



# CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

## Beyond subhalos: Probing the collective effect of the Universe's small-scale structure with gravitational lensing

Francis-Yan Cyr-Racine, Charles R. Keeton, and Leonidas A. Moustakas

Phys. Rev. D **100**, 023013 — Published 25 July 2019

DOI: [10.1103/PhysRevD.100.023013](https://doi.org/10.1103/PhysRevD.100.023013)

# Beyond subhalos: Probing the collective effect of the Universe’s small-scale structure with gravitational lensing

Francis-Yan Cyr-Racine,<sup>1,\*</sup> Charles R. Keeton,<sup>2</sup> and Leonidas A. Moustakas<sup>3</sup>

<sup>1</sup>*Department of Physics, Harvard University, Cambridge, Massachusetts 02138 USA*

<sup>2</sup>*Department of Physics and Astronomy, Rutgers,*

*The State University of New Jersey, Piscataway, New Jersey 08854, USA*

<sup>3</sup>*Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California 91109, USA*

(Dated: June 7, 2019)

Gravitational lensing has emerged as a powerful probe of the matter distribution on subgalactic scales, which itself may contain important clues about the fundamental origins and properties of dark matter. Broadly speaking, two different approaches have been taken in the literature to map the small-scale structure of the Universe using strong lensing, with one focused on measuring the position and mass of a small number of discrete massive subhalos appearing close in projection to lensed images, and the other focused on detecting the collective effect of all the small-scale structure between the lensed source and the observer. In this paper, we follow the latter approach and perform a detailed study of the sensitivity of galaxy-scale gravitational lenses to the ensemble properties of small-scale structure. As in some previous studies, we adopt the language of the substructure power spectrum to characterize the statistical properties of the small-scale density field. We present a comprehensive theory that treats lenses with extended sources as well as those with time-dependent compact sources (such as quasars) in a unified framework for the first time. Our approach uses mode functions to provide both computational advantages and insights about couplings between the lens and source. The goal of this paper is to develop the theory and gain the intuition necessary to understand how the sensitivity to the substructure power spectrum depends on the source and lens properties, with the eventual aim of identifying the most promising targets for such studies.

## I. INTRODUCTION

The distribution of matter on kiloparsec scales and smaller may hold important clues about the fundamental nature of dark matter. Within the standard cold dark matter (CDM) paradigm, for instance, we expect dark matter at these scales to be distributed among a large number of low-mass subhalos that are largely devoid of gas and stars [1]. On the other hand, theories in which dark matter has a significant free-streaming length [2–6] or interacts with a relativistic species at early times (see e.g. Refs. [7–22]) predict a dearth of low-mass dark matter subhalos. Moreover, the presence of dark matter self-interaction [23–29] or dissipative dynamics [30–34] can modify the inner structure of subhalos, hence providing another potential probe of dark matter microphysics. Ultra-light axions [35, 36] might also lead to interesting phenomenology on small scales (see e.g. Refs. [37, 38]).

Several observational techniques can be used to probe the dark matter distribution on subgalactic scales. In addition to kinematical studies of ultra-faint dwarf galaxies within the Local Group [39–44], detailed analyses of potential perturbations to local stellar streams [45–49], to the Milky Way’s galactic disk [50, 51], or to the timing of distant pulsars [52, 53] have been proposed as methods to put constraints on the abundance of low-mass subhalos. Analyses of stellar wakes [54] could also be used to detect starless dark matter subhalos locally.

Both within and beyond the Local Group, gravitational lensing has emerged as a powerful probe of the dark matter distribution on the smallest scales. Indeed, nearby dark matter subhalos devoid of stars could be detected through lensing using subtle time-dependent astrometric perturbations of more distant sources [55–57]. Further out, galaxy-scale strong lensing systems in which a background high-redshift source is multiply imaged by a massive foreground galaxy are promising laboratories to study the dark matter distribution on the smallest scales. For example, the presence of flux-ratio anomalies in quasar lenses [58–64] has led to a measurement of the typical abundance of mass substructures within lens galaxies [65, 66], with future measurements appearing promising (see e.g. Refs. [67, 68]) if potential systematics can be properly accounted for [69–74]. Quasar flux ratio measurements have also been used to constrain the position and mass of potential individual subhalos within the lens galaxies [75–77]. In a different regime, the strong magnification near the caustic of galaxy clusters acting as strong lenses could also be used to probe subhalos [78].

Using images of lensed extended sources, gravitational imaging [79, 80] has led to the statistically-significant detection of a few subhalos with masses above  $\sim 10^8 M_\odot$  [81–83]. A somewhat similar technique using spatially-resolved spectroscopic observations of gravitational lenses [84, 85] has also led to the direct detection of a  $\sim 10^9 M_\odot$  subhalo [86]. Taken together, these measurements can in principle be used to put constraints on the subhalo mass function (see e.g. Refs. [87–90]), although interpreting these detections in terms of physical subhalo masses might be more subtle than initially thought [91].

---

\* fcyracci@physics.harvard.edu

In the near future, transdimensional techniques [92, 93] will allow a better understanding of various lensing degeneracies associated with subhalo detection within gravitational lenses.

Within the CDM paradigm, typical lens galaxies are expected to contain a large amount of mass substructure, with the subhalo mass function rapidly rising toward small halo mass [1, 94]. While the smallest subhalos are likely to be individually too light to be detected with standard gravitational imaging techniques, their collective effect might be detectable, in the spirit of the original work by Dalal and Kochanek [65]. Such a collective approach based on Approximate Bayesian Computation was used in Ref. [95] to put constraints on the abundance of substructure within the lens RX J1131-1231.

A different approach brought forward by Ref. [96] (see also [97]) aims at characterizing the mass substructure through their power spectrum. Interestingly, the substructure power spectrum contains information about the inner profile, mass function, and abundance of subhalos within lens galaxies [98–100], hence packing a lot of information within a single function. More generally, the power spectrum allows one to capture the effects of overdensities that cannot be easily described within the traditional language of the halo model (i.e. pancakes, streams, etc.) [101]. Since the density field on such small scales is expected to be nonGaussian, the power spectrum does not in general capture all its properties. The power spectrum can nonetheless contain important clues about the behavior of dark matter on sub-kiloparsec scales, just like measuring the nonlinear matter power spectrum of the large-scale structure of the Universe can lead to important information about the sum of neutrino masses. Recently, Ref. [102] provided the first upper limits on the amplitude of the substructure power spectrum in the lens SDSS J0252+0039.

In this work, we perform an in-depth analysis of the sensitivity of gravitational lens images to the substructure power spectrum, focusing on data taken with charged coupled devices (CCD), which are common in optical astronomy. We present here a likelihood-based mathematical framework necessary to extract the substructure power spectrum directly from pixel-based images. While our general approach follows a similar philosophy to that of Ref. [96], our computational technique differs at several levels, especially in our use of a mode function-based approach. Our approach also differs significantly from Ref. [97] since we make no a priori assumption about the statistical symmetries of the lensed source. Importantly, we extend our power spectrum mathematical framework to include compact time-varying sources such as quasars, hence opening substructure power spectrum measurements to a broader range of gravitational lenses. Since the goal of this paper is to present the framework necessary to extract measurements of the substructure power spectrum from lensed images and develop some intuition about their sensitivity to this latter quantity, we focus here on simple paramet-

ric source and lens models. Also, to avoid the complexity related to multi-plane lensing (see e.g. Refs. [103–105]), we shall concentrate in this work on substructure within the lens galaxy. We note however that most of the machinery developed here is likely applicable in the presence of line-of-sight structure with only minor modifications, as long as the effects of the substructure (both subhalos and those along the line of sight) on the lensed image are small.

Given the unique potential of this technique in probing sub-kiloparsec structure within galaxies and along the line-of-sight at cosmological distances from the Milky Way, we aim this paper at an audience that is not necessarily familiar with galaxy-scale strong lensing. As such, we carefully review the different ingredients and assumptions entering our analysis. Readers with expertise in cosmic microwave background or large-scale structure analyses will find several familiar concepts and techniques throughout this paper. Expert readers could skip directly to Sec. IV for details about our method to extract the substructure power spectrum from images of gravitationally lensed sources.

This paper is organized as follows. In Sec. II, we review the mass decomposition of the lens galaxy into macrolens and substructure, and then introduce the substructure convergence power spectrum. In Sec. III, we review the impact of mass substructure on observed images of galaxy-scale gravitational lenses, focusing on extended sources. In Sec. IV, we present the derivation of our likelihood for the substructure power spectrum in the case of extended lensed images. The numerical implementation of this likelihood is discussed in Sec. V, and simple Fisher forecasts are presented in Sec. VI. We present in Sec. VII complete Markov Chain Monte Carlo analyses of mock lensed images of extended sources to assess sensitivity to the substructure power spectrum. In Sec. VIII, we generalize our likelihood computation to include time-dependent compact sources such as quasars and present simple Fisher forecasts. We finally conclude in Sec. IX.

Throughout this paper, we assume a Planck 2015 cosmology [106]. We also take the redshift of the source to  $z_{\text{src}} = 0.6$  and that of the lens to be  $z_{\text{lens}} = 0.25$ , which results in a critical density for lensing  $\Sigma_{\text{crit}} = 5.998 \times 10^{10} M_{\odot}/\text{arcsec}^2 = 3.686 \times 10^9 M_{\odot}/\text{kpc}^2$  in the lens plane. A useful number to keep in mind is that for these choices of cosmology and redshifts,  $1 \text{ arcsec} \approx 4 \text{ kpc}$  in the lens plane.

## II. SMALL-SCALE STRUCTURE AND LENS GALAXIES

We begin this paper by reviewing the distinction between the so-called macro lens mass model and the small-scale substructure contained within the lens galaxies or along the line of sight. We then review the relevant statistical properties of the substructure that are most interesting from a gravitational lensing point of view.

### A. Mass decomposition for galaxy-scale lenses

In this work, we specialize to the case of galaxy-scale strong gravitational lenses, in which multiple images of a background source are generated. In general, the exact structure of the gravitational potential  $\phi_{\text{lens}}$  responsible for the lensing is the result of the complex assembly history of the lens galaxy as well as its subsequent dynamical evolution. In addition, structures along the line of sight can also contribute to the richness of the projected gravitational potential. Despite this apparent complexity, many observed galaxy-scale gravitational lenses can be reasonably fitted with relatively simple mass models, such as power-law ellipsoids (see however Ref. [107]).

A typical lens galaxy contains structure on a variety of scales, with the larger scale features responsible for the broad morphology of the observed lensed images, while the small-scale structure (e.g. satellite galaxies, giant molecular clouds, globular clusters, etc.) give rise to small corrections to the lensed observables. This suggests that we can decompose the lensing convergence (i.e. the two-dimensional projected mass density divided by the critical surface density for lensing) into a dominant macro component  $\kappa_0(\mathbf{y})$ , and a small contribution  $\kappa_{\text{sub}}(\mathbf{y})$  parametrizing the difference between the actual projected mass distribution and the dominant component  $\kappa_0$ , that is,

$$\kappa_{\text{lens}}(\mathbf{y}) = \kappa_0(\mathbf{y}) + \kappa_{\text{sub}}(\mathbf{y}). \quad (1)$$

Note that we have absorbed the mean convergence in the substructure (denoted  $\bar{\kappa}_{\text{sub}}$ ) within  $\kappa_0$  such that the  $\kappa_{\text{sub}}$  field as defined above has zero expectation value,  $\langle \kappa_{\text{sub}} \rangle = 0$ . We note that in the absence of lensing time-delay observations, stellar kinematic measurements, or strong priors on the brightness and size of the source, it is difficult to constrain  $\bar{\kappa}_{\text{sub}}$  due to the mass-sheet degeneracy [108]. We shall refer to  $\kappa_0$  (and  $\phi_0$ ) as the macro lens (or component) since it is responsible for determining the broad configuration of the lens. In general, it contains the contributions from the smooth dark matter halo, the dominant baryonic structure (disk, bulge, or otherwise), and possibly from single massive subhalos significantly affecting the configuration of the lens (such as those identified in Refs. [83, 86, 88]).

On the other hand, the substructure convergence  $\kappa_{\text{sub}}$  (and its related lensing potential  $\phi_{\text{sub}}$ ) contains contributions from the usual dark matter subhalos and satellite galaxies orbiting the main lens galaxy, as well as from other astrophysical structures such as tidal streams, debris, dense gas clouds, and globular clusters. The structure of the line-of-sight density field also contributes to  $\kappa_{\text{sub}}$ , although we do not explicitly take it into account here since this would require multiplane lensing. The crucial point is that the perturbations encoded in  $\kappa_{\text{sub}}$  are subdominant<sup>1</sup> compared to  $\kappa_0$ , hence resulting in subtle

disturbances in the lensed images. Detecting the presence of these small perturbations and characterizing their statistical properties could yield important clues about the nature of dark matter.

### B. Two-point statistics for substructure convergence

The typical length scales probed by strong lensing observations are deep into the nonlinear regime of cosmological structure formation where both baryon- and dark matter-dominated objects contribute significantly to the overall projected mass density. Taken at face value, this should imply that the statistics of the projected mass density fluctuations encoded in  $\kappa_{\text{sub}}$  are highly non-Gaussian and difficult to compute from first principles without resorting to expensive numerical simulations. However, in the cases where  $\kappa_{\text{sub}}$  and  $\phi_{\text{sub}}$  receive contributions from a large number of dark matter and baryonic mass substructures in the lens galaxy, it is possible to invoke the central limit theorem to argue that the statistics of  $\kappa_{\text{sub}}$  and  $\phi_{\text{sub}}$  are approximately Gaussian [96, 109, 110]. This occurs for instance in cold dark matter models where the subhalo mass function rises steeply towards small halo mass. For such approximately Gaussian scenarios, the statistical properties of the substructure convergence field are almost entirely captured by its two-point correlation function, or its Fourier transform, the power spectrum.

In a realistic lens galaxy, the mass substructure convergence and potential will never be exactly Gaussian. As was argued in Ref. [96], the non-Gaussian signatures are dominated by massive subhalos or line-of-sight structure appearing close in projection to lensed images. In high signal-to-noise images, these massive halos could in principle be directly detected through gravitational imaging [80, 85, 111] and incorporated into the main lens model described by  $\kappa_0$ , hence leaving the statistics of the mass substructures contributing to  $\kappa_{\text{sub}}$  roughly Gaussian. But even if the statistics of the substructure convergence are not entirely Gaussian, it is nonetheless interesting to measure the substructure convergence power spectrum since it still contains important information about small-scale structure within lens galaxies.

In contrast to the more familiar case of cosmological large-scale structure, we do not expect the substructure density field within a single lens galaxy to be either homogeneous or isotropic. However, given the relatively small projected area probed by the strong lensing region, neglecting the spatial variation of the substructure power spectrum across the lensed image is likely a good approximation. The relevant quantity is then the lens

<sup>1</sup> By construction, if the  $\kappa_{\text{sub}}$  perturbations were large, they would

lead to easily detectable effects, implying that they should have been absorbed in  $\kappa_0$ .

plane-averaged substructure power spectrum, which for a single lens plane is defined as [98]

$$P_{\text{sub}}(\mathbf{k}) = \int d^2\mathbf{r} e^{-i\mathbf{k}\cdot\mathbf{r}} \int_A d^2\mathbf{s} \frac{\langle \kappa_{\text{sub}}(\mathbf{s}) \kappa_{\text{sub}}(\mathbf{s} + \mathbf{r}) \rangle}{A}, \quad (2)$$

where  $\mathbf{k}$  is the Fourier wavenumber and  $A$  is the area of the strong lensing region where we perform the spatial average. This power spectrum should be interpreted as an ensemble average over multiple realization of a given lens, or for practical purposes, as an ensemble average over multiple galaxy-scale lenses. As such, we expect that it will be on average statistically isotropic, and the monopole power spectrum should carry thus most of the signal. The latter is simply given by

$$P_{\text{sub}}^{(0)}(k) = \frac{1}{2\pi} \int_0^{2\pi} d\theta_k P_{\text{sub}}(\mathbf{k}), \quad (3)$$

where  $\theta_k$  is the polar angle of the Fourier wavenumber  $\mathbf{k}$ , and  $k = |\mathbf{k}|$ . In this work, we shall focus our attention on the monopole contribution to the substructure convergence power spectrum but we note that the formalism developed below can easily accommodate anisotropic contributions<sup>2</sup> to the power spectrum.

A detailed study of the substructure convergence power spectrum for a dark matter-only population of subhalos was presented in Ref. [98]. We illustrate in the upper panel of Fig. 1 examples of the monopole convergence power spectrum for this idealized scenario. We note that we only display the one-subhalo term and we thus neglect here subhalo-subhalo clustering. The black dashed line shows the convergence power spectrum for an unclustered distribution of point masses, which is essentially a white noise (flat) power spectrum. The amplitude of this power spectrum is simply given by [96, 98, 109, 110]

$$P_{\text{sub}}^{(0)}(k) = \frac{\bar{\kappa}_{\text{sub}} \langle m^2 \rangle}{\Sigma_{\text{crit}} \langle m \rangle}, \quad (\text{point masses}) \quad (4)$$

where  $\langle m^n \rangle$  is the  $n$ th moment of the subhalo mass function [110]. We also display the convergence power spectrum for two other populations of subhalos, a truncated Navarro-Frenk-White (NFW) [113] subhalo population inspired by standard CDM, and a cored (modeled with a Burkert [114] profile) subhalo population inspired by self-interacting dark matter. In both cases, we see that the convergence power spectrum follows that of the point masses on scales larger than the size of the largest subhalos. Once that scale is reached, however, the convergence power spectrum for both the NFW and cored population begins to rapidly decay away from the point mass case.

<sup>2</sup> Since lens galaxies are generally not spherically symmetric due, e.g., to the presence of a disk or an elliptical baryonic distribution, it would be interesting, though likely difficult, to look at the possible dependence of the convergence power spectrum on the angle between Fourier modes and the baryonic major axis.

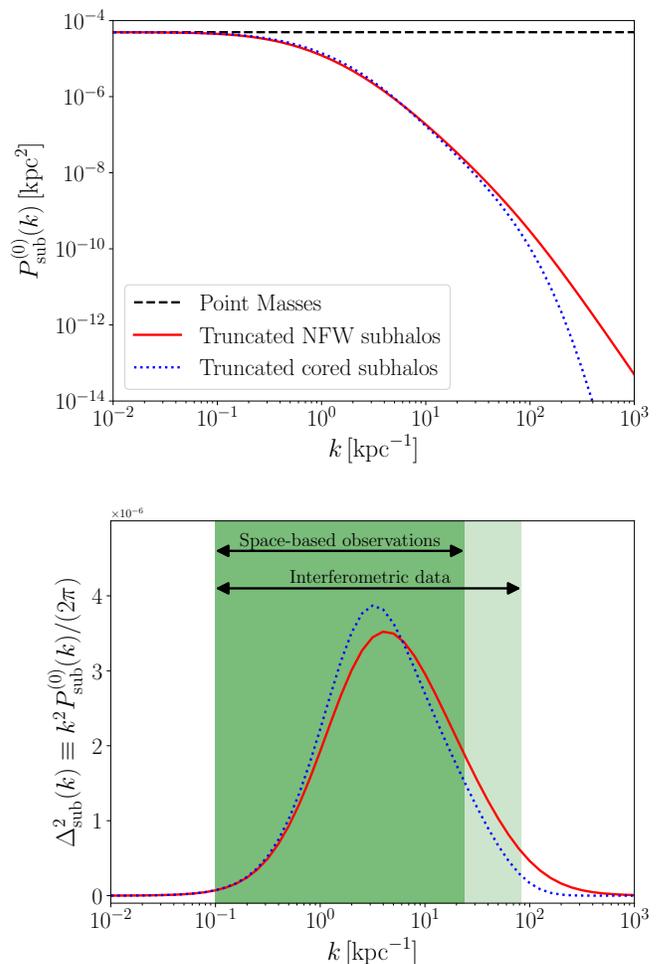


FIG. 1. The upper panel shows the substructure convergence power spectra for a population of point masses (dashed black), of truncated Navarro-Frenk-White (NFW) subhalos (solid red), and of truncated cored (Burkert) subhalos (dotted blue). Here, we have assumed a mean substructure abundance of  $\bar{\kappa}_{\text{sub}} = 0.01$ , a power law subhalo mass function  $dN/dm \propto m^\beta$  with  $\beta = -1.9$  and  $10^5 M_\odot < m < 10^8 M_\odot$ . For the truncated NFW, we take the scale radius to scale with the subhalo mass as  $r_s = 0.11 (m/10^6 M_\odot)^{1/2}$  kpc. The subhalos are tidally truncated at a radius  $r_t = (m/10^6 M_\odot)^{1/3} (r_{3D}/100 \text{ kpc})^{2/3}$  kpc [112], where  $r_{3D}$  is the three-dimensional distance of the subhalo from the center of the host. We average over all possible  $r_{3D}$  locations of a subhalo up to a radius of 409 kpc, as described in Ref. [98]. We adopt similar relations for the truncated cored subhalos but with a core radius  $r_c = 0.7 r_s$ . The lower panel shows the dimensionless power per log interval in wavenumber,  $\Delta_{\text{sub}}^2(k) \equiv k^2 P_{\text{sub}}^{(0)}(k)/(2\pi)$ , for both the truncated NFW and cored subhalo population. The two green bands display the range of scales that could be probed by space-based observations (dark green) and interferometric data (light green).

On small enough scales, the convergence power spectrum starts to probe the inner slope of the subhalos' density profile, with the cored subhalo case displaying less power

on these scales than the more cuspy NFW subhalos. We refer the reader to the figure caption for more details on the parameter used to generate these power spectra.

To understand on which scales the mass substructures matter the most, it is instructive to consider the dimensionless power spectrum

$$\Delta_{\text{sub}}^2(k) \equiv \frac{k^2 P_{\text{sub}}^{(0)}(k)}{2\pi}, \quad (5)$$

which describes the amount of substructure convergence power per logarithmic interval in wavenumber. We illustrate this dimensionless power spectrum in the lower panel of Fig. 1 for both the truncated NFW and truncated cored subhalo population. Since we observe that most of the power lies on scales 0.03 – 10 kpc for these plausible models, we naturally expect this range of scales to dominate the constraints that we can put on substructure using lensing data. Remarkably, such length scales correspond to the typical scales probed by galaxy-size strong lens observations. Assuming Einstein radii of 1–2'' and redshifts of  $z_{\text{lens}} = 0.1$ –1 which are typical for galaxy-scale lenses (see e.g. Ref. [115]), the range of scales that can in principle<sup>3</sup> be probed with space-based lensing observations (assuming a 70 mas image resolution) is indicated by the dark green band in the lower panel of Fig. 1. The higher image resolution achievable with interferometric data (such as those from the Atacama Large Millimeter/sub-millimeter Array) can significantly extend this range to higher wavenumbers (pale green band), hence allowing measurements on most of the length scales where mass substructure matters the most. Thus, galaxy-scale strong lenses constitute ideal laboratory to study small-scale dark matter structures.

In the remainder of this paper, we shall use both the point mass and the truncated NFW power spectra shown in Fig. 1 as benchmark fiducial models to generate simulated data and test their sensitivity to the amplitude and shape of the convergence power spectrum.

### III. GRAVITATIONAL LENSING IN THE PRESENCE OF MASS SUBSTRUCTURES

In this section, we review the impact of small-scale structure on lensing residuals for the case of extended, time-independent sources. Time-dependent compact sources such as quasars will be treated separately in Sec. VIII.

