, but for select other 2.1 and later OpSims, with baseline metrics for reference. All OpSims are plotted in reference to `baseline_v2.2_10yrs` in this plot. Three representative `twilight_neo` runs are plotted. This has the OpSims on the y-axis and the metrics on the x-axis for ease of plotting.

Table 2
Summary of OpSims Alphabetical by Family that Are Relevant to Microlensing and Milky Way Science with Descriptions of Pertinent Aspects

OpSim (Family) Name	Description
baseline	Baseline survey strategies.
baseline_v3.0_10yrs	Includes high-priority areas across the Galactic plane and bulge.
baseline_v2.2_10yrs	Optimizations to code and DDF strategy change.
baseline_v2.1_10yrs	Includes Virgo Cluster and acquisition of good seeing images in <i>r</i> and <i>i</i> bands.
baseline_v2.0_10yrs	Added in the Galactic bulge, LMC, and SMC.
baseline_retrofoot_v2.0_10yrs	Uses 2018 footprint and v2.0 baseline strategy.
retro_baseline_v2.0_10yrs	Uses 2018 footprint and strategy.
ddf_*	Varies survey strategy of DDFs.
Galactic plane	Simulations that explore the Galactic plane survey strategy.
plane_priority_priorityX_pbf_*	Includes regions of Galactic plane with priority $\geq X$ (priorityX) not including pencil beam fields (pbf = pencil beams false).
plane_priority_priorityX_pbt_*	Includes regions of Galactic plane with priority $\geq X$ (priorityX) including pencil beam fields (pbt = pencil beams true).
pencil_fsX_10yrs	Varies size/number of pencil beams where $X = 1$ (fs1) is 20 smaller fields and $X = 2$ is four larger ones (fs2).
vary_gp_gpfracX_*	Spends $X\%$ (gpfracX) of survey time on areas of the Galactic plane not including in the WFD.
good_seeing_*	Adds requirement of at least three good seeing images per year per pointing.
microsurveys	Surveys requiring $<3\%$ of LSST time.
roman_v2.0_10yrs	Adds microsurvey of Roman GBTDS field.
smc_movie_v2.0_10yrs	Adds two nights of observing of the SMC.
rolling	Varies strategy to alternate high-cadence coverage of areas of the sky.
noroll_v2.0_10yrs	No rolling (noroll).
rolling_nsX_rwY_*	Splits the sky into X regions (nsX) with $Y\%$ (rwY) strength of rolling.
rolling_nsX_strengthY_*	Splits the sky into X regions (nsX) with $Y\%$ (strengthY) strength of rolling.
rolling_bulge_nsX_rwY_*	Splits the Galactic bulge into X regions (nsX) with $Y\%$ (rwY) strength of rolling.
rolling_early_v2.0_10yrs	Rolling cadence beginning in year 1 of LSST.
six_rolling_*	Splits sky into six regions and performs a rolling cadence.
rolling_bulge_6_v2.0_10yrs	Splits Galactic bulge into six regions and performs a rolling cadence.
rolling_with_const_*	Intersperses rolling cadence with constant years.
rolling_flipped_*	Flips the order of the two-band 90% strength rolling cadence.
rolling_all_sky_*	Rolls the whole sky with the same cadence.
triplets	Includes triplet observations in a single night.
presto_gapX	Triplets spaced X hours (gapX) apart.
presto_half_gapX	Triplets spaced X hours (gapX) apart every other night.
long_gaps_nightsoffX_delayed-1	Triplets every X nights (nightsoffX).
long_gaps_nightsoffX_delayed1827	Triplets every X nights (nightsoffX) starting after year 5.
twilight NEO	Survey added in twilight to observe near-Earth objects.
vary_nes_nesfracX_*	Survey strategy spends $X\%$ (nesfracX) of survey time on the NES.

Note. Those that end in an underscore and star (*) indicate that there are multiple OpSims related to that entry. Indented entries belong to the family listed above. See R. L. Jones et al. (2020) and The Rubin Observatory Survey Cadence Optimization Committee (2023) for more detailed descriptions.

