

Probing dark matter structure down to 10^7 solar masses: flux ratio statistics in gravitational lenses with line of sight halos

Daniel Gilman¹ ^{*}, Simon Birrer¹, Tommaso Treu¹, Anna Nierenberg², Andrew Benson³

¹*Department of Physics and Astronomy, University of California, Los Angeles, CA 90095, USA*

²*Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr, Pasadena, CA 91109, USA*

³*Carnegie Observatories, 813 Santa Barbara Street, Pasadena, CA 91101, USA*

Accepted . Received

ABSTRACT

Strong lensing provides a powerful means of investigating the nature of dark matter as it probes the mass function and density profiles of halos on sub-galactic scales. We present an extension of a forward modeling framework that uses flux ratios from quadruply imaged quasars (quads) to measure the shape and amplitude of the halo mass function, including line of sight (LOS) halos and main deflector subhalos. We apply this machinery to 50 mock lenses — roughly the number of known quads — with mass functions exhibiting a free-streaming cutoff parameterized by the half-mode mass m_{hm} . Assuming cold dark matter (CDM), we forecast bounds on m_{hm} and the corresponding thermal relic particle masses for scenarios with a range of tidal destruction severity. With significant tidal destruction, at 2σ we constrain $m_{\text{hm}} < 10^{7.9}$ ($10^{8.4}$) M_{\odot} , or a 4.4 (3.1) keV thermal relic, with image flux uncertainties from measurements and lens modeling of 2% (6%). With less severe tidal destruction we constrain $m_{\text{hm}} < 10^7$ ($10^{7.4}$) M_{\odot} , or an 8.2 (6.2) keV thermal relic. If dark matter is warm, with $m_{\text{hm}} = 10^{7.7} M_{\odot}$ (5.1 keV), we would favor WDM with $m_{\text{hm}} > 10^{7.7} M_{\odot}$ over CDM with relative likelihoods of 22:1 and 8:1 with flux uncertainties of 2% and 6%, respectively. These bounds improve over those obtained by modeling only main deflector subhalos because LOS objects produce additional flux perturbations, especially for high redshift systems. These results indicate that ~ 50 quads can conclusively differentiate between warm and cold dark matter.

Key words: [gravitational lensing: strong - cosmology: dark matter - galaxies: structure - methods: statistical]

1 INTRODUCTION

Theories of particle dark matter predict that the enigmatic particle(s) collect in gravitationally bound halos. The mass function and density profiles of these objects depend on the particle nature of dark matter itself. For example, theories with cold dark matter predict an abundance of low mass halos, and cuspy r^{-1} central density profiles (Moore et al. 1999; Springel et al. 2008; Fiacconi et al. 2016). In warm dark matter (WDM) models, diffusion of dark matter particles in the early universe wipes out density fluctuations below a characteristic scale that depends on the production mechanism of the WDM particle candidate (Kusenko 2009; Shoemaker & Kusenko 2009; Abazajian 2017). Suppression of small-scale

power in WDM models results in a turnover in the halo mass function and a dearth of small-scale structure at later times (Bode et al. 2001; Schneider et al. 2012; Lovell et al. 2014). In self-interacting dark matter (SIDM) theories, scattering between dark matter particles produces cored density profiles in individual halos (Spergel & Steinhardt 2000; Rocha et al. 2013; Vogelsberger et al. 2016; Kamada et al. 2017; Tulin & Yu 2018). Finally, in ‘fuzzy’ dark matter scenarios the kpc-scale de Broglie wavelength of ultra-light dark matter particles results in quantum mechanical phenomena on galactic scales, which produces large soliton cores (Hui et al. 2017; Robles et al. 2019). To date, the strongest constraints on WDM come from the Lyman- α forest (Viel et al. 2013; Iršič et al. 2017), while cosmological probes on large scales (Cyr-Racine et al. 2014; Bringmann et al. 2017) and

* gilmanda@ucla.edu

in galaxy clusters (Kim et al. 2017b; Andrade et al. 2019) constrain the interaction cross section in SIDM models.

Two challenges to CDM in particular spur interest in alternative theories. First, natural CDM particle candidates have not yet been detected, despite decades of experimental searches (Aprile et al. 2018). Second, the suppression of small scale structure in WDM, and cored density profiles associated with SIDM, possibly alleviate tension between observations and the predictions on sub-galactic scales, dubbed the ‘Small-Scale Crisis’ of CDM (see Bullock & Boylan-Kolchin 2017, and references therein).

Traditional astrophysical challenges to the CDM model, however, are predicated on assumptions related to baryonic physics. This has the undesirable consequence of propagating uncertainties from sub-galactic astrophysics onto inferences of dark matter properties, and results in covariance between baryonic astrophysics and dark matter physics. Supernova and stellar feedback inside halos, for instance, and the tidal destruction of subhalos by their host galaxy, mimic the observable signatures of SIDM and WDM models, respectively (Tollet et al. 2016; Read et al. 2018; Despali & Vegetti 2017; Garrison-Kimmel et al. 2017; Kim et al. 2017a; Despali et al. 2018a). Moreover, in some cases, the uncertainties related to baryonic astrophysical processes can be larger than the differences between CDM, WDM, and SIDM (e.g. Nierenberg et al. 2016). To isolate dark matter physics from sub-galactic astrophysics, and to differentiate between CDM, WDM, and SIDM, one must look to masses below $10^8 M_\odot$, where subhalos are expected to be devoid of stars and completely dark in the case of CDM, or absent in the case of WDM.

Gravitational lensing offers a direct probe of this elusive, low-mass regime. It circumvents the complications associated with using luminous matter to trace the dark matter by enabling the direct measurement of the distribution of matter across cosmological distance, and is sensitive to mass scales where astrophysical effects are thought to be too weak to significantly alter the structure of halos. It also complements other probes of dark matter, such as the Lyman- α forest, since lensing depends on different systematics and measures the halo mass function directly.

Ultimately, analysis of strong lenses hinges on separating mass distributions that vary on large scales (the lensing galaxy and its parent dark matter halo) from small scale structure in the main lens plane and along the line of sight. In strong lens systems with luminous arcs, the analysis consists of iteratively fitting a smooth model to the flux in pixels of an image while simultaneously reconstructing the background source. This process can reveal the presence of small scale structure in the arcs (see e.g. Vegetti et al. 2014; Hezaveh et al. 2016b; Vegetti et al. 2018; Ritondale et al. 2018). Birrer et al. (2017b) performed this analysis, and placed constraints on the free streaming length of dark matter. Recently, several authors have proposed using the surface brightness residuals from lens models fit to luminous arcs and to infer the power spectrum of dark matter in strong lenses (Hezaveh et al. 2016a; Diaz Rivero et al. 2018; Cyr-Racine et al. 2018), and Bayer et al. (2018) applied this method to a strong lens system.

In addition to extended arcs, some strong gravitational lenses produce four images (quads) of an unresolved background source, such as a quasar. The magnification ratios

(flux ratios) between multiple images of unresolved sources have long been recognized as powerful probes of small scale structure near lensed images (Mao & Schneider 1998; Metcalf & Madau 2001), and have been used to test the predictions of CDM and to detect structure near individual objects (Dalal & Kochanek 2002; Amara et al. 2006; Xu et al. 2012; Nierenberg et al. 2014; Xu et al. 2015; Nierenberg et al. 2017). Recently, Nierenberg et al. (2014, 2017) used image flux ratios measured from narrow line emission, a method first proposed by (Moustakas & Metcalf 2003), to study substructure in strong lenses. The significance of this advance derives from the fact that the magnification of a lensed image is a function of background source diameter (Dobler & Keeton 2006); the narrow-line region, which typically subtends angles on scales of a few tens of milliarcseconds, is resilient to contaminating effects of microlensing by stars, while still being sensitive to the milliarcsecond perturbations sourced by dark matter halos above $10^6 M_\odot$ with current astrometric precision of a few m.a.s. (Nierenberg et al. 2017).

In this work, we extend the formalism presented by (Gilman et al. 2018) to include the contribution from dark matter halos along the line of sight. Since field halos do not orbit in a steep galactic potential with star formation, stellar feedback, and other complications, they constitute an ideal laboratory for studying the intrinsic structure of dark matter halos. Several studies investigate the role of the line of sight halos on flux ratio perturbations in strong lenses (Chen et al. 2003; Metcalf 2005; Xu et al. 2012; Inoue & Takahashi 2012), and Despali et al. (2018b) address the line of sight contribution in the context of gravitational imaging with luminous arcs. The consensus from these works is that the line of sight halos affect lensing observables, possibly becoming the dominant source of perturbation to smooth lens models for lenses at high redshift.

The analysis presented here builds on previous analysis of multiple image lenses in several ways. First, we quantify the signal from non-linear multi-plane lensing effects on flux ratios with finite-size background sources, and combine this multi-plane lensing machinery with a forward-generative model to measure the shape and amplitude of the halo mass function by combining flux ratio statistics from a sample of lenses. We also marginalize over parameters such as the background source size and the power law profile of the main deflector, both of which can affect the flux ratios between images. We demonstrate how well this method constrains the free-streaming length of dark matter in the presence of uncertainties associated with measurements and lens modeling, and apply the machinery to a set of 50 simulated quads. The number 50 is chosen since it is roughly the size of the current sample of known quads (Shajib et al. 2018, HST GO-15652) with a similar distribution of lens and source redshifts.

This paper is organized as follows: First in Section 2, we describe our prescription for modeling the line of sight halo mass function, and the subhalo mass function in the main lens plane. In Section 3, we discuss the impact of line of sight halos on flux ratio observables. In Section 4, we describe the forward modeling procedure implemented in the simulations, and the process for creating mock datasets. Finally, in Section 5 we present the results of simulations run with a mock data set in which we infer dark matter and lens model parameters with a Bayesian framework. Finally,

Section 6 summarizes our main results. All lensing computations performed in this work utilize the open-source gravitational lensing software LENSSTRONOMY (Birrer et al. 2015; Birrer & Amara 2018). Cosmological calculations, in particular the line of sight halo mass function and two-halo term, are computed with the software package COLOSSUS (Diemer 2017). We assume a flat cosmology with $\sigma_8 = 0.82$, $\Omega_m = 0.28$ and $h = 0.7$. When quoting halo masses, we refer to M_{200} computed with respect to the critical density of the universe at $z=0$.

2 MODELING THE LINE OF SIGHT AND SUBHALO MASS FUNCTIONS

This section describes the parameterization of the subhalo mass function in the main deflector, and the halo mass function along the line of sight, as well as the density profile for individual halos. We then describe our parameterization of free-streaming effects in WDM models, both on the mass functions and the mass-concentration relation. The forward model, described in Section 4, will use these parameterizations to render realistic populations of dark matter structure for lensing computations.

2.1 Mass profile of individual halos

We model the density profiles of dark matter halos using truncated NFW profiles (Baltz et al. 2009)

$$\rho(r, r_s, r_t) = \frac{\rho_0}{x(1+x)^2} \frac{\tau^2}{x^2 + \tau^2} \quad (1)$$

where $\tau = \frac{r_t}{r_s}$ and $x = \frac{r}{r_s}$.

In the main lens plane, we tidally truncate subhalos through a Roche limit approximation, assuming a roughly isothermal mass profile for the main lens halo mass distribution. This truncation corresponds to a scaling $r_t \propto (M_{200} r_{3D}^2)^{\frac{1}{3}}$ (Tormen et al. 1998; Cyr-Racine et al. 2016). We truncate according to this scaling using the expression

$$r_t = 0.68 \left(\frac{M_{200}}{10^6 M_\odot} \right)^{\frac{1}{3}} \left(\frac{r_{3D}}{100 \text{ kpc}} \right)^{\frac{2}{3}} [\text{kpc}]. \quad (2)$$

This results in a skewed distribution of τ with mean $\langle \tau \rangle \sim 6$, and a tail extending to $\tau \sim 20$.

We truncate line of sight halos at r_{50} , or the radius where the mean enclosed density is $50\rho_{\text{crit}}$ ¹. Finally, we adopt the mass-concentration-redshift relation for CDM halos presented by (Diemer & Joyce 2018) with a scatter of 0.13 dex (Dutton & Macciò 2014). We render subhalos in the range $10^6 - 10^{10} M_\odot$, which captures perturbations from the smallest subhalos that affect image magnifications, given the source sizes we model.

¹ We introduce this truncation to keep the total mass per unit volume along the line of sight finite, since the mass of an NFW profile diverges. Since r_{50} is much larger than the scale radius of an NFW halo, this truncation negligibly impacts observables.

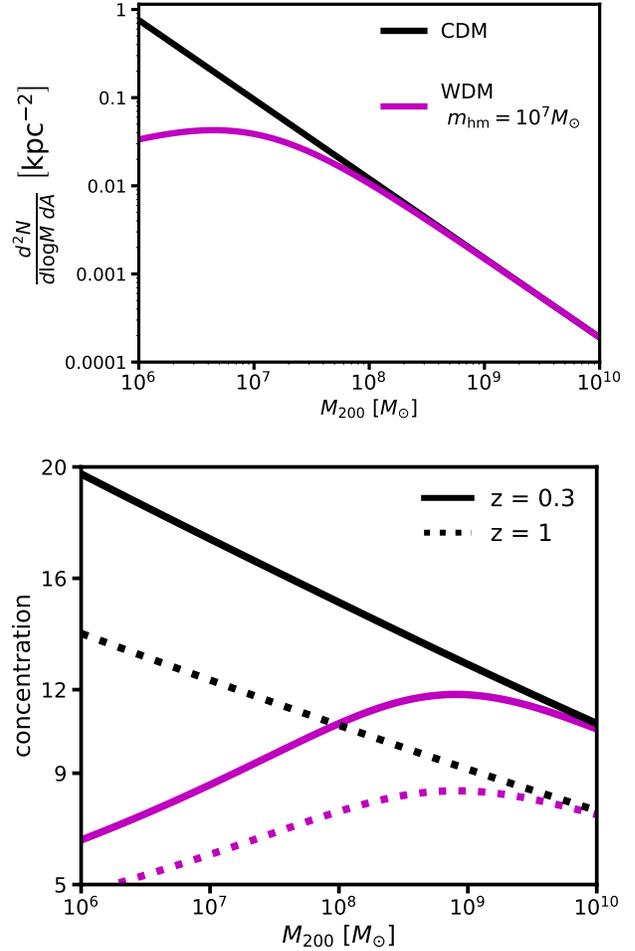


Figure 1. Top: The subhalo mass function for CDM and a WDM scenario with a half-mode mass of $10^7 M_\odot$. The line of sight halo mass function looks qualitatively similar, but evolves slightly with redshift. **Bottom:** The mass concentration relation for the same CDM and WDM models as in the upper panel. The effects of free-streaming on the mass-concentration relation alter the properties of halos two orders of magnitude above the half-mode mass.