#### A. Preliminaries

Let us take a source at redshift  $z_{\text{src}}$  with a surface brightness profile at wavelength  $\lambda_{\text{src}} \equiv \lambda/(1 + z_{\text{src}})$

<sup>3</sup> We say ‘‘in principle’’ since the exact length scales that can actually be probed depend on the structure of the lensed source.

and time  $t_s$  described by  $\tilde{S}(\mathbf{u}, t_s, \lambda_{\text{src}})$ , where  $\lambda$  is the wavelength of the observation. Here,  $\mathbf{u}$  stands for two-dimensional angular coordinates in the source plane. We also take a gravitational lens at redshift  $z_{\text{lens}}$  specified by the projected mass density (convergence)  $\kappa_{\text{lens}}(\mathbf{y})$  and its related lensing potential  $\phi_{\text{lens}}(\mathbf{y})$ , which we assume to be static in time<sup>4</sup>. Again,  $\mathbf{y}$  stands for two-dimensional angular coordinates, but in the image plane. The lensing convergence and potential are related to each other via the usual Poisson equation  $\nabla^2 \phi_{\text{lens}} = 2\kappa_{\text{lens}}$ . Let us denote the lensed image of the source detected at time  $t$  and wavelength  $\lambda$  by  $\tilde{O}(\mathbf{y}, t, \lambda)$ . Formally, the relation between the source and the lensed image is [116]

$$\tilde{O}(\mathbf{y}, t, \lambda) = \int d\mathbf{u} \int dt_s \tilde{S}(\mathbf{u}, t_s, \lambda_{\text{src}}) \delta(t - t_s - \tau(\mathbf{y}, \mathbf{u})) \times \delta\left(\mathbf{u} - \mathbf{y} + \vec{\nabla} \phi_{\text{lens}}(\mathbf{y})\right), \quad (6)$$

where  $\vec{\nabla}$  denotes the gradient with respect to  $\mathbf{y}$ ,  $\delta$  is the Dirac delta function, and where the excess time delay is

$$\tau(\mathbf{y}, \mathbf{u}) = t_0 \left[ \frac{1}{2} |\mathbf{y} - \mathbf{u}|^2 - \phi_{\text{lens}}(\mathbf{y}) \right], \quad (7)$$

with

$$t_0 \equiv \frac{1 + z_{\text{lens}}}{c} \frac{D_1 D_s}{D_{\text{ls}}}, \quad (8)$$

where  $c$  is the speed of light,  $D_1$ ,  $D_s$ , and  $D_{\text{ls}}$  are the angular diameter distance between the observer and the lens, the observer and the source, and the lens and the source, respectively. Performing the integrals in Eq. (6), we obtain

$$\tilde{O}(\mathbf{y}, t, \lambda) = \tilde{S}\left(\mathbf{y} - \vec{\nabla} \phi_{\text{lens}}(\mathbf{y}), t - \tau(\mathbf{y}), \lambda_{\text{src}}\right), \quad (9)$$

where we use the shorthand notation  $\tau(\mathbf{y}) = \tau(\mathbf{y}, \mathbf{y} - \vec{\nabla} \phi_{\text{lens}}(\mathbf{y}))$ . In a typical observational scenario, the lens galaxy itself, the sky background, and possibly other objects along or close to the line of sight can also contribute to the observed photon flux in addition to the image of the lensed source. We take this into account by adding an external surface brightness contribution  $\tilde{L}(\mathbf{y}, t, \lambda)$  to the lensed image  $\tilde{O}(\mathbf{y}, t, \lambda)$ . Furthermore, the image is usually observed through a filter  $F_{\lambda_i}(\lambda)$  centered at a characteristic wavelength  $\lambda_i$ . We thus define the two following quantities

$$S_{\lambda_i}(\mathbf{u}, t_s) = \int d\lambda F_{\lambda_i}(\lambda) \tilde{S}(\mathbf{u}, t_s, \lambda_{\text{src}}), \quad (10)$$

$$L_{\lambda_i}(\mathbf{y}, t) = \int d\lambda F_{\lambda_i}(\lambda) \tilde{L}(\mathbf{y}, t, \lambda), \quad (11)$$

<sup>4</sup> On the typical timescale of a single astronomical observation, this is indeed a very good approximation. However, for a multi-epoch observational campaign of a given lens system, it is possible to notice changes in the gravitational potential due to stellar microlensing. In this work, we treat microlensing separately from  $\phi_{\text{lens}}$  (see Sec. VIII A) and we can therefore take the latter to be static.

which are the wavelength-integrated lensed source and external surface brightness, respectively. In addition, the light from all sources will be processed by the optics of the instrument used to observe it, as well by potential atmospheric disturbances. We take this effect into account by convolving the image with a point-spread function (PSF)  $W_{\lambda_i}(\mathbf{y}, t)$ , which in general depends on the wavelength  $\lambda_i$  and time of the observation. The actual model for the observed surface brightness  $\hat{O}_{\lambda_i}(\mathbf{x}, t)$  is thus given by

$$\hat{O}_{\lambda_i}(\mathbf{x}, t) = \int d\mathbf{y} \left[ S_{\lambda_i}(\mathbf{y} - \vec{\nabla}\phi_{\text{lens}}(\mathbf{y}), t - \tau(\mathbf{y})) \right. \\ \left. + L_{\lambda_i}(\mathbf{y}, t) \right] W_{\lambda_i}(\mathbf{x} - \mathbf{y}, t). \quad (12)$$

Finally, the light is usually collected for an exposure of length  $T_{\text{exp}}$  on a detector made of an array of two-dimensional pixels, and then converted to counts per pixel. Taking  $P_j(\mathbf{x})$  to be the pixel response function of the  $j$ th pixel<sup>5</sup>, and  $\mathcal{S}_{\text{inv}}^{(\lambda_i)}$  to be the inverse sensitivity of the detector for the filter labeled by  $\lambda_i$ , the number of counts in the  $j$ th pixel for the  $k$ th exposure is

$$O_{\lambda_i}(\mathbf{x}_j, t_k) = \frac{1}{\mathcal{S}_{\text{inv}}^{(\lambda_i)}} \int_{t_k}^{t_k + T_{\text{exp}}} dt \int d^2\mathbf{x} P_j(\mathbf{x}) \hat{O}_{\lambda_i}(\mathbf{x}, t) \\ \approx \frac{A_{\text{pix}} T_{\text{exp}}}{\mathcal{S}_{\text{inv}}^{(\lambda_i)}} \hat{O}_{\lambda_i}(\mathbf{x}_j, t_k), \quad (13)$$

where we have assumed in going from the first to the second line that the lensed source, foregrounds, and the PSF are static on the exposure timescale, and that the pixel response function is uniform across the detector and given by a 2D rectangular function. Here,  $A_{\text{pix}}$  is the area of a pixel, and  $\mathbf{x}_j$  is the position of the  $j$ th pixel.

## B. Lensing residuals for extended sources

In this section, we review (see e.g. Refs. [80, 111, 117]) the structure of the residuals between a strongly-lensed image of an extended source (e.g. a galaxy) created using only a macro lens  $\kappa_{\text{lens}} = \kappa_0$  from an image generated by a mass model that includes substructure,  $\kappa_{\text{lens}} = \kappa_0 + \kappa_{\text{sub}}$ . By causality, an extended source can be considered static on the typical time scales associated with astronomical observations. We can thus neglect the time dependence of the source in Eq. (12). However, the PSF and foreground light can vary from observations to observations, and we keep their time dependence explicit. Since the  $\phi_{\text{sub}}$  potential causes only small distortions to the observed image, we can perform a perturbative anal-

ysis in  $\phi_{\text{sub}}$  and expand  $\hat{O}_{\lambda_i}(\mathbf{x}, t)$  from Eq. (12) as

$$\hat{O}_{\lambda_i}(\mathbf{x}, t) \approx \int d\mathbf{y} W_{\lambda_i}(\mathbf{x} - \mathbf{y}, t) \left[ S_{\lambda_i}(\mathbf{y} - \vec{\nabla}\phi_0(\mathbf{y})) \right. \\ \left. - \vec{\nabla}_{\mathbf{u}} S_{\lambda_i}(\mathbf{u}) \Big|_{\mathbf{u}=\mathbf{y}-\vec{\nabla}\phi_0(\mathbf{y})} \cdot \vec{\nabla}\phi_{\text{sub}}(\mathbf{y}) + L_{\lambda_i}(\mathbf{y}, t) \right], \quad (14)$$

where  $\vec{\nabla}_{\mathbf{u}}$  denotes the source-plane gradient, and where  $\vec{\nabla}\phi_{\text{sub}}(\mathbf{y})$  is the deflection vector field created by the substructure. We note that the source-plane gradient arising in Eq. (14) can be translated to the image plane via the relation

$$\vec{\nabla}_{\mathbf{u}} S_{\lambda_i}(\mathbf{u}) \Big|_{\mathbf{u}=\mathbf{y}-\vec{\nabla}\phi_0(\mathbf{y})} = \left( \frac{\partial \mathbf{y}}{\partial \mathbf{u}} \right) \cdot \vec{\nabla} S_{\lambda_i}(\mathbf{y} - \vec{\nabla}\phi_0(\mathbf{y})), \quad (15)$$

where we recognize that the prefactor is nothing more than the magnification tensor  $\mathbf{M}_0 \equiv \partial \mathbf{y} / \partial \mathbf{u}$  generated by the macro lens component. Denoting by  $\hat{O}_{\lambda_i}^{(0)}(\mathbf{x}, t)$  the image of the source lensed purely by the macro potential (as well as potential foregrounds), the residuals between the image of a source lensed by the total lens potential  $\phi_{\text{lens}}$ , and that of the same source lensed only by the macro component  $\phi_0$  is

$$\delta \hat{O}_{\lambda_i, \text{sub}}(\mathbf{x}, t) \equiv \hat{O}_{\lambda_i}(\mathbf{x}, t) - \hat{O}_{\lambda_i}^{(0)}(\mathbf{x}, t) \\ \approx - \int d\mathbf{y} W_{\lambda_i}(\mathbf{x} - \mathbf{y}, t) \\ \times \left[ \vec{\nabla}_{\mathbf{u}} S_{\lambda_i}(\mathbf{u}) \Big|_{\mathbf{u}=\mathbf{y}-\vec{\nabla}\phi_0(\mathbf{y})} \cdot \vec{\nabla}\phi_{\text{sub}}(\mathbf{y}) \right]. \quad (16)$$

We thus obtain the well-known result [80, 111, 117] that the lensed image residuals of an extended source are proportional to the gradient of that source evaluated in the image plane. Residuals are largest when this latter gradient is either aligned or anti-aligned with the deflection field created by substructures. The above expression automatically captures the well-known fact that large, smoothly varying sources lead to lensed images that are largely insensitive to short-scale variations in the substructure potential.

## C. Validity of the linear approximation

We illustrate in Fig. 2 how the image residuals predicted by Eq. (16) (middle left panel) compare with the actual residuals (middle right panel) gotten by simply taking the difference  $\hat{O}_{\lambda_i}(\mathbf{x}, t) - \hat{O}_{\lambda_i}^{(0)}(\mathbf{x}, t)$  between an actual mock image generated with a nonvanishing  $\kappa_{\text{sub}}$  and a model image generated using only a macro lens model  $\kappa_0$ . The details of the lens, source, foreground, noise, and PSF models used to generate these images can be found in Appendix B below. Here, the substructures is modeled as truncated NFW subhalos with a convergence power spectrum given by the red solid line of Fig. 1. Despite the noise, we see that the predicted residuals match very closely the actual image residuals, implying that the

<sup>5</sup> For an ideal pixel, this function should be unity within the area spanned by the pixel, and zero elsewhere.

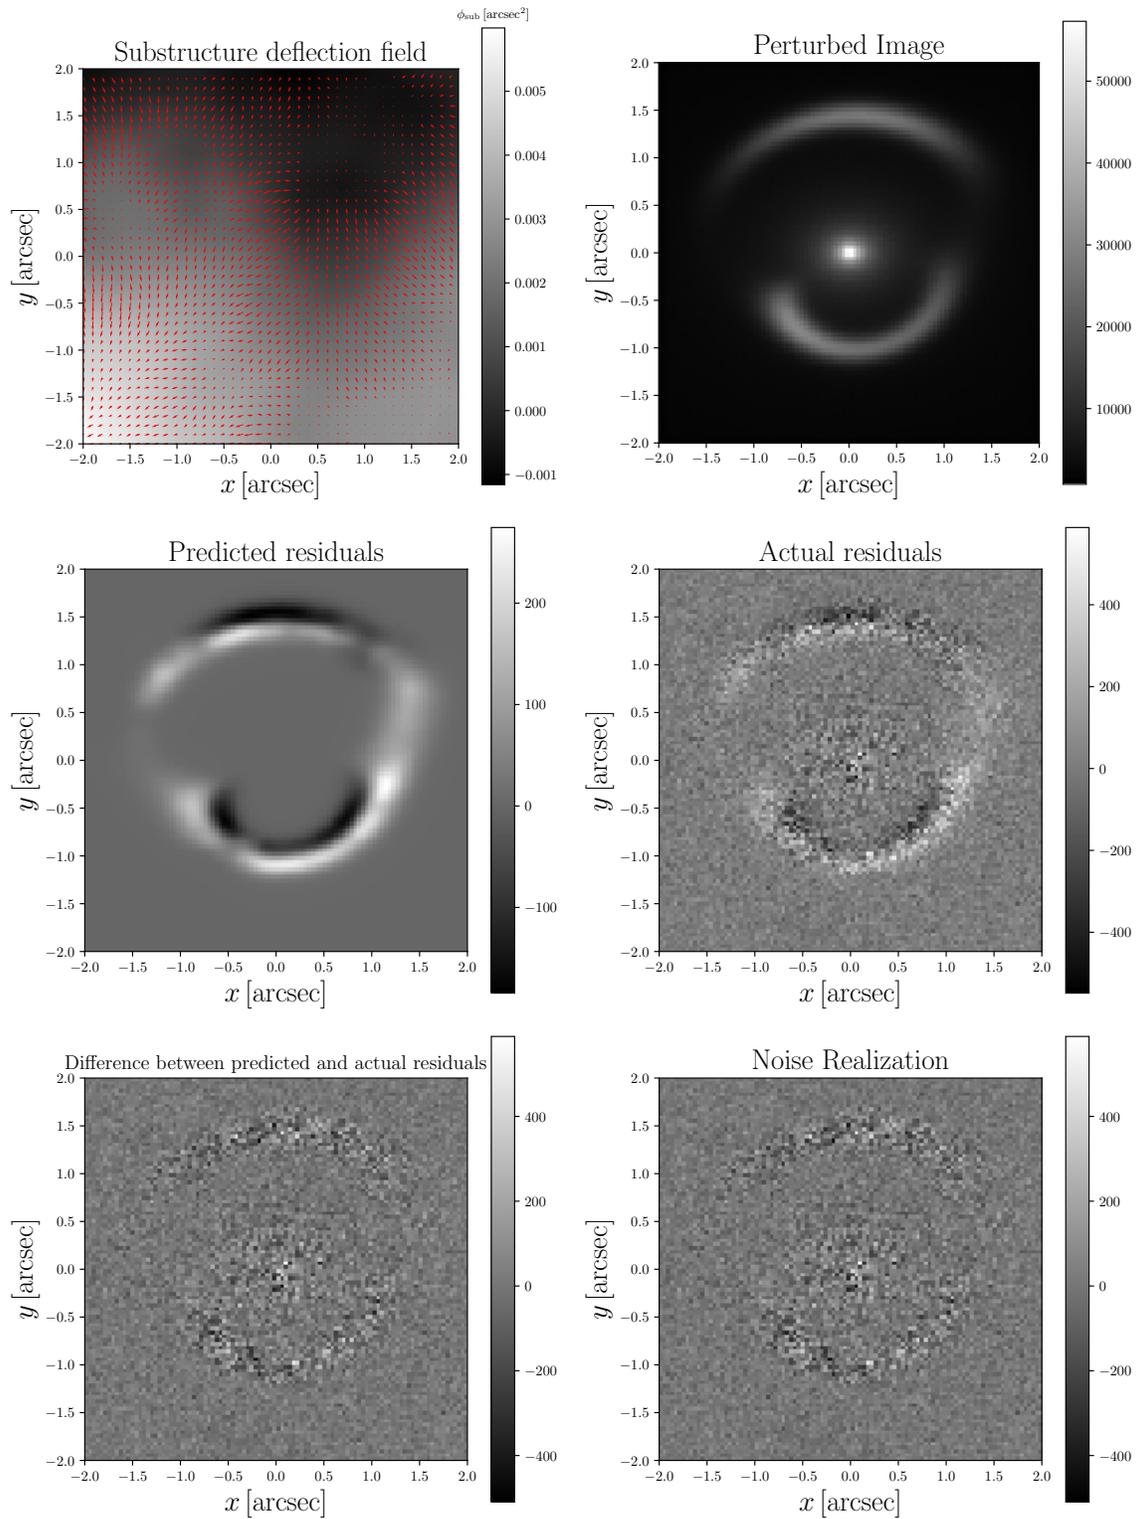


FIG. 2. Illustration of the lensing residuals in the presence of substructure. The upper left panel shows the substructure deflection field (red arrows) from a population of truncated NFW subhalos with a convergence power spectrum given by the red solid line of Fig. 1. The upper right panel shows a mock image of a background galaxy (modeled as a  $n = 1/2$  Sérsic profile) that is strongly lensed by a foreground galaxy (modeled as an isothermal ellipsoid plus external shear). The light from the lens galaxy (modeled as a  $n = 4$  Sérsic), sky background, and Poisson noise are added to the image. The middle left panel shows the predicted image residuals as given by Eq. (16) for the specific substructure realization used to generate the lensed image. The middle right panel shows the actual image residuals gotten by subtracting from the mock image a model image generated with the same lens galaxy, source parameters, environment, and observational conditions, but without the substructure. Finally, the lower left panel shows the difference between the residuals predicted by Eq. (16) and the actual image residuals which is entirely consistent with the Poisson noise realization used to generate the mock image (lower right panel).

most of the substructure information is captured by the first-order expansion performed in Eq. (16). This is also illustrated in the lower left panel of Fig. 2 where we display the difference between the predicted and actual image residuals, which is entirely consistent with the Poisson noise realization (lower right panel) used to generate the mock image.

Mathematically, one can assess the validity of the first-order approximation used to derive Eq. (16) by looking at the next order contribution to the lensing residuals, which takes the form

$$\frac{1}{2} \nabla \phi_{\text{sub}}(\mathbf{y}) \cdot \mathcal{H}_{S_{\lambda_i}} \cdot \nabla \phi_{\text{sub}}(\mathbf{y}), \quad (17)$$

where  $\mathcal{H}_{S_{\lambda_i}}$  is the Hessian matrix containing the second derivatives of the source surface brightness profile. Let us estimate the magnitude of this second-order term compared to the leading-order contribution  $\nabla_{\mathbf{u}} S_{\lambda_i} \cdot \nabla \phi_{\text{sub}}$ . For a source of typical size  $r_{\text{src}}$ , the ratio of the determinant of the Hessian matrix to that of the norm of the gradient of the source is  $\sim 1/r_{\text{src}}$ , which implies that the second-order corrections to Eq. (16) are negligible if

$$\frac{|\mathcal{H}_{S_{\lambda_i}}| |\nabla \phi_{\text{sub}}|}{|\nabla_{\mathbf{u}} S_{\lambda_i}|} \sim \frac{|\nabla \phi_{\text{sub}}|}{r_{\text{src}}} \ll 1, \quad (18)$$

that is, the typical magnitude of the substructure deflection must be much smaller than the size of the source. Since the typical magnitude of the substructure deflection field is [109, 110]

$$|\nabla \phi_{\text{sub}}| \sim \left( \frac{\bar{\kappa}_{\text{sub}} \langle m^2 \rangle}{\Sigma_{\text{crit}} \langle m \rangle} \right)^{1/2} \sim 10^{-3} \text{ arcsec}, \quad (19)$$

the first-order approximation should be valid for a broad range of galaxy-scale lenses with extended sources where the source sizes are typically in the range 0.1 – 1 arcsec (see e.g. Ref. [115]). However, the first-order calculation obviously breaks down for compact sources such as quasars for which a different treatment is necessary (see Sec. VIII A).

#### D. Degeneracy with the source brightness profile

When confronted with the image residuals given in Eq. (16), the immediate question that comes to mind is: *Can these residuals be reabsorbed by an appropriate modification to the source brightness profile?* To answer this question, let us imagine that we add a small contribution  $\delta S_{\lambda_i}(\mathbf{u})$  to the source surface brightness profile, that is,

$$S_{\lambda_i}(\mathbf{u}) \rightarrow S_{\lambda_i}(\mathbf{u}) + \delta S_{\lambda_i}(\mathbf{u}). \quad (20)$$

We want  $\delta S_{\lambda_i}(\mathbf{u})$  to absorb the image residuals introduced by the presence of substructures within the lens. To do so, the small source correction once projected to the image plane must have the form

$$\delta S_{\lambda_i}(\mathbf{y} - \nabla \phi_0(\mathbf{y})) = \nabla_{\mathbf{u}} S_{\lambda_i} \Big|_{\mathbf{u}=\mathbf{y}-\vec{\nabla}\phi_0(\mathbf{y})} \cdot \nabla \phi_{\text{sub}}(\mathbf{y}), \quad (21)$$

where we only kept terms that are first order in perturbation variables  $\{\nabla \phi_{\text{sub}}, \delta S_{\lambda_i}\}$ . To actually compute the source correction, we must project this expression back to the source plane. First, in the weak lensing regime, the lens equation  $\mathbf{u} = \mathbf{y} - \nabla \phi_0(\mathbf{y})$  has a unique solution  $\mathbf{y}(\mathbf{u})$  and projecting Eq. (21) back to the source plane is a well-defined procedure. It is then always possible to add a source correction of the form

$$\delta S_{\lambda_i}(\mathbf{u}) = \nabla_{\mathbf{u}} S_{\lambda_i}(\mathbf{u}) \cdot \nabla \phi_{\text{sub}}(\mathbf{y}(\mathbf{u})). \quad (22)$$

to “gauge” away the image residuals caused by substructures<sup>6</sup>. We note that the same logic applies in the case where  $\nabla \phi_{\text{sub}}$  is a constant vector.

However, in the strong lensing regime where the lens equation  $\mathbf{u} = \mathbf{y} - \nabla \phi_0(\mathbf{y})$  has multiple solutions and  $\nabla \phi_{\text{sub}}$  is a random nonconstant vector, the mapping between lens and source planes is no longer one-to-one. It is then no longer possible to define a unique  $\delta S_{\lambda_i}(\mathbf{u})$  that can entirely absorb the image residuals given in Eq. (16). For instance, let us imagine the case where the lens equation has two distinct solutions  $\mathbf{y}_1(\mathbf{u})$  and  $\mathbf{y}_2(\mathbf{u})$ . We can still define a source brightness correction  $\delta S_{\lambda_i}^{(1)}(\mathbf{u})$  using the first solution  $\mathbf{y}_1(\mathbf{u})$ . Adding this correction to the overall source model will indeed nullify the image residuals in the neighborhood of the first image, but will also introduce extra residuals near the second image. This argument is easily generalized to a larger number of images. The main message here is that *the redundancy introduced by having multiple images of the lensed source ensures that substructure effects can never be entirely absorbed by adding complexity to the source.*

Thus, in the strong lensing regime, the image residuals caused by the stochastic substructure in the lens (or along the line of sight) are in general not degenerate with the source surface brightness profile. Of course, for a realistic noisy image with imperfect knowledge of the PSF, it will be always possible to absorb some of the residuals by a modification of the source brightness profile (due to the specifics of the noise realization and the PSF side lobes), but the argument presented above shows that it is never possible to completely eliminate the substructure-caused residuals by changing the source.

#### E. Degeneracy with foregrounds

One might also worry that the image residuals caused by the substructure could be degenerate with foreground light, either from the lens galaxy itself, or other faint objects along the line of sight. Unfortunately, this possibility appears difficult to eliminate since there is no

<sup>6</sup> We note that if the structure of the source is well-known—as is the case for the cosmic microwave background—it is not necessarily possible to just “gauge” away the effect of the substructure, even in the weak lensing regime. See Ref. [118] for details.

a priori reason for why foregrounds could not mimic substructure-caused residuals. However, Eq. (16) has a very specific functional form linking the structure of the source and of the macro lens to that of the residuals. Since the broad structure of the lens provides information about the general configuration of the source and of the macro lens, it should in principle be possible to distinguish image residuals caused by substructures from mis-modeled foregrounds. In any case, it is unlikely in our opinion that foregrounds could exactly reproduce the mathematical structure given in Eq. (16), so they are unlikely to be completely degenerate with the impact of substructures.

We thus conclude that it is in principle possible to extract information about mass substructure within gravitational lenses by examining image residuals such as those illustrated in Fig. 2. We now turn our attention to how exactly one could extract that information from lensed images of extended sources.

#### IV. LIKELIHOOD ANALYSIS FOR THE SUBSTRUCTURE POWER SPECTRUM: EXTENDED SOURCE

In this section, we analyze how gravitationally lensed images of an extended source can be used to extract constraints on the properties of the substructure inside lens galaxies. As noted in Ref. [96], the fact that the leading-order image residuals for a lensed extended source are linearly proportional to the gradient of the substructure potential provides a straightforward way to write down a likelihood for the two-point correlation function of the substructure's projected mass density. We derive a general expression for this likelihood below.