Table 3
Metric Results for `baseline_v2.0_10yrs` Breaking Down Metric Results According to Location Since the t_E Distributions Change as a Function of Location

Metric	t_E (days)	Galactic Bulge	Galactic Plane	Magellanic Clouds
Discovery	1–5	220	810	51
	5–10	486	1618	113
	10–20	752	2497	162
	20–30	926	3123	209
	30–60	1116	3768	249
	60–90	1267	4317	291
	100–200	1542	5259	343
	200–500	1862	6445	421
	500–1000	1839	6490	430
Npts	10–20	707	2251	159
	20–30	1319	4199	257
	30–60	1761	5683	342
	200–500	2800	9855	655
Fisher	10–20	224	792	26
	20–30	512	1738	80
	30–60	822	2753	146
	200–500	2403	8252	529

Note. Events were sorted into these locations using the Healpix maps from R. A. Street et al. (2023b). The Discovery Metric is the number of events (out of 10,000 in the entire footprint) with two points on the rising side of the light curve with $\geq 3\sigma$ difference in magnitude, the Npts Metric has the number of light curves with at least 10 points in $t_0 \pm t_E$ (see Section 2.2), and the Fisher Metric has the number of light curves with $\frac{\sigma_{t_E}}{t_E} < 0.1$ (see Section 2.2.1).

Table 4
Summary Statistics of OpSims in This Paper

OpSim	Nvisits	Area with > 825 Visits	Median WFD Inter- night Gaps
baseline_v3.0_10yrs	2,086,079	2921.58	2.95
baseline_v2.2_10yrs	2,074,975	10,569.22	2.98
baseline_v2.1_10yrs	2,081,749	12,434.14	3.02
baseline_v2.0_10yrs	2,086,980	12,893.23	3.00
light_roll_v2.99_10yrs	2,082,888	650.45	3.07
draft2_rw0.5_v2.99_10yrs	2,082,209	1545.14	3.02
draft2_rw0.9_v2.99_10yrs	2,085,283	2368.49	2.95
draft2_rw0.9_uz_v2.99_10yrs	2,086,079	2921.58	2.95

Note. Nvisits corresponds to the “Nvisits All visits” metric and is the total number of visits through the entire survey. Area with > 825 visits corresponds to the “fOArea fO All sky HealpixSlicer” metric and is the area in square degrees that received over 825 visits. Median WFD Internight Gaps corresponds to the “Median Median Internight Gap WFD all bands HealpixSubsetSlicer” metric and is the median gap in days between observations in the WFD region including all filters. Results for these metrics and OpSims, as well as additional metrics and OpSims, can be found on GitHub at https://github.com/lsst-pst/survey_strategy/blob/main/fbs_2.0/summary_2023_01_01.csv.

(This table is available in its entirety in machine-readable form in the [online article](#).)

Table 5
Discovery Metric Values of OpSims in This Paper

OpSim	Discovery Metric (t_E Range in Days)								
	1–5	5–10	10–20	20–30	30–60	60–90	100–200	200–500	500–1000
baseline_v3.0_10yrs	779	1618	2518	3239	3917	4559	5563	6674	6693
baseline_v2.2_10yrs	860	1656	2539	3212	3900	4462	5462	6710	6765
baseline_v2.1_10yrs	751	1535	2440	3111	3764	4312	5290	6639	6773
baseline_v2.0_10yrs	761	1580	2427	3116	3747	4284	5234	6551	6600
light_roll_v2.99_10yrs	854	1727	2726	3457	4120	4708	5707	6798	6851
draft2_rw0.5_v2.99_10yrs	882	1717	2655	3413	4128	4732	5803	6867	6843
draft2_rw0.9_v2.99_10yrs	831	1657	2540	3290	3969	4528	5546	6740	6841
draft2_rw0.9_uz_v2.99_10yrs	779	1618	2518	3239	3917	4559	5563	6674	6693

Note. Each entry has the number of light curves detected (defined in Section 2.2) out of 10,000 simulated light curves per entry. Each column has the t_E range of events simulated. Results for these metrics and OpSims, as well as additional metrics and OpSims, can be found on GitHub at https://github.com/lsst-pst/survey_strategy/blob/main/fbs_2.0/summary_2023_01_01.csv.

(This table is available in its entirety in machine-readable form in the [online article](#).)