2.2 The line of sight halo mass function

We model line of sight structure using the mass function of Sheth and Tormen (Sheth et al. 2001), plus a boost from the 2-halo term at a distance r from the main deflector $\xi_{2\text{halo}}(r, M, z)$, where M denotes the halo mass of the parent dark matter halo. The two-halo term accounts for the correlated structure in the vicinity of the main lens halo. To leading order, this term rescales the background density and the amplitude of the halo mass function. The inclusion of $\xi_{2\text{halo}}$ results in a roughly 5 – 15% boost in the number of halos located at approximately the main lens redshift, depending on the normalization of the subhalo mass function and the lens redshift. We review the form of the two-halo term and its implementation in lensing simulations in Appendix A.

We introduce a rescaling factor δ_{los} to account for theoretical uncertainty regarding the amplitude of the halo mass function. This term accounts for statistical fluctua-

tions around the mean density of the universe, which may lead to modestly over-dense or under-dense lines of sight to individual lenses. We note, however, that due to the vast cosmological distances probed by strong lenses (order Gpc, versus kpc-scale dark matter halos and filament diameters) the dark matter structure in these volumes should be well represented by the average halo mass function in the universe², which corresponds to $\delta_{\text{los}} = 1$, modulo uncertainties in parameters such as σ_8 and Ω_m .

With these modifications, the line of sight halo mass function takes the form

$$\frac{d^2 N_{\text{los}}}{dm dV} = \delta_{\text{los}} (1 + \xi_{2\text{halo}}) \left. \frac{d^2 N}{dm dV} \right|_{\text{ShethTormen}}. \quad (3)$$

This mass function yields accurate counts of isolated halos over a wide mass range. We do not model the subhalos of these objects along the line of sight, subsuming the possible effects of these small perturbers into the marginalization over δ_{los} . We render line of sight halos in the range $10^6 - 10^{10} M_{\odot}$. Line of sight halos are rendered in a double-cone geometry with opening angle $3R_{\text{Ein}}$, where R_{Ein} is the Einstein radius of a given lens, and a closing angle behind the main lens plane such that the cone closes at the source redshift.

2.3 The subhalo mass function of the main deflector

We parameterize the subhalo mass function in terms of a projected number density per unit area. In principle, the abundance and spatial distribution of substructure depends on the total mass of the parent dark matter halo and redshift (Gao et al. 2011; Han et al. 2016), and tidal stripping, which can dramatically reduce the subhalo content of galactic halos (Despali & Vegetti 2017; Han et al. 2016; Garrison-Kimmel et al. 2017; Jiang & van den Bosch 2017; Richings et al. 2018). We may therefore write the subhalo mass function as

$$\frac{d^2 N}{dm dA} = \frac{\Sigma_{\text{sub}}}{m_0} \left(\frac{m}{m_0} \right)^{-\alpha} F(M) H(z) \quad (4)$$

where F and H encode dependence on the parent halo mass M and redshift, respectively.

We render subhalos out to a maximum projected radius of $R_{\text{max}} = 3R_{\text{Ein}}$, and render the subhalo z -coordinates in three dimensions out to the virial radius of the parent halo. In the semi-cylindrical volume defined by the virial radius and R_{max} , we assume the spatial distribution of subhalos follows the mass profile of the parent NFW halo outside $r_{3D} = 0.5R_s$, where R_s is the scale radius of the parent halo, and assume the spatial distribution (per unit volume) is constant inside $0.5R_s$. This reflects the impact of tidal stripping, which tends to preferentially destroy subhalos in the central regions of halos (Jiang & van den Bosch 2017). This procedure sets the distribution of subhalo z -coordinates, which affects the truncation of subhalos through Equation 2. When we render halo populations from this mass function and the line of sight halo mass function, we draw from a Poisson

² To maintain the correct matter density in the universe, we add negative convergence sheets at each lens plane such that the total mass added in halos at each lens plane equals zero, on average (see Birrer et al. 2017a).

distribution with mean $\langle N \rangle$ obtained by normalizing and integrating Equation 4 (see Section 4).

2.4 Modeling free-streaming effects in WDM

Diffusion of dark matter particles in the early universe suppresses small scale power in the matter power spectrum below a characteristic ‘free-streaming length’ that depends on the WDM particle mass and formation mechanism. For a more detailed review of WDM theory, see Benson et al. (2013); Schneider et al. (2013).

We parameterize free-streaming effects on the mass function through the half-mode mass m_{hm} , defined with respect to the length scale where the WDM transfer function is damped with respect to the CDM transfer function by one-half. In WDM models, the number of halos below m_{hm} is strongly suppressed with respect to CDM. We adopt the functional form for this effect given by Lovell et al. (2014)

$$\frac{dN_{\text{wdm}}}{dm} = \frac{dN_{\text{cdm}}}{dm} \left(1 + \frac{m_{\text{hm}}}{m} \right)^{-1.3}. \quad (5)$$

We note that other parameterizations for the turnover in the mass function differ slightly from Equation 5 (see Schneider et al. 2012; Benson et al. 2013). For instance, the WDM mass functions by Benson et al. (2013) exhibit a harder turnover than the parameterization in Equation 5 due to physical effects, namely, the presence of dark matter velocity dispersion at early times. Other (non-physical) variables, including the different algorithms for identifying and assigning mass to halos, and the choice of window function used to compute the matter power spectrum, can yield different mass functions for the same dark matter model. We do not explicitly address these complications in this work.

Thermally produced dark matter particles (thermal relics), assuming they comprise the entirety of the dark matter, admit a one-to-one mapping between the half-mode mass m_{hm} and the mass of the dark matter particle m_{DM} . To translate between these two quantities, we use the scaling $m_{\text{hm}} \sim m_{\text{DM}}^{-3.33}$ (see Schneider et al. 2012), and normalize this relation using the $2 \times 10^8 M_{\odot} h^{-1} \sim 3.3 \text{ keV}$ constraint from the Lyman- α forest (Viel et al. 2013). This yields

$$m_{\text{hm}}(m) = 10^{10} \left(\frac{m_{\text{DM}}}{1 \text{ keV}} \right)^{-3.33} M_{\odot} h^{-1}. \quad (6)$$

In addition to a suppressed mass function below the free streaming scale, free streaming alters the concentration-mass relation of WDM halos (Schneider et al. 2012; Macciò et al. 2013; Bose et al. 2016; Ludlow et al. 2016). We model this suppression using the parameterization given by (Bose et al. 2016)

$$\frac{c_{\text{wdm}}(m, z)}{c_{\text{cdm}}(m, z)} = (1+z)^{\beta(z)} \left(1 + 60 \frac{m_{\text{hm}}}{m} \right)^{-0.17}. \quad (7)$$

with $\beta(z) = 0.026z - 0.04$.³ We plot the subhalo mass function and the mass concentration relation in Figure 1. Due to the factor of 60 in Equation 7, the effect on halo concentrations affects the central densities of objects with masses significantly above the half-mode mass.

³ We remind the reader that we use the $c_{\text{cdm}}(m, z)$ relation presented in Diemer & Joyce (2018).

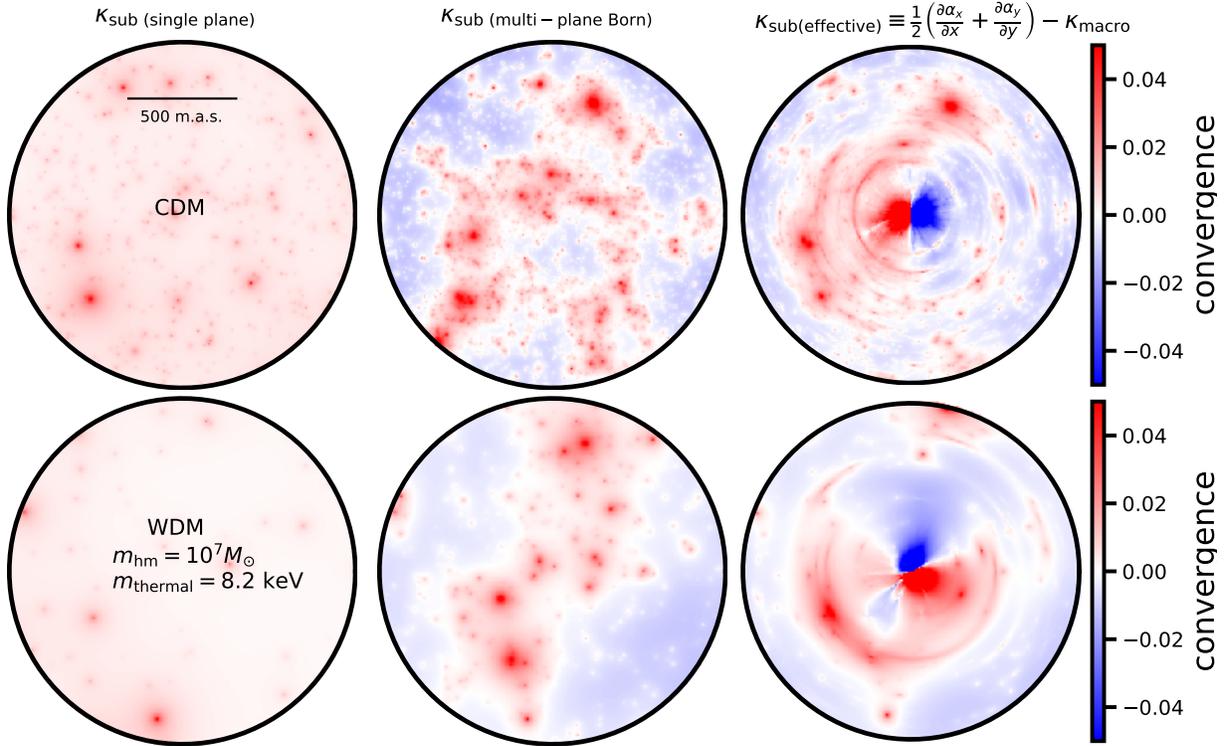


Figure 2. A visualization of the mass distributions that affect observables in single-plane and multi-plane lensing. The top and bottom rows show a single realization of CDM and WDM structure, respectively. **Left:** The convergence map from subhalos of the main deflector only, with $\Sigma_{\text{sub}} = 0.012 \text{ kpc}^{-2}$, which corresponds to a projected mass fraction in substructure at the Einstein radius of 1% at $z = 0.5$. **Center:** The full line of sight realization viewed in projection. Computing deflection angles with respect to these mass distribution effectively employs the Born approximation, in which the deflection angles from halos at different redshifts are computed by assuming light travels along an unperturbed path. There are blue regions with negative mass due to the inclusion of negative convergence sheets at each lens plane, to ensure that the mean density along the line of sight is the background density of the universe. **Right:** The *effective multi-plane convergence* for these realizations. The deflection angles corresponding to these convergence maps, after subtracting off the convergence from the main deflector, include the non-linear effects present in multi-plane lensing not captured by the Born approximation (see Appendix A).

3 EFFECT OF LINE OF SIGHT STRUCTURE ON IMAGE FLUX RATIOS

In order to constrain different dark matter models, we must accurately predict image flux ratios in the presence of perturbing dark matter halos in the main lens plane and along the line of sight. To this end, in this exploratory section we investigate the effect of halos at multiple redshifts on flux ratio observables. First, we present visualizations of the non-linear effects present in multi-plane lensing by defining an effective single plane mass distribution for a multi-plane lens system. We then quantify the signal in flux ratios from line of sight structures using a summary statistic, and compare the contributions from subhalos in the main deflector to the signal from line of sight objects for lenses at different redshifts.

3.1 Multi-plane lensing

As photons traverse the cosmos from a background source to an observer, they experience numerous deflections by dark matter halos along the line of sight. One formulation of the equation describing these deflections and the path of de-

flected light rays is given by (Schneider 1997)

$$\beta_{\mathbf{S}} = \boldsymbol{\theta} - \frac{1}{D_s} \sum_{n=1}^{S-1} D_{ns} \boldsymbol{\alpha}_n (D_n \boldsymbol{\beta}_n). \quad (8)$$

where $\beta_{\mathbf{S}}$ and $\boldsymbol{\theta}$ denote angular coordinates in the source plane and on the sky, respectively, and where D_n and D_{ns} denote angular diameter distances to the n th lens plane, and between the n th lens plane and the source plane.

In the case of a single lens plane, the deflection field from multiple halos is a linear superposition of the deflections from each individual halo. In the case of multiple lens planes, however, Equation 8 becomes a recursive equation for the $\boldsymbol{\beta}_n$, coupling the deflections from halos at different redshifts. Equation 8 describes a physical process akin to looking through a magnifying glass through the lens of another magnifying glass (or in the case of substructure lensing, through thousands of other magnifying glasses). For additional details on multi-plane lensing, see Schneider et al. (1992).