In the following, we denote the array of parameters describing the source as  $\mathbf{q}_{\text{src}}$ , those describing the macro lens galaxy as  $\mathbf{q}_{\text{gal}}$ , those describing its environment by  $\mathbf{q}_{\text{env}}$ , those describing the foreground light by  $\mathbf{q}_{\text{fore}}$ , and those describing the properties of the instrument used to make the observations by  $\mathbf{q}_{\text{inst}}$ . For notational convenience, we gather these different sets of parameters into a single array  $\mathbf{q} = \{\mathbf{q}_{\text{src}}, \mathbf{q}_{\text{gal}}, \mathbf{q}_{\text{env}}, \mathbf{q}_{\text{fore}}, \mathbf{q}_{\text{inst}}\}$ . We also take the statistical properties of the mass substructures to be described by an array of parameters  $\mathbf{q}_{\text{sub}}$ , which may include parameters describing the amplitude and shape of the substructure convergence power spectrum. In general, the data for a given lens system will consist of a time series of pixel counts taken with different filters centered at wavelength  $\lambda_i$ , that is,

$$O_{\text{obs}} = \{O_{\text{obs}, \lambda_i}(\mathbf{x}_j, t_k)\}, \quad (23)$$

where  $t_k$  denotes the epoch of the  $k$ th observations, and  $\mathbf{x}_j$  the position of the  $j$ th pixel.

We first obtain the image residuals by subtracting from the data a model for the lensed image  $O_{\lambda}^{(0)}(\mathbf{x}_j, t_k; \mathbf{q})$  for a given choice of source, macro lens, environment, fore-

ground light, and instrumental configuration. The observed residuals at the  $j$ th pixel are then given by

$$\delta O_{\text{obs}, \lambda}(\mathbf{x}_j, t_k; \mathbf{q}) \equiv O_{\text{obs}, \lambda}(\mathbf{x}_j, t_k) - O_{\lambda}^{(0)}(\mathbf{x}_j, t_k; \mathbf{q}). \quad (24)$$

For an appropriate choice of source structure and macro lens, the observed residuals should be caused by the effect of mass substructures and instrumental noise

$$\delta O_{\text{obs}, \lambda}(\mathbf{x}_j, t_k; \mathbf{q}) = \delta O_{\text{sub}, \lambda}(\mathbf{x}_j, t_k; \mathbf{q}, \mathbf{q}_{\text{sub}}) + N_{\lambda}(\mathbf{x}_j, t_k), \quad (25)$$

where  $N_{\lambda}(\mathbf{x}_j, t_k)$  is the instrumental noise in the  $j$ th pixel and where we have explicitly written that the substructure residuals depend on the choice of source, macro lens and its environment, as well as on parameters  $\mathbf{q}_{\text{sub}}$  describing the substructure population.

#### A. Basis function expansion

To extract information from lensed images about the small-scale structure of the matter density field, it is useful to expand the substructure deflection field in terms of a set of orthonormal basis functions  $\{\nabla\varphi_l\}$ . In general, this basis should be chosen as to both maximally simplify the analysis and facilitate comparison between measurements and theoretical predictions. There are many possible choices of basis function, including including a Fourier basis, a polar harmonic expansion, or a shapelet basis [119]. To retain generality, we refrain at this point from specifying an actual orthonormal basis and write the deflection field as

$$\nabla\phi_{\text{sub}}(\mathbf{x}) = \sum_{l=1}^{N_{\text{modes}}} \mathcal{A}_l \nabla\varphi_l(\mathbf{x}) \quad (26)$$

where  $\mathcal{A}_l$  is the (usually complex) amplitude of the  $l$ th mode of the substructure deflection field. The basis functions satisfy the following orthonormality condition

$$\frac{1}{A_{\text{img}}} \int_{A_{\text{img}}} d^2\mathbf{x} \nabla\varphi_l(\mathbf{x}) \cdot \nabla\varphi_{l'}^*(\mathbf{x}) = \delta_{ll'}, \quad (27)$$

where  $A_{\text{img}}$  is the area of the sky spanned by the data. This orthogonality condition can be used to invert Eq. (26) for the mode amplitudes

$$\begin{aligned} \mathcal{A}_l &= \frac{1}{A_{\text{img}}} \int_{A_{\text{img}}} d^2\mathbf{x} \nabla\varphi_l^*(\mathbf{x}) \cdot \nabla\phi_{\text{sub}}(\mathbf{x}) \\ &= -\frac{2}{A_{\text{img}}} \int_{A_{\text{img}}} d^2\mathbf{x} \varphi_l^*(\mathbf{x}) \kappa_{\text{sub}}(\mathbf{x}), \end{aligned} \quad (28)$$

where we have used integration by parts<sup>7</sup> and the Poisson equation  $\nabla^2\phi_{\text{sub}} = 2\kappa_{\text{sub}}$  to write the last line. The

<sup>7</sup> We define our mode functions  $\varphi_l$  such that they vanish on the boundary of the integration domain, ensuring that the surface term does not contribute to Eq. (28).

residuals can then be written as

$$\delta O_{\text{sub},\lambda}(\mathbf{x}_j, t_k) = \sum_{l=1}^{N_{\text{modes}}} \mathcal{A}_l \mathcal{W}_l^\lambda(\mathbf{x}_j, t_k), \quad (29)$$

where  $\mathcal{W}_l^\lambda(\mathbf{x}, t_k)$  is the gradient of the source projected into the  $l^{\text{th}}$  mode and convolved with the PSF. For an extended source, it takes the form

$$\mathcal{W}_l^\lambda(\mathbf{x}, t_k) = -\frac{A_{\text{pix}} T_{\text{exp}}}{S_{\text{inv}}^{(\lambda)}} \int d\mathbf{y} \mathcal{W}_\lambda(\mathbf{x} - \mathbf{y}, t_k) \times \nabla \varphi_l(\mathbf{y}) \cdot \nabla_{\mathbf{u}} S_\lambda \Big|_{\mathbf{u}=\mathbf{y} - \vec{\nabla} \phi_0(\mathbf{y})}, \quad (30)$$

where we have used Eq. (13) and assumed a uniform pixel response function (the above kernel can easily be generalized for a more complex pixel response). While the  $\nabla \varphi_l(\mathbf{x})$  modes are orthogonal by construction, we note that the kernels  $\mathcal{W}_l^\lambda(\mathbf{x}, t_k)$  do not generally form an orthogonal basis of the lensing residuals due to the presence of the gradient of the source.

## B. Likelihood

It is useful at this point to introduce a matrix and vector notation that will streamline the likelihood derivation. We shall assume here a single observation at a given wavelength  $\lambda_i$ , and will therefore drop the time and wavelength indices for now. We will restore them in the next subsection. Let us gather the observed lensing residuals into a vector  $\delta \mathbf{O}_{\text{obs}}$  of length  $N_{\text{pix}}$ . Similarly, we denote the image residuals caused by substructure as  $\delta \mathbf{O}_{\text{sub}}$ . We also gather the mode basis  $\mathcal{W}_l(\mathbf{x}_j)$  into an  $N_{\text{pix}} \times N_{\text{modes}}$  matrix  $\mathbf{W}_E$ , such that  $(\mathbf{W}_E)_{jl} = \mathcal{W}_l(\mathbf{x}_j)$ . We note that  $\mathbf{W}_E$  is a nothing more than a change-of-basis matrix from the mode space to the pixel space. We finally gather the mode amplitudes  $\mathcal{A}_l$  into a vector  $\mathbf{a} \equiv \{\mathcal{A}_l\}$  of length  $N_{\text{modes}}$ . With these definitions, the image residuals caused by substructure can be simply written as

$$\delta \mathbf{O}_{\text{sub}} = \mathbf{W}_E \mathbf{a}. \quad (31)$$

Since our data consist of photon counts on pixels, we expect the noise in each pixel to have a Poisson contribution in addition to other instrumental sources such as readout noise. Since our aim is to detect the subtle effects of mass substructures, we are primarily interested in high signal-to-noise images in which the photon counts per pixel will be large, implying that we can approximate the Poisson shot noise by a Gaussian contribution. We thus assume that the noise has statistical properties entirely given by

$$\langle N(\mathbf{x}_i) N(\mathbf{x}_j) \rangle_N = \mathbf{C}_{N,ij}, \quad (32)$$

where  $\langle \dots \rangle_N$  denotes ensemble averaging over noise realizations. The likelihood for the parameters  $\mathbf{q}$  and  $\mathbf{q}_{\text{sub}}$

marginalized over the unknown amplitudes  $\mathcal{A}_l$  is then

$$\mathcal{L}(\mathbf{q}, \mathbf{q}_{\text{sub}}) \propto \int d\mathbf{a} d\mathbf{a}^\dagger \mathcal{P}_{\text{sub}}(\mathbf{a} | \mathbf{q}_{\text{sub}}) \times \frac{e^{-\frac{1}{2}(\delta \mathbf{O}_{\text{obs}} - \mathbf{W}_E \mathbf{a})^\dagger \mathbf{C}_N^{-1} (\delta \mathbf{O}_{\text{obs}} - \mathbf{W}_E \mathbf{a})}}{\sqrt{|\mathbf{C}_N|}}. \quad (33)$$

In general, the conditional probability distribution for the  $\mathcal{A}_l$  coefficients given a choice of substructure parameters  $\mathbf{q}_{\text{sub}}$ ,  $\mathcal{P}_{\text{sub}}(\mathbf{a} | \mathbf{q}_{\text{sub}})$ , is difficult to determine since it depends on rather complex galaxy formation physics. However, as we argued in Sec. II B the statistics of the substructure convergence field could be approximated as Gaussian, implying that the statistics of the  $\mathcal{A}_l$  coefficients can be approximately captured by their two-point functions. In this approximation, we have

$$\mathcal{P}_{\text{sub}}(\mathbf{a} | \mathbf{q}_{\text{sub}}) \simeq \frac{e^{-\frac{1}{2} \mathbf{a}^\dagger \mathbf{C}_{\text{sub}}^{-1} \mathbf{a}}}{\sqrt{(2\pi)^{N_{\text{modes}}} |\mathbf{C}_{\text{sub}}|}}, \quad (34)$$

where we have written the variance of the  $\mathcal{A}_l$  coefficients as

$$(\mathbf{C}_{\text{sub}})_{ll'} \equiv \langle \mathcal{A}_l \mathcal{A}_{l'}^* \rangle, \quad (35)$$

and where  $|\mathbf{C}_{\text{sub}}| = |\det \mathbf{C}_{\text{sub}}|$ . In the presence of small nonGaussianities, we note that Eq. (34) could be generalized by performing an Edgeworth expansion similar to that performed in Ref. [110]. Given the basis functions  $\nabla \varphi_l(\mathbf{x})$ , the  $\mathbf{C}_{\text{sub}}$  takes the form

$$(\mathbf{C}_{\text{sub}})_{ll'} = \frac{4}{A_{\text{img}}^2} \int \frac{d^2 \mathbf{k}}{(2\pi)^2} P_{\text{sub}}(\mathbf{k}) \tilde{\varphi}_l(\mathbf{k}) \tilde{\varphi}_{l'}^*(\mathbf{k}), \quad (36)$$

where  $\tilde{\varphi}_l(\mathbf{k})$  is the Fourier transform of  $\varphi_l(\mathbf{x})$ . It is then useful to define the following matrix

$$\mathbf{G} \equiv \mathbf{W}_E^\dagger \mathbf{C}_N^{-1} \mathbf{W}_E, \quad (37)$$

which has the convenient property of being Hermitian,  $\mathbf{G}^\dagger = \mathbf{G}$ . We note that the matrix  $\mathbf{G}_{ll'}$  is essentially the noise covariance matrix projected into the  $l$  and  $l'$  modes. We also introduce  $\mathbf{g}$ , the noise-weighted data vector projected into the mode space

$$\mathbf{g} \equiv \mathbf{W}_E^\dagger \mathbf{C}_N^{-1} \delta \mathbf{O}_{\text{obs}}, \quad (38)$$

as well as the standard  $\chi^2$  in the absence of substructure

$$\chi^2 \equiv \delta \mathbf{O}_{\text{obs}}^\dagger \mathbf{C}_N^{-1} \delta \mathbf{O}_{\text{obs}}. \quad (39)$$

With the simplifying choice given in Eq. (34) the likelihood is Gaussian in the  $\mathbf{a}$  variables, and we can thus analytically marginalize over these coefficients. Equation (33) then becomes

$$\mathcal{L}(\mathbf{q}, \mathbf{q}_{\text{sub}}) \propto \frac{e^{-\frac{1}{2}[\chi^2 - \mathbf{g}^\dagger \mathbf{D}^{-1} \mathbf{g}]}{\sqrt{|\mathbf{C}_N| |\mathbf{C}_{\text{sub}}| |\mathbf{D}|}}, \quad (40)$$

where

$$\mathbf{D} = \mathbf{G} + \mathbf{C}_{\text{sub}}^{-1}. \quad (41)$$

The matrix  $\mathbf{D}$  and the vector  $\mathbf{g}$  depend on both the source and smooth lens parameters  $\mathbf{q}$ , while the data enter through  $\chi^2$  and  $\mathbf{g}$ . The substructure parameters  $\mathbf{q}_{\text{sub}}$  only enter through the covariance matrix  $\mathbf{C}_{\text{sub}}$ .

Equation (40) is the likelihood written in the “mode” basis. It could be cast into the perhaps more familiar pixel basis by using the Woodbury matrix identity to write

$$\chi^2 - \mathbf{g}^\dagger \mathbf{D}^{-1} \mathbf{g} = \delta \mathbf{O}_{\text{obs}}^\top \mathbf{V}^{-1} \delta \mathbf{O}_{\text{obs}}, \quad (42)$$

and

$$|\mathbf{C}_N| |\mathbf{C}_{\text{sub}}| |\mathbf{D}| = |\mathbf{V}|, \quad (43)$$

where

$$\mathbf{V} \equiv \mathbf{C}_N + \mathbf{W}_E \mathbf{C}_{\text{sub}} \mathbf{W}_E^\dagger. \quad (44)$$

With this definition, the likelihood admits the simple form

$$\mathcal{L}(\mathbf{q}, \mathbf{q}_{\text{sub}}) \propto \frac{e^{-\frac{1}{2} \delta \mathbf{O}_{\text{obs}}^\top \mathbf{V}^{-1} \delta \mathbf{O}_{\text{obs}}}}{\sqrt{|\mathbf{V}|}}. \quad (45)$$

While the matrix  $\mathbf{V}$  has a simple interpretation as the “noise + signal” covariance matrix and the form of the likelihood given in Eq. (45) is rather intuitive, we note that it is often computationally advantageous to use the mode basis likelihood (Eq. (40)) since it usually involves lower dimensional matrices and vectors. Indeed, due to noise, the number of measurable modes with non-vanishing signal-to-noise ratio is usually much smaller than the number of image pixels and it is more efficient to first perform the projection into the mode basis (Eqs. (37) and (38)) before computing the likelihood. Furthermore, depending on the exact choice of mode basis, the matrix  $\mathbf{G}$  and vector  $\mathbf{g}$  can have important symmetries that significantly simplify their computation (see e.g. Sec. VB below).

### C. Generalization to an ensemble of observations

Let us now turn our attention to the case where we have a sequence of observations of the same gravitational lens taken at different time stamps  $t_k$  and/or with different filters centered at wavelength  $\lambda$ . Generalizing Eq. (33) to the case of a time series of independent observations, the likelihood marginalized over the coefficients  $\mathcal{A}_l$  for an extended source takes the form

$$\begin{aligned} \mathcal{L}(\mathbf{q}, \mathbf{q}_{\text{sub}}) &\propto \int d\mathbf{a} d\mathbf{a}^\dagger \mathcal{P}_{\text{sub}}(\mathbf{a}|\mathbf{q}_{\text{sub}}) \\ &\times \frac{e^{-\frac{1}{2} \sum_{k,\lambda} \Delta_\lambda(t_k)^\dagger \mathbf{C}_{N,\lambda}^{-1}(t_k) \Delta_\lambda(t_k)}}{\sqrt{|\tilde{\mathbf{C}}_N|}}, \end{aligned} \quad (46)$$

where

$$\Delta_\lambda(t_k) \equiv \delta \mathbf{O}_{\text{obs},\lambda}(t_k) - \mathbf{W}_{E,\lambda}(t_k) \mathbf{a}, \quad (47)$$

and where the exact form of the  $\tilde{\mathbf{C}}_N$  matrix will be given below. Note that the argument of the exponent is now summed over epochs and wavelengths, which is valid if the noise of observations taken at different epoch is uncorrelated. It is understood that each image could have its own source parameters (if the source has a different morphology at different wavelength, for instance), foreground parameters, pixelization, and PSF. On the other hand, the macro lens, its environment, and the substructure contribution to the lensing deflection are taken to be the same across all images. Taking  $\mathcal{P}_{\text{sub}}(\mathbf{a}|\mathbf{q}_{\text{sub}})$  as given in Eq. (34), we can marginalize over the  $\mathcal{A}_l$  coefficients to obtain

$$\mathcal{L}(\mathbf{q}, \mathbf{q}_{\text{sub}}) \propto \frac{e^{-\frac{1}{2} [\tilde{\chi}^2 - \tilde{\mathbf{g}}^\dagger \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{g}}]}}{\sqrt{|\tilde{\mathbf{C}}_N| |\mathbf{C}_{\text{sub}}| |\tilde{\mathbf{D}}|}}, \quad (48)$$

where

$$\tilde{\chi}^2 = \sum_\lambda \sum_{k=1}^{N_{\text{obs}}^\lambda} \chi_\lambda^2(t_k), \quad \tilde{\mathbf{g}} = \sum_\lambda \sum_{k=1}^{N_{\text{obs}}^\lambda} \mathbf{g}_\lambda(t_k), \quad (49)$$

$$\tilde{\mathbf{D}} = \sum_\lambda \sum_{k=1}^{N_{\text{obs}}^\lambda} \mathbf{G}_\lambda(t_k) + \mathbf{C}_{\text{sub}}^{-1}, \quad (50)$$

$$|\tilde{\mathbf{C}}_N| = \prod_\lambda \prod_{k=1}^{N_{\text{obs}}^\lambda} |\mathbf{C}_{N_\lambda}(t_k)|, \quad (51)$$

and where  $N_{\text{obs}}^\lambda$  is the number of exposures with a filter centered at wavelength  $\lambda$ . We note that these expressions are very similar to the single-observation case, except that the relevant quantities are now summed over all available observations.

## V. NUMERICAL IMPLEMENTATION IN A DISCRETE FOURIER BASIS

In this section, we specialize the general framework presented in the previous section to a discrete Fourier basis. We first present the specifics of the Fourier case before discussing the details of our numerical implementation.

### A. Discrete Fourier basis

We define our discrete orthonormal Fourier basis functions such that

$$\varphi_l(\mathbf{x}) = \begin{cases} \frac{e^{i\mathbf{k}_l \cdot \mathbf{x}}}{k_l} & \text{if } \mathbf{x} \in A_{\text{img}}, \\ 0 & \text{otherwise,} \end{cases} \quad (52)$$

where  $k_l = |\mathbf{k}_l|$ , and  $A_{\text{img}}$  is the sky area spanned by the data. For a  $N_{\text{pix}} \times N_{\text{pix}}$  image with side-length  $R$  ( $A_{\text{img}} = R^2$ ), the orthonormality condition given in Eq. (27) implies that the wavenumber  $\mathbf{k}_l$  must take discrete values

$$\mathbf{k}_l = \left( \frac{2\pi l_x}{R}, \frac{2\pi l_y}{R} \right), \quad (53)$$

with

$$l_x, l_y = \begin{cases} -\frac{N_{\text{pix}}-1}{2}, \dots, \frac{N_{\text{pix}}-1}{2} & (N_{\text{pix}} \text{ odd}) \\ -\frac{N_{\text{pix}}}{2} - 1, \dots, \frac{N_{\text{pix}}}{2} - 1 & (N_{\text{pix}} \text{ even}), \end{cases} \quad (54)$$

but where the zero mode with  $l_x = l_y = 0$  is not included. Here, it is understood that the mode index  $l$  is a shorthand notation for the doublet  $\{l_x, l_y\}$  characterizing the  $x$  and  $y$  coordinates of the Fourier mode. In this basis, the  $\mathcal{W}_l^\lambda(\mathbf{x}, t_k)$  kernel takes the form

$$\mathcal{W}_l^\lambda(\mathbf{x}, t_k) = -\frac{iA_{\text{pix}}T_{\text{exp}}}{\mathcal{S}_{\text{inv}}^{(\lambda)}} \int d\mathbf{y} W_\lambda(\mathbf{x} - \mathbf{y}, t_k) e^{i\mathbf{k}_l \cdot \mathbf{y}} \quad (55)$$

$$\times \hat{\mathbf{k}}_l \cdot \nabla_{\mathbf{u}} S_\lambda \Big|_{\mathbf{u}=\mathbf{y}-\vec{\nabla}\phi_0(\mathbf{y})},$$

which has the symmetry property  $\mathcal{W}_l^{\lambda*} = \mathcal{W}_{-l}^\lambda$ , a consequence of the image residuals being real. This implies that the total number of independent Fourier modes is for an  $N_{\text{pix}} \times N_{\text{pix}}$  image is

$$N_{\text{modes,ind}} = \begin{cases} \frac{N_{\text{pix}}^2 - 1}{2} & N_{\text{pix}} \text{ odd} \\ \frac{N_{\text{pix}}(N_{\text{pix}} - 1)}{2} - 1 & N_{\text{pix}} \text{ even}. \end{cases} \quad (56)$$

In practice, the size of the source and the finite width of the PSF put an upper bound on the largest  $k$ -mode that can be measured from a given lensed image, and it is therefore sometime unnecessary to consider the full range of Fourier modes given in Eq. (54). We illustrate examples of the  $\mathcal{W}_l^\lambda$  kernel in Fig. 3 for four different choices of Fourier modes. The source and lens model used to generate these are the same as in Fig. 2. Not surprisingly, we see that the various kernels pick out different features of the image residuals, with the lower wavenumbers picking structures stretching across the image, and the higher wavenumbers selecting small-scale brightness fluctuations along the lensed arc.

The Fourier transform of the mode functions given in Eq. (52) is

$$\tilde{\varphi}_l(\mathbf{k}) = \frac{A_{\text{img}}}{k_l} j_0\left(\frac{1}{2}(k_x R - 2\pi l_x)\right) j_0\left(\frac{1}{2}(k_y R - 2\pi l_y)\right), \quad (57)$$

where  $j_0(x)$  is the spherical Bessel function of order zero, and where we have written  $\mathbf{k} = (k_x, k_y)$ . Since we are focusing our attention here on the monopole of the substructure convergence power spectrum (Eq. (3)), it is useful to define the following window function

$$\Phi_{ll'}(k) = \frac{k_l k_{l'}}{(2\pi)^2 A_{\text{img}}} \int_0^{2\pi} d\theta_k \tilde{\varphi}_l(\mathbf{k}) \tilde{\varphi}_{l'}^*(\mathbf{k}), \quad (58)$$

which has the convenient normalization

$$\int_0^\infty dk k \Phi_{ll}(k) = 1. \quad (59)$$

This window function encodes the fact that we have access to only a small region of the sky, and thus can measure a limited number of Fourier modes. Now, compared to large-scale structure surveys, the window function plays a subdued role here since the gradient of the source appearing in  $\mathcal{W}_l^\lambda$  kernel (Eq. (55)) already limits the sensitivity of the data to Fourier modes with wavelength on the order of the lens' Einstein radius or smaller, independently of the size of  $A_{\text{img}}$ . Furthermore, for the modes given in Eq. (53),  $\Phi_{ll}(k)$  is strongly peaked at  $k = k_l$  while  $\Phi_{ll'}(k)$  is oscillatory for  $l \neq l'$ , hence leading to strong cancellation<sup>8</sup> of the off-diagonal elements. We can thus approximate the window function as

$$\Phi_{ll'}(k) \approx \frac{\delta(k - k_l)}{k_l} \delta_{ll'}, \quad (60)$$

which yields a  $\mathbf{C}_{\text{sub}}$  covariance matrix of the form

$$(\mathbf{C}_{\text{sub}})_{ll'} = \frac{4}{A_{\text{img}} k_l k_{l'}} \int dk k P_{\text{sub}}^{(0)}(k) \Phi_{ll'}(k)$$

$$\approx \frac{4P_{\text{sub}}^{(0)}(k_l)}{A_{\text{img}} k_l^2} \delta_{ll'}. \quad (61)$$

We note that for a constant  $P_{\text{sub}}^{(0)}(k)$  (as in the case of a population of point masses), Eq. (61) becomes exact for the diagonal elements of  $\mathbf{C}_{\text{sub}}$ . In general, as long as the value of the convergence power spectrum does not vary rapidly over the width of the window function, we find Eq. (61) to be an excellent approximation. For the remainder of this paper, we adopt for simplicity the approximation given in Eq. (61) for the substructure covariance matrix, but note that it is straightforward to generalize our calculation to also include off-diagonal elements of  $\mathbf{C}_{\text{sub}}$ .