Table 6
Npts Metric and Fisher Metric Values of OpSims in This Paper

OpSim	Npts Metric (t_E Range in Days)				Fisher Metric (t_E Range in Days)			
	10–20	20–30	30–60	200–500	10–20	20–30	30–60	200–500
OpSim	10–20	20–30	30–60	200–500	10–20	20–30	30–60	200–500
baseline_v3.0_10yrs	0.24	0.41	0.56	1	0.05	0.12	0.22	0.80
baseline_v2.2_10yrs	0.27	0.44	0.58	1	0.08	0.17	0.28	0.85
baseline_v2.1_10yrs	0.27	0.44	0.59	1	0.07	0.16	0.26	0.82
baseline_v2.0_10yrs	0.26	0.44	0.59	1	0.07	0.16	0.27	0.82
light_roll_v2.99_10yrs	0.26	0.45	0.60	1	0.06	0.14	0.25	0.82
draft2_rw0.5_v2.99_10yrs	0.25	0.44	0.59	1	0.05	0.14	0.25	0.82
draft2_rw0.9_v2.99_10yrs	0.24	0.41	0.56	1	0.05	0.13	0.22	0.80
draft2_rw0.9_uz_v2.99_10yrs	0.24	0.41	0.56	1	0.05	0.12	0.22	0.80

Note. The first four columns have the fraction of light curves with at least 10 points in $t_0 \pm t_E$ (see Section 2.2). The second four columns have the fraction of light curves with $\frac{\sigma t_E}{t_E} < 0.1$ (see Section 2.2.1). Each column has the t_E range of events simulated. Results for these metrics and OpSims, as well as additional metrics and OpSims, can be found on GitHub at https://github.com/lsst-pst/survey_strategy/blob/main/fbs_2.0/summary_2023_01_01.csv.

(This table is available in its entirety in machine-readable form in the [online article](#).)

ORCID iDs

Natasha S. Abrams  <https://orcid.org/0000-0002-0287-3783>
 Markus P. G. Hundertmark  <https://orcid.org/0000-0003-0961-5231>
 Somayeh Khakpash  <https://orcid.org/0000-0002-1910-7065>
 Rachel A. Street  <https://orcid.org/0000-0001-6279-0552>
 R. Lynne Jones  <https://orcid.org/0000-0001-5916-0031>
 Jessica R. Lu  <https://orcid.org/0000-0001-9611-0009>
 Etienne Bachelet  <https://orcid.org/0000-0002-6578-5078>
 Yiannis Tsapras  <https://orcid.org/0000-0001-8411-351X>
 Marc Moniez  <https://orcid.org/0000-0001-8716-6561>
 Rosanne Di Stefano  <https://orcid.org/0000-0003-0972-1376>
 Martin Makler  <https://orcid.org/0000-0003-2206-2651>
 Markus Rabus  <https://orcid.org/0000-0003-2935-7196>

References

Albrecht, A., Bernstein, G., Cahn, R., et al. 2006, arXiv:astro-ph/0609591
 Alcock, C., Allsman, R. A., Alves, D. R., et al. 2000, *ApJ*, 542, 281
 Beaulieu, J. P., Bennett, D. P., Fouqué, P., et al. 2006, *Natur*, 439, 437
 Bennett, D. P., Batista, V., Bond, I. A., et al. 2014, *ApJ*, 785, 155
 Bianco, F. B., Drout, M. R., Graham, M. L., et al. 2019, *PASP*, 131, 068002
 Bianco, F. B., Ivezić, Ž., Jones, R. L., et al. 2022, *ApJS*, 258, 1
 Blaineau, T., Moniez, M., & Dagoret-Campagne, S. 2021, Microlensing towards the Magellanic Clouds: searching for long events, LSST Cadence Note 37634, https://docushare.lsst.org/docushare/dsweb/Get/Document-37634/LMC_SMC.pdf
 Connolly, A. J., Angeli, G. Z., Chandrasekharan, S., et al. 2014, *Proc. SPIE*, 9150, 915014
 Dal Tio, P., Pastorelli, G., Mazzi, A., et al. 2022, *ApJS*, 262, 22
 Gaudi, B. S. 2012, *ARA&A*, 50, 411
 Gould, A. 1992, *ApJ*, 392, 442
 Gould, A. 1994, *ApJL*, 421, L75
 Gould, A. 2013, arXiv:1304.3455
 Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, 585, 357
 Hundertmark, M., Tsapras, Y., Street, R., et al. 2021, Alerting transient phenomena in the Galactic Plane in time to coordinate follow-up, LSST Cadence Note 37638, https://docushare.lsst.org/docushare/dsweb/Get/Document-37638/Galactic_Plane_Transients.pdf
 Hunter, J. D. 2007, *CSE*, 9, 90
 Ivezić, Z. 2022, Survey Cadence Optimization Committee’s Phase 1 Recommendation PSTN-053, Vera C. Rubin Observatory, <https://pstn-053.lsst.io/>
 Jones, R. L., Yoachim, P., Ivezić, Z., Neilsen, E. H., & Ribeiro, T. 2020, Survey Strategy and Cadence Choices for the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST), v1.2, Zenodo, doi: 10.5281/zenodo.4048838
 Jones, R. L., Yoachim, P., Chandrasekharan, S., et al. 2014, *Proc. SPIE*, 9149, 91490B
 Jungman, G., Kamionkowski, M., Kosowsky, A., & Spergel, D. N. 1996, *PhRvD*, 54, 1332
 Kim, D. J., Kim, H. W., Hwang, K. H., et al. 2018, *AJ*, 155, 76
 Lam, C. Y., Abrams, N., Andrews, J., et al. 2023, arXiv:2306.12514
 Lam, C. Y., & Lu, J. R. 2023, *ApJ*, 955, 116