The number of halos along the line of sight often outnumber main lens plane subhalos, to a degree that depends on the lens and source redshifts, and the normalization of the subhalo mass function. However, number counts do not accurately reflect the effects of these line of sight objects on

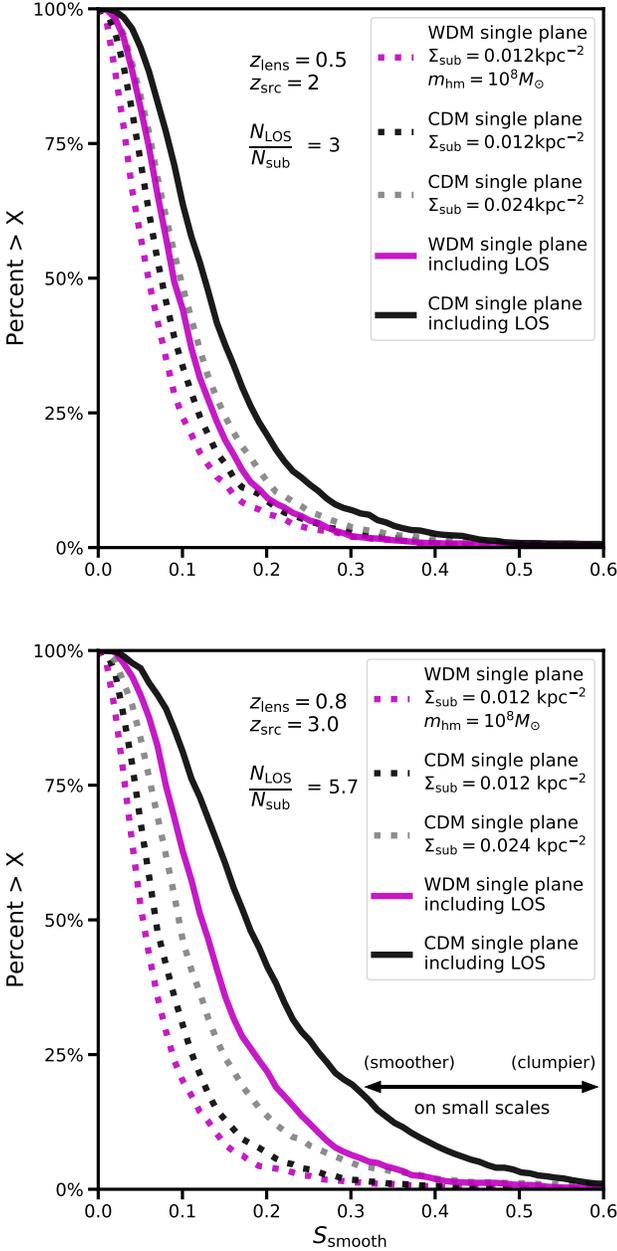


Figure 3. Distributions of the summary statistic in Equation 10 for different dark matter mass functions, and lens and source redshifts. Dotted curves represent realizations of main deflector subhalos only, while the solid curves include both subhalos and line of sight halos. Black and grey curves denote CDM mass functions with normalizations $\Sigma_{\text{sub}} = 0.012 \text{ kpc}^{-2}$ and 0.024 kpc^{-2} , respectively, while magenta curves correspond to WDM mass functions with $\Sigma_{\text{sub}} = 0.012 \text{ kpc}^{-2}$ and $m_{\text{hm}} = 10^8 M_{\odot}$. Mass functions with more small scale structure produce more frequent flux ratio anomalies with respect to smooth lens models, which results in longer tails in the cumulative distribution of these statistics. The boost in the frequency and magnitude of flux ratio anomalies is much stronger for configurations with higher lens/source redshifts.

lensing observables. First, the geometry defined by the lens and source redshifts results in different lensing efficiencies for halos at different redshifts. Second, the coupling between deflections by halos at different redshifts results in non-linear effects that impact the deflection angles.

To glean some physical intuition of the lensing effects at play in a multi-plane system, we adopt a definition of the lensing surface mass density for multi-plane systems that encodes redshift-dependent lensing efficiency, and non-linear coupling between different lens planes. We define $\kappa_{\text{effective}}$, the *effective multi-plane convergence*, as

$$\kappa_{(\text{effective})} \equiv \frac{1}{2} \nabla \cdot \boldsymbol{\alpha} \quad (9)$$

where $\boldsymbol{\alpha}$ is the deflection field of the lens system, or the mapping from a coordinate on the sky to a position in source plane through multi-plane ray-tracing.

This definition expresses the convergence of a multi-plane realization in terms of deflections angles (α_x, α_y) rather than a lensing potential, but is equivalent to the usual definition of convergence in the case of a single lens plane.⁴ We compute these deflection angles by ray-tracing through the line of sight according to Equation 8. To obtain an effective substructure convergence $\kappa_{\text{sub}(\text{effective})}$, we simply subtract the convergence profile of the main deflector κ_{macro} (the macromodel), from the full $\kappa_{(\text{effective})}$.

The definition of κ in Equation 9 permits a comparison between single plane and multi-plane ‘convergence’ maps. For illustrative purposes, in Figure 2, we render a full multi-plane realization of NFW halos between $10^{5.7}$ and $10^{10} M_{\odot}$, for a CDM and WDM scenario. The far left panels show only the single-plane realizations of the subhalo mass function, as would be present in a typical strong lens halo. The central panels show the single plane realizations plus the a full line of sight realization viewed in projection, with coupling between the multiple lens planes turned off. The lensing properties of this convergence map correspond to adopting the Born approximation in lensing, in which lensing quantities are computed by assuming the light rays follow unperturbed paths through the lens planes in front of and behind the main deflector. The far right panels show the *effective multi-plane convergence* for these realizations. In Appendix B, we compare flux ratios computed with the Born approximation to those computed with full ray-tracing, and find the two approaches yield significantly different observables.

Comparing the mass distribution in the far left panels with those on the far right suggests the inclusion of line of sight objects will dramatically affect the statistics of flux ratio distributions in strong lenses caused by small scale density fluctuations in the projected mass density. In the following sections, we will show that this is indeed the case.

⁴ Convergence is equivalent to the projected surface mass density in units of the critical density for lensing $\Sigma_{\text{crit}} = \frac{c^2}{4\pi G} \frac{D_s}{D_{ds} D_d}$ in single plane lensing, where subscripts d and s denote the lens and source redshifts. For multiple lens planes, we express κ as a vector-field derived quantity.

3.2 Flux ratio statistics with line of sight halos

We perform a simple experiment to build intuition for the impact of line of sight halos on flux ratio observables. First, we compute a set of image positions \boldsymbol{x} and flux ratios $\boldsymbol{f}_{\text{reference}}$ for a smooth lens mass distribution, which for simplicity we model as an elliptical isothermal-ellipsoid with external shear (SIE+Shear). Next, given a dark matter model with fixed Σ_{sub} and m_{hm} (with $\delta_{\text{los}} = 1$ and a background source size of 40 pc FWHM), we render 1,000 realizations of halos this model from Equations 3 and 4. For each of these realizations, we optimize a smooth model to fit the image positions, and compute the model flux ratios \boldsymbol{f}' with respect to this optimized lens model. We then compute the summary statistic⁵

$$S_{\text{smooth}}(\boldsymbol{f}', \boldsymbol{f}_{\text{reference}}) \equiv \sqrt{\sum_{i=1}^3 (f'_i - f_{\text{reference}(i)})^2}. \quad (10)$$

The statistic S_{smooth} encodes the amount of flux ratio anomaly with respect to a smooth lens model induced by the presence of dark matter halos. In principle, the distributions of this statistic depend on the reference smooth lens model used to compute $\boldsymbol{f}_{\text{reference}}$, but since we construct these distributions merely for visualization purposes the choice of smooth model is not crucial. These complications notwithstanding, we note that the SIE+Shear profile used to compute S_{smooth} reasonably describes the large-scale mass profile of a typical deflector (Auger et al. 2010; Gilman et al. 2017).

Figure 3 shows distributions of S_{smooth} for different lens (source) redshifts of 0.5 (2) and 0.8 (3) with different dark matter models. The addition of line of sight halos increases the frequency of a flux ratio anomaly with respect to a smooth lens model, and the boost is substantially higher for configurations with higher lens and source redshifts. The inclusion of line of sight structure also increases the difference in relative amplitudes between the CDM and WDM (solid black and magenta curves) relative to models with lens plane subhalos only. Finally, the distribution of summary statistics for a CDM mass function with a high normalization (grey dotted curve) resembles the statistics produced in a WDM model with a lower value of Σ_{sub} . This reflects a degeneracy between the amplitude of the subhalo mass function in the main lens plane, and the turnover scale in the mass function.

In the next Section, we amend the definition of the summary statistic in Equation 10 slightly, replacing $\boldsymbol{f}_{\text{reference}}$ with a set of observed fluxes from a strong lens $\boldsymbol{f}_{\text{obs}}$. We write this new statistic S_{lens} as

$$S_{\text{lens}}(\boldsymbol{f}', \boldsymbol{f}_{\text{obs}}) \equiv \sqrt{\sum_{i=1}^3 (f'_i - f_{\text{obs}(i)})^2}. \quad (11)$$

Through the forward model, we will attempt to minimize this statistic by computing flux ratios \boldsymbol{f}' with different dark matter mass functions. In Appendix C, we describe the implementation of a fast algorithm for lens model optimiza-

tions with many line of sight halos, which we use to compute the statistic in Equation 11.

4 SIMULATIONS OF SUBSTRUCTURE LENSING: SETUP AND METHODOLOGY

In this section, we describe the setup of simulations designed to project the constraining power of flux ratios on a WDM mass function. We first outline the physical assumptions imposed in the simulations, and the priors on the parameters sampled in the forward model. Next, we walk through the forward modeling procedure. The subsequent section describes our implementation of flux uncertainties, both from measurement errors and lens modeling. We then describe how, after accounting for uncertainty in the image fluxes, we construct posterior distributions for the model parameters. Finally, we describe the procedure for creating simulated datasets we will use to test this machinery and make forecasts.

4.1 Physical assumptions and priors

The methodology we present is flexible, and accommodates any parameterization for the quantities such as the subhalo mass function, line of sight halo mass function, main deflector mass profile, etc. However, for the purpose of making forecast statements and presenting the methodology, we make several simplifying assumptions regarding the implementation of dark matter physics, mass models, and lensing quantities.

First, we do not marginalize over the mass, concentration, or ellipticity of the host dark matter halo. We assume a halo mass of $10^{13}M_{\odot}$, which is typical for a lensing galaxy (Gavazzi et al. 2007), when distributing halos spatially and evaluating the two-halo term in Equation 3. We do not expect the ellipticity of the parent dark matter halo to affect the lens model predictions for image fluxes, since the ellipticity of the lensing galaxy and external shear dominate the quadrupole moment of the mass distribution (Keeton et al. 1997). We also ignore any redshift dependence in the subhalo mass function, although we evolve the line of sight halo mass function evolve with redshift. With these simplifications, the subhalo mass function in Equation 4 takes the form

$$\frac{d^2 N_{(13)}}{dm dA} = \frac{\Sigma_{\text{sub}}}{m_0} \left(\frac{m}{m_0} \right)^{-\alpha} \quad (12)$$

where the subscript (13) refers to the assumed halo mass of $10^{13}M_{\odot}$. We assume $\alpha = 1.9$ (Springel et al. 2008; Fiacconi et al. 2016). We derive a projected mass density in subhalos by integrating Equation 12 over mass, and find values of $\Sigma_{\text{sub}} \sim 0.01 - 0.02 \text{ kpc}^{-2}$ yield surface mass densities in substructure similar to those derived in simulations of early-type galaxy halos of $10^7 M_{\odot} \text{ kpc}^{-2}$ with a pivot mass of $m_0 = 10^8 M_{\odot}$ (Fiacconi et al. 2016). This normalization in principle depends on the severity of tidal stripping, the host halo mass, the halo redshift, and the halo formation time. Rather than modeling all of these effects from first principles, we subsume them in the normalization Σ_{sub} , and impose a wide (flat) prior on this parameter between $0 - 0.045 \text{ kpc}^{-2}$. Gilman et al. (2018) demonstrate that the

⁵ The summation i runs over the three flux ratios derived from the four image fluxes.

mean normalization in the lens sample effectively scales the information content available per lens; we perform the same analysis in this work, examining how the constraints on dark matter respond to different values of Σ_{sub} .

Regarding the implementation of WDM mass functions, we assume that the parameterization of the mass function turnover near m_{hm} (Equation 5) applies to both halos along the line of sight, and for subhalos in the main lens halo. As we vary the half-mode mass m_{hm} between $10^{4.8} - 10^{10} M_{\odot}$, none of the models considered are truly ‘cold’ in the sense of GeV-scale WIMPS with free-streaming masses of order an Earth mass. However, provided $m_{\text{hm}} \ll m_{\text{low}} = 10^6 M_{\odot}$, the halo populations rendered result in the same observables as those in a CDM universe.⁶ We therefore interpret inferences that favor models with $m_{\text{hm}} < 10^6 M_{\odot}$ as consistent with CDM, even though the true half-mode mass may be in fact be much lower than the value we recover. Finally, while we implement scatter and redshift dependence in the mass concentration relation in Equation 7, we do not marginalize over the parameters describing the turnover for WDM models.

We vary the rescaling parameter for the line of sight halo mass function between 0.7 and 1.3. This accounts for uncertainties the prediction of the halo mass function which is typically at the 10 – 30% level (Despali et al. 2016), and accounts for variance in the average density along the line of sight to strong lenses.

The background source size enters the forward model because the perturbation to image magnifications depends on the source size relative to the deflection angle of a perturber (Dobler & Keeton 2006). Upper limits on the size of the narrow-line region from (Nierenberg et al. 2017) correspond to physical sizes of ~ 50 pc, which agrees with the surface brightness profiles seen in low redshift AGN (Müller-Sánchez et al. 2011). We therefore allow the source size to vary between 25 and 50 pc.