## B. Numerical implementation

To implement and test the likelihood presented in Secs. IV B and IV C in the Fourier basis, we have developed the software package `PkLens`<sup>9</sup>. Written in pure `Python 3`, `PkLens` uses just-in-time compilation and automatic parallelization from the `numba` [120] package to accelerate key parts of the computation.

The reality condition  $\mathcal{W}_{-l}^\lambda = \mathcal{W}_l^{\lambda*}$  implies that the  $\mathbf{G}$  matrix defined in Eq. (37) can be written in the following block structure

$$\mathbf{G} = \begin{pmatrix} \mathbf{X} & \mathbf{Y} \\ \mathbf{Y}^* & \mathbf{X}^* \end{pmatrix}, \quad (62)$$

<sup>8</sup> For instance, we find that  $\int_0^\infty dk k \Phi_{ll+1} \sim 10^{-3}$ .

<sup>9</sup> `PkLens` will be made publicly available upon publication of this manuscript.

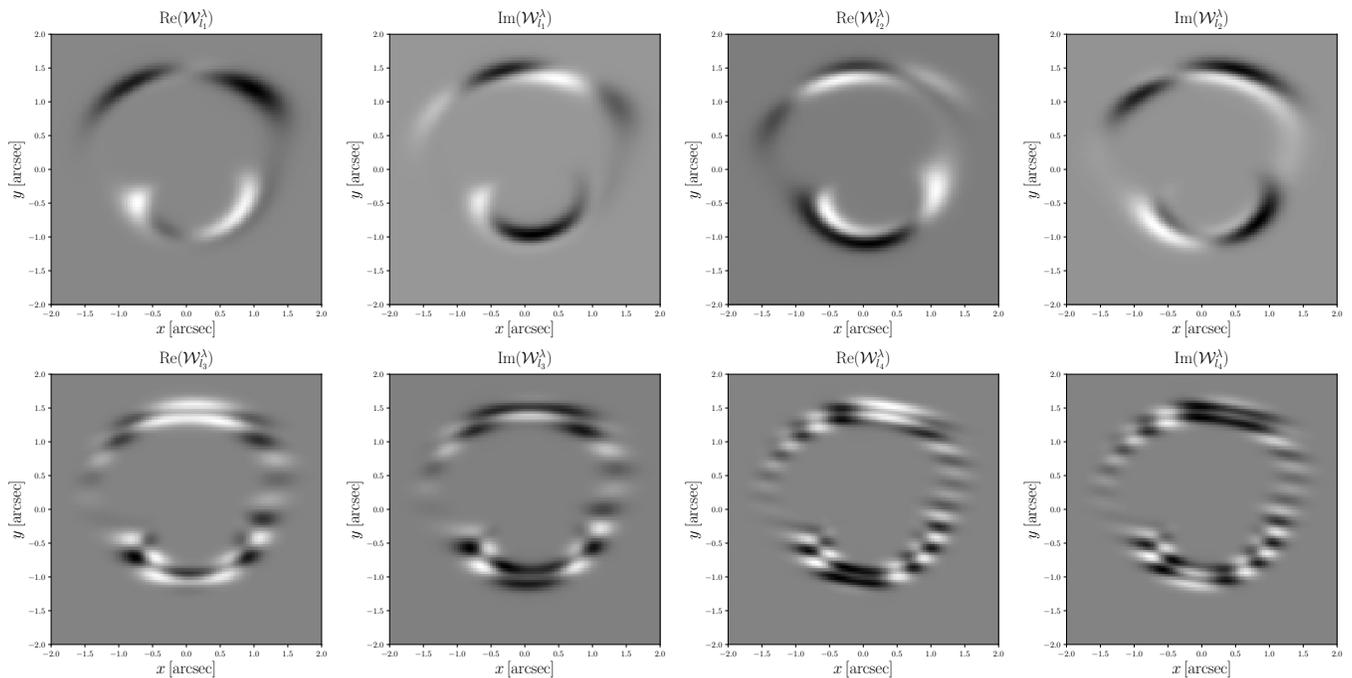


FIG. 3. Real and imaginary parts of the  $\mathcal{W}_l^\lambda$  kernel for four different Fourier modes. The kernels are ordered from long wavelength modes (top left corner) to short wavelength modes (lower right corner) with  $\{k_{l_1}, k_{l_2}, k_{l_3}, k_{l_4}\} = \{1.57, 2.22, 11.0, 19.1\}$  arcsec $^{-1}$ . The source and lens model used here are the same as in Fig. 2.

where  $\mathbf{X} = \mathbf{X}^\dagger$  is an Hermitian block and  $\mathbf{Y} = \mathbf{Y}^T$  is a symmetric block, both of size  $N_{\text{modes,ind}} \times N_{\text{modes,ind}}$ . We thus need to compute only half the elements of  $\mathbf{X}$  and half that of  $\mathbf{Y}$  (for a total of  $N_{\text{modes,ind}}$  entries) to fully characterize the matrix  $\mathbf{G}$ . This structure of the  $\mathbf{G}$  matrix allows us to use blockwise inversion in order to compute the matrix  $\mathbf{D}^{-1}$  appearing in the likelihood given in Eq. (40), hence significantly speeding up the linear algebra. Similarly, only half of the  $\mathbf{g}_l$  vector entries need to be computed since  $\mathbf{g}_{-l} = \mathbf{g}_l^*$ .

## VI. FISHER ANALYSIS

To develop some intuition about the sensitivity of different lens configurations and observational scenarios to the substructure convergence power spectrum, it is instructive to first carry out a simple Fisher analysis of the likelihood given in Eq. (48). We adopt a binned substructure convergence power spectrum as our fitting model, and the relevant parameters here are thus the logarithm of the amplitude of  $P_{\text{sub}}(k)$  within each bin,  $\mathbf{q}_{\text{sub}} = \{\ln P_{\text{sub},i}\}_{i=1,\dots,N_{\text{bins}}}$ . For the analysis shown in this section, we divide the range of scales probed by a given lensed image into 4 wavenumber bins that are evenly spaced in  $\log_{10}(k)$ . In the following, for each filter centered at wavelength  $\lambda$ , we assume that we have  $N_{\text{obs}}^\lambda$  observations of the same lens.

### A. Fisher matrix and sensitivity function

The Fisher matrix for the binned log amplitude of the power spectrum takes the form

$$F_{ij} \equiv - \left\langle \frac{\partial^2 \ln \mathcal{L}}{\partial \ln P_{\text{sub},i} \partial \ln P_{\text{sub},j}} \right\rangle = \frac{P_{\text{sub},i} P_{\text{sub},j}}{2} \text{Tr} \left[ \mathbf{\Gamma} \frac{\partial \mathbf{C}_{\text{sub}}}{\partial P_{\text{sub},i}} \mathbf{\Gamma} \frac{\partial \mathbf{C}_{\text{sub}}}{\partial P_{\text{sub},j}} \right], \quad (63)$$

where

$$\mathbf{\Gamma} \equiv (\mathbf{G}^{-1} + \mathbf{C}_{\text{sub}})^{-1} = \mathbf{G} - \mathbf{G}\mathbf{D}^{-1}\mathbf{G}. \quad (64)$$

To understand how the Fisher matrix scales with the observational parameters, it is instructive to consider a simple example where we neglect the off-diagonal entries of the  $\mathbf{G}$  matrix. In this case, the diagonal entries of the Fisher matrix admit the form

$$F_{ii} = \frac{1}{2} \sum_{l \in i} \frac{(\mathcal{S}_l P_{\text{sub},i})^2}{(1 + \mathcal{S}_l P_{\text{sub},i})^2}, \quad (65)$$

where the sum runs over all Fourier modes whose magnitude falls within the range of the  $i$ th bin. We have assumed here that the noise for each observation is Poissonian with  $\mathbf{C}_{N_\lambda, ij} = \delta_{ij} \sigma_1 O_\lambda(\mathbf{x}_i)$  (where  $O_\lambda(\mathbf{x}_i)$  is given in Eq. (13) and  $\sigma_1 = 1$  for pure Poisson noise). In Eq. (65), we have introduced the sensitivity  $\mathcal{S}_l$  of a given gravitational lens observation to the  $l$ th mode of the substructure convergence field. It is defined as the product

of a mode-independent prefactor  $Q_{\text{obs}}^\lambda$  that depends on the depth and quality of the observation and of a mode-dependent function  $\Lambda_l^\lambda$  that only depends on the spatial structure of the macrolens, source, and PSF,

$$S_l = \sum_\lambda Q_{\text{obs}}^\lambda \times \Lambda_l^\lambda, \quad (66)$$

where

$$Q_{\text{obs}}^\lambda \equiv N_{\text{obs}}^\lambda \frac{T_{\text{exp}} \mathcal{F}_\lambda}{\sigma_1 \mathcal{S}_{\text{inv}}^{(\lambda)}}, \quad (67)$$

and

$$\Lambda_l^\lambda \equiv \frac{4}{k_l^2 N_{\text{pix}}} \sum_m \frac{|(W_\lambda * \nabla \varphi_l \cdot \nabla_{\mathbf{u}} \hat{S}_\lambda)(\mathbf{x}_m)|^2}{W_\lambda * \hat{S}_\lambda(\mathbf{x}_m)}, \quad (68)$$

where the sum runs over all the pixels in the image, and the “\*” symbol stands for the convolution operation. Note that we have written the source surface brightness as  $S_\lambda(\mathbf{u}) = \mathcal{F}_\lambda \hat{S}_\lambda(\mathbf{u})$ , where  $\mathcal{F}_\lambda$  is the total source flux within the bandpass of the filter and  $\int d^2\mathbf{u} \hat{S}_\lambda(\mathbf{u}) = 1$ . For simplicity, we have omitted the foreground contribution when writing Eq. (68).

Since  $\Lambda_l^\lambda$  describes the intrinsic sensitivity of a given lens configuration to the  $l$ th mode of the substructure density field, it is a useful figure of merit to rapidly assess whether a given lens can provide competitive constraints on the substructure convergence power spectrum. The dimensionless prefactor  $Q_{\text{obs}}^\lambda$  simply captures how the sensitivity  $S_l$  scales with exposure time, number of observations, source flux, noise level, and detector sensitivity. Not surprisingly, the sensitivity is improved for a longer total exposure, a brighter source, a lower noise level, and by lowering the value of  $\mathcal{S}_{\text{inv}}^{(\lambda)}$  (which could be done by using a larger telescope and/or a more sensitive camera).

For very large value of the sensitivity,  $P_{\text{sub},i} S_l \gg 1$ , Eq. (65) implies that the measurement uncertainty  $\delta \ln P_{\text{sub},i} = \sqrt{(F_{ii}^{-1})}$  on the amplitude of the binned power spectrum becomes sample variance dominated with

$$\delta \ln P_{\text{sub},i} \simeq \frac{\sqrt{2}}{\sqrt{N_i}}, \quad (P_{\text{sub},i} S_l \gg 1) \quad (69)$$

where  $N_i$  is the number of modes within the  $i$ th bin. This is a familiar result that arises for instance in the study of cosmological large-scale structure. On the other hand, for low sensitivity  $P_{\text{sub},i} S_l \ll 1$ , we can Taylor expand Eq. (65) to obtain

$$\delta \ln P_{\text{sub},i} \simeq \frac{\sqrt{2}}{P_{\text{sub},i} \sqrt{\sum_{l \in i} S_l^2}}, \quad (P_{\text{sub},i} S_l \ll 1) \quad (70)$$

from which we obtain  $\delta P_{\text{sub},i} \propto 1/(N_{\text{obs}} T_{\text{exp}})$ .

To gain some intuition about the structure of the sensitivity function as a function of wavenumber  $k_l$ , let us

consider a Gaussian source of width  $\sigma_s$  lensed by a singular isothermal sphere lens model into an Einstein ring with radius  $b_{\text{ein}}$ . Also, let us consider a series of observations with a single filter and a Gaussian PSF of size  $\sigma_{\text{PSF}}$  (we shall drop the wavelength index  $\lambda$  in the following). For this simple system, Eq. (68) admits as leading behavior in the limit that  $\sigma_{\text{PSF}}, \sigma_s < b_{\text{ein}}$

$$\Lambda_l \propto b_{\text{ein}} \frac{\sigma_{\text{PSF}}^2 + \sigma_s^2 + k_l^2 \sigma_{\text{PSF}}^4}{k_l^2 (\sigma_{\text{PSF}}^2 + \sigma_s^2)^{5/2} A_{\text{img}}} e^{-\frac{1}{2} k_l^2 \sigma_{\text{eff}}^2}, \quad (71)$$

where

$$\sigma_{\text{eff}} = \sqrt{2} \frac{\sigma_{\text{PSF}} \sigma_s}{\sqrt{\sigma_{\text{PSF}}^2 + \sigma_s^2}}. \quad (72)$$

We compare this analytical estimate of  $\Lambda_l$  to exact numerical computations for three different choices of PSF size in Fig. 4. There, the solid lines shows the numerical results, while the dashed lines display the approximate expression given in Eq. (71). The dotted vertical line shows the wavenumber  $k_s$  corresponding to the size of the Gaussian source.

On scales larger than  $\sigma_{\text{eff}}$  ( $k_l \lesssim \pi/\sigma_{\text{eff}}$ ), the sensitivity decays as  $\Lambda_l \propto k_l^{-2}$ . This scaling is simply the result of the Poisson equation linking the deflection field probed by the lensing observations to the substructure convergence whose power spectrum we are trying to measure, that is,  $|\tilde{\alpha}_{\text{sub}}(k)|^2 \sim 4|\tilde{\kappa}_{\text{sub}}(k)|^2/k^2$ . Since the substructure deflection field couples to the gradient of the source’s surface brightness profile convolved with the PSF (see Eq. (16)), the sensitivity becomes strongly suppressed for Fourier modes probing scales smaller than either the size of the source or the PSF, whichever is smallest. This can be seen in Fig. 4 where the examples with  $\sigma_{\text{PSF}} > \sigma_s$  (purple and green lines) have rapidly decaying sensitivity for  $k_l \gtrsim k_s$ , whereas the example with  $\sigma_{\text{PSF}} < \sigma_s$  (red lines) roughly retains the  $k_l^{-2}$  scaling until wavenumbers corresponding to the PSF size. While this sensitivity cutoff is exponential for our example with a Gaussian source and PSF, we generally expect it to be milder for more realistic choices of source and PSF models.

In addition to the sensitivity cutoff for  $k_l \gtrsim \pi/\sigma_{\text{eff}}$ , Eq. (71) admits the general scaling  $\Lambda_l \propto b_{\text{ein}}/\sigma_{\text{PSF}}^3$  for PSF size larger than the source. The leading factor of  $b_{\text{ein}}/\sigma_{\text{PSF}}$  essentially counts the number of independent sections of the Einstein ring that are available for the analysis. This is quite intuitive: at fixed PSF and image size, a larger Einstein ring contains more information about the substructure density field than a smaller one. The further  $1/\sigma_{\text{PSF}}^2$  factor stems from the reduced signal to noise per pixel as the PSF size increases. Indeed, for larger PSFs, the light from a given part of the source gets spread to a larger area of the focal plane, resulting in a smaller photon count in each pixel and thus to a reduced sensitivity. Of course, this scaling breaks down once the PSF size becomes smaller than the typical size of the source, as can be seen in Eq. (71). As a general rule of thumb, high-resolution images will always provide better

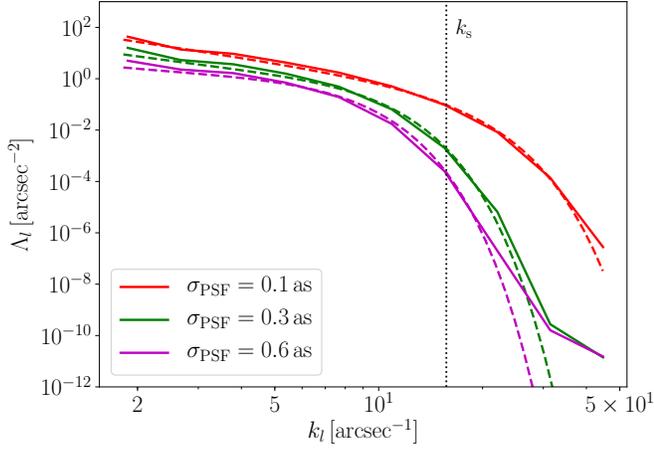


FIG. 4. Sensitivity function  $\Lambda_l$  for a Gaussian source of size  $\sigma_s = 0.2$  arcsec lensed by a singular isothermal sphere lens model with Einstein radius  $b_{\text{ein}} = 1.2$  arcsec. The solid lines show the results of exact numerical computations for three different sizes of Gaussian PSF, while the dashed lines illustrate the approximate expression given in Eq. (71). The vertical dotted line shows the approximate wavenumber corresponding to the size of the source  $k_s = \pi/\sigma_s$ . Note that higher values of  $\Lambda_l$  means that the mock lens has greater sensitivity to the substructure convergence power spectrum for wavenumber  $k_l$ .

constraining power on the substructure power spectrum than low-resolution images, and lenses with larger Einstein radii (or more complete Einstein rings) will generally have display greater sensitivity to the effect of substructure.

### B. Simple Fisher forecast

We now use the Fisher matrix from Eq. (63) to quantitatively estimate the error on the binned substructure convergence power spectrum for two different observational scenarios. As usual, these Fisher forecasts should interpreted with caution, especially since we neglect here possible covariances between the effect of substructures and changes to the macrolens and source parameters. These possible degeneracies will be explored in the next section once we performed complete Markov Chain Monte Carlo analyses of lensed images. *The results shown here should thus be taken as illustrative of the best-case sensitivity to the substructure power spectrum that could be achieved within the observational scenarios we consider below.* In the following, we consider images on a  $50 \times 50$  pixel grid, and adopt the truncated NFW substructure power spectrum shown in Fig. 1 as our fiducial model for the Fisher forecast.

Figure 5 shows the Fisher forecast for a high-resolution observations with a PSF full width at half maximum (FWHM) of 0.07 arcsec, similar to what is achievable with the Hubble Space Telescope (HST) for optical wavelengths. We display Fisher estimates of the error bars

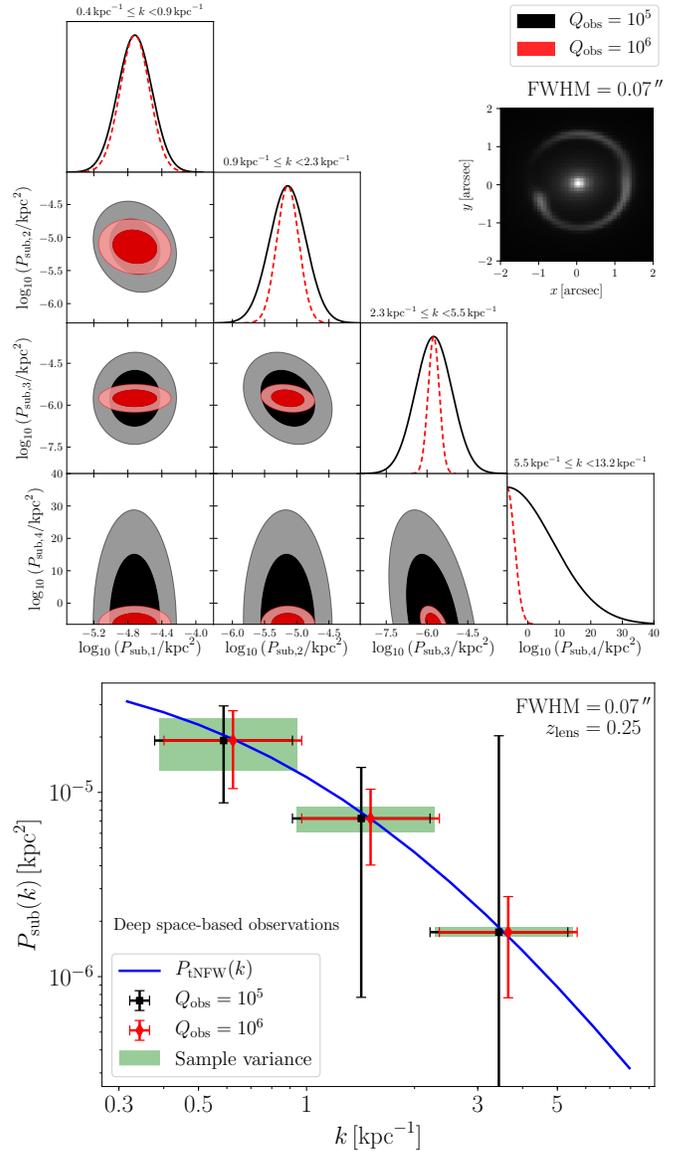


FIG. 5. Fisher forecast for the substructure convergence power spectrum in four logarithmic wavenumber bins. The top panel shows the different projection of the inverse Fisher matrix for an HST-like observations with a PSF FWHM of 0.07 arcsec for two values of the quality factor  $Q_{\text{obs}}$  (defined in Eq. (67)). The darkly and lightly shaded areas show the 68% and 95% confidence regions, respectively. The inset shows the lens configuration used. The lower panel shows the resulting Fisher error bars on  $P_{\text{sub}}(k)$  for the three lowest wavenumber bins. The blue solid line shows the fiducial power spectrum model used in the forecast, which corresponds to the truncated NFW model shown in Fig. 1. The error bars show the  $1\text{-}\sigma$  regions, while the green rectangles display the sample variance contribution within each bin. For clarity, the wavenumber bin center for each observational scenario shown have been offset by 6%, with the green rectangle showing the true wavenumber bin used in the analysis.

for two choices of the observation quality factor  $Q_{\text{obs}}$ .

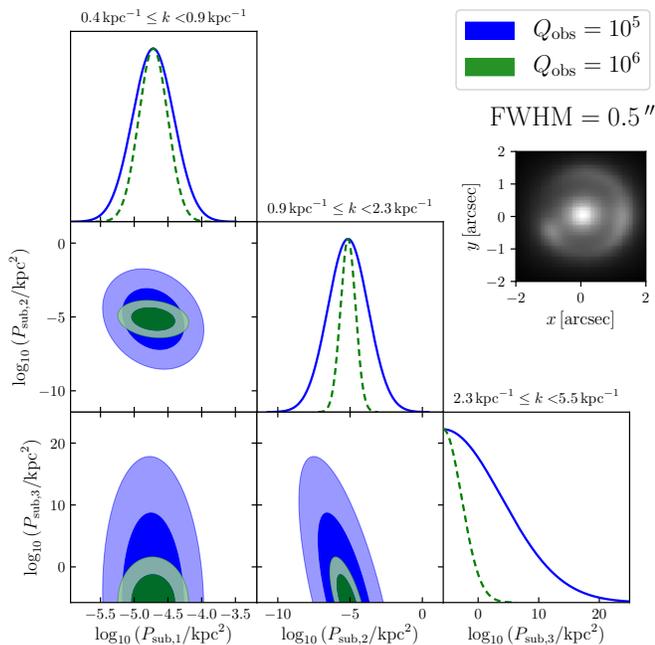


FIG. 6. Same as the upper panel of Fig. 5 but for seeing-limited ground-based observations with  $\text{FWHM} = 0.5$  arcsec (modeled here as a Moffat profile). Here, we only display the results for the wavenumber range  $0.4\text{kpc}^{-1} \leq k \leq 5.5\text{kpc}^{-1}$  since there is very little sensitivity to higher wavenumbers given the poor image resolution. Note the larger range power spectrum value shown on the axes as compared to Fig. 5 due to the much lower sensitivity here. The inset shows the lens configuration used.

As a concrete example,  $Q_{\text{obs}} = 10^5$  could be achieved by combining 10 observations of a source with total flux  $\mathcal{F}_\lambda = 10^{-18}$  erg/cm<sup>2</sup>/s/Å (approximately corresponding to an unlensed AB magnitude of 24), each observed through the FW555 filter for  $T_{\text{exp}} = 2000$  seconds with the UVIS detector on Wide-Field Camera 3 (WFC3) aboard HST.