Lam, C. Y., Lu, J. R., Hosek, M. W. J., Dawson, W. A., & Golovich, N. R. 2020, *ApJ*, **889**, 31

Lam, C. Y., Lu, J. R., Udalski, A., et al. 2022, *ApJL*, **933**, L23

LSST Science Collaboration, Marshall, P., Anguita, T., et al. 2017, arXiv:1708.04058

Lu, J. R., Sinukoff, E., Ofek, E. O., Udalski, A., & Kozłowski, S. 2016, *ApJ*, **830**, 41

Medford, M. S., Abrams, N. S., Lu, J. R., Nugent, P., & Lam, C. Y. 2023, *ApJ*, **947**, 24

Meurer, A., Smith, C. P., Paprocki, M., et al. 2017, *PeerJ*, 3, e103

Moniez, M. 2010, *GReGr*, **42**, 2047

Moniez, M., Sajadian, S., Karami, M., Rahvar, S., & Ansari, R. 2017, *A&A*, **604**, A124

Mróz, P., Udalski, A., & Gould, A. 2022, *ApJL*, **937**, L24

Mróz, P., Udalski, A., Skowron, J., et al. 2017, *Natur*, **548**, 183

Mróz, P., Udalski, A., Szymański, M. K., et al. 2020, *ApJS*, **249**, 16

Naghib, E., Yoachim, P., Vanderbei, R. J., Connolly, A. J., & Jones, R. L. 2019, *AJ*, **157**, 151

Paczynski, B. 1986, *ApJ*, **304**, 1

Penny, M. T., Gaudi, B. S., Kerins, E., et al. 2019, *ApJS*, **241**, 3

Poindexter, S., Afonso, C., Bennett, D. P., et al. 2005, *ApJ*, **633**, 914

Poleski, R. 2016, *MNRAS*, **455**, 3656

Refsdal, S. 1966, *MNRAS*, **134**, 315

Rodriguez, A. C., Mróz, P., Kulkarni, S. R., et al. 2022, *ApJ*, **927**, 150

The Rubin Observatory Survey Cadence Optimization Committee 2023, Survey Cadence Optimization Committee's Phase 2 Recommendations PSTN-055, Vera C. Rubin Observatory, <https://pstn-055.lsst.io/>

Sahu, K. C., Anderson, J., Casertano, S., et al. 2022, *ApJ*, **933**, 83

Sajadian, S., & Poleski, R. 2019, *ApJ*, **871**, 205

Spergel, D., Gehrels, N., Baltay, C., et al. 2015, arXiv:1503.03757

Street, R., Lund, M., Khakpash, S., et al. 2018, arXiv:1812.03137

Street, R. A., Gough-Kelly, S., Lam, C., et al. 2023a, arXiv:2306.13792

Street, R. A., Li, X., Khakpash, S., et al. 2023b, *ApJS*, **267**, 15

Sumi, T., Abe, F., Bond, I. A., et al. 2003, *ApJ*, **591**, 204

Tsapras, Y. 2018, *Geosc*, **8**, 365

Tsapras, Y., Hundertmark, M., Wyrzykowski, Ł., et al. 2016, *MNRAS*, **457**, 1320

Udalski, A., Szymański, M. K., & Szymański, G. 2015, *AcA*, **65**, 1

Wyrzykowski, Ł., Kruszyńska, K., Rybicki, K. A., et al. 2023, *A&A*, **674**, A23

Yee, J. C. 2013, *ApJL*, **770**, L31

Yee, J. C., Albrow, M., Barry, R. K., et al. 2014, arXiv:1409.2759

Yee, J. C., Shvartzvald, Y., Gal-Yam, A., et al. 2012, *ApJ*, **755**, 102

Zhai, R., Rodriguez, A. C., Lam, C. Y., et al. 2023, arXiv:2311.18627