We model the main deflector as a power-law ellipsoid plus external shear. This is a generalization of the widely applied, physically motivated (e.g. Treu et al. 2006) singular isothermal sphere (SIE) profile used to model lensing galaxies. Studies of early-type deflectors find mass profiles $\rho(r) \sim r^{-\gamma_{\text{macro}}}$ modestly steeper than r^{-2} (Treu et al. 2009; Auger et al. 2010; Shankar et al. 2017), so we allow the power-law profile γ_{macro} to vary between 2 and 2.2. We assume deflectors with complex morphologies, including features like stellar disks, have been identified and removed from our sample, and describe residual baryonic effects by adding perturbations to the forward model image fluxes, a process we describe in Section 4.3. We marginalize over uncertainties in image positions by rendering Gaussian astrometric uncertainties of ± 3 m.a.s. in the forward model.

As a final note, we point out that many of the simplifying assumptions we impose in our forecasts effectively ignore relevant information that could be used to inform a prior. For example, the velocity dispersion of the lensing galaxy

could inform a prior on the halo mass and the normalization Σ_{sub} , and possibly the macromodel profile γ_{macro} . Since Σ_{sub} is somewhat correlated with m_{hm} (see Section 5), this could improve constraints on the free-streaming length of the dark matter. Similarly, modeling redshift dependence in the normalization of the subhalo mass function could break the covariance between Σ_{sub} and δ_{los} (see Section 5). This information would therefore improve the precision on the inferred dark matter properties, and it is possible that we overestimate uncertainties by omitting it.

4.2 Forward modeling procedure

To constrain the halo mass function, we adopt a forward modeling approach. This consists of generating mock data sets by simulating the physical processes that affect lensing observables, including the size of the background source, dark matter halos in the main lens halo and along the line of sight, the mass profile of the main deflector, and statistical measurement errors. This approach handles complicated degeneracies between model parameters - for example, between halo redshift and halo mass (e.g. Despali et al. 2018b) - by building these features directly into the forward-generated data sets. In effect, we exchange the task of computing a complicated likelihood function with the challenge of simulating the relevant physics in strong lensing.

This first step in the forward model is to sample all parameters from their respective prior probability densities, summarized in Table 1. For convenience, for the i th realization, we denote the collection of the model parameters \mathbf{M}_i . Using the parameters describing the dark matter ($\Sigma_{\text{sub}}, \delta_{\text{los}}, m_{\text{hm}}$), we render a the full population of line of sight halos and lens plane subhalos, as described in Section 2.

Next, using the observed image positions⁷ and fluxes from a strong lens, we optimize a power-law plus external shear lens model with power law slope γ_{macro} to fit the observed image positions in the presence of the full population of dark matter halos, and ray-trace to compute the flux ratios with background source modeled as a Gaussian with a FWHM of σ_{src} . While optimizing the macromodel to fit image positions, we allow the lens Einstein radius, centroid, ellipticity, ellipticity angle, shear, and shear angle to vary, while keeping the power-law slope γ_{macro} fixed for each optimization. If necessary, we may extend the forward modeling of γ_{macro} to additional mass profile parameters to add complexity in the lens macromodel.

At this stage, we have a set of observed flux ratios and a set of flux ratios simulated in the forward model. We use the model-predicted flux ratios \mathbf{f}' with the observed flux ratios \mathbf{f}_{obs} to compute the summary statistic in Equation 11, which we then assign to the set of parameters \mathbf{M}_i . We repeat this entire procedure 600,000 times for each lens (see the convergence test in Appendix D).

⁶ This is only true if the signal in flux ratio saturates at m_{low} , otherwise we would miss part of the signal from halos with mass $< m_{\text{low}}$. We verify that halos of mass below $10^6 M_{\odot}$ do not significantly affect the flux ratio signal for the background source sizes 25-50 pc.

⁷ We add random statistical measurement errors of ± 3 m.a.s. to the image positions for each realization.

Table 1. Parameters sampled in the forward model

parameter	definition	prior
$\Sigma_{\text{sub}} [\text{kpc}^{-2}]$	normalization of subhalo mass function (Equation 12) (rendered between $10^6 - 10^{10} M_{\odot}$)	uniform: [0, 0.045]
$m_{\text{hm}} [M_{\odot}]$	half-mode mass (Equations 5 and 7) \propto to free streaming length and thermal relic mass m_{DM}	log-uniform: [4.8, 10]
δ_{los}	rescaling factor for the line of sight Sheth-Tormen mass function (Equation 3, rendered between $10^6 - 10^{10} M_{\odot}$)	uniform: [0.7, 1.3]
$\sigma_{\text{src}} [\text{pc}]$	source size parameterized as FWHM of a Gaussian	uniform: [25, 50]
γ_{macro}	logarithmic slope of main deflector mass model	uniform: [2, 2.2]
$\delta_{xy} [\text{m.a.s.}]$	image position uncertainties	$\mathcal{N}(0, 3)$

4.3 Accounting for uncertainty in image fluxes

We introduce uncertainties in the image fluxes by adding perturbations to the fluxes in the mock data, and by rendering these perturbations in the model fluxes. Explicitly, we modify each model-predicted image flux f_i as

$$f_i \rightarrow f_i + \mathcal{N}(0, \delta). \quad (13)$$

The most straightforward interpretation of this procedure is the incorporation of statistical measurement errors. For reference, current measurements of narrow-line fluxes achieve precision of 3 – 6% (Nierenberg et al. 2014, 2017). These perturbations also simulate the role of unknown sources of uncertainty, or simply those we do not explicitly model. For example, in cases where a more complex macromodel is required, the additional degrees of freedom that must be marginalized over result in a larger variation in image fluxes at fixed image positions, which effectively introduces an additional source of flux uncertainty.

We will explicitly consider flux perturbations of 2%, 4%, 6%, and 8%. The intermediate values of 4% and 6% represent current measurement precision (Nierenberg et al. 2017) and modeling uncertainties (Gilman et al. 2017). The 2% value represents a best-case scenario with precise measurements — perhaps with observations from future telescopes such as JWST — and a sample of morphologically simple deflectors that do not require complex macromodels. The 8% value corresponds to a scenario where the majority of the systems in the lens sample require marginalizations over complex macromodels.

4.4 Bayesian Inference

To construct posterior probability densities for the parameters \mathbf{M} listed in Table 1, we rank the 600,000 \mathbf{M}_i by their summary statistics, with those that minimize the statistic ranked highest. A subset of these models (we use the top 1,500) form a probability density $p'(\mathbf{M}|\text{data})$, which becomes an increasingly good approximation of the true posterior distribution $p(\mathbf{M}|\text{data})$ as the number of forward model samples increases. This procedure falls in the category of Approximate Bayesian Computing methods (for a review, see (Lintusaari et al. 2017)), and is widely applied to

problems with intractable likelihood functions (Akeret et al. 2015; Hahn et al. 2017; Birrer et al. 2017b; Davies et al. 2018). We apply a kernel density estimator to the 1,500 sample that form $p'(\mathbf{M}|\text{data})$, and multiply the resulting probability densities to obtain the final posterior. We test for convergence in this algorithm in Appendix D.

We acknowledge that, formally, a marginalization of the macromodel, rather than an optimization of the macromodel, yields the desired posterior distribution of dark matter parameters. We avoid this computationally prohibitive step⁸ with two justifications: First, the volume of macromodel parameter space is typically tightly constrained by the requirement that the macromodel fit the image positions. For macromodels parameterized as power-law ellipsoids, the image fluxes do not vary significantly over this volume, and the variation in image fluxes induced by marginalizing over the macromodel is negligible compared to other sources of uncertainty⁹. Second, we note that for each of the 600,000 realizations rendered in the forward model, each macromodel re-optimization is independent. Thus, over the course of many realizations, covariance between macromodel parameters and the parameters describing the dark matter content is reflected in the summary statistics.

4.5 Creating simulated data sets

To create mock data sets, we parameterize the lens macromodel as a power-law ellipsoid, and generate mock lenses by sampling the Einstein radii, ellipticity, and external shears, as well as lens and source redshifts, from the distributions of these quantities used by Oguri & Marshall (2010). We plot the lens and source redshifts of the 50 quads in our mock lens sample in Figure 4. We sample power law slopes drawn from a distribution centered at 2.05 ± 0.04 , consistent with

⁸ This is computationally prohibitive because the vast majority of macromodel parameter configurations do not fit the image positions, and therefore consume computation time without contributing to the desired posterior distribution.

⁹ We test this by re-sampling a once-optimized macromodel around the peak of the likelihood, and computing the variation in image fluxes.

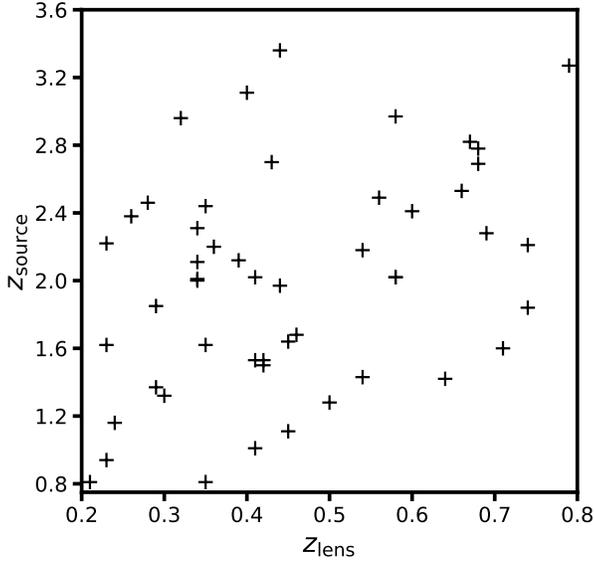


Figure 4. The lens and source redshifts for the 50 quads in our mock lens sample. We draw these parameters, along with the lens velocity dispersion, ellipticity, and shear from the distributions used by Oguri & Marshall (2010).

the morphological properties of the early-type galaxies that dominate the strong lensing cross section (Auger et al. 2010; Shankar et al. 2017). The background source is parameterized by a circular Gaussian with a FWHM, which we specify within the range 25 – 50 pc, consistent with the upper limits on the size inferred by Nierenberg et al. (2017), and comparable to the luminous extent of the narrow line region of quasars (Müller-Sánchez et al. 2011). Finally, we choose background source positions to produce roughly equal numbers of cross, fold, and cusp image configurations. When generating the mock data sets, we add measurement errors to the image positions of 3 m.a.s., and model statistical measurement errors by adding perturbations to the image fluxes, as described in Section 4.3.

5 SIMULATIONS OF SUBSTRUCTURE LENSING: RESULTS

This section presents the results of our analysis, in which we test the forward modeling machinery described in the previous section to constrain dark matter properties. We discuss how measurement and modeling uncertainties affect the precision of constraints on both CDM and WDM mass functions, and make projections for the constraints on the half-mode mass. We explicitly consider 4 models: Two CDM cases with a different normalization of the subhalo mass function Σ_{sub} , and two WDM cases with half-mode masses of $10^{7.7}M_{\odot}$, and 10^7M_{\odot} .

5.1 Joint inference on model parameters

Beginning with the CDM mass functions, in Figure 5 we show posterior distributions for all the parameters sampled in the forward model for a CDM mass function with a normalization of $\Sigma_{\text{sub}} = 0.018 \text{ kpc}^{-2}$. As described in Section

4, we add flux perturbations of 2%, 4%, and 6% the mock data and model fluxes to simulate measurement errors, and additional sources of flux uncertainty that stem from lens modeling. We marginalize over ten realizations of these flux perturbations to reduce shot noise in the posterior distributions.

The boost in signal from the line of sight halos permits 2σ bounds on the half-mode mass that range between $m_{\text{hm}} < 10^{7.1}M_{\odot}$, or a 7.9 keV thermal relic particle, to $m_{\text{hm}} < 10^{8.8}M_{\odot}$ (2.4 keV) as statistical measurement errors and modeling uncertainties in image fluxes increase from 2% to 6%. This rapid erosion of constraining power underscores the necessity of accurately measuring image fluxes, and accurate lens model predictions for these observables.

The most visibly striking covariance in Figure 5 exists between Σ_{sub} and m_{hm} (see also Figure 6). Physically, this feature corresponds to adding more substructure by increasing the normalization, and subsequently removing some of the subhalos by raising the half-mode mass such that the total amount of flux perturbation remains relatively constant. Thus, above a sensitivity threshold of roughly 10^6M_{\odot} , Σ_{sub} and m_{hm} are positively correlated. The opposite is true for Σ_{sub} and δ_{los} : the additional source of flux perturbation from extra line of sight structure is partially offset by reducing the number of lens plane subhalos, and these parameters are anti-correlated. Finally, there is weak evidence (notice the slightly tilted 2σ contours) for a positive correlation between the power-law slope of the macromodel γ_{macro} and the source size σ_{src} . Without a priori knowledge of the true source size, the focusing power of a lens with a steeper mass profile makes larger background sources look smaller. Thus, a more extended background source is focused to the same size image by a steeper mass profile and these parameters are positively correlated. We emphasize that despite the covariance between parameters such as Σ_{sub} and m_{hm} , the data still constrains these parameters independently. The covariance affects the precision of the inference, but it does not result in completely unconstrained posterior distributions.

The normalization of the subhalo mass function Σ_{sub} plays an important role in the constraints on WDM and CDM models. Systems with more substructure are effectively weighted more than systems with fewer subhalos, and the strength of the constraints on m_{hm} reflect this weighting. We illustrate this effect in Figure 6, through comparison with Figure 5. The former has $\Sigma_{\text{sub}} = 0.01 \text{ kpc}^{-2}$, while the latter has nearly twice as many lens plane subhalos with $\Sigma_{\text{sub}} = 0.018 \text{ kpc}^{-2}$. The constraints on m_{hm} are weaker for the simulation with less substructure, because the data contains less signal. Due to the covariance between Σ_{sub} and m_{hm} , a significant portion of the volume of the posterior lies in high Σ_{sub} , high m_{hm} parameter space, which results in a peak in the marginalized constraint on m_{hm} . Stronger theoretical priors on Σ_{sub} , which take into account the role of halo mass, redshift, and tidal stripping, may improve constraints on m_{hm} by breaking this covariance.