For  $Q_{\text{obs}} = 10^6$ , the substructure convergence power spectrum error bars approach the sample-variance limit within the lowest wavenumber bin, as shown in the lower panel of Fig. 5. For the second and third wavenumber bins, the error bars grow modestly according to Eq. (65) for both image depths shown. There is then a significant decrease in sensitivity within the highest wavenumber bin due to the limits imposed by the source size and PSF (corresponding here to  $k_s \simeq 8\text{kpc}^{-1}$  and  $k_{\text{FWHM}} \simeq 11\text{kpc}^{-1}$ , respectively), in accordance with our discussion above. For the wavenumber range  $2.3\text{kpc}^{-1} \leq k \leq 13.2\text{kpc}^{-1}$ , the deeper observations ( $Q_{\text{obs}} = 10^6$ ) lead to a significant gain in sensitivity as compared to the shallower one ( $Q_{\text{obs}} = 10^5$ ) due to these modes being in the noise-dominated regime described by Eq. (70). Within the highest wavenumber bin, only coarse upper limits appears possible in either case, which is why we do not illustrate them in the lower panel of Fig. 5.

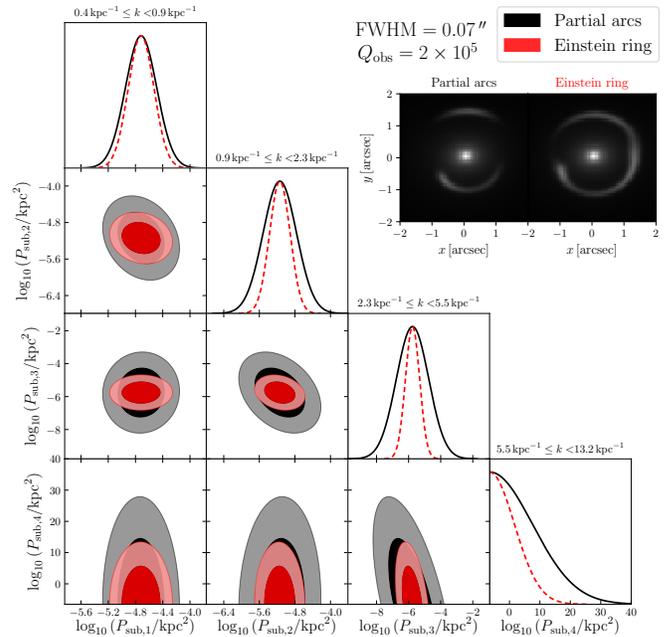


FIG. 7. Similar to the upper panel of Fig. 5 but comparing the Fisher forecasts for two different lens configurations (shown in the inset). Here, the structure of the source, the macrolens, its environment and the foregrounds are all kept the same, and only the source position is changed. The nearly complete Einstein ring displays greater sensitivity to the substructure power spectrum than the partial arcs.

Figure 6 displays Fisher forecast for seeing-limited ground-based observations with PSF  $\text{FWHM} = 0.5$  arcsec for the same two values of  $Q_{\text{obs}}$ . The error bars are significantly larger in this case, with the two highest wavenumber bins only yielding very coarse limits. While a measurement of the power spectrum amplitude for the lowest wavenumber bin appears possible, we caution that the degeneracies between the effects of substructure and the structure of the lens, source, and foregrounds are likely to be quite severe for the low resolution ground-based data. This will likely degrade the constraining power of the low-resolution data compared to the simple Fisher forecast illustrated in Fig. 6, unless a high-resolution image is also available to help break the various degeneracies.

It is important to realize that, at equal value of the observational quality factor, distinct lens configurations will display different sensitivity to the substructure power spectrum. An example of this is illustrated in Fig. 7 where we show the figure forecast for two different lens configurations: a set of partial arcs and a nearly complete Einstein ring (the latter being the configuration used in the forecast above). For these images, the structure of the source, macrolens, its environment and foregrounds are all kept fixed and only the position of the source is modified. Due to its greater coverage of the image plane that allows it to probe more substructure modes, the

Einstein ring displays greater sensitivity to the substructure power spectrum in all wavenumber bins shown. As discussed in Sec. VIA, at a given image resolution and depth, lensed images covering a larger area of the image plane will generally display greater sensitivity to the power spectrum.

## VII. ANALYSIS OF SIMULATED IMAGES

Having developed some intuition about the various factors affecting the sensitivity of a gravitational lens observation to the substructure convergence power spectrum, we now turn our attention to more realistic analyses of mock images. We explore in this section how the posterior distribution of the binned substructure power spectrum amplitudes is affected by degeneracies with the macro lens, source, foreground, and noise parameters for a few representative observational scenarios. Details on how we generate our mock observations of gravitational lenses are given in Appendix B. Null tests to confirm the accuracy of our numerical implementation, and a study of possible degeneracies between the substructure power spectrum amplitudes and the noise parameters are presented in Appendix C. We present here the results of our Markov Chain Monte Carlo (MCMC) analyses exploring the complete degeneracies between lens, source, foreground, and noise parameters on the one hand, and substructure power spectrum on the other.

### A. Substructure power spectrum inference

As in the case of our Fisher analysis, we adopt the logarithm of the binned substructure convergence power spectrum amplitudes as our fitting model. We divide the range of scales probed by a given lensed image into 3 wavenumber bins that are evenly spaced in  $\log_{10}(k)$  within the range<sup>10</sup>  $0.4 \text{ kpc}^{-1} \leq k \leq 5.5 \text{ kpc}^{-1}$ . We note that the choice of binning presented in this work is arbitrary and driven only by convenience and simplicity. We leave a thorough study of the optimal binning strategy to future work, but note that it is very likely to depend on the specifics of each dataset. We adopt a broad scale-free (log-uniform) prior on the amplitude within each bin,  $\log_{10}(P_{\text{sub},i}/\text{arcsec}^2) \in [-10, -1]$ .

We use the affine invariant sampler `emcee` [121] to sample the likelihood given in Eq. (40), allowing the macrolens, source, foreground, and noise parameters to vary freely within broad flat priors (see Appendix B for details on the different components entering the mock

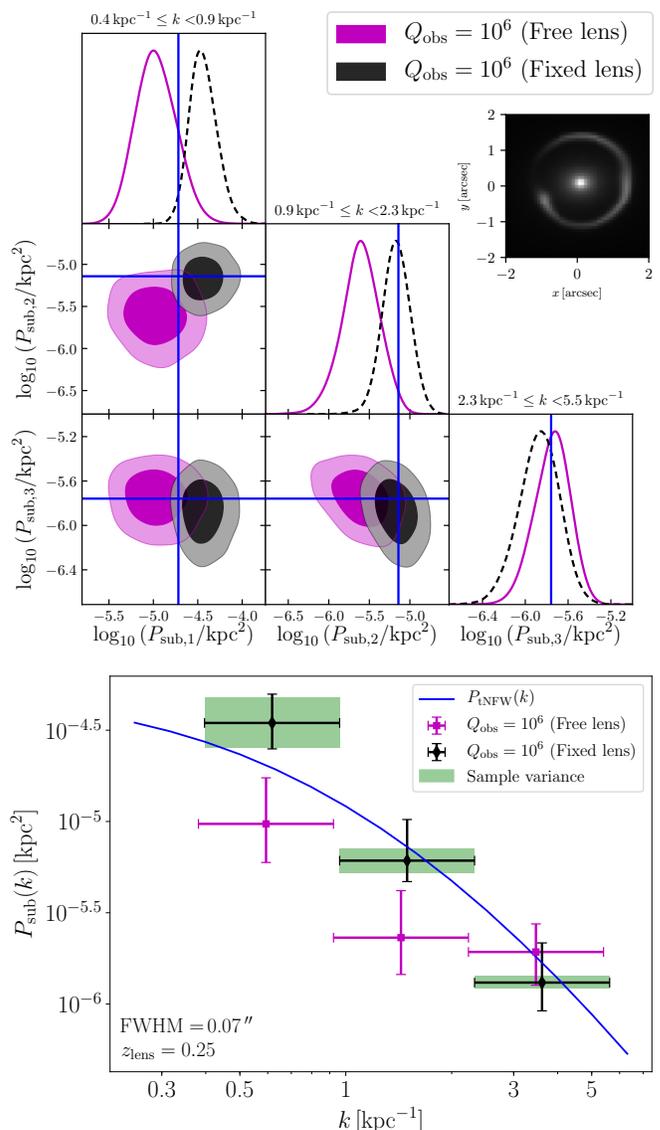


FIG. 8. Posterior summary for the binned substructure convergence power spectrum amplitudes for the case of a nearly complete Einstein ring observed with quality factor  $Q_{\text{obs}} = 10^6$  and an HST-like PSF with a FWHM of 0.07 arcsecond. The top panel shows the marginalized confidence intervals when the macro lens, source, foreground, and noise parameters are either held fixed (dashed black) or allowed to vary freely in the MCMC chains (magenta). The dark (light) contours show the 68% (95%) confidence intervals. The solid blue lines show the average values of the substructure power spectrum used to generate the mock image within each bin. The inset shows the mock image configuration used in the analysis. The lower panel shows the corresponding error bars in the substructure power spectrum space. The blue solid line shows the substructure power spectrum (corresponding to the truncated NFW model shown in Fig. 1) used to generate the mock image. The two-sided error bars show the 68% highest posterior density intervals, and the green rectangles show the sample contribution to the error bars within each bin.

<sup>10</sup> Since we are considering lenses at redshift  $z_{\text{lens}} = 0.25$  with a Planck 2015 Cosmology [106], this interval corresponds to angular scales in the range  $1.6 \text{ arcsec}^{-1} \leq k \leq 22.2 \text{ arcsec}^{-1}$ .

lensed images). To determine how much the macro-model, source, noise, and foregrounds can somewhat reabsorb the effects of the substructure, we also run MCMC chains that keep the parameters of these latter components fixed to their true values. For each mock image and case considered, we use twice as many `emcee` walkers as the number of free parameters in the computation. As illustrative examples, we perform the substructure power spectrum inference for the two lens configurations used for the Fisher forecast shown in Fig. 7: the nearly complete Einstein ring, and the two partial lensed arcs, each observed with a quality factor of  $Q_{\text{obs}} = 10^6$  and an HST-like PSF with FWHM = 0.07 arcsec.

The resulting posterior distribution summary for the case of the nearly complete Einstein ring is shown in Fig. 8. The top panel shows the marginalized confidence intervals of the substructure power spectrum for the three wavenumber bins. The black lines and contours denote the confidence intervals when the macrolens, source, foreground, and noise parameters are held fixed at their true values, while the magenta contours and lines show the intervals when all parameters are allowed to freely vary. The solid blue lines denote the average values of the substructure convergence power spectrum within each wavenumber bin of the model used to generate the mock images. The lower panel shows the resulting 68% highest posterior density (HPD) intervals for each bin, on top of the truncated NFW power spectrum used to generate the mock data. The green rectangles illustrate the contribution from sample variance within each bin.

We first note that even in the case of a fixed macro lens, source, and foregrounds (black contours), there are fluctuations of the highest posterior values with respect to the true input values, with the scatter being larger at low wavenumbers. This is expected given that we are looking at a single realization of the substructure convergence field. The relatively small number of substructure modes within the lower wavenumber bins gives rise to a significant lens-to-lens variation on the inferred value of the power spectrum at larger scales. However, our test shows that averaging over a sufficient number of substructure realizations leads to an unbiased estimate of the substructure power spectrum amplitude within each bin. Of course, while this is easy to do for mock data, performing this average over substructure realizations with real data will require carefully combining the measurements from different lens systems, taking into account how their respective substructure population depends on the lens galaxy's properties.

We see in Fig. 8 that even after the macro lens, source, and foregrounds are allowed to reabsorb some of the surface brightness features caused by the substructure, the mock data still retain significant sensitivity to the underlying convergence power spectrum. To some extent, the results shown here are still somewhat optimistic since the model used to fit the data is the same as the model used to generate the mock data. Nevertheless, it is encouraging that the inferred confidence regions on the substructure

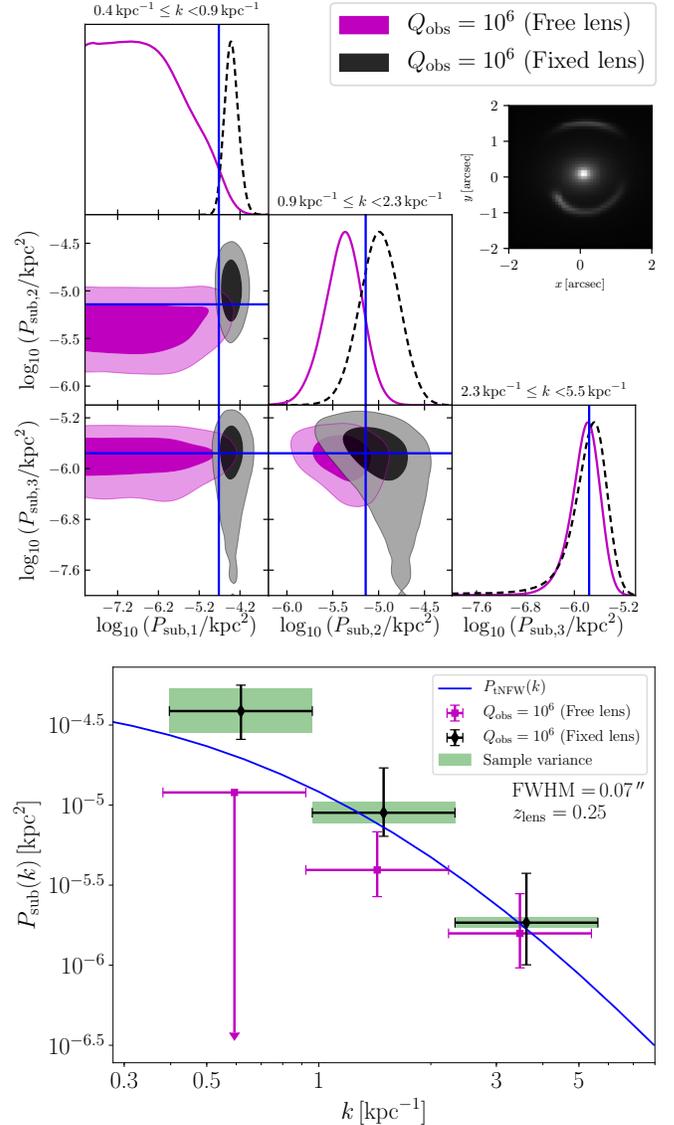


FIG. 9. Similar to Fig. 8 but for a different lens configuration made of two partial lensed arcs (see inset in top panel). In the lower panel, the black diamond error bars show the highest posterior density (HPD) intervals for  $Q_{\text{obs}} = 10^6$  while keeping the macro lens, source, foreground, and noise parameters fixed. The magenta square error bars show the HPD intervals for  $Q_{\text{obs}} = 10^6$  allowing all macro lens, source, foreground, and noise parameters to vary within the MCMC chains. The one-sided error bars show the 95% HDP upper limits.

structure convergence power spectrum do not change dramatically in size once the possible degeneracies with the macro lens, source, and foregrounds are fully taken into account. The highest posterior values of the binned power spectrum amplitudes do shift with respect to the true input values, implying that the macro lens, source, and foregrounds are absorbing some of the substructure effects. This reabsorption is more important on the larger scales (low wavenumber) where it is easier for these components

to compensate for small surface brightness fluctuations caused by the substructure. In any case, the resulting biases on the binned power spectrum amplitude are modest, and the true input values are recovered within the 95% confidence intervals. Overall, we observe little covariances among the different binned power spectrum amplitudes, in agreement with our Fisher estimates presented in Sec. VIB.

Figure 9 shows the corresponding posterior summary for the case of the partial lensed arcs. Much of the same discussion as above applies here except that in this case, the macro lens, source, and foregrounds can reabsorb a significant fraction of the substructure effect within the lowest wavenumber bin. This is caused by the less redundant nature of this lens configuration (which essentially shows two distinct images of the extended source) as compared to the Einstein ring, which makes it more difficult to break the degeneracy between the macro lens, source, and foregrounds on the one hand, and the effects of the substructure on the other.

### VIII. GENERALIZATION TO COMPACT TIME-DEPENDENT SOURCES

We now turn our attention to time-dependent compact (unresolved) sources, such as quasars, and consider how they can be used to probe the substructure within lens galaxies. As outlined in Sec. III C, the first-order perturbative approach used to compute the lensing residuals for extended sources fails in the case of compact sources since the surface brightness gradient of such object can become very large, hence leading to large second-order corrections to the lensing residuals. A different approach is thus required to treat these compact sources, which we summarize here.

#### A. Lensing of time-dependent compact sources

We assume that the compact source admits the following structure in the source plane

$$S_\lambda(\mathbf{u}, t_s) = s_\lambda(t_s) \delta(\mathbf{u} - \mathbf{u}_c), \quad (73)$$

where  $s_\lambda(t_s)$  is the luminosity of the compact source at source-plane time  $t_s$  and wavelength  $\lambda/(1+z_{\text{src}})$ , and  $\mathbf{u}_c$  is the position of the compact source. If the compact source is not exactly point-like, the delta function in Eq. (73) could be replaced with the actual source profile, with only minor modifications to the remainder of the calculation. We can substitute the above into Eq. (12) to obtain the model for the lensed compact source

$$\hat{O}_\lambda(\mathbf{x}, t) = \int d\mathbf{y} W_\lambda(\mathbf{x} - \mathbf{y}, t) s_\lambda(t - \tau(\mathbf{y})) \times \delta(\mathbf{y} - \vec{\nabla} \phi_{\text{lens}}(\mathbf{y}) - \mathbf{u}_c). \quad (74)$$

We recognize the argument of the Dirac delta function as the lens equation for the compact source, which allows us to immediately write

$$\delta(\mathbf{y} - \vec{\nabla} \phi_{\text{lens}}(\mathbf{y}) - \mathbf{u}_c) = \sum_{j=1}^{N_{\text{img}}} \mu_j \delta(\mathbf{y} - \mathbf{y}_j), \quad (75)$$

where  $\mathbf{y}_j$  and  $\mu_j = |\det \mathbf{M}(\mathbf{y}_j)|$  are the location and the magnification of the  $j$ th image, respectively. Using the delta functions to perform the  $\mathbf{y}$  integration in Eq. (74), our model for the lensed image of the time-dependent compact source takes the form

$$\hat{O}_\lambda(\mathbf{x}, t) = \sum_{j=1}^{N_{\text{img}}} \mu_j s_\lambda(t - \tau(\mathbf{y}_j)) W_\lambda(\mathbf{x} - \mathbf{y}_j, t). \quad (76)$$

The substructure potential  $\phi_{\text{sub}}$  enters Eq. (76) through the magnification  $\mu_j$ , the image position  $\mathbf{y}_j$ , and the arrival time  $\tau(\mathbf{y}_j)$ . In order to write down an explicit expression for Eq. (76), one would need to solve the full lens equation  $\mathbf{y} - \vec{\nabla} \phi_{\text{lens}}(\mathbf{y}) - \mathbf{u}_c = 0$  to determine the image positions for every possible choice of source position  $\mathbf{u}_c$ . This is a daunting task since the solution depends non linearly on the substructure potential  $\phi_{\text{sub}}$ , whose statistics we are trying to probe in the first place.

Our strategy to handle this difficulty is to treat the quantities  $\{\mu_j, \mathbf{y}_j, \Delta t_{j1} \equiv \tau(\mathbf{y}_1) - \tau(\mathbf{y}_j)\}$  as free nuisance parameters whose values are entirely driven by the data and an appropriate choice of priors. As we discuss below, the impact of the substructure on the image positions and arrival time delays can be *exactly* captured by a series of constraint equations. It is important to note, however, that the impact of the substructure on the magnification of point images is often obscured by the presence of stellar microlensing and dust absorption. While magnification information of point images has been used in the literature to constrain the presence of substructure within lens galaxies, we adopt here a conservative point-of-view and treat the  $\mu_j$ s as free parameters independent of  $\phi_{\text{sub}}$ .

The remaining ingredient entering Eq. (76) is the time dependence of the source  $s_\lambda(t_s)$ . Here, we choose to characterize the source variability as a Gaussian process. Within this framework, the source function  $s_\lambda(t - \tau(\mathbf{y}_j))$  is explicitly written in terms of a fluctuation vector  $\delta \mathbf{s}_\lambda = \{\{s_{\lambda,j}(t_k) - \hat{s}_\lambda\}_{j=1 \dots N_{\text{img}}}\}_{k=1 \dots N_{\text{obs}}^\lambda}$  describing the brightness of the source at each observation time  $t_k$  and for each lensed image. Here,  $\hat{s}_\lambda$  is the time-averaged compact source luminosity. Essentially,  $\delta \mathbf{s}_\lambda$  contains the fluctuation light curve for each observed lensed image of the compact source. This vector is drawn from a Gaussian distribution given by

$$\mathcal{P}_{\text{src}}(\delta \mathbf{s}_\lambda) = \frac{1}{\sqrt{(2\pi)^{N_s} |\mathbf{C}_{\text{src}}|}} e^{-\frac{1}{2} \delta \mathbf{s}_\lambda^T \mathbf{C}_{\text{src}}^{-1} \delta \mathbf{s}_\lambda}, \quad (77)$$

where  $N_s$  is the total number of elements in the vector  $\delta \mathbf{s}_\lambda$ , which is  $N_s = N_{\text{img}} \times N_{\text{obs}}^\lambda$  because we reconstruct

the source flux from every image at every epoch. The exact structure of the covariance matrix  $\mathbf{C}_{\text{src}}$  depends on the type of Gaussian process used to describe the variability of the compact source. A few possible choices have been considered in the literature [122–126] including the damped random walk, the powered exponential, and the Matérn covariance function. Our substructure analysis derived in the following sections is general enough to accommodate any choice of covariance function. We note that in general the covariance matrix  $\mathbf{C}_{\text{src}}$  itself depends on the values of the time delays between images. Taking into account the conversion to counts within pixels (Eq. (13)), our model for the residuals of the lensed images of a time-dependent compact source observed at time  $t_k$  is thus

$$\delta O_\lambda(\mathbf{x}, t_k) = \frac{A_{\text{pix}} T_{\text{exp}}}{\mathcal{S}_{\text{inv}}^{(\lambda)}} \sum_{j=1}^{N_{\text{img}}} \mu_j \delta s_{\lambda,j}(t_k) W_\lambda(\mathbf{x} - \mathbf{y}_j, t_k), \quad (78)$$

where  $\delta s_{\lambda,j}(t_k) \equiv s_{\lambda,j}(t_k) - \hat{s}_\lambda$ .

