


Testing self-interacting dark matter with galaxy warps

K. Pardo^{1,2,*}, H. Desmond^{2,†} and P. G. Ferreira^{2,‡}

¹*Astrophysical Sciences, Princeton University, Princeton, New Jersey 08544, USA*

²*Astrophysics, University of Oxford,
Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, United Kingdom*

 (Received 26 July 2019; published 10 December 2019)

Self-interacting dark matter (SIDM) is an able alternative to collisionless dark matter. If dark matter does have self-interactions, we would expect this to cause a separation between the collisionless stars and the dark matter halo of a galaxy as it falls through a dark matter medium. For stars arranged in a disk, this would generate a U-shaped warp. The magnitude of this warping depends on the SIDM cross section, type of self-interaction, relative velocity of galaxy and background, halo structure, and density of the dark matter medium. In this paper we set constraints on long-range (light mediator) dark matter self-interaction by means of this signal. We begin by measuring U-shaped warps in 3,213 edge-on disk galaxies within the Sloan Digital Sky Survey. We then forward-model the expected warp from SIDM on a galaxy-by-galaxy basis by combining models of halo structure, density and velocity field reconstructions, and models for the dark matter interactions. We find no evidence for a contribution to the warps from SIDM. Our constraints are highly dependent on the uncertain velocities of our galaxies: we find $\bar{\sigma}/m_{\text{DM}} \lesssim 3 \times 10^{-13} \text{ cm}^2/\text{g}$ at fixed velocity $v = 300 \text{ km/s}$, a bound that scales roughly linearly with increasing v . We also consider galaxy velocities from the CosmicFlows-3 catalogue. These limits are stronger than those from dwarf galaxy evaporation, and we show that they scale well with additional data from the next generation of photometric galaxy surveys. Finally, we forecast constraints for contact and intermediate-range interactions that could be achieved with a similar sample of galaxies in cluster environments, where multistreaming and the fluid approximation are satisfied.

DOI: [10.1103/PhysRevD.100.123006](https://doi.org/10.1103/PhysRevD.100.123006)

I. INTRODUCTION

Most of the matter in our Universe is composed of dark matter (DM). In particular, noninteracting cold dark matter (CDM) seems to fit the observations of the cosmic microwave background, large scale structure, and rotation curves well (e.g., [1,2]). However, we have yet to detect a DM particle to determine its properties directly (e.g., [3–5]). There are also possible discrepancies between observations and CDM predictions, stemming mainly from overprediction of power on small scales (see Ref. [6] for a recent review).

All particles in the standard model have nongravitational interactions, which makes it reasonable to consider self-interactions in the dark sector as well. In addition, self-interacting dark matter (SIDM) could alleviate the possible small-scale CDM issues by redistributing dark matter out of the centers of halos and suppressing small-scale structure formation [7]. For velocity-independent interactions, SIDM cross sections per unit DM mass of $\sigma/m \sim 0.1\text{--}1 \text{ cm}^2/\text{g}$

would be needed to fit the current observations [8]. However, there are constraints on SIDM from a wide variety of systems and experiments (for a comprehensive review, see Ref. [9]). For example, SIDM would lead to the evaporation of halos due to high-momentum-transfer collisions. Thus, the existence of DM halos in dwarf galaxies places constraints on the cross section [10,11]. SIDM would also allow for the spherical relaxation of cluster halos. The observation of elliptical cluster halos places strong limits on the SIDM cross section from cluster ellipticities [12], although these are disputed [13].

SIDM also modifies the distribution of DM in galaxy and galaxy cluster collisions. In the canonical CDM picture, the DM halos do not interact but pass through one another without collision, while the gas shock-heats and decelerates. If DM has self-interactions, then we would expect the DM to experience a drag as well, with a magnitude depending on the interaction cross section. Thus, the centroid of the DM compared to that of the gas could be used to constrain the SIDM cross section. This method has been employed successfully for galaxy cluster collisions, most famously the Bullet Cluster, which disfavors interaction cross sections $\sigma/m > 0.7\text{--}1 \text{ cm}^2/\text{g}$ [14,15] (although some simulations find weaker constraints [16]).

*kpardo@caltech.edu

†harry.desmond@physics.ox.ac.uk

‡pedro.ferreira@physics.ox.ac.uk

We can also expect this effect to leave imprints in the subhalos of clusters—from galaxies falling into clusters. Specifically, we can look for the separation between the centroid of the stars and the DM. The centroid separation technique has been successfully used in simulations [17]. Unfortunately, a clear detection of this effect in data is challenging due to the weak-lensing accuracy required, as well as other systematics [18]. However, recent work has shown that the infall of galaxies into clusters can leave signatures at larger scales [19].

Instead, we can try to look for other markers of this centroid separation. Reference [20] recently considered the SIDM dynamics of a disk galaxy falling into a large galaxy cluster. Using numerical simulations, they found that the separation between the DM and stellar centroids should also produce a warp in the stellar disk of the galaxy. This would be a U-shaped warp facing in the direction of motion—a signature difficult to mimic with baryonic effects. The largest warps should occur in galaxies on first infall into galaxy clusters. The dark matter densities are highest in galaxy clusters and the first infall allows for ample time to form the warp before the direction of the drag force changes at periapsis. There have been no observational searches for this signal in data thus far.

Warps are most easily measured in disk galaxies, which are not typically found in galaxy clusters. Although at lower magnitude, warping should occur in *any* galaxy moving in a dark matter medium. Unfortunately, in this case contact interactions would not be expected to leave an appreciable signature. The interaction timescale at low background density is much longer than the time it takes for the background particles to relax into the halo potential. In the contact case, the relevant relative velocity is between coincident DM streams, which are not present in halos that are relaxed with their environments. Only in clusters does such multistreaming occur. Neither of these restrictions apply to light-mediator models, where interactions are frequent and operate over large distances.