It is possible that by extending the range of the prior on Σ_{sub} to higher values, the covariance between m_{hm} and Σ_{sub} would result in weaker constraints on the half-mode mass. However, extending the prior in this manner would imply a degree of ignorance surrounding the parameter Σ_{sub} that would likely be exaggerated given the current state of numerical simulations of dark matter halos and their substructure.

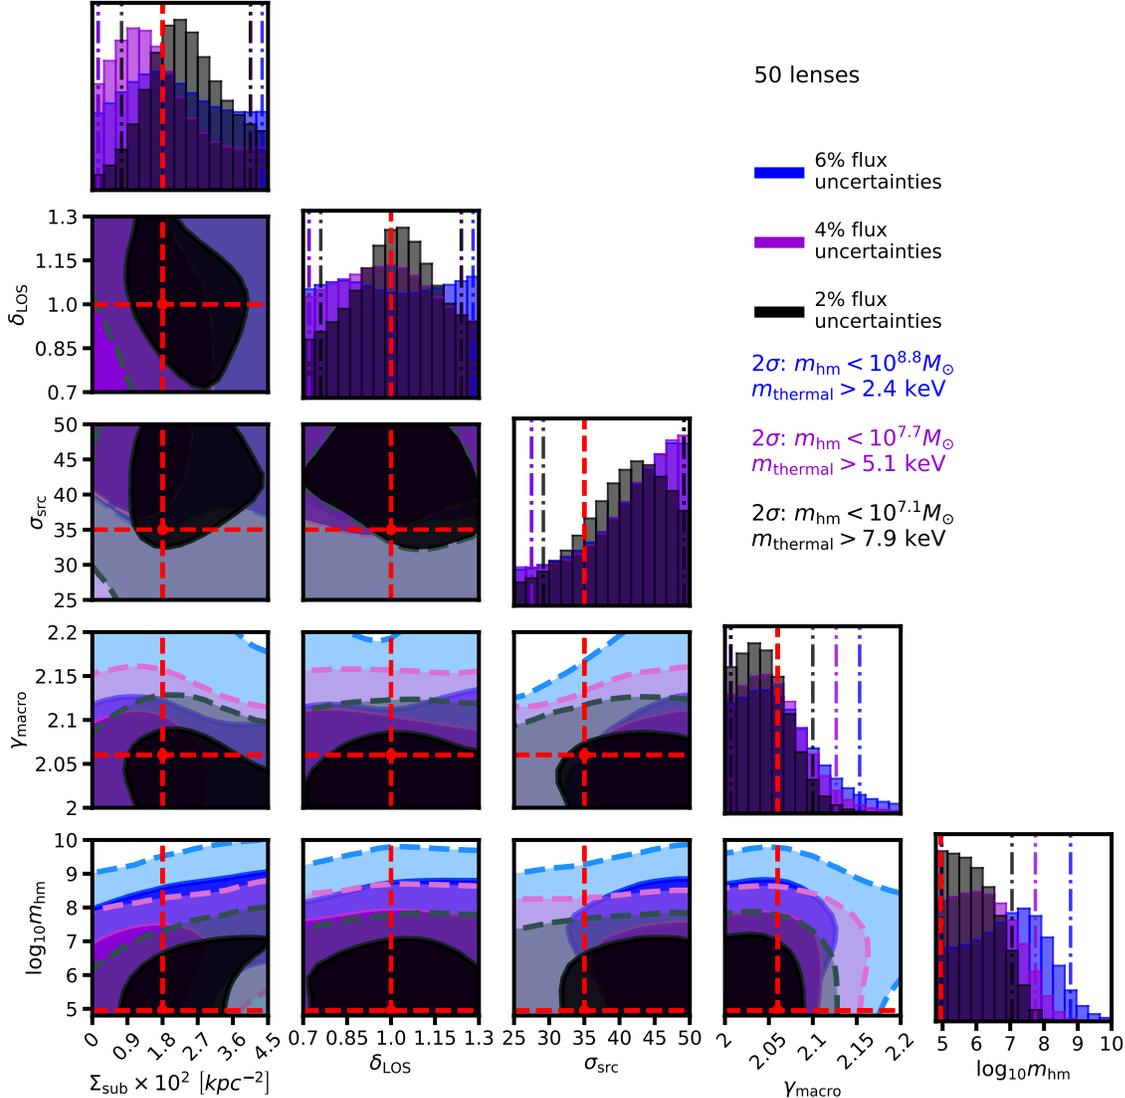


Figure 5. The posterior distributions resulting from the forward modeling analysis of a sample of 50 lenses, with flux uncertainties stemming from measurement errors and lens modeling controlled at the 2%, 4%, and 6% level. Vertical bars in the marginal distribution indicate 2σ confidence interval, while dashed (solid) lines in the panels denote 2σ (1σ) contours. The marginalized constraints on m_{hm} range from $10^{7.2}M_{\odot}$ for the case of 2% flux uncertainties, to $10^{8.8}M_{\odot}$ for uncertainties of 6%.

ture (Benson 2012; Wheeler et al. 2018; Bozek et al. 2018; Lovell et al. 2018). Keeping the width of the prior fixed, we implicitly assume that one may predict Σ_{sub} for each lens halo to within the width a factor of 4.5, or the width of the prior on Σ_{sub} .

In Figures 7 and 8, we show the constraints on WDM mass functions with m_{hm} of $10^{7.7}M_{\odot}$ and 10^7M_{\odot} , which correspond to thermal relic dark matter particles of 5.1 and 8.2 keV, respectively. Both datasets have $\Sigma_{\text{sub}} = 0.012 \text{ kpc}^{-2}$. As in Figures 5 and 6, we marginalize over every parameter listed in Table 1, but focus only on the joint distribution of Σ_{sub} and m_{hm} . We see evidence for a turnover in the mass function, even though it lies below 10^8M_{\odot} . When interpreting the marginalized posteriors for m_{hm} in cases where there

is a clear peak in WDM territory, we use the relative likelihood between the lowest m_{hm} bin (at $10^{4.8}M_{\odot}$) and the peak of the posterior as a summary statistic, since the statement regarding the 2σ confidence interval depends on the width of the prior.¹⁰

In the case of $m_{\text{hm}} = 10^{7.7}M_{\odot}$, with flux uncertainties of 2%, 4%, and 6%, we favor WDM mass functions with

¹⁰ Sometimes, inference on CDM mass functions results in a posterior distribution peaked around some value of m_{hm} , due to the covariance between m_{hm} and other parameters. This effect is visible in Figure 6. In the case of Figure 6, the maximum likelihood ratio between WDM and CDM with uncertainties of 2% equals two.

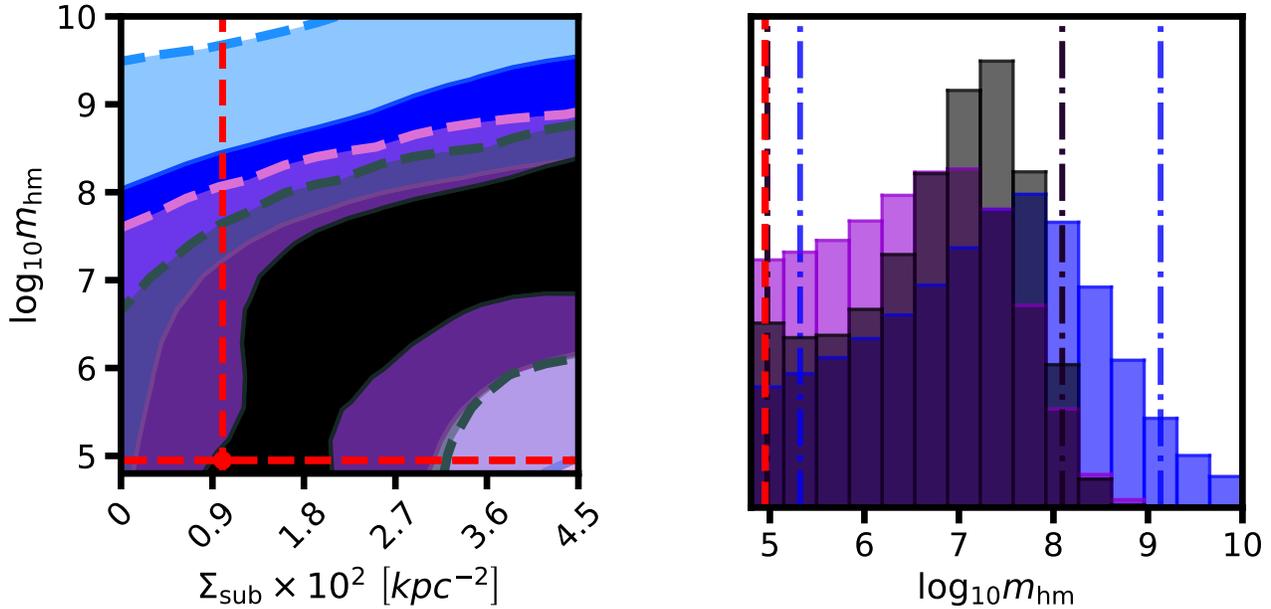


Figure 6. Inference on a CDM mass function with a normalization of the subhalo mass function $\Sigma_{\text{sub}} = 0.01$, roughly half the value of the normalization assumed in Figure 5. The color scheme is the same as in Figure 5, with black, magenta, and blue representing flux uncertainties of 2%, 4%, and 6%, respectively. In this case, the marginalized constraints on m_{hm} are $10^{9.1} M_{\odot}$, $10^{8.1} M_{\odot}$, and $10^{8.1} M_{\odot}$ for flux uncertainties of 6%, 4%, and 2% (for the 4% and 2% flux uncertainties, the 2σ confidence interval happen to be the same). These constraints are weaker by roughly an order of magnitude in mass over the bounds quoted in Figure 5, which illustrates the role of the normalization of the subhalo mass function on the possible constraints on m_{hm} .

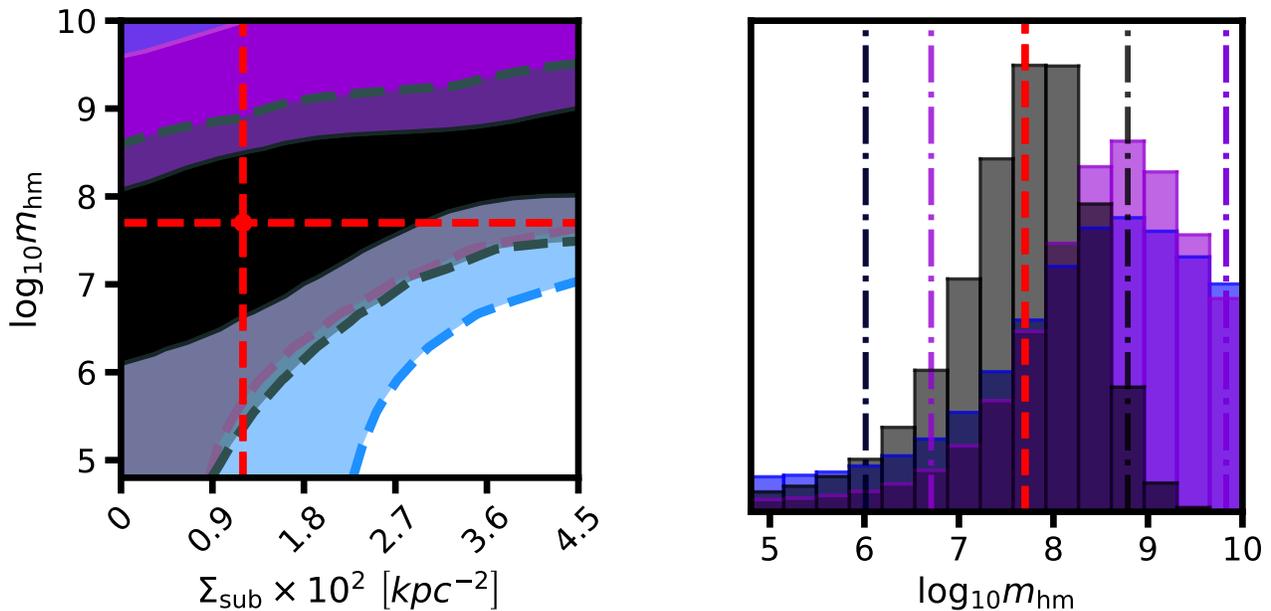


Figure 7. Inference on a WDM mass function with a half-mode mass of $10^{7.7} M_{\odot}$ ($m_{\text{thermal}} = 5.4 \text{ keV}$), with the same color scheme as Figure 5. As in Figure 5, we marginalize over the parameters listed in 1 and over various degrees of flux uncertainty, and the color scheme is the same as in Figures 5 and 6. For flux uncertainties of 2%, 4%, and 6%, we favor WDM with $m_{\text{hm}} > 10^{7.7} M_{\odot}$ over CDM with likelihood ratios of 22, 30, and 8, respectively.