### B. Constraint equations for compact sources

As discussed in the previous section, the time-dependent source introduces nuisance parameters corresponding to the usual “reduced observables” (image position, magnification, and relative time delay) that are often used to characterize lensed images of quasars. The key point here is that these nuisance parameters depend on the substructure coefficients  $\mathcal{A}_l$  through the lens equation. This leads to additional constraints on the relative amplitudes of the different substructure modes. To see this, we start by writing the lens equation for the compact source in the presence of mass substructures

$$\mathbf{u}_c = \mathbf{y}_j - \nabla \phi_0(\mathbf{y}_j) - \sum_{l=1}^{N_{\text{modes}}} \mathcal{A}_l \nabla \varphi_l(\mathbf{y}_j), \quad (79)$$

where  $\mathbf{y}_j$  is the position of the  $j$ th image. We can use the  $j = 1$  image to eliminate the source position in Eq. (79) in order to obtain  $2(N_{\text{img}} - 1)$  constraint equations for the  $\mathcal{A}_l$  coefficients

$$\left\{ \begin{aligned} & \sum_{l=1}^{N_{\text{modes}}} \mathcal{A}_l [\nabla \varphi_l(\mathbf{y}_j) - \nabla \varphi_l(\mathbf{y}_1)] \\ & = \mathbf{y}_j - \mathbf{y}_1 + \nabla \phi_0(\mathbf{y}_1) - \nabla \phi_0(\mathbf{y}_j) \end{aligned} \right\}_{j=2, \dots, N_{\text{img}}} \quad (80)$$

An additional set of constraint equations can be obtained by looking at the arrival time delay. Using Eq. (7), the

arrival time delay for the  $j$ th image is

$$\begin{aligned} \tau_j &= t_0 \left[ \frac{1}{2} |\mathbf{y}_j - \mathbf{u}_c|^2 - \phi_0(\mathbf{y}_j) - \sum_{l=1}^{N_{\text{modes}}} \mathcal{A}_l \varphi_l(\mathbf{y}_j) \right] \\ &= t_0 \left[ \frac{1}{2} (|\mathbf{u}_c|^2 - |\mathbf{y}_j|^2) - \mathcal{D} \phi_0(\mathbf{y}_j) \right. \\ & \quad \left. - \sum_{l=1}^{N_{\text{modes}}} \mathcal{A}_l \mathcal{D} \varphi_l(\mathbf{y}_j) \right], \end{aligned} \quad (81)$$

where we have defined the projection operator

$$\mathcal{D} \equiv 1 - \mathbf{y} \cdot \nabla, \quad (82)$$

and where we have used Eq. (79) to eliminate the term  $\mathbf{u}_c \cdot \mathbf{y}_j$  in going from the first to the second line. We note that the operator  $\mathcal{D}$  effectively ensures that only the gauge invariant part of the projected potential contributes to the arrival time delay<sup>11</sup>. As before, we can use the  $j = 1$  image to eliminate the source position from Eq. (81) in order to obtain  $N_{\text{img}} - 1$  additional constraints on the amplitudes  $\mathcal{A}_l$

$$\left\{ \begin{aligned} & \sum_{l=1}^{N_{\text{modes}}} \mathcal{A}_l \mathcal{D} [\varphi_l(\mathbf{y}_j) - \varphi_l(\mathbf{y}_1)] = \frac{1}{2} (|\mathbf{y}_1|^2 - |\mathbf{y}_j|^2) \\ & + \mathcal{D} \phi_0(\mathbf{y}_1) - \mathcal{D} \phi_0(\mathbf{y}_j) - \frac{\Delta t_{1j}}{t_0} \end{aligned} \right\}_{j=2, \dots, N_{\text{img}}} \quad (83)$$

For each choice of nuisance parameters  $\{\mathbf{y}_1, \{\mathbf{y}_j, \Delta t_{1j}\}_{j=2, \dots, N_{\text{img}}}\}$  describing the images of the compact source, Eqs. (80) and (83) form  $3(N_{\text{img}} - 1)$  linear equations for  $N_{\text{modes}}$  unknown amplitudes  $\mathcal{A}_l$ . Schematically, this system takes the form

$$\mathbf{L} \mathbf{a} = \mathbf{b}, \quad (84)$$

where again  $\mathbf{a} = \{\mathcal{A}_l\}$ , and where the structure of the matrix  $\mathbf{L}$  and vector  $\mathbf{b}$  can be read off Eqs. (80) and (83). It is important to emphasize that the constraints enforced by Eq. (84) are *exact* and do not necessitate a perturbative expansion in the small substructure potential  $\phi_{\text{sub}}$ . Essentially, the presence of a compact time-dependent source restricts the values of the  $\{\mathcal{A}_l\}$  coefficients to lie on a hyperplane in the parameter space. In general the number of constraints is much smaller than the number of measurable modes and the matrix  $\mathbf{L}$  is therefore not invertible. In the following, we enforce this constraint by multiplying the likelihood by a Gaussian factor

$$\mathcal{P}_c(\mathbf{a}|\mathbf{b}) \propto \frac{e^{-\frac{1}{2}(\mathbf{L} \mathbf{a} - \mathbf{b})^\dagger \boldsymbol{\Sigma}^{-1} (\mathbf{L} \mathbf{a} - \mathbf{b})}}{\sqrt{|\boldsymbol{\Sigma}|}}, \quad (85)$$

and taking the limit  $\boldsymbol{\Sigma} \rightarrow 0$  at the end of the calculation. Conveniently, this also allows us to relax the constraints by taking  $\boldsymbol{\Sigma} \rightarrow \infty$ .

<sup>11</sup> Indeed, the term  $\mathcal{D} \phi$  is invariant under the gauge transformation  $\phi \rightarrow \phi + \mathbf{c} \cdot \mathbf{x}$ , where  $\mathbf{c}$  is a constant vector.

### C. Likelihood

Having established the structure of these constraints, we can now write down the joint likelihood for the parameters  $\mathbf{q}$ ,  $\mathbf{q}_{\text{sub}}$ , and the nuisance parameters  $\mathbf{p}_{\mathbf{Q}} \equiv \{\{\mu_j, \mathbf{y}_j, \Delta t_{1j}\}_{j=1, \dots, N_{\text{img}}}\}$  describing the properties of the lensed images of the compact source. As in Sec. IV B, it is useful to use a matrix and vector notation to simplify the derivation of the likelihood. We denote the data vector as  $\delta \mathbf{O}_{\text{obs}}$ , which we take to contain the image residuals for all pixels from the  $N_{\text{obs}}$  observations. It thus has a total length of  $N_d \equiv N_{\text{pix}} \times N_{\text{obs}}$ . As before, the image residuals caused by substructure are denoted as  $\delta \mathbf{O}_{\text{sub}} = \mathbf{W}_{\text{E}} \mathbf{a}$ , but where the  $\mathbf{W}_{\text{E}}$  matrix now has dimensions of  $N_d \times N_{\text{modes}}$ . The contribution from the lensed compact source is denoted as  $\delta \mathbf{O}_{\text{C}} = \mathbf{W}_{\text{C}} \delta \mathbf{s}$ , where  $\delta \mathbf{s}$  is the compact source vector defined in Sec. VIII A of length  $N_s = N_{\text{img}} \times N_{\text{obs}}$ , and  $\mathbf{W}_{\text{C}}$  is a matrix of size  $N_d \times N_s$  which maps  $\mathbf{s}$  to the observable space as given in Eq. (78).

With these definitions, our model for the observed image residuals takes the form

$$\delta \mathbf{O}_{\text{obs}} = \mathbf{W}_{\text{E}} \mathbf{a} + \mathbf{W}_{\text{C}} \delta \mathbf{s} + \mathbf{N}, \quad (86)$$

where  $\mathbf{N}$  is a noise vector of length  $N_d$ . As before, we are mainly interested in high signal-to-noise observations and we thus take the noise to have Gaussian statistical properties specified by

$$\langle \mathbf{N} \mathbf{N} \rangle_N = \mathbf{C}_N, \quad (87)$$

where  $\mathbf{C}_N$  is now the  $N_d \times N_d$  total noise covariance matrix for all observations. Since we expect the noise from different observations to be uncorrelated, the  $\mathbf{C}_N$  matrix will generally have a block diagonal structure. The likelihood then takes the form

$$\begin{aligned} \mathcal{L} \propto & \int d\mathbf{a} d\mathbf{a}^\dagger d\delta \mathbf{s} \mathcal{P}_{\text{c}}(\mathbf{a}|\mathbf{b}) \mathcal{P}_{\text{sub}}(\mathbf{a}) \mathcal{P}_{\text{src}}(\delta \mathbf{s}) \\ & \times \frac{e^{-\frac{1}{2}(\delta \mathbf{O}_{\text{obs}} - \mathbf{W}_{\text{E}} \mathbf{a} - \mathbf{W}_{\text{C}} \delta \mathbf{s})^\dagger \mathbf{C}_N^{-1} (\delta \mathbf{O}_{\text{obs}} - \mathbf{W}_{\text{E}} \mathbf{a} - \mathbf{W}_{\text{C}} \delta \mathbf{s})}}{\sqrt{|\mathbf{C}_N| |\mathbf{C}_{\text{src}}| |\mathbf{C}_{\text{sub}}| |\boldsymbol{\Sigma}|}}, \end{aligned} \quad (88)$$

where we have explicitly written down the marginalization over the substructure amplitudes specified by the vector  $\mathbf{a}$ , the source light curve fluctuations  $\delta \mathbf{s}$ . The constraints from Eq. (84) are implemented using the  $\mathcal{P}_{\text{c}}$  factor. Using the expressions for  $\mathcal{P}_{\text{src}}(\delta \mathbf{s})$  and  $\mathcal{P}_{\text{sub}}(\mathbf{a}|\mathbf{q}_{\text{sub}})$  given in Eqs. (77) and (34), respectively, and defining the vectors

$$\mathbf{v} = \begin{bmatrix} \mathbf{a} \\ \delta \mathbf{s} \end{bmatrix}, \quad \mathbf{d} = \begin{bmatrix} \delta \mathbf{O}_{\text{obs}} \\ \mathbf{b} \end{bmatrix}, \quad (89)$$

and the block matrices

$$\mathbf{C}_s^{-1} = \begin{bmatrix} \mathbf{C}_{\text{sub}}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{C}_{\text{src}}^{-1} \end{bmatrix}, \quad (90)$$

$$\mathbf{A} = \begin{bmatrix} \mathbf{C}_N^{-1} & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Sigma}^{-1} \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} \mathbf{W}_{\text{E}} & \mathbf{W}_{\text{C}} \\ \mathbf{L} & \mathbf{0} \end{bmatrix}, \quad (91)$$

we can write the likelihood as

$$\begin{aligned} \mathcal{L} \propto & \int d\mathbf{v} d\mathbf{v}^\dagger \times \\ & \frac{e^{-\frac{1}{2} \left( \begin{bmatrix} \mathbf{d}^\dagger & \mathbf{v}^\dagger \end{bmatrix} \begin{bmatrix} \mathbf{A} & -\mathbf{A}\mathbf{B} \\ -\mathbf{B}^\dagger \mathbf{A} & \mathbf{C}_s^{-1} + \mathbf{B}^\dagger \mathbf{A}\mathbf{B} \end{bmatrix} \begin{bmatrix} \mathbf{d} \\ \mathbf{v} \end{bmatrix} \right)}}{\sqrt{|\mathbf{C}_N| |\mathbf{C}_{\text{sub}}| |\mathbf{C}_{\text{src}}| |\boldsymbol{\Sigma}|}}. \end{aligned} \quad (92)$$

Using standard partial Gaussian integration, the marginalization over the  $\mathbf{v}$  and  $\mathbf{v}^\dagger$  vectors can be performed to yield

$$\mathcal{L} \propto \frac{e^{-\frac{1}{2} \mathbf{d}^\dagger \mathbf{M}^{-1} \mathbf{d}}}{\sqrt{|\mathbf{M}|}}, \quad (93)$$

where

$$\begin{aligned} \mathbf{M} & \equiv (\mathbf{A} - \mathbf{A}\mathbf{B}(\mathbf{C}_s^{-1} + \mathbf{B}^\dagger \mathbf{A}\mathbf{B})^{-1} \mathbf{B}^\dagger \mathbf{A})^{-1} \\ & = \mathbf{A}^{-1} + \mathbf{B}\mathbf{C}_s\mathbf{B}^\dagger, \end{aligned} \quad (94)$$

where we have used the Woodbury matrix identity in the last line. Since only  $\boldsymbol{\Sigma}$  (and not its inverse) appears within the covariance matrix  $\mathbf{M}$  determining the likelihood given in Eq. (93), it is well-defined to take the limit  $\boldsymbol{\Sigma} \rightarrow 0$  to enforce the constraints from the image positions and time delays of the compact time-dependent source.

Equation (93) is the likelihood written in the ‘‘pixel’’ basis. As explained in Sec. IV B, it is often computationally advantageous to first perform the projection into the basis formed by the modes of the substructure deflection field since it reduces the dimensions of the matrices that need to be inverted. In this ‘‘mode’’ basis, the likelihood takes the form (after taking the limit  $\boldsymbol{\Sigma} \rightarrow 0$ )

$$\mathcal{L} \propto \frac{e^{-\frac{1}{2}(\tilde{\chi}^2 - \mathbf{w}^\dagger \mathbf{U}^{-1} \mathbf{w} + (\mathbf{b} - \boldsymbol{\omega})^\dagger \boldsymbol{\Upsilon}^{-1} (\mathbf{b} - \boldsymbol{\omega}))}}{\sqrt{|\mathbf{C}_N| |\mathbf{C}_{\text{sub}}| |\mathbf{C}_{\text{src}}| |\mathbf{U}| |\boldsymbol{\Upsilon}|}}, \quad (95)$$

where

$$\mathbf{U} = \begin{bmatrix} \tilde{\mathbf{D}} & \mathbf{K} \\ \mathbf{K}^\dagger & \mathbf{F} \end{bmatrix}, \quad \mathbf{w} = \begin{bmatrix} \tilde{\mathbf{g}} \\ \mathbf{h} \end{bmatrix}, \quad (96)$$

and where  $\tilde{\chi}^2$ ,  $\tilde{\mathbf{g}}$ , and  $\tilde{\mathbf{D}}$  are defined in Eqs. (49) and (50). The other quantities appearing in Eq. (95) are

$$\mathbf{F} = \mathbf{C}_{\text{src}}^{-1} + \mathbf{W}_{\text{C}}^\dagger \mathbf{C}_N^{-1} \mathbf{W}_{\text{C}}, \quad (97)$$

$$\mathbf{K} = \mathbf{W}_{\text{E}}^\dagger \mathbf{C}_N^{-1} \mathbf{W}_{\text{C}}, \quad (98)$$

$$\mathbf{h} = \mathbf{W}_{\text{C}}^\dagger \mathbf{C}_N^{-1} \delta \mathbf{O}_{\text{obs}}, \quad (99)$$

and

$$\boldsymbol{\omega} = \mathbf{L}\mathbf{S}^{-1}\tilde{\mathbf{g}} \quad \text{and} \quad \boldsymbol{\Upsilon} = \mathbf{L}\mathbf{S}^{-1}\mathbf{L}^\dagger, \quad (100)$$

where  $\mathbf{S}$  is the Schur complement of the matrix  $\mathbf{F}$

$$\mathbf{S} = \tilde{\mathbf{D}} - \mathbf{K}\mathbf{F}\mathbf{K}^\dagger. \quad (101)$$

The likelihood given in Eq. (95) has a similar structure as that occurring for an extended source only (Eq. (48)) except for the term in the exponent proportional to  $\Upsilon^{-1}$  which encodes the constraints from the compact source. Both  $\omega$  and  $\Upsilon$  have sizes determined by the number of linear compact source constraints, which for an  $N_{\text{img}}$  lens will be at most  $3(N_{\text{img}} - 1)$ , a small number indeed.

#### D. Compact sources and power spectrum sensitivity

We wrap up this section by considering how the linear constraints imposed by the presence of a compact source can improve the constraints on the substructure convergence power spectrum. For this purpose, it is useful to isolate the part of the covariance matrix  $\mathbf{M}$  that depends on  $\mathbf{C}_{\text{sub}}$  by writing it as

$$\mathbf{M} = \mathbf{M}_0 + \mathbf{Q}\mathbf{C}_{\text{sub}}\mathbf{Q}^\dagger, \quad (102)$$

where

$$\mathbf{M}_0 = \begin{bmatrix} \mathbf{C}_N + \mathbf{W}_C \mathbf{C}_{\text{src}} \mathbf{W}_C^\dagger & \mathbf{0} \\ \mathbf{0} & \mathbf{\Sigma} \end{bmatrix}, \quad \mathbf{Q} = \begin{bmatrix} \mathbf{W}_E \\ \mathbf{L} \end{bmatrix}. \quad (103)$$

Interestingly, we observe that the intrinsic variability of the compact source encoded in  $\mathbf{C}_{\text{src}}$  essentially acts as an additional source of noise for the pixels where its images form. On the other hand, the constraints on the modes of the substructure deflection field encoded in the projection operator  $\mathbf{L}$  only appears in the second term of Eq. (102). Given the Gaussian structure of the likelihood given in Eq. (93), the Fisher matrix for the binned logarithmic amplitude of the substructure power spectrum takes the form

$$\begin{aligned} \mathbf{F}_{ij} &= \frac{1}{2} \text{Tr} \left[ \mathbf{M}^{-1} \frac{\partial \mathbf{M}}{\partial \ln P_{\text{sub},i}} \mathbf{M}^{-1} \frac{\partial \mathbf{M}}{\partial \ln P_{\text{sub},j}} \right] \\ &= \frac{1}{2} \text{Tr} \left[ \mathbf{Q}^\dagger \mathbf{M}^{-1} \mathbf{Q} \frac{\partial \mathbf{C}_{\text{sub}}}{\partial \ln P_{\text{sub},i}} \mathbf{Q}^\dagger \mathbf{M}^{-1} \mathbf{Q} \frac{\partial \mathbf{C}_{\text{sub}}}{\partial \ln P_{\text{sub},j}} \right] \\ &= \frac{P_{\text{sub},i} P_{\text{sub},j}}{2} \text{Tr} \left[ \mathbf{\Gamma} \frac{\partial \mathbf{C}_{\text{sub}}}{\partial P_{\text{sub},i}} \mathbf{\Gamma} \frac{\partial \mathbf{C}_{\text{sub}}}{\partial P_{\text{sub},j}} \right], \end{aligned} \quad (104)$$

where

$$\mathbf{\Gamma} = \left[ (\mathbf{Q}^\dagger \mathbf{M}_0^{-1} \mathbf{Q})^{-1} + \mathbf{C}_{\text{sub}} \right]^{-1}. \quad (105)$$

We first note that Eq. (105) exactly reduces to Eq. (64) in the limit where the source has no compact component ( $\mathbf{C}_{\text{src}} \rightarrow 0$  and  $\mathbf{\Sigma} \rightarrow \infty$ ).

In the remainder of this section, we will focus our attention on the case where a compact source is present, but where we have only a single deep observations of the lens, which corresponds to the limit  $\mathbf{C}_{\text{src}} \rightarrow 0$  and  $\mathbf{\Sigma} \rightarrow 0$ . In that case, the  $\mathbf{\Gamma}$  matrix takes the form

$$\mathbf{\Gamma} = \left[ \tilde{\mathbf{G}}^{-1} + \mathbf{C}_{\text{sub}} - \tilde{\mathbf{G}}^{-1} \mathbf{L}^\dagger (\mathbf{L} \tilde{\mathbf{G}}^{-1} \mathbf{L}^\dagger)^{-1} \mathbf{L} \tilde{\mathbf{G}}^{-1} \right]^{-1}. \quad (106)$$

The two first terms in the square brackets are the same as in the purely extended source case, while the third term encodes the constraints from the presence of the compact source.

We compare in Fig. 10 the Fisher forecasts on the binned substructure power spectrum amplitudes with and without the additional constraints encoded in the  $\mathbf{L}$  operator, for the case of a single deep observation of a lens (i.e. no time delay constraints). We assume a cusp-like lens configuration where 4 images of the lensed quasars are present in addition to a partial Einstein ring (as illustrated in the insets of Fig. 10). These highly-idealized forecasts keep the structure of the source, foregrounds, the macrolens and its environment fixed to their true values. As such, they should be cautiously interpreted as a best-case scenario for the additional sensitivity that lensed images of compact sources can bring to the substructure power spectrum measurements, in the absence of time delay measurements.

The left panel displays the Fisher forecast comparison for a high-resolution image (FWHM = 0.07 arcsec) where we observe that the added quasar constraints do not significantly modify the constraints, except in the highest wavenumber bin where the improvement is modest. The right panel shows the Fisher forecast comparison for a seeing-limited deep observation (FWHM = 0.5 arcsec) of the same lens. In this case, it is assumed that we have an independent measurement of the quasar image positions (from, e.g., a shallow HST observation) since the low resolution nature of the image used for the power spectrum analysis likely precludes such a measurement. The improvement to the binned power spectrum constraints are slightly more significant in this case, with the largest gain in the two highest wavenumber bins. This reflects the fact that lensed images of compact sources are more sensitive to smaller-scale perturbations than the extended image studied in the rest of the paper.

In general, the presence of a compact source can improve the bounds on the substructure convergence power spectrum at large wavenumbers. The small relative number of constraints arising from a typical four-image quasar lens (6 from image positions, with 3 additional constraints possible if time delays are known) compared to the total number of distinct substructure modes that can be probed (usually several hundreds) implies that the gain in sensitivity is modest. Nevertheless, adding the constraints from the quasar images could help break potential degeneracies between the macro lens and the effects of the substructure, so a further investigation is required to determine their full impact. Also, for lenses with no extended source component, precise time delays

Here, we have treated the magnifications of the observed compact source images as free parameters. If microlensing or dust extinction (if relevant) could be properly understood and modeled, the substructure information contained in the brightness of the quasar images could also be harnessed by multiplying the likelihood given in Eq. (93) by an extra (nonGaussian) function

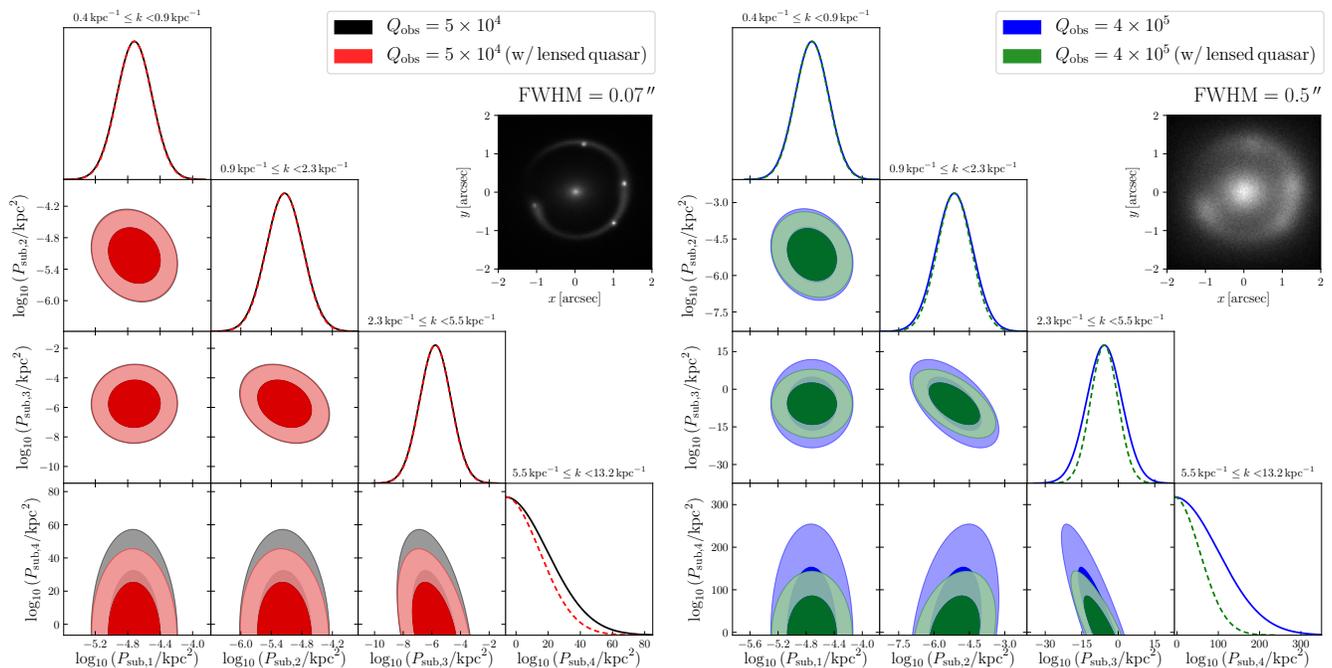


FIG. 10. Fisher forecast comparison for the binned substructure convergence power spectrum with and without the additional constraints from having a quadruply-imaged lensed quasar within the image. Here, we use 4 wavenumber bins in the range  $0.4 \text{ kpc}^{-1} \leq k \leq 13.2 \text{ kpc}^{-1}$ , and only use the 6 constraints from the quasar image positions (no time delay constraints). We adopt the truncated NFW substructure power spectrum shown in Fig. 1 as our fiducial model. The left panel shows the comparison for a high-resolution image (FWHM = 0.07 arcsec) with an observational quality factor of  $Q_{\text{obs}} = 5 \times 10^4$ . The right panel shows the Fisher forecast comparison for seeing-limited (FWHM = 0.5 arcsec) observations with a quality factor of  $Q_{\text{obs}} = 4 \times 10^5$ . The lensed image configuration used in each case are shown in the insets.

describing the impact of substructure on the  $\mu_j$  parameters. Given the highly nonlinear nature of the magnification perturbations caused by substructure [109], deriving such a function likely requires extensive forward simulations [65, 67, 75].

## IX. CONCLUSIONS

In this paper, we have derived a mode function-based approach to extract statistical information about the projected substructure density field in proximity to strongly lensed images of high redshift sources. Focusing on two-point statistics, we have derived a likelihood for the substructure convergence power spectrum, given pixelated images of gravitationally-lensed extended sources. We have implemented this likelihood within the software package `PkLens` and have performed simple Fisher forecasts to assess the sensitivity of different observational scenarios. Using simple lens, source, foreground models, we have explored the possible degeneracies these latter components and the collective effect of the substructure. We have finally generalized our power spectrum likelihood to take into account the presence of compact time-dependent lensed sources such as quasars within the observed images.