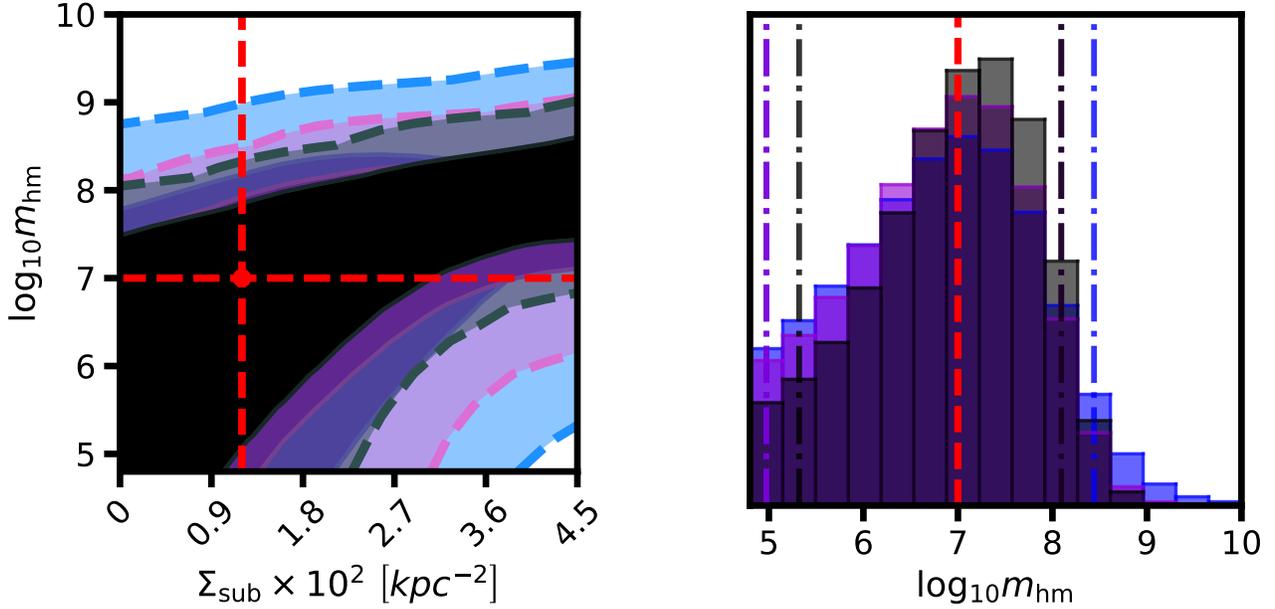


Figure 8. Inference on a WDM mass function with a half-mode mass of $10^7 M_\odot$ ($m_{\text{thermal}} = 8.2 \text{ keV}$), marginalized over the parameters listed in Table 1, with the same color scheme as Figure 5. For each degree of uncertainty in image fluxes, the peak of the posterior coincides with the location of the the turnover at $10^7 M_\odot$, but the width of the distributions increases. With uncertainties of 2%, 4%, and 6% we favor WDM mass functions with $m_{\text{hm}} > 10^7 M_\odot$ over CDM with likelihood ratios of 4:1, 3:1, and 2:1, respectively.

$m_{\text{hm}} > 10^{7.7} M_\odot$ over CDM with relative likelihoods of 22:1, 30:1, and 8:1, respectively¹¹. With uncertainties of 4% and 6%, the posterior distributions of m_{hm} shift towards higher masses, and the posteriors no longer resolve the position of the turnover in the mass function and mass-concentration relation. The shift to higher values of m_{hm} is a consequence of the weak signal produced by very warm mass functions with a paucity of small-scale structure. Increased flux uncertainties wash out the information from the ‘weak signal’ regime of parameter space with $m_{\text{hm}} > 10^{7.7} M_\odot$, and the constraints on this region of parameter space deteriorate because the data itself lies in this ‘weak signal’ regime. This reasoning is similar to the interpretation of Σ_{sub} as an information scaling parameter for CDM mass function: like a CDM mass function with a high normalization, a ‘colder’ WDM mass function produces more significant flux perturbation events, and is more resilient to increased uncertainties in image fluxes. If this reasoning is correct, we should expect the posteriors on m_{hm} for ‘colder’ WDM mass function to remain relatively stationary, modulo an increased variance, after adding perturbations to the image fluxes.

This effect is seen in Figure 8, which has $m_{\text{hm}} = 10^7 M_\odot$. The shift of the posterior distributions towards higher masses as flux uncertainties increase does not happen in this case because the WDM mass function with $m_{\text{hm}} = 10^7 M_\odot$ produces stronger perturbations in the data than the warmer, ‘weak signal’ model with $m_{\text{hm}} = 10^{7.7} M_\odot$. This is because the halos are both more numerous and more concentrated that the WDM model with $m_{\text{hm}} = 10^{7.7} M_\odot$. In turn, the stronger signal survives additional flux uncer-

tainties, and is sufficient to constrain very warm mass functions. The locations of the peaks of the posteriors coincide with the true value of m_{hm} , but the width of the distributions widen. In this case, we favor WDM mass functions with $m_{\text{hm}} > 10^7 M_\odot$ over CDM mass functions with relative likelihoods of 4:1, 3:1, and 2:1 with flux uncertainties of 2%, 4%, and 6%, respectively. The fact that we statistically favor WDM models over CDM models suggests that we could infer a turnover in the mass function at $m_{\text{hm}} = 10^7 M_\odot$ (or an 8.2 keV WDM particle) at higher significance with a larger sample of quads.

5.2 Marginalized constraints on the free-streaming length

The posterior distributions in Figures 5 and 6 give a sense for how the constraints on the half-mode mass in WDM models depends on the precision with which one measures image fluxes and predicts them with lens models, and on parameters such as the normalization of the subhalo mass function. To take into account sample variance, in Figure 9 we plot the marginalized constraints on the half-mode mass as a function of the number of lenses, Σ_{sub} , and flux measurement uncertainties of 2%, 4%, 6%, and 8%. We plot the bounds on m_{hm} for both a high ($\Sigma_{\text{sub}} = 0.02 \text{ kpc}^{-2}$) and low ($\Sigma_{\text{sub}} = 0.008 \text{ kpc}^{-2}$) normalization of the subhalo mass function. To produce these curves, we compute 200 bootstraps of 50 lenses, and average over many realizations of flux uncertainties.

With a sample of 50 lenses it will be possible to probe below $10^8 M_\odot$ in the halo mass function, to a degree that depends on the amount of substructure in the main deflector, measurement precision of image fluxes, and precise lens

¹¹ The increase from 22 to 30 is likely due to shot noise.

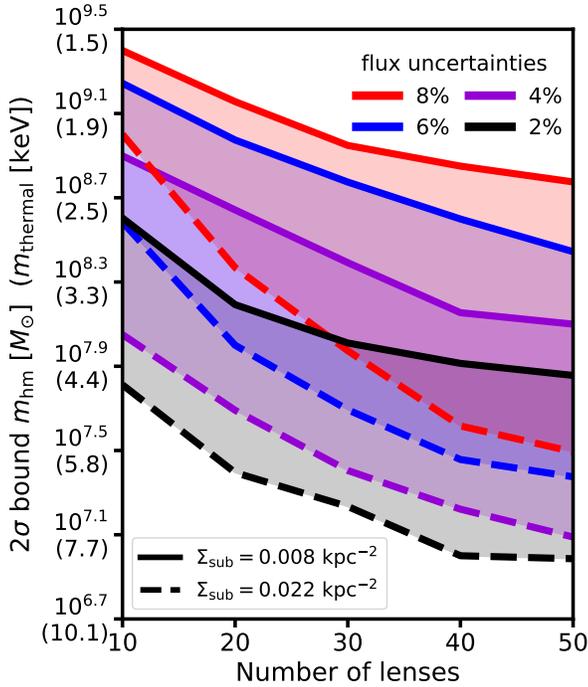


Figure 9. Forecasts for the constraints on the half-mode mass as a function of the number of lenses, including line of sight halos and subhalos of the main deflector. Black, purple, blue, and red colors denote flux uncertainties of 2%, 4%, 6%, and 8%. The solid line corresponds to a normalization $\Sigma_{\text{sub}} = 0.008 \text{kpc}^{-2}$, while dashed lines correspond to $\Sigma_{\text{sub}} = 0.022 \text{kpc}^{-2}$. The y-axis labels represent the 2σ bound on m_{hm} , with the mass of the corresponding thermal relic dark matter particle in parentheses. Models with more subhalos (dashed lines), and hence more signal, are more resilient to flux uncertainties than models with fewer lens planes subhalos (solid lines) and produce stronger constraints on m_{hm} .

model predictions for this observable. With control over image fluxes at the level for 4%, routinely achieved at present (Nierenberg et al. 2014, 2017), the bounds on m_{hm} with 50 quads range between $10^{7.1} - 10^{8.1} M_{\odot}$ for values of Σ_{sub} of 0.01 and 0.022kpc^{-2} , respectively. With more precise predictions of Σ_{sub} made on a lens-by-lens basis, these bounds may improve. We also note that future surveys, such as LSST, WFIRST, and Euclid, will discover hundreds of quads (Oguri & Marshall 2010), so the sample of available quads will eventually be much larger than 50.

6 SUMMARY AND CONCLUSIONS

We have presented a method to perform Bayesian inference on the halo mass function through a forward modeling analysis of image flux ratios in quadruply imaged quasars. We model the contribution from line of sight halos, which boost the signal per lens and permit stronger constraints on the properties of dark matter with fewer systems. We demonstrate the method with a sample of 50 quads, comparable in number to the currently observed sample size, and project the constraints on the free streaming length of a WDM particle under different degrees of flux measurement and lens modeling uncertainties, while marginalizing over parameters

describing the size of the background source, the lens macro-model, and the amplitude of the line of sight halo mass function.

Our key results can be summarized as follows:

- With a sample of 50 quads, we are able to constrain the free streaming length of dark matter on scales below $10^8 M_{\odot}$. Assuming CDM, with mean subhalo mass function normalizations $\Sigma_{\text{sub}} = 0.022 \text{kpc}^{-2}$ (0.008kpc^{-2}) we forecast bounds on the half-mode mass of 10^7 ($10^{7.9}$) M_{\odot} , $10^{7.1}$ ($10^{8.1}$) M_{\odot} , $10^{7.4}$ ($10^{8.4}$) M_{\odot} , $10^{7.5}$ ($10^{8.8}$) M_{\odot} for flux uncertainties of 2%, 4%, 6%, and 8%, respectively. These m_{hm} limits translate to bounds on the mass of thermal relic particles of 8.2 (4.4), 7.7, (3.8), 6.2 (3.1), 5.8 (2.4) keV.
- Line of sight halos contribute substantially to the signal in flux ratios, even dominating the signal in lens systems with higher lens and source redshifts. However, the normalization of the subhalo mass function still plays a key role in scaling the information content per lens, with higher values of this parameter translating into tighter constraints on the mass function. The half-mode mass is also covariant with the normalization, which affects the marginalized constraints on this parameter. These features underscore the importance of theoretical work to predict the projected surface mass density of substructure inside galactic halos with accurate models of baryonic feedback and tidal stripping.
- In the case that dark matter is warm, we are able to infer the location of the turnover in the mass function with 50 quads, even if it lies below $10^8 M_{\odot}$. With a half-mode mass of $10^{7.7} M_{\odot}$, which corresponds to a 5.1 keV thermal relic particle, we favor WDM mass functions with $m_{\text{hm}} > 10^{7.7} M_{\odot}$ over CDM with relative likelihoods of 22:1, 30:1 and 8:1 for flux uncertainties of 2%, 4%, and 6%, respectively. With the same set of flux uncertainties and a half-mode mass of $10^7 M_{\odot}$, we favor WDM with $m_{\text{hm}} > 10^7 M_{\odot}$ over CDM with relative likelihoods of 4:1, 3:1, and 2:1. These constraints will likely improve with additional lenses, which suggests that a future large sample of quads could be used to infer a turnover in the halo mass function at $10^7 M_{\odot}$ at high statistical significance.

Our work is broadly consistent with other studies of the line of sight contribution in substructure lensing. For instance, by ray tracing through N-body simulations, (Xu et al. 2012) compare the frequency of flux anomalies induced by line of sight versus main lens halos, and reach the conclusion that line of sight halos contribute at the same level as subhalos. More recently, (Despali et al. 2018b) analyze the role of line of sight halos in the context of gravitational imaging. This method differs somewhat from this analysis in that it aims to detect individual halos along the line of sight, and in the main lens plane, but the authors reach a similar conclusion: the line of sight contribution substantially boosts the signal per lens.

The strength of the constraints on WDM models depend sensitively on the normalization of the subhalo mass function. This is partly due to the interpretation of the normalization as scaling the information content per lens, and also due to the covariance between the normalization and the half-mode mass, although we stress that despite this covariance both parameters can be constrained independently. This highlights the importance of refining theoretical predictions for the value of the normalization, accounting for halo

mass, redshift, and the destruction of subhalos by tidal stripping. To this end, observables from each lens system, such as the central velocity dispersion, half-light radius, redshift, etc. may be used to inform the prior on the normalization and thus further improve the inferred posterior with actual data.

The macromodel used to describe the mass profile of the main deflector plays a key role in this analysis. Several studies demonstrate that simple parameterizations sometimes fail to fit the flux ratios of substructure-less mass profiles, leading to ‘artificial’ flux ratio anomalies in the sense that they do not derive from dark matter substructure (Gilman et al. 2017; Hsueh et al. 2018). However, we note that these cases are dominated by the presence of undetected stellar disks, which are rare in the early-type galaxies that dominate the lensing cross section (Auger et al. 2010; Shankar et al. 2017). Also, we point out that identifying morphologically complexity in the main deflector and modeling it can remove these ‘artificial’ anomalies (Hsueh et al. 2016). While we do not explicitly account morphologically complex deflectors in this work, we do allow some freedom in the macromodel by marginalizing over the power law slope, and account for additional variations in the image fluxes as high as 8% that would result from marginalizing over additional macromodel parameters in the forward model.

Finally, we note that the formalism we present naturally accommodates other parameterizations of the halo mass function, and density profile for individual objects.

ACKNOWLEDGMENTS

We thank Francis-Yan Cyr-Racine, Giulia Despali, Chuck Keeton, Stacy Kim, Alex Kusenko, Leonidas Moustakas, Annika Peter, and Simona Vegetti for helpful suggestions and interesting discussions throughout the course of this project.

DG, TT, and SB acknowledge support by the US National Science Foundation through grant AST-1714953. DG, TT, SB and AN acknowledge support from HST-GO-15177. Support for Program number GO-15177 was provided by NASA through a grant from the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Incorporated, under NASA contract NAS5-26555. AN acknowledges support from the NASA Postdoctoral Program Fellowship.

This work used computational and storage services associated with the Hoffman2 Shared Cluster provided by the UCLA Institute for Digital Research and Education’s Research Technology Group. This work also used computational and storage services associated with the Aurora and Halo super computers. These resources were provided by funding from the JPL Office of the Chief Information Officer.

REFERENCES

- Abazajian K. N., 2017, *Phys. Rept.*, 711, 1
- Akeret J., Refregier A., Amara A., Seehars S., Hasner C., 2015, *JCAP*, 8, 043
- Amara A., Metcalf R. B., Cox T. J., Ostriker J. P., 2006, *MNRAS*, 367, 1367
- Andrade K. E., Minor Q., Nierenberg A., Kaplinghat M., 2019, arXiv e-prints, p. arXiv:1901.00507
- Aprile E., et al., 2018, *Phys. Rev. Lett.*, 121, 111302
- Auger M. W., Treu T., Bolton A. S., Gavazzi R., Koopmans L. V. E., Marshall P. J., Moustakas L. A., Burles S., 2010, *ApJ*, 724, 511
- Baltz E. A., Marshall P., Oguri M., 2009, *JCAP*, 1, 015
- Bayer D., Chatterjee S., Koopmans L. V. E., Vegetti S., McKean J. P., Treu T., Fassnacht C. D., 2018, preprint, (arXiv:1803.05952)
- Benson A. J., 2012, *NewA*, 17, 175
- Benson A. J., et al., 2013, *MNRAS*, 428, 1774
- Birrer S., Amara A., 2018, *Physics of the Dark Universe*, 22, 189
- Birrer S., Amara A., Refregier A., 2015, *ApJ*, 813, 102
- Birrer S., Welschen C., Amara A., Refregier A., 2017a, *JCAP*, 4, 049
- Birrer S., Amara A., Refregier A., 2017b, *JCAP*, 5, 037
- Bode P., Ostriker J. P., Turok N., 2001, *ApJ*, 556, 93
- Bose S., Hellwing W. A., Frenk C. S., Jenkins A., Lovell M. R., Helly J. C., Li B., 2016, *MNRAS*, 455, 318
- Bozek B., et al., 2018, *MNRAS*,
- Bringmann T., Kahlhoefer F., Schmidt-Hoberg K., Walia P., 2017, *Physical Review Letters*, 118, 141802
- Bullock J. S., Boylan-Kolchin M., 2017, *Annual Review of Astronomy and Astrophysics*, 55, 343
- Chen J., Kravtsov A. V., Keeton C. R., 2003, *ApJ*, 592, 24
- Cyr-Racine F.-Y., de Putter R., Raccanelli A., Sigurdson K., 2014, *PhysRevD*, 89, 063517
- Cyr-Racine F.-Y., Moustakas L. A., Keeton C. R., Sigurdson K., Gilman D. A., 2016, *PhysRevD*, 94, 043505
- Cyr-Racine F.-Y., Keeton C. R., Moustakas L. A., 2018, preprint, (arXiv:1806.07897)
- Dalal N., Kochanek C. S., 2002, *ApJ*, 572, 25
- Davies F. B., Hennawi J. F., Eilers A.-C., Lukić Z., 2018, *ApJ*, 855, 106
- Despali G., Vegetti S., 2017, *MNRAS*, 469, 1997
- Despali G., Giocoli C., Angulo R. E., Tormen G., Sheth R. K., Baso G., Moscardini L., 2016, *MNRAS*, 456, 2486
- Despali G., Sparre M., Vegetti S., Vogelsberger M., Zavala J., Marinacci F., 2018a, arXiv e-prints,
- Despali G., Vegetti S., White S. D. M., Giocoli C., van den Bosch F. C., 2018b, *MNRAS*, 475, 5424
- Diaz Rivero A., Cyr-Racine F.-Y., Dvorkin C., 2018, *PhysRevD*, 97, 023001
- Diemer B., 2017, preprint, (arXiv:1712.04512)
- Diemer B., Joyce M., 2018, preprint, (arXiv:1809.07326)
- Dobler G., Keeton C. R., 2006, *MNRAS*, 365, 1243
- Dutton A. A., Macciò A. V., 2014, *MNRAS*, 441, 3359
- Fiacconi D., Madau P., Potter D., Stadel J., 2016, *ApJ*, 824, 144
- Gao L., Frenk C. S., Boylan-Kolchin M., Jenkins A., Springel V., White S. D. M., 2011, *MNRAS*, 410, 2309
- Garrison-Kimmel S., et al., 2017, *MNRAS*, 471, 1709
- Gavazzi R., Treu T., Rhodes J. D., Koopmans L. V. E., Bolton A. S., Burles S., Massey R. J., Moustakas L. A., 2007, *ApJ*, 667, 176
- Gilman D., Agnello A., Treu T., Keeton C. R., Nierenberg A. M., 2017, *MNRAS*, 467, 3970
- Gilman D., Birrer S., Treu T., Keeton C. R., Nierenberg A., 2018, *MNRAS*, 481, 819

- Hahn C., Vakili M., Walsh K., Hearin A. P., Hogg D. W., Campbell D., 2017, *MNRAS*, 469, 2791
- Han J., Cole S., Frenk C. S., Jing Y., 2016, *MNRAS*, 457, 1208
- Hezaveh Y., Dalal N., Holder G., Kisner T., Kuhlen M., Perreault Levasseur L., 2016a, *JCAP*, 11, 048
- Hezaveh Y. D., et al., 2016b, *ApJ*, 823, 37
- Hsueh J.-W., Fassnacht C. D., Vegetti S., McKean J. P., Spingola C., Auger M. W., Koopmans L. V. E., Lagattuta D. J., 2016, *MNRAS*, 463, L51
- Hsueh J.-W., Despali G., Vegetti S., Xu D., Fassnacht C. D., Metcalf R. B., 2018, *MNRAS*, 475, 2438
- Hui L., Ostriker J. P., Tremaine S., Witten E., 2017, *PhysRevD*, 95, 043541
- Inoue K. T., Takahashi R., 2012, *MNRAS*, 426, 2978
- Iršič V., et al., 2017, *PhysRevD*, 96, 023522
- Jiang F., van den Bosch F. C., 2017, *MNRAS*, 472, 657
- Kamada A., Kaplinghat M., Pace A. B., Yu H.-B., 2017, *Phys. Rev. Lett.*, 119, 111102
- Keeton C. R., Kochanek C. S., Seljak U., 1997, *ApJ*, 482, 604
- Kim S. Y., Peter A. H. G., Hargis J. R., 2017a, preprint, ([arXiv:1711.06267](https://arxiv.org/abs/1711.06267))
- Kim S. Y., Peter A. H. G., Wittman D., 2017b, *MNRAS*, 469, 1414
- Kusenko A., 2009, *Phys. Rept.*, 481, 1
- Lintusaari J., Gutmann M. U., Dutta R., Kaski S., Corander J., 2017, *Systematic Biology*, 66, e66
- Lovell M. R., Frenk C. S., Eke V. R., Jenkins A., Gao L., Theuns T., 2014, *MNRAS*, 439, 300
- Lovell M. R., et al., 2018, *MNRAS*, 481, 1950
- Ludlow A. D., Bose S., Angulo R. E., Wang L., Hellwing W. A., Navarro J. F., Cole S., Frenk C. S., 2016, *MNRAS*, 460, 1214
- Macciò A. V., Ruchayskiy O., Boyarsky A., Muñoz-Cuartas J. C., 2013, *MNRAS*, 428, 882
- Mao S., Schneider P., 1998, *MNRAS*, 295, 587
- Metcalf R. B., 2005, *ApJ*, 629, 673
- Metcalf R. B., Madau P., 2001, *ApJ*, 563, 9
- Moore B., Ghigna S., Governato F., Lake G., Quinn T., Stadel J., Tozzi P., 1999, *ApJL*, 524, L19
- Moustakas L. A., Metcalf R. B., 2003, *MNRAS*, 339, 607
- Müller-Sánchez F., Prieto M. A., Hicks E. K. S., Vives-Arias H., Davies R. I., Malkan M., Tacconi L. J., Genzel R., 2011, *ApJ*, 739, 69
- Nierenberg A. M., Treu T., Wright S. A., Fassnacht C. D., Auger M. W., 2014, *MNRAS*, 442, 2434
- Nierenberg A. M., Treu T., Menci N., Lu Y., Torrey P., Vogelsberger M., 2016, *MNRAS*, 462, 4473
- Nierenberg A. M., et al., 2017, *MNRAS*, 471, 2224
- Oguri M., Marshall P. J., 2010, *MNRAS*, 405, 2579
- Read J. I., Walker M. G., Steger P., 2018, preprint, ([arXiv:1808.06634](https://arxiv.org/abs/1808.06634))
- Richings J., Frenk C., Jenkins A., Robertson A., 2018, *arXiv e-prints*,
- Ritondale E., Vegetti S., Despali G., Auger M. W., Koopmans L. V. E., McKean J. P., 2018, *arXiv e-prints*, p. [arXiv:1811.03627](https://arxiv.org/abs/1811.03627)
- Robles V. H., Bullock J. S., Boylan-Kolchin M., 2019, *MNRAS*, 483, 289
- Rocha M., Peter A. H. G., Bullock J. S., Kaplinghat M., Garrison-Kimmel S., Oñorbe J., Moustakas L. A., 2013, *MNRAS*, 430, 81
- Schneider P., 1997, *MNRAS*, 292, 673
- Schneider P., Ehlers J., Falco E. E., 1992, *Gravitational Lenses*, doi:10.1007/978-3-662-03758-4.
- Schneider A., Smith R. E., Macciò A. V., Moore B., 2012, *MNRAS*, 424, 684
- Schneider A., Smith R. E., Reed D., 2013, *MNRAS*, 433, 1573
- Shajib A. J., et al., 2018, *MNRAS*,
- Shankar F., et al., 2017, *ApJ*, 840, 34
- Sheth R. K., Mo H. J., Tormen G., 2001, *MNRAS*, 323, 1
- Shoemaker I. M., Kusenko A., 2009, *PhysRevD*, 80, 075021
- Spergel D. N., Steinhardt P. J., 2000, *Physical Review Letters*, 84, 3760
- Springel V., et al., 2008, *MNRAS*, 391, 1685
- Tinker J. L., Robertson B. E., Kravtsov A. V., Klypin A., Warren M. S., Yepes G., Gottlöber S., 2010, *ApJ*, 724, 878
- Tollet E., et al., 2016, *MNRAS*, 456, 3542
- Tormen G., Diaferio A., Syer D., 1998, *MNRAS*, 299, 728
- Treu T., Koopmans L. V., Bolton A. S., Burles S., Moustakas L. A., 2006, *ApJ*, 640, 662
- Treu T., Gavazzi R., Gorecki A., Marshall P. J., Koopmans L. V. E., Bolton A. S., Moustakas L. A., Burles S., 2009, *ApJ*, 690, 670
- Tulin S., Yu H.-B., 2018, *Phys. Rept.*, 730, 1
- Vegetti S., Koopmans L. V. E., Auger M. W., Treu T., Bolton A. S., 2014, *MNRAS*, 442, 2017
- Vegetti S., Despali G., Lovell M. R., Enzi W., 2018, preprint, ([arXiv:1801.01505](https://arxiv.org/abs/1801.01505))
- Viel M., Becker G. D., Bolton J. S., Haehnelt M. G., 2013, *PhysRevD*, 88, 043502
- Vogelsberger M., Zavala J., Cyr-Racine F.-Y., Pfrommer C., Bringmann T., Sigurdson K., 2016, *MNRAS*, 460, 1399
- Wheeler C., et al., 2018, *arXiv e-prints*,
- Xu D. D., Mao S., Cooper A. P., Gao L., Frenk C. S., Angulo R. E., Helly J., 2012, *MNRAS*, 421, 2553
- Xu D., Sluse D., Gao L., Wang J., Frenk C., Mao S., Schneider P., Springel V., 2015, *MNRAS*, 447, 3189

APPENDIX A: IMPLEMENTING THE TWO-HALO TERM

The two-halo term describes an excess of matter (relative to the mean density of the universe) near a large halo, or a peak in the density field. It is evaluated using the software package COLOSSUS (Diemer 2017), and takes the form

$$\xi_{2\text{halo}}(r, M, z) = b(M, z) \xi_{\text{lin}}(r, z) \quad (\text{A1})$$

where $b(M, z)$ is the halo bias around a mass M , computed with the model presented by Tinker et al. (2010), and

$$\xi_{\text{lin}}(r, z) = \frac{1}{2\pi^2} \int_0^\infty k^2 P(k, z) \frac{\sin(kr)}{kr} dk \quad (\text{A2})$$

is the linear matter-matter correlation function at a distance r . While in principle WDM free-streaming should affect the linear power spectrum $P(k, z)$, we do not model this effect.

We define a boost parameter β in terms of $\xi_{2\text{halo}}$ as

$$\beta(M, z) = \frac{2}{r_{\text{max}} - r_{\text{min}}} \int_{r_{\text{min}}}^{r_{\text{max}}} \xi_{2\text{halo}}(r', M, z) dr' \quad (\text{A3})$$

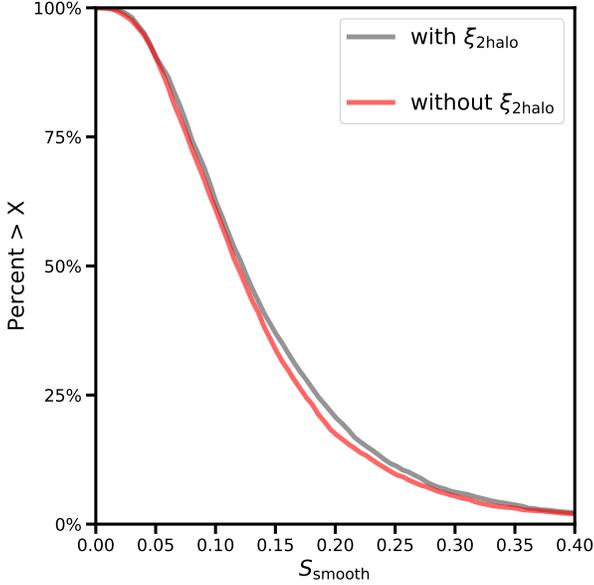


Figure A1. Distributions of the summary statistic S_{smooth} , which represents the amount of flux ratio anomaly with respect to a smooth lens model (see the discussion in Section 3). The grey curve is computed with the two-halo contribution, and the red curve is computed without it. Both models include only line of sight halos to isolate the contribution from $\xi_{2\text{halo}}$. The largest difference between the curves, an offset of 4%, lies at $S_{\text{smooth}} \sim 0.2$.

where M denotes the parent halo mass, and the factor of 2 accounts for symmetry around the parent halo. We choose $r_{\text{min}} = 0.5\text{Mpc}$ and $r_{\text{max}} = 10\text{Mpc}$, which captures most of the contribution from the correlation function while omitting the contribution from regions inside the virial radius of the parent halo. Defining $A_0(z)$ as the normalization of the halo mass function in the lens plane closest to the main lens halo, we incorporate the two halo term by taking $A_0(z) \rightarrow (1 + \beta) \times A_0(z)$, and add these halos at the main lens redshift.