Not too surprisingly, our results indicate that deep high-resolution images provide the best sensitivity to the substructure power spectrum, up to scales approximately corresponding to the smallest observable feature of the background source. At fixed image resolution and pixelization, lenses with larger Einstein radii or displaying a larger fraction of a complete Einstein ring provide better sensitivity to the collective effect of the substructure. We generally find that substructure perturbations with smaller wavenumbers are more likely to be reabsorbed by changes in the macro lens and source models, compared to those with larger wavenumbers. This can bias low the inferred amplitude of the substructure power spectrum on these scales. We leave to future work a detailed study of this potential bias and of possible techniques to address it. Finally, we find that for lenses containing both extended images and multiply-imaged quasars (such as RX J1131-1231 [127]), the extra constraints provided by the quasar image positions provide a modest improvement on the substructure power spectrum constraints, with most of the gain in sensitivity being concentrated at the smallest scales. While we have not considered the additional constraints coming from the relative time delays [128], it is likely that they could provide a slight improvement to the substructure sensitivity as compared to those shown here, while at the same time helping break

some degeneracies in the lens modeling.

Since one of the goals of this paper was to develop intuition about the different factors influencing the sensitivity to the substructure power spectrum, and not to perform detailed lens modeling (see e.g. Ref. [129]), we have focused here on simple parametric lens and source models. The next step is to embed the numerical infrastructure developed here within a lens modeling framework that allows for more flexible source and foreground structures (see e.g. Refs. [130–132]), as well as dynamical PSF reconstruction. Given that the deflections caused by the substructure couple primarily to the gradient of the source, it could be argued that the forecasts based on simple sources presented here are conservative since more complex sources are likely to display enhanced sensitivity to substructure due to their greater spatial variability. It is however possible that this improvement is somewhat offset by the larger number of parameters needed to accurately describe the source. We study this tradeoff between source complexity and substructure constraints in an upcoming publication, in which we apply our formalism to actual observational data.

In this work, we have used a binned substructure power spectrum as our fitting model, but the framework we developed here (see Eq. (36)) is completely general and can easily accommodate parametric models of  $P_{\text{sub}}(\mathbf{k})$  such as a power laws [102], as well as anisotropic power spectra. While we have used the substructure power spectrum for a population of truncated NFW subhalos [98] within the main lens as our fiducial model throughout this manuscript, we note that our derived likelihood is sensitive to the overall collective effect of the substructure between the lensed source and the observer, including line-of-sight structures either in front or behind the main lens, and baryonic structures such as globular clusters and giant molecular clouds. In order to eventually extract constraints on dark matter physics from substructure power spectrum measurements, robust predictions about the contribution from these latter objects to the power spectrum would have to be computed. Similarly, a more thorough study of the impact of nonGaussianities on the inferred power spectrum would have to be performed.

Eventually, deriving robust constraints on the Universe’s small-scale structure from lensing power spectrum measurements will require combining the results from multiple lens systems. This will require a detailed study of how the inferred substructure power spectrum (including the line-of-sight contribution) depends on the properties (redshift, mass, concentration, environment, etc.) of the lens galaxy (see e.g. Ref. [133]). Alternatively, given the large number of gravitational lenses that will become known in the next decade [134], one could imagine building carefully selected samples that share similar properties and should thus have comparable substructure populations. Combining lens systems will allow the reduction of sample variance uncertainties on the substructure power spectrum at larger scales and, ultimately, provide

a key test of the cold dark matter paradigm.

## ACKNOWLEDGMENTS

F.-Y. C.-R. wishes to thank the Kavli Institute for Theoretical Physics for their hospitality during the completion of this work. We thank the Aspen Center for Physics, which is supported by National Science Foundation (NSF) grant PHY-1607611, for their hospitality during the initial stages of this work. This research was supported in part by the NSF under Grant No. PHY-1748958. F.-Y. C.-R. acknowledges the support of the National Aeronautical and Space Administration (NASA) ATP grant NNX16AI12G at Harvard University. C. R. K. acknowledges the support of NSF grant AST-1716585. Part of the research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with NASA.

*Software:* The following software packages were used in this work: `Matplotlib` [135], `Scipy` [136], `Numpy` [137], `astropy` [138], `numba` [120], `emcee` [121], `Getdist` [139], `h5py` [140], and `corner` [141].

## Appendix A: Creating a random realization with a given power spectrum

To create a realization of the convergence field with a given power spectrum, we first generate a map where each pixel is drawn from a standard normal distribution, that is,

$$\langle \kappa_{ij} \kappa_{kl} \rangle = \delta_{ik} \delta_{jl}, \quad \langle \kappa_{ij} \rangle = 0, \quad (\text{A1})$$

where  $\kappa_{ij}$  is the value of the convergence at the pixel with indices  $ij$ , with the first index referring to the  $x$ -axis and the second referring to the  $y$ -axis. We note that this is a white noise map. The discrete Fourier transform of  $\kappa_{kl}$  is given by

$$\tilde{\kappa}_{mn} = A_{\text{pix}} \sum_{k=0}^{N_x-1} \sum_{l=0}^{N_y-1} \kappa_{kl} e^{-2\pi i \frac{km}{N_x}} e^{-2\pi i \frac{ln}{N_y}}, \quad (\text{A2})$$

where  $A_{\text{pix}}$  is the area of one pixel,  $N_x$  is the number of pixels along the  $x$ -axis, and  $N_y$  is the number of pixels along the  $y$ -axis. The power spectrum of  $\tilde{\kappa}_{mn}$  is given by

$$\begin{aligned} \langle \tilde{\kappa}_{mn}^* \tilde{\kappa}_{m'n'} \rangle &= A_{\text{pix}}^2 \sum_{k=0}^{N_x-1} \sum_{l=0}^{N_y-1} \sum_{k'=0}^{N_x-1} \sum_{l'=0}^{N_y-1} \langle \kappa_{kl} \kappa_{k'l'} \rangle \\ &\quad \times e^{2\pi i \frac{km}{N_x}} e^{2\pi i \frac{ln}{N_y}} e^{-2\pi i \frac{k'm'}{N_x}} e^{-2\pi i \frac{l'n'}{N_y}} \\ &= A_{\text{pix}}^2 \sum_{k=0}^{N_x-1} \sum_{l=0}^{N_y-1} e^{2\pi i \frac{k(m-m')}{N_x}} e^{2\pi i \frac{l(n-n')}{N_y}} \\ &= A_{\text{pix}}^2 N_x N_y \delta_{mm'} \delta_{nn'}, \end{aligned} \quad (\text{A3})$$

where we used Eq. (A1) and the definition of the Kronecker delta in terms of Fourier series. Now, remember that we want to create a map of a convergence field with an input monopole power spectrum  $P_{\text{sub}}^{(0)}(k)$ . The discrete Fourier transform of such a convergence field is given by

$$\langle |\tilde{\kappa}_{mn}^{\text{sub}}|^2 \rangle = A_{\text{pix}} N_x N_y P_{\text{sub}}^{(0)}(k_{mn}). \quad (\text{A4})$$

Thus, to convert from the white noise Fourier variables to the desired substructure  $\tilde{\kappa}_{mn}^{\text{sub}}$  Fourier mode, we write

$$\tilde{\kappa}_{mn}^{\text{sub}} = \sqrt{\frac{P_{\text{sub}}^{(0)}(k_{mn})}{A_{\text{pix}}}} \tilde{\kappa}_{mn}. \quad (\text{A5})$$

Note that  $P_{\text{sub}}$  has units of area so the ratio  $P_{\text{sub}}^{(0)}(k_{mn})/A_{\text{pix}}$  is dimensionless. Finally, to compute the actual convergence map, we perform the inverse discrete Fourier transform

$$\kappa_{kl}^{\text{sub}} = \frac{1}{N_x N_y A_{\text{pix}}} \sum_{m=0}^{N_x-1} \sum_{n=0}^{N_y-1} \tilde{\kappa}_{mn}^{\text{sub}} e^{2\pi i \frac{mk}{N_x}} e^{2\pi i \frac{nl}{N_y}}, \quad (\text{A6})$$

where we note that the factor  $N_x N_y A_{\text{pix}}$  is nothing more than the area of the whole region where we are computing the convergence field. The substructure deflection field is generated in Fourier space as

$$\tilde{\alpha}_{\text{sub}}(\mathbf{k}) = \left( \frac{2ik_x \tilde{\kappa}_{\text{sub}}(\mathbf{k})}{k^2}, \frac{2ik_y \tilde{\kappa}_{\text{sub}}(\mathbf{k})}{k^2} \right). \quad (\text{A7})$$

An inverse discrete Fourier transform similar to Eq. (A6) is then performed to create the configuration space deflection field used to create the lensed image.

## Appendix B: Image Simulation

We explain here how we generate mock lensed images that are perturbed by a random realization of a substructure deflection field. Our procedure to generate a random realization of a substructure convergence and deflection field from a given input power spectrum is reviewed in Appendix A. To avoid periodicity effects due to the use of numerical Fast Fourier Transform (FFT), we generate a substructure deflection field spanning an area several times larger than the strong lensing region where we have sensitivity to substructures. Our actual mock images are generated by combining a model image gotten by applying Eq. (13) with a realistic noise realization. We summarize below the different ingredients entering our simulated images.

- **Macro lens:** We take the macro lens to consist of an isothermal ellipsoidal mass distribution. This model has 5 free parameters: the  $(x, y)$  lens position, the Einstein radius, the ellipticity, and the direction of the latter in the plane of the sky.
- **Lens environment:** We model the lens environment as an external shear (2 parameters).

- **Source:** The source is taken to be an elliptical Sérsic profile, which is described by 6 parameters: the 2D source position, the source flux, its half-light radius, the source ellipticity, and the angle of the semi-major axis in the plane of the sky.
- **Foregrounds:** We take the light from the lens galaxy to also be given by an elliptical  $n = 4$  Sérsic profile (6 parameters). We allow for a possible offset between the centroid of the lens light and that of its mass model (see e.g. Ref. [142]). We also add a uniform sky background (1 parameter) to take into account zodiacal light, Earth shine, etc. Unless otherwise mentioned, we take the sky background surface brightness to be  $L_{\text{sky}} = 4.5 \times 10^{-18}$  erg/cm<sup>2</sup>/Å/s/arcsec<sup>2</sup>.
- **Point Spread Function (PSF):** In most cases, we convolve our images with actual Hubble Space Telescope PSFs<sup>12</sup> as measured by the UVIS detector of the Wide Field Camera 3 (WFC3) using the F555W filter. However, to generate ground-based seeing-limited observations, we adopt a simple Moffat profile with power law index 2 whose FWHM is a free parameter. PSF convolution is handled with standard FFT techniques.
- **Noise:** We assume uncorrelated noise across the pixel array, and adopt the following model for the noise within each pixel

$$\mathbf{C}_{N_\lambda, ij} = \delta_{ij} (\sigma_0^2 + \sigma_1 O_\lambda(\mathbf{x}_i) + s(\mathbf{x}_i)), \quad (\text{B1})$$

where the first term describes a constant noise contribution (1 parameter) across each pixel (such as readout noise), the second term (1 parameter) is proportional to the observed count in each pixel and mimic Poisson (shot) noise, while the last term is introduced to take into account bad pixels or cosmic rays, and can be used to mask certain pixels by giving them large uncertainties. For simplicity, we set  $s(\mathbf{x}_i) = 0$  here. When generating a mock image, we draw a random zero-mean Gaussian noise realization whose variance is given by Eq. (B1) and add it to the image. Here, we conservatively take  $\sigma_0 = 1$  count/pixel, which is significantly larger than the read noise of the UVIS detector of WFC3, and take  $\sigma_1 = 1$ , which makes the second term of Eq. (B1) be exactly Poissonian.

- **Pixelization:** Throughout this work, we consider images of size  $4'' \times 4''$ , which is large enough to capture the relevant features of most galaxy-scale lenses. We restrict ourselves to images with  $50 \times 50$  pixels, which ensures that the likelihood from Eq. (40) can be evaluated in a few seconds on a

<sup>12</sup> See <http://www.stsci.edu/hst/wfc3/analysis/PSF>.

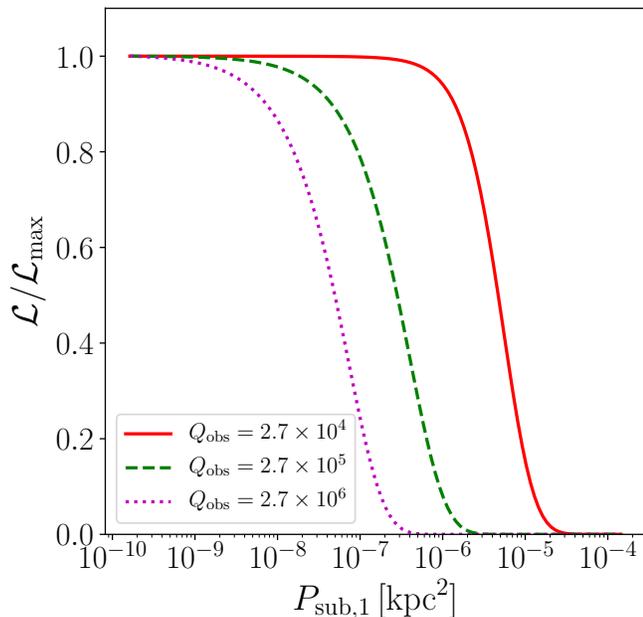


FIG. 11. Posterior distributions for the amplitude  $P_{\text{sub},1}$  of the substructure convergence power spectrum within the wavenumber range  $0.4 \text{ kpc}^{-1} \leq k \leq 5.5 \text{ kpc}^{-1}$ . The mock lensed images used in this analysis (similar in configuration to that shown in Fig. 2 with an HST-like PSF) do not contain substructure perturbations, and the resulting posterior distributions, which are consistent with a vanishing power spectrum amplitude, provide an important null test of our inference framework. The three different curves display different of the quality factor  $Q_{\text{obs}}$ , with the upper bound on the substructure power spectrum amplitude getting more stringent for the deeper images, as expected.

modest computer. This implies a linear pixel size of  $0.08''$ , which is about twice the pixel size of the UVIS detector of WFC3. When computing a model image using Eq. (13), we use 64 light rays per pixel to estimate the photon count within each pixel. We have checked that this number is large enough to ensure that the photon count within the pixels is converged to better than 0.1%.

Throughout the simulated images used in this work, we use the quantity  $Q_{\text{obs}}$  defined in Eq. (67) to characterize the quality of the mock observations. As an example, for the UVIS detector of WFC3 with the F555W filter on (for which  $\mathcal{S}_{\text{inv}}^{(\lambda_i)} = 1.8 \times 10^{-19} \text{ erg/cm}^2/\text{\AA}/\text{count}$ )<sup>13</sup>, we obtain  $Q_{\text{obs}} = 4.4 \times 10^5$  for a combination of 10 45-minute exposures of an unlensed source of AB magnitude 23.

<sup>13</sup> See the technical document <http://www.stsci.edu/hst/wfc3/documents/ISRs/WFC3-2017-14.pdf>.

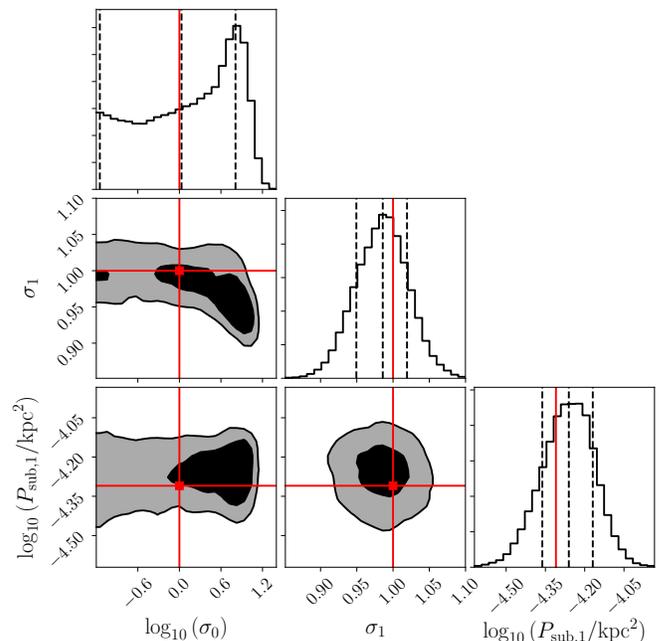


FIG. 12. Posterior distribution for the noise parameters given in Eq. (B1) and the amplitude of the substructure convergence power spectrum. The black and gray regions represent the 68% and 95% confidence regions, respectively. The dashed vertical lines shown in the one-dimensional posteriors shown along the diagonal correspond to the 16%, 50%, and 84% quartiles. The solid red lines and squares represent the true values used to generate the mock data. Here, the fiducial data is generated assuming the point-mass substructure power spectrum shown in Fig. 1, with an image configuration similar to that shown in Fig. 2 with a quality factor  $Q_{\text{obs}} = 2.7 \times 10^5$  and an HST-like PSF.

### Appendix C: Null test and degeneracy with noise parameters

To test the sensitivity of gravitationally lensed extended sources to the substructure convergence power spectrum, we feed mock images generated using the procedure outlined in the Appendix B to the likelihood given in Eq. (48) and generate Markov Chain Monte Carlo (MCMC) samples from the posterior using affine-invariant sampler `emcee` [121]. We adopt broad uniform priors for all macro lens, source, and foreground parameters, except for the surface brightness parameters for which log-uniform priors are taken. The constant noise contribution  $\sigma_0$  also receives a log-uniform prior while we adopt a uniform prior on  $\sigma_1$ . This latter choice has very little impact on our results since the posterior for  $\sigma_1$  is generally very sharply peaked around unity (consistent with Poisson noise).

As in the case of our Fisher analysis, we adopt the logarithm of the binned substructure convergence power spectrum amplitudes as our fitting model, and divide the range of scales probed by a given lensed image into 1 to 4

wavenumber bins that are evenly spaced in  $\log_{10}(k)$ . We adopt a broad log-uniform prior on the amplitude within each bin,  $\log_{10}(P_{\text{sub},i}/\text{arcsec}^2) \in [-11, -1]$ .

We first test the validity of our framework by performing null tests in which we sample the likelihood given in Eq. (48) for lensed images that have not been perturbed by substructure. For simplicity, we take the substructure fitting model to consist of a single power spectrum bin spanning the range  $0.4 \text{ kpc}^{-1} \leq k \leq 5.5 \text{ kpc}^{-1}$ . Using a lensed image configuration similar to that shown in Fig. 2 and considering three different values of the quality factor  $Q_{\text{obs}}$ , we obtain the posterior distributions shown in Fig. 11 for the convergence power spectrum amplitude  $P_{\text{sub},1}$  within that single bin. These posterior distributions are all consistent with a vanishing amplitude of the substructure perturbations, with the upper limits on the power spectrum amplitude improving for images with higher  $Q_{\text{obs}}$ , as should be expected. This indicates that our formalism does not appear to interpret random noise features as spurious substructure within the lens galaxy. While this result is not surprising given the idealized nature of our mock images, it does provide an important consistency check that our numerical implementation of the likelihood given in Eq. (48) is sound.

We then turn our attention to the potential degeneracy between image noise and the amplitude of the substructure power spectrum. For simplicity, we consider here mock lenses that are perturbed by a random realization of point-mass substructures with a power spectrum given by the dashed black line of Fig. 1. While the point-mass power spectrum is likely unrealistic, it has the advantage

and being described by a single parameter: its amplitude. We show the joint posterior distribution for the noise parameters and the amplitude of the substructure power spectrum (denoted  $P_{\text{sub},1}$  here) in Fig. 12. To perform the inference, we use all Fourier modes spanning the range  $0.4 \text{ kpc}^{-1} \leq k \leq 5.5 \text{ kpc}^{-1}$  (for a total of 180 independent Fourier modes<sup>14</sup>).

We observe no degeneracy between the substructure power spectrum amplitude and the noise parameters, and note that we successfully recover the true amplitude of the input power spectrum (solid red line) to within less than one standard deviation. The marginalized posterior for the parameter  $\log_{10}(\sigma_0)$  is quite broad and extends to significantly larger (and lower) values than that used to generate the mock image. This is a consequence of the Poisson term (that proportional to  $\sigma_1$  in Eq. (B1)) dominating the noise budget in the mock image. Only when the contribution from  $\sigma_0^2$  becomes comparable in magnitude to  $\sigma_1 O_\lambda$  can the mock data start displaying sensitivity to  $\sigma_0$ , explaining the long flat tail of the posterior at low values of this noise parameter.

These findings are broadly consistent with those of Ref. [96] which found no degeneracy between uncorrelated image noise and the substructure power spectrum amplitude. Again, this is a consequence of the highly nonrandom structure of the image residuals caused by substructure (Eq. (16)) which is not easily mimicked by pure random noise. In a realistic image where adjacent pixels may have correlated noise due, for example, to drizzling, it is possible however for the noise structure to be somewhat degenerate with the effect of the substructure.

- 
- [1] V. Springel, J. Wang, M. Vogelsberger, A. Ludlow, A. Jenkins, *et al.*, *Mon. Not. R. Astron. Soc.* **391**, 1685 (2008), arXiv:0809.0898 [astro-ph].
- [2] J. R. Bond and A. S. Szalay, *Astrophys. J.* **274**, 443 (1983).
- [3] P. Bode, J. P. Ostriker, and N. Turok, *Astrophys. J.* **556**, 93 (2001), arXiv:astro-ph/0010389 [astro-ph].
- [4] J. J. Dalcanton and C. J. Hogan, *Astrophys. J.* **561**, 35 (2001), arXiv:astro-ph/0004381 [astro-ph].
- [5] D. Boyanovsky, H. J. de Vega, and N. Sanchez, *Phys. Rev. D* **78**, 063546 (2008), 0807.0622.
- [6] D. Boyanovsky and J. Wu, *Phys. Rev. D* **83**, 043524 (2011), 1008.0992.
- [7] E. D. Carlson, M. E. Machacek, and L. J. Hall, *Astrophys. J.* **398**, 43 (1992).
- [8] C. Boehm, A. Riazuelo, S. H. Hansen, and R. Schaeffer, *Phys. Rev. D* **66**, 083505 (2002), arXiv:astro-ph/0112522 [astro-ph].
- [9] L. Ackerman, M. R. Buckley, S. M. Carroll, and M. Kamionkowski, *Phys. Rev. D* **79**, 023519 (2009), arXiv:0810.5126 [hep-ph].
- [10] J. L. Feng, M. Kaplinghat, H. Tu, and H.-B. Yu, *JCAP* **0907**, 004 (2009), arXiv:0905.3039 [hep-ph].
- [11] D. E. Kaplan, G. Z. Krnjaic, K. R. Rehermann, and C. M. Wells, *Journal of Cosmology and Astroparticle Physics* **1005**, 021 (2010), arXiv:0909.0753 [hep-ph].
- [12] L. G. van den Aarsen, T. Bringmann, and C. Pfrommer, *Phys. Rev. Lett.* **109**, 231301 (2012), arXiv:1205.5809 [astro-ph.CO].
- [13] F.-Y. Cyr-Racine and K. Sigurdson, *Phys. Rev. D* **87**, 103515 (2013), 1209.5752.
- [14] F.-Y. Cyr-Racine, R. de Putter, A. Raccanelli, and K. Sigurdson, *Phys. Rev. D* **89**, 063517 (2014), 1310.3278.
- [15] M. R. Buckley, J. Zavala, F.-Y. Cyr-Racine, K. Sigurdson, and M. Vogelsberger, *Phys. Rev. D* **90**, 043524 (2014), arXiv:1405.2075.
- [16] C. Boehm, J. A. Schewtschenko, R. J. Wilkinson, C. M. Baugh, and S. Pascoli, *Mon. Not. R. Astron. Soc.* **445**, L31 (2014), arXiv:1404.7012.

<sup>14</sup> Given the resolution of our mock images (0.08 arcsec pixel size), we could technically extend this range to  $0.4 \text{ kpc}^{-1} \leq k \lesssim 13.2 \text{ kpc}^{-1}$ . However, the small-scale modes with  $k > 5.5 \text{ kpc}^{-1}$  are noise dominated for our choice of exposure, and adding them does not change the posterior shown in Fig. 12.