In Figure A1 we plot the distribution of summary statistics S_{smooth} for a CDM mass function that includes the boost from the two-halo term, and one that does not. In both cases, we set to $\Sigma_{\text{sub}} = 0$ to isolate the impact of the two-halo term. The lens and source redshifts are set at 0.6 and 2, respectively. The largest differences between the curves occurs at $S_{\text{smooth}} \sim 0.2$, and is equal to 4%. We conclude that the contribution from $\xi_{2\text{halo}}$ is at most at the level of a few percent, although this may increase if a larger halo mass than $10^{13}M_{\odot}$ is used to evaluate Equation A3.

APPENDIX B: THE BORN APPROXIMATION IN SUBSTRUCTURE LENSING

The Born approximation computes the deflection at each subsequent plane along an unperturbed path. This speeds up lensing computations since a full backwards ray-tracing routine is not required. In Figure B1, we compare the distribution of flux ratio anomalies computed with respect to a smooth lens model (see the discussion in Section 3.2) using the Born approximation, and through full multi-plane ray-tracing. The difference between the solid and dotted curves

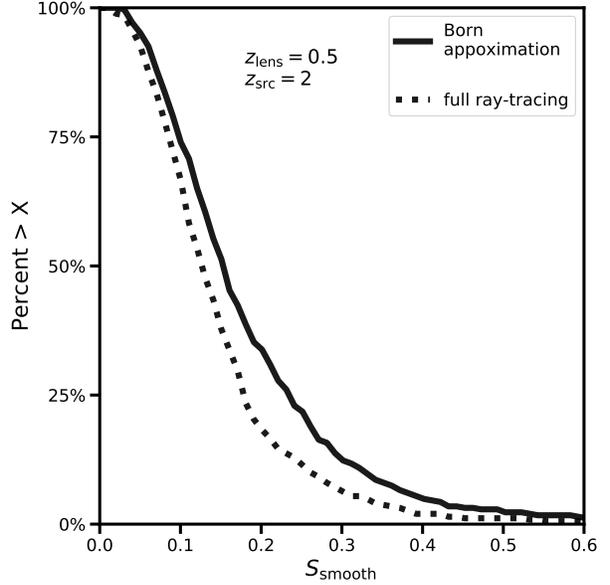


Figure B1. The two curves show distributions of summary statistics computed with respect to a smooth lens model. The curves are computed for the same CDM mass function, with and without the use of the Born approximation. The disagreement between the two curves suggests that the Born approximation does not predict image flux ratios accurately enough to differentiate between dark matter models.

in the figure, which represent flux ratios computed with and without the Born approximation, respectively, is comparable to the difference of WDM and CDM mass functions in Figure 3. Thus, we conclude that full multi-plane ray-tracing approach is required to accurately predict image flux ratios and probe dark matter on small scales.

APPENDIX C: A FAST ALGORITHM FOR MULTI-PLANE LENSING COMPUTATIONS

For each observed lens, our forward modeling approach requires finding a set of macromodel parameters that cast the four light rays in a quadrupole image system to the same location in the source plane. For a single realization, this typically requires hundreds to thousands of backwards ray-tracing computations.

This task is computationally light for models with halos only in front of and at the same redshift as the main deflector because the path through the foreground field of halos is not coupled to the deflections produced in the main lens plane (owing to the recursive nature of Equation 8). Put differently, as soon as one specifies image positions on the sky and draws a realization of dark matter halos, the path through the foreground field is fully determined. In contrast, the path through the field of background halos is coupled to the deflections produced by the macromodel. The path through the background field therefore changes for each new proposal of macromodel parameters. This necessitates repeated computations of the potentially thousands of deflection angles of halos behind the main lens plane, which requires hundreds to thousands as many function evaluations as those needed in single plane lensing computations.

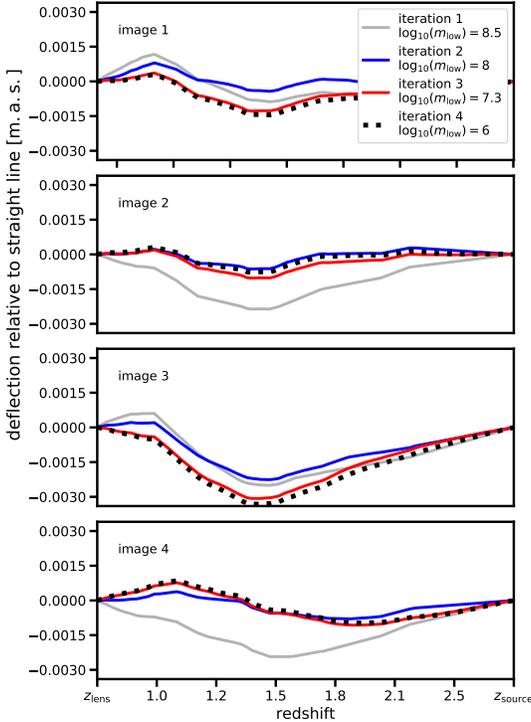


Figure C1. A visualization of the perturbative ray tracing algorithm we use to optimize lens models with potentially thousands of line of sight halos. The panels show the path through the background field relative to a straight line for multiple iterations of the algorithm, in which progressively smaller halos are rendered in progressively smaller apertures around the path of the light rays. This procedure speeds up optimizations of lens models with line of sight halos by at least an order of magnitude.

We address this computational challenge by implementing a perturbative approach to lens model optimizations. First, we optimize the macromodel to fit image positions with only foreground halos and main deflector subhalos present. We denote this optimized lens model \vec{m}^* . This proceeds quickly, since the macromodel deflection angles are not coupled to those from foreground and main lens plane halos. Next, we add the largest background halos with $m > 10^8 M_\odot$, and *re-optimize* \vec{m}^* . Even though the deflections from these massive halos are coupled to those of the macromodel and need to be continuously re-evaluated during the optimization, since there are relatively few of them this proceeds fairly quickly. Next, we add halos in the range $10^{7.5} - 10^8 M_\odot$, but only in 300 m.a.s. apertures around the path of the rays computed with respect to \vec{m}^* . Since the area in which we render these smaller halos is relatively small, and since the macromodel solution \vec{m}^* is already close to the true solution, this optimization also proceeds quickly. We iterate this process for progressively smaller halos until we reach $10^6 M_\odot$.

A visual representation of this process is presented in Figure C1, where we plot the path through the background halos relative to a straight line for subsequent iterations of the perturbative approach. After adding the $10^8 M_\odot$ background halos, the path through the background lens planes

changes only slightly, which reflects the fact that these massive objects dominate the deflection field.

This procedure accomplishes the optimization of a macromodel with background halos 10-50 times faster than a naive optimization with all background halos included simultaneously. We test that the flux ratio statistics are identical to those obtained by ray tracing through full realizations without the perturbative approach implemented. We note that this algorithm is reminiscent of the Born approximation in that it initially neglects the presence of small deflections from subhalos along the line of sight, but differs fundamentally from the Born approximation in that the full non-linear coupling between every subhalo is eventually accounted for.

APPENDIX D: CONVERGENCE OF POSTERIOR DISTRIBUTIONS

We approximate the true posterior distributions for model parameters by retaining the top 1,500 samples (ranked by their summary statistics) out of the 600,000 realizations computed per lens. To test whether this procedure yields an accurate approximation to the true posterior distribution, we appeal to a certain feature of Approximate Bayesian Computing algorithms, namely, that the approximation to the true posterior distribution converges as the number of samples increases. We can therefore test for convergence by applying the same cut on the top 1,500 samples to an ‘under-sampled’ model with only 400,000 realizations per lens, and check that the posterior distribution stays approximately fixed in place. We generate the sample of 400,000 by drawing the realizations randomly from the computed set of 600,000.

We perform this test and plot the results in Figure D1. While there is some movement in the 1σ contours, the 2σ contours trace each other closely. Importantly the constraints on the half-mode mass are the same between the two inferences, which is the most important criterion for our purpose of forecasting bounds on dark matter warmth. Finally, we note that ABC routines tend to yield conservative approximations to the true posterior distributions, in the sense that with more samples the volume of the resulting posterior distribution shrinks. This explains why black contours (400,000 samples) tend to cover more area than the red contours (600,000 samples). As additional forward model samples improve the precision of the inference, the constraints we present would only improve by computing additional realizations.

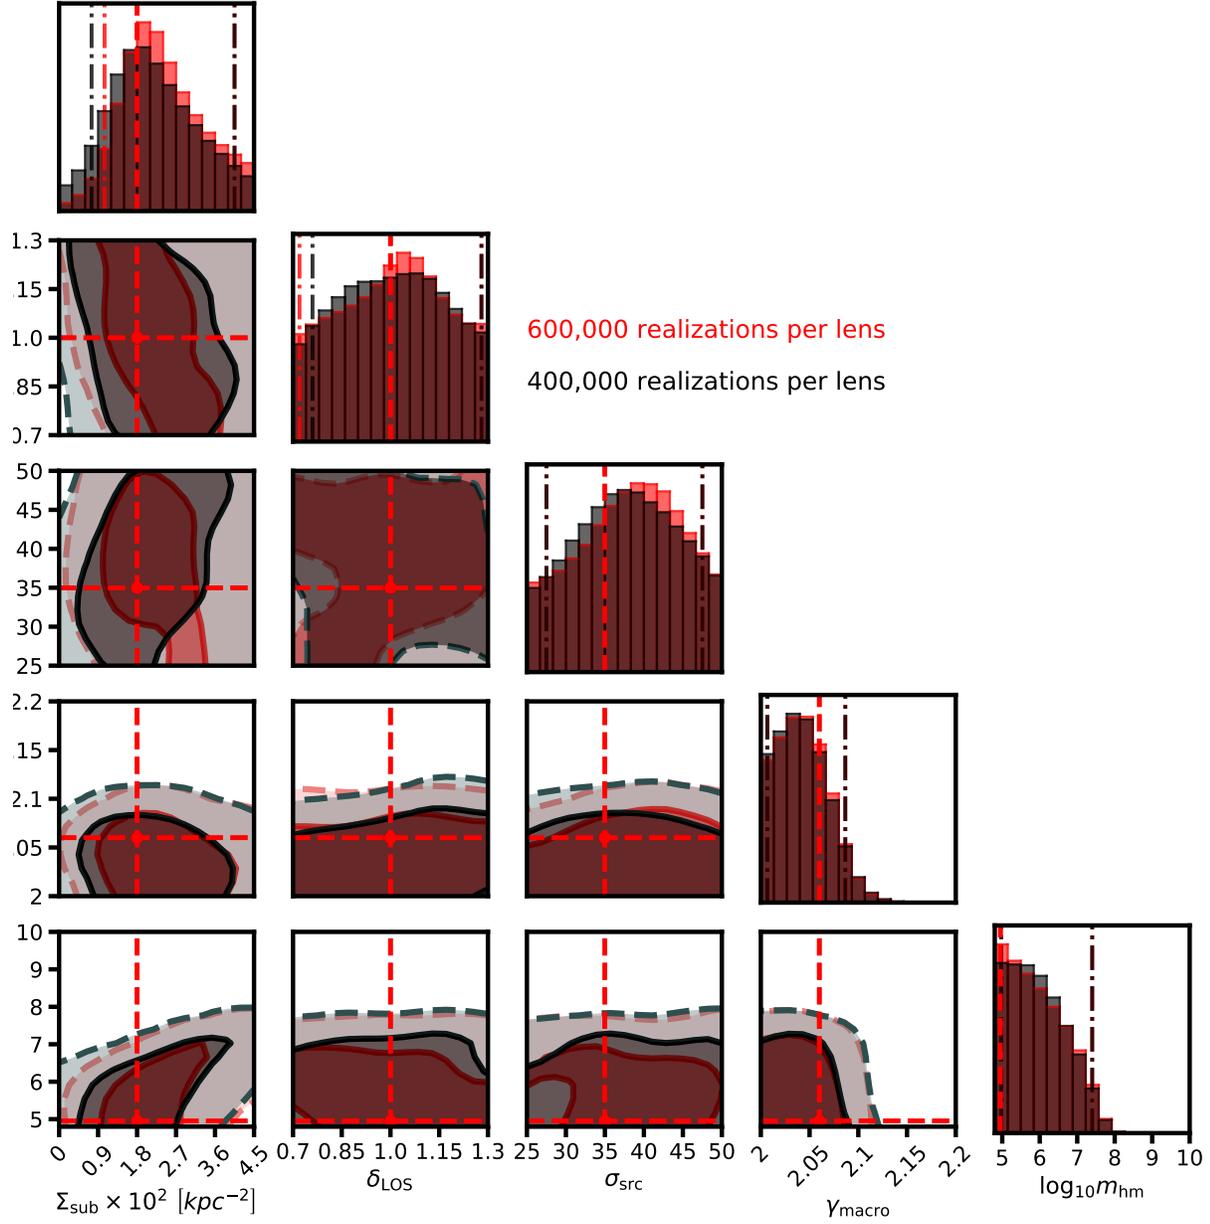


Figure D1. A convergence test for the forward model simulations. The overall agreement between the black and red distributions indicates that the posteriors we derive, and the numerical operations involved to produce them including the kernel density estimation, are robust to changes in the number of forward model samples per lens.