- [17] J. A. Schewtschenko, R. J. Wilkinson, C. M. Baugh, C. Boehm, and S. Pascoli, *Mon. Not. R. Astron. Soc.* **449**, 3587 (2015), arXiv:1412.4905 [astro-ph.CO].
- [18] X. Chu and B. Dasgupta, *Phys. Rev. Lett.* **113**, 161301 (2014), arXiv:1404.6127 [hep-ph].
- [19] M. A. Buen-Abad, G. Marques-Tavares, and M. Schmaltz, *Phys. Rev.* **D92**, 023531 (2015), arXiv:1505.03542 [hep-ph].
- [20] F.-Y. Cyr-Racine, K. Sigurdson, J. Zavala, T. Bringmann, M. Vogelsberger, and C. Pfrommer, *Phys. Rev. D* **93**, 123527 (2016), arXiv:1512.05344.
- [21] Z. Chacko, Y. Cui, S. Hong, T. Okui, and Y. Tsai, *JHEP* **12**, 108 (2016), arXiv:1609.03569 [astro-ph.CO].
- [22] P. Ko and Y. Tang, *Phys. Lett.* **B762**, 462 (2016), arXiv:1608.01083 [hep-ph].
- [23] M. Vogelsberger, J. Zavala, and A. Loeb, *Mon. Not. R. Astron. Soc.* **423**, 3740 (2012), arXiv:1201.5892 [astro-ph.CO].
- [24] M. Rocha, A. H. Peter, J. S. Bullock, M. Kaplinghat, S. Garrison-Kimmel, *et al.*, *Mon. Not. R. Astron. Soc.* **430**, 81 (2013), arXiv:1208.3025 [astro-ph.CO].
- [25] A. H. G. Peter, M. Rocha, J. S. Bullock, and M. Kaplinghat, *Mon. Not. R. Astron. Soc.* **430**, 105 (2013), arXiv:1208.3026 [astro-ph.CO].
- [26] J. Zavala, M. Vogelsberger, and M. G. Walker, *Mon. Not. R. Astron. Soc.* **431**, L20 (2013), arXiv:1211.6426 [astro-ph.CO].
- [27] M. Kaplinghat, R. E. Keeley, T. Linden, and H.-B. Yu, *Phys. Rev. Lett.* **113**, 021302 (2014), arXiv:1311.6524 [astro-ph.CO].
- [28] M. Kaplinghat, S. Tulin, and H.-B. Yu, *Phys. Rev. Lett.* **116**, 041302 (2016), arXiv:1508.03339 [astro-ph.CO].
- [29] M. Vogelsberger, J. Zavala, F.-Y. Cyr-Racine, C. Pfrommer, T. Bringmann, and K. Sigurdson, *Mon. Not. R. Astron. Soc.* **460**, 1399 (2016), arXiv:1512.05349 [astro-ph.CO].
- [30] J. J. Fan, A. Katz, L. Randall, and M. Reece, *Phys. Rev. Lett.* **110**, 211302 (2013), arXiv:1303.3271 [hep-ph].
- [31] J. Fan, A. Katz, L. Randall, and M. Reece, *Phys. Dark Univ.* **2**, 139 (2013), arXiv:1303.1521 [astro-ph.CO].
- [32] M. R. Buckley and A. DiFranzo, *Phys. Rev. Lett.* **120**, 051102 (2018), arXiv:1707.03829 [hep-ph].
- [33] P. Agrawal, F.-Y. Cyr-Racine, L. Randall, and J. Scholtz, *JCAP* **1708**, 021 (2017), arXiv:1702.05482 [astro-ph.CO].
- [34] P. Agrawal and L. Randall, *JCAP* **1712**, 019 (2017), arXiv:1706.04195 [hep-ph].
- [35] W. Hu, R. Barkana, and A. Gruzinov, *Phys. Rev. Lett.* **85**, 1158 (2000), arXiv:astro-ph/0003365 [astro-ph].
- [36] L. Hui, J. P. Ostriker, S. Tremaine, and E. Witten, *Phys. Rev.* **D95**, 043541 (2017), arXiv:1610.08297 [astro-ph.CO].
- [37] X. Du, C. Behrens, and J. C. Niemeyer, *Mon. Not. Roy. Astron. Soc.* **465**, 941 (2017), arXiv:1608.02575 [astro-ph.CO].
- [38] P. Mocz, M. Vogelsberger, V. H. Robles, J. Zavala, M. Boylan-Kolchin, A. Fialkov, and L. Hernquist, *Mon. Not. Roy. Astron. Soc.* **471**, 4559 (2017), arXiv:1705.05845 [astro-ph.CO].
- [39] M. G. Walker and J. Peñarrubia, *Astrophys. J.* **742**, 20 (2011), arXiv:1108.2404.
- [40] P. Salucci, M. I. Wilkinson, M. G. Walker, G. F. Gilmore, E. K. Grebel, A. Koch, C. F. Martins, and R. F. G. Wyse, *Mon. Not. Roy. Astron. Soc.* **420**, 2034 (2012), arXiv:1111.1165 [astro-ph.CO].
- [41] C. F. P. Laporte, M. G. Walker, and J. Peñarrubia, *Mon. Not. Roy. Astron. Soc.* **433**, 54 (2013), arXiv:1303.1534 [astro-ph.GA].
- [42] V. Bonnivard *et al.*, *Mon. Not. Roy. Astron. Soc.* **453**, 849 (2015), arXiv:1504.02048 [astro-ph.HE].
- [43] N. Caldwell, M. G. Walker, M. Mateo, E. W. Olszewski, S. Kaposov, V. Belokurov, G. Torrealba, A. Geringer-Sameth, and C. I. Johnson, *Astrophys. J.* **839**, 20 (2017), arXiv:1612.06398 [astro-ph.GA].
- [44] N. Banik, G. Bertone, J. Bovy, and N. Bozorgnia, *JCAP* **1807**, 061 (2018), arXiv:1804.04384 [astro-ph.CO].
- [45] W. H. W. Ngan and R. G. Carlberg, *Astrophys. J.* **788**, 181 (2014), arXiv:1311.1710.
- [46] R. G. Carlberg, *Astrophys. J.* **820**, 45 (2016), arXiv:1512.01620.
- [47] J. Bovy, *Physical Review Letters* **116**, 121301 (2016), arXiv:1512.00452.
- [48] D. Erkal, V. Belokurov, J. Bovy, and J. L. Sanders, *Mon. Not. R. Astron. Soc.* **463**, 102 (2016), arXiv:1606.04946.
- [49] J. Bovy, D. Erkal, and J. L. Sanders, *Mon. Not. R. Astron. Soc.* **466**, 628 (2017), arXiv:1606.03470.
- [50] L. M. Widrow, S. Gardner, B. Yanny, S. Dodelson, and H.-Y. Chen, *The Astrophysical Journal Letters* **750**, L41 (2012), arXiv:1203.6861.
- [51] R. Feldmann and D. Spolyar, *Mon. Not. Roy. Astron. Soc.* **446**, 1000 (2015), arXiv:1310.2243 [astro-ph.GA].
- [52] H. A. Clark, G. F. Lewis, and P. Scott, *Mon. Not. Roy. Astron. Soc.* **456**, 1402 (2016), [Erratum: *Mon. Not. Roy. Astron. Soc.* 464, no.1, 955 (2017)], arXiv:1509.02941 [astro-ph.CO].
- [53] H. A. Clark, G. F. Lewis, and P. Scott, *Mon. Not. Roy. Astron. Soc.* **456**, 1394 (2016), [Erratum: *Mon. Not. Roy. Astron. Soc.* 464, no.2, 2468 (2017)], arXiv:1509.02938 [astro-ph.CO].
- [54] M. Buschmann, J. Kopp, B. R. Safdi, and C.-L. Wu, *Phys. Rev. Lett.* **120**, 211101 (2018), arXiv:1711.03554 [astro-ph.GA].
- [55] A. L. Erickcek and N. M. Law, *Astrophys. J.* **729**, 49 (2011), arXiv:1007.4228 [astro-ph.CO].
- [56] F. Li, A. L. Erickcek, and N. M. Law, *Phys. Rev.* **D86**, 043519 (2012), arXiv:1202.1284 [astro-ph.CO].
- [57] K. Van Tilburg, A.-M. Taki, and N. Weiner, *JCAP* **1807**, 041 (2018), arXiv:1804.01991 [astro-ph.CO].
- [58] S. Mao and P. Schneider, *Mon. Not. R. Astron. Soc.* **295**, 587 (1998), astro-ph/9707187.
- [59] R. B. Metcalf and P. Madau, *Astrophys. J.* **563**, 9 (2001), astro-ph/0108224.
- [60] M. Chiba, *Astrophys. J.* **565**, 17 (2002), astro-ph/0109499.
- [61] R. B. Metcalf and H. Zhao, *The Astrophysical Journal Letters* **567**, L5 (2002), astro-ph/0111427.
- [62] N. Dalal and C. S. Kochanek, (2002), astro-ph/0202290.
- [63] C. R. Keeton, B. S. Gaudi, and A. O. Petters, *Astrophys. J.* **598**, 138 (2003), astro-ph/0210318.
- [64] R. B. Metcalf and A. Amara, *Mon. Not. R. Astron. Soc.* **419**, 3414 (2012), arXiv:1007.1599.
- [65] N. Dalal and C. Kochanek, *Astrophys. J.* **572**, 25 (2002), astro-ph/0111456.
- [66] J.-W. Hsueh, W. Enzi, S. Vegetti, M. Auger, C. D. Fassnacht, G. Despali, L. V. E. Koopmans, and J. P. McK-

- ean, (2019), arXiv:1905.04182 [astro-ph.CO].
- [67] D. Gilman, S. Birrer, T. Treu, C. R. Keeton, and A. Nierenberg, *Mon. Not. Roy. Astron. Soc.* **481**, 819 (2018), arXiv:1712.04945 [astro-ph.CO].
- [68] D. Gilman, S. Birrer, T. Treu, A. Nierenberg, and A. Benson, (2019), arXiv:1901.11031 [astro-ph.CO].
- [69] D. D. Xu, S. Mao, A. P. Cooper, L. Gao, C. S. Frenk, R. E. Angulo, and J. Helly, *Mon. Not. R. Astron. Soc.* **421**, 2553 (2012), arXiv:1110.1185 [astro-ph.CO].
- [70] D. D. Xu, S. Mao, A. P. Cooper, J. Wang, L. Gao, C. S. Frenk, and V. Springel, *Mon. Not. R. Astron. Soc.* **408**, 1721 (2010), arXiv:1004.3094.
- [71] D. Xu, D. Sluse, L. Gao, J. Wang, C. Frenk, S. Mao, P. Schneider, and V. Springel, *Mon. Not. R. Astron. Soc.* **447**, 3189 (2015), arXiv:1410.3282.
- [72] J. W. Hsueh, C. D. Fassnacht, S. Vegetti, J. P. McKean, C. Spingola, M. W. Auger, L. V. E. Koopmans, and D. J. Lagattuta, *Mon. Not. Roy. Astron. Soc.* **463**, L51 (2016), arXiv:1601.01671 [astro-ph.CO].
- [73] J.-W. Hsueh, G. Despali, S. Vegetti, D. Xu, C. D. Fassnacht, and R. B. Metcalf, *Mon. Not. Roy. Astron. Soc.* **475**, 2438 (2018), arXiv:1707.07680 [astro-ph.GA].
- [74] J. W. Hsueh, L. Oldham, C. Spingola, S. Vegetti, C. D. Fassnacht, M. W. Auger, L. V. E. Koopmans, J. P. McKean, and D. J. Lagattuta, *Mon. Not. Roy. Astron. Soc.* **469**, 3713 (2017), arXiv:1701.06575 [astro-ph.GA].
- [75] R. Fadely and C. R. Keeton, *Mon. Not. R. Astron. Soc.* **419**, 936 (2012), arXiv:1109.0548 [astro-ph.CO].
- [76] A. M. Nierenberg, T. Treu, S. A. Wright, C. D. Fassnacht, and M. W. Auger, *Mon. Not. R. Astron. Soc.* **442**, 2434 (2014), arXiv:1402.1496.
- [77] A. M. Nierenberg, T. Treu, G. Brammer, A. H. G. Peter, C. D. Fassnacht, C. R. Keeton, C. S. Kochanek, K. B. Schmidt, D. Sluse, and S. A. Wright, *Mon. Not. Roy. Astron. Soc.* **471**, 2224 (2017), arXiv:1701.05188 [astro-ph.CO].
- [78] L. Dai, T. Venumadhav, A. A. Kaurov, and J. Miralda-Escudé, *Astrophys. J.* **867**, 24 (2018), arXiv:1804.03149 [astro-ph.CO].
- [79] L. V. E. Koopmans, *Mon. Not. Roy. Astron. Soc.* **363**, 1136 (2005), arXiv:astro-ph/0501324 [astro-ph].
- [80] S. Vegetti and L. V. E. Koopmans, *Mon. Not. R. Astron. Soc.* **392**, 945 (2009), arXiv:0805.0201.
- [81] S. Vegetti, O. Czoske, and L. V. E. Koopmans, *Mon. Not. R. Astron. Soc.* **407**, 225 (2010), arXiv:1002.4708.
- [82] S. Vegetti, L. V. E. Koopmans, A. Bolton, T. Treu, and R. Gavazzi, *Mon. Not. R. Astron. Soc.* **408**, 1969 (2010), arXiv:0910.0760.
- [83] S. Vegetti, D. J. Lagattuta, J. P. McKean, M. W. Auger, C. D. Fassnacht, and L. V. E. Koopmans, *Nature (London)* **481**, 341 (2012), arXiv:1201.3643.
- [84] L. A. Moustakas and R. B. Metcalf, *Mon. Not. R. Astron. Soc.* **339**, 607 (2003), astro-ph/0206176.
- [85] Y. Hezaveh, N. Dalal, G. Holder, M. Kuhlen, D. Marrone, *et al.*, *Astrophys. J.* **767**, 9 (2013), arXiv:1210.4562 [astro-ph.CO].
- [86] Y. D. Hezaveh *et al.*, *Astrophys. J.* **823**, 37 (2016), arXiv:1601.01388 [astro-ph.CO].
- [87] S. Vegetti and L. V. E. Koopmans, *Mon. Not. R. Astron. Soc.* **400**, 1583 (2009), arXiv:0903.4752 [astro-ph.CO].
- [88] S. Vegetti, L. V. E. Koopmans, M. W. Auger, T. Treu, and A. S. Bolton, *Mon. Not. R. Astron. Soc.* **442**, 2017 (2014), arXiv:1405.3666.
- [89] R. Li, C. S. Frenk, S. Cole, L. Gao, S. Bose, and W. A. Hellwing, *Mon. Not. Roy. Astron. Soc.* **460**, 363 (2016), arXiv:1512.06507 [astro-ph.CO].
- [90] E. Ritondale, S. Vegetti, G. Despali, M. W. Auger, L. V. E. Koopmans, and J. P. McKean, *Mon. Not. Roy. Astron. Soc.* **485**, 2179 (2019), arXiv:1811.03627 [astro-ph.CO].
- [91] Q. E. Minor, M. Kaplinghat, and N. Li, *Astrophys. J.* **845**, 118 (2017), arXiv:1612.05250 [astro-ph.GA].
- [92] B. J. Brewer, D. Huijser, and G. F. Lewis, *Mon. Not. Roy. Astron. Soc.* **455**, 1819 (2016), arXiv:1508.00662 [astro-ph.IM].
- [93] T. Daylan, F.-Y. Cyr-Racine, A. Diaz Rivero, C. Dvorkin, and D. P. Finkbeiner, *Astrophys. J.* **854**, 141 (2018), arXiv:1706.06111 [astro-ph.CO].
- [94] G. Despali and S. Vegetti, *Mon. Not. Roy. Astron. Soc.* **469**, 1997 (2017), arXiv:1608.06938 [astro-ph.GA].
- [95] S. Birrer, A. Amara, and A. Refregier, *Journal of Cosmology and Astroparticle Physics* **5**, 037 (2017), arXiv:1702.00009.
- [96] Y. Hezaveh, N. Dalal, G. Holder, T. Kisner, M. Kuhlen, and L. Perreault Levasseur, *Journal of Cosmology and Astroparticle Physics* **1611**, 048 (2016), arXiv:1403.2720 [astro-ph.CO].
- [97] S. Chatterjee and L. V. E. Koopmans, *Mon. Not. R. Astron. Soc.* **474**, 1762 (2018), arXiv:1710.03075.
- [98] A. Diaz Rivero, F.-Y. Cyr-Racine, and C. Dvorkin, *Phys. Rev.* **D97**, 023001 (2018), arXiv:1707.04590 [astro-ph.CO].
- [99] S. Brennan, A. J. Benson, F.-Y. Cyr-Racine, C. R. Keeton, L. A. Moustakas, and A. R. Pullen, (2018), arXiv:1808.03501 [astro-ph.GA].
- [100] A. Díaz Rivero, C. Dvorkin, F.-Y. Cyr-Racine, J. Zavala, and M. Vogelsberger, *Phys. Rev.* **D98**, 103517 (2018), arXiv:1809.00004 [astro-ph.CO].
- [101] R. E. Angulo, O. Hahn, and T. Abel, *Mon. Not. Roy. Astron. Soc.* **434**, 3337 (2013), arXiv:1304.2406 [astro-ph.CO].
- [102] D. Bayer, S. Chatterjee, L. V. E. Koopmans, S. Vegetti, J. P. McKean, T. Treu, and C. D. Fassnacht, (2018), arXiv:1803.05952 [astro-ph.GA].
- [103] C. R. Keeton, *Astrophys. J.* **584**, 664 (2003), astro-ph/0209040.
- [104] C. McCully, C. R. Keeton, K. C. Wong, and A. I. Zabludoff, *Monthly Notices of the Royal Astronomical Society* **443**, 3631 (2014).
- [105] G. Despali, S. Vegetti, S. D. M. White, C. Giocoli, and F. C. van den Bosch, *Mon. Not. Roy. Astron. Soc.* **475**, 5424 (2018), arXiv:1710.05029 [astro-ph.CO].
- [106] P. A. R. Ade *et al.* (Planck), *Astron. Astrophys.* **594**, A13 (2016), arXiv:1502.01589 [astro-ph.CO].
- [107] A. G. Woldesenbet and L. L. R. Williams, *Mon. Not. Roy. Astron. Soc.* **420**, 2944 (2012), arXiv:1110.6857 [astro-ph.CO].
- [108] E. E. Falco, M. V. Gorenstein, and I. I. Shapiro, *The Astrophysical Journal Letters* **289**, L1 (1985).
- [109] C. R. Keeton, ArXiv e-prints (2009), arXiv:0908.3001 [astro-ph.CO].
- [110] F.-Y. Cyr-Racine, L. A. Moustakas, C. R. Keeton, K. Sigurdson, and D. A. Gilman, *Phys. Rev.* **D94**, 043505 (2016), arXiv:1506.01724 [astro-ph.CO].
- [111] L. V. E. Koopmans, *Mon. Not. R. Astron. Soc.* **363**, 1136 (2005), astro-ph/0501324.
- [112] J. Binney and S. Tremaine, *Galactic Dynamics*, 2nd ed. (Princeton University Press, 2008) p. 920.

- [113] J. F. Navarro, C. S. Frenk, and S. D. M. White, *Astrophys. J.* **490**, 493 (1997), arXiv:astro-ph/9611107 [astro-ph].
- [114] A. Burkert, *IAU Symposium 171: New Light on Galaxy Evolution Heidelberg, Germany, June 26-30, 1995*, IAU Symp. **171**, 175 (1996), [Astrophys. J.447,L25(1995)], arXiv:astro-ph/9504041 [astro-ph].
- [115] A. S. Bolton, S. Burles, L. V. E. Koopmans, T. Treu, R. Gavazzi, L. A. Moustakas, R. Wayth, and D. J. Schlegel, *The Astrophysical Journal* **682**, 964 (2008), arXiv:0805.1931.
- [116] P. Schneider, J. Ehlers, and E. Falco, *Gravitational Lenses*, Astronomy and Astrophysics Library (Springer, 1999).
- [117] R. Blandford, G. Surpi, and T. Kundić, in *Gravitational Lensing: Recent Progress and Future Goals*, Astronomical Society of the Pacific Conference Series, Vol. 237, edited by T. G. Brainerd and C. S. Kochanek (2001) p. 65, astro-ph/0001496.
- [118] H. N. Nguyen, N. Sehgal, and M. Madhavacheril, *Phys. Rev.* **D99**, 023502 (2019), arXiv:1710.03747 [astro-ph.CO].
- [119] A. Refregier, *Mon. Not. Roy. Astron. Soc.* **338**, 35 (2003), arXiv:astro-ph/0105178 [astro-ph].
- [120] <https://github.com/numba/numba>.
- [121] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, *Publ. Astron. Soc. Pac.* **125**, 306 (2013), arXiv:1202.3665 [astro-ph.IM].
- [122] B. C. Kelly, J. Bechtold, and A. Siemiginowska, *Astrophys. J.* **698**, 895-910 (2009), arXiv:0903.5315 [astro-ph.CO].
- [123] C. L. MacLeod, K. Brooks, Ž. Ivezić, C. S. Kochanek, R. Gibson, A. Meisner, S. Kozłowski, B. Sesar, A. C. Becker, and W. H. de Vries, *Astrophys. J.* **728**, 26 (2011), arXiv:1009.2081 [astro-ph.CO].
- [124] R. F. Mushotzky, R. Edelson, W. Baumgartner, and P. Gandhi, *The Astrophysical Journal Letters* **743**, L12 (2011), arXiv:1111.0672.
- [125] A. Hojjati, A. G. Kim, and E. V. Linder, *Phys. Rev. D* **87**, 123512 (2013), 1304.0309.
- [126] Y. Zu, C. S. Kochanek, S. Kozłowski, and A. Udalski, *Astrophys. J.* **765**, 106 (2013), arXiv:1202.3783 [astro-ph.CO].
- [127] D. Sluse, J. F. Claeskens, B. Altieri, R. A. Cabanac, O. Garcet, D. Hutsemekers, C. Jean, A. Smette, and J. Surdej, *Astron. Astrophys.* **449**, 539 (2006), arXiv:astro-ph/0512473 [astro-ph].
- [128] C. R. Keeton and L. A. Moustakas, *Astrophys. J.* **699**, 1720 (2009), 0805.0309.
- [129] J. Nightingale, S. Dye, and R. Massey, *Mon. Not. R. Astron. Soc.* (2018), 10.1093/mnras/sty1264, arXiv:1708.07377 [astro-ph.CO].
- [130] S. Warren and S. Dye, *Astrophys. J.* **590**, 673 (2003), arXiv:astro-ph/0302587 [astro-ph].
- [131] S. H. Suyu, P. J. Marshall, M. P. Hobson, and R. D. Blandford, *Mon. Not. Roy. Astron. Soc.* **371**, 983 (2006), arXiv:astro-ph/0601493 [astro-ph].
- [132] S. Birrer, A. Amara, and A. Refregier, *Astrophys. J.* **813**, 102 (2015), arXiv:1504.07629 [astro-ph.CO].
- [133] Y.-Y. Mao, M. Williamson, and R. H. Wechsler, *Astrophys. J.* **810**, 21 (2015), arXiv:1503.02637 [astro-ph.CO].
- [134] M. Oguri and P. J. Marshall, *Mon. Not. Roy. Astron. Soc.* **405**, 2579 (2010), arXiv:1001.2037 [astro-ph.CO].
- [135] J. D. Hunter, *Computing in Science & Engineering* **9**, 90 (2007).
- [136] E. Jones, T. Oliphant, P. Peterson, *et al.*, “SciPy: Open source scientific tools for Python,” (2001).
- [137] T. E. Oliphant, *Computing in Science & Engineering* **9**, 10 (2007).
- [138] A. M. Price-Whelan *et al.* (The Astropy Collaboration), *ArXiv e-prints* (2018), arXiv:1801.02634 [astro-ph.IM].
- [139] <https://github.com/cmbant/getdist>.
- [140] <https://github.com/h5py/h5py>.
- [141] D. Foreman-Mackey, *The Journal of Open Source Software* **24** (2016), 10.21105/joss.00024.
- [142] Y. Shu, A. S. Bolton, L. A. Moustakas, D. Stern, A. Dey, J. R. Brownstein, S. Burles, and H. Spinrad, *Astrophys. J.* **820**, 43 (2016), arXiv:1602.02927.