In this paper, we place constraints on the long-range SIDM cross section by measuring the warps of stellar disks. Section II summarizes the theory of dark matter self-interaction and the relevant physical effects that it induces. Section III describes our methods, including our forward-modeling of SIDM warps and measurements of real galaxies. Section IV gives our results, while Sec. V discusses and concludes.

II. THEORY

DM self-interactions generally induce a drag force on the DM halo of a galaxy traveling through some background overdense region. The form of the drag force depends on the type of self-interaction. For a long-range interaction (velocity and angle dependent), we expect a drag force $\propto \rho_{\text{bg}}/v^2$ (e.g., [11]). Were the fluid approximation to hold,

a contact (velocity-independent) interaction would generate a drag force $\propto \rho_{\text{bg}}v^2$, where ρ_{bg} is the density of the background dark matter and v is the relative velocity between the halo and the background. For intermediate-range interactions (i.e., where the mass of the mediator is close to the mass of the DM particle), we expect a force law between these two cases.

Other physics also affect the final force law. For any one collision between particles, there is a probability of the halo particle being ejected. Over time, this leads to an evaporation of the halo, which damps the drag force. Finally, we expect some velocity dispersion in both the halo and the background. This causes a distribution of incoming particle velocity directions, further damping the drag force.

In this section, we develop the equations for the expected stellar warp produced by self-interactions between the DM in a galactic halo and a background overdensity. We begin by finding the drag force per particle mass for the three different types of DM self-interactions (long range, contact, and intermediate range), along with the modifications due to evaporation and velocity dispersion of the halo. We then describe the warp this produces within the galaxy's stellar disk.

A. Halo deceleration from DM self-interactions

Consider a halo moving through some background medium with relative velocity \vec{v} . We would like to find the force per unit mass on the halo in the direction of \vec{v} from DM self-interactions between particles in the halo and particles in the medium. This drag force law depends on several factors, such as the angular and velocity dependencies of the self-interaction and the effects of evaporation and velocity dispersion.

1. Long-range interactions

Let us first consider interactions arising from a long-range force. For now, focus on a two-particle interaction: one particle from the halo and one from the background overdensity. In the center of mass (COM) frame, the velocity of the halo particle in the direction of the relative velocity changes by

$$\delta v_{\parallel} = |\vec{v}|(\cos\theta - 1), \quad (1)$$

where θ is the scattering angle in the COM frame. Note that $\delta v_{\parallel}/|\vec{v}| \leq 0$ always.

The total number of interactions is given by

$$dN = \frac{\rho_{\text{bg}}}{m_{\text{DM}}} \frac{d\sigma}{d\Omega} |\vec{v}| dt d\Omega, \quad (2)$$

where ρ_{bg} is the density of the background overdensity and $d\sigma/d\Omega$ is the differential cross section.

The total drag acceleration is given by integrating over all interactions, which can be written as

$$\vec{a}_{\text{drag}} = \frac{\vec{F}_{\text{drag}}}{m_{\text{DM}}} = \frac{\rho_{\text{bg}}}{m_{\text{DM}}} |\vec{v}|^2 \int \frac{d\sigma}{d\Omega} (\cos \theta - 1) d\Omega \hat{v}. \quad (3)$$

Long-range interactions describe DM that interacts via a massless mediator. This introduces angle and velocity dependencies in the cross section, which can be written as [9,21]

$$\frac{d\sigma}{d\Omega} = \frac{\sigma_0 \sin \theta}{\left(\frac{v}{c}\right)^4 \sin^4\left(\frac{\theta}{2}\right)}, \quad (4)$$

where σ_0 describes the coupling strength of the interaction. This is the well-known Rutherford scattering formula.

Using Eq. (3), we find that the drag acceleration from long-range interactions is

$$\vec{a}_{\text{drag}} = -\frac{1}{4} \left(\frac{\tilde{\sigma}}{m_{\text{DM}}} \right) \rho_{\text{bg}} \frac{c^4}{|\vec{v}|^2} \hat{v}, \quad (5)$$

where we define an effective cross section $\tilde{\sigma} \equiv -32\pi(2+\pi)\sigma_0$ (see Ref. [11] for a similar approach). We use this form so as to match the contact interactions equation (see below).

The drag force from long-range interactions is maximized for *small* relative velocities. Although evaporation is important for the contact interaction case, the lack of high-momentum-transfer collisions (for suitably small $\tilde{\sigma}/m_{\text{DM}}$) means that evaporation is negligible for long-range interactions (see, e.g., [21]).

We have assumed that all of the particles in the halo are traveling with velocity \vec{v} . More realistically, the particles in the halo have some velocity dispersion. The authors of Ref. [21] find that, for a Maxwellian velocity distribution, this leads to a suppression of the drag force, which is well approximated by

$$\chi_p = \frac{|\vec{v}|^3}{|\vec{v}|^3 + |\vec{v}_{\text{disp}}|^3}, \quad (6)$$

where v_{disp} is again the velocity dispersion of the particles. The background should also have a velocity dispersion, but we ignore this for the purposes of this paper since it is small compared to the dispersion of the halo. Our final equation for the long-range drag acceleration is then

$$\vec{a}_{\text{drag}} = -\frac{1}{4} \chi_p \left(\frac{\tilde{\sigma}}{m_{\text{DM}}} \right) \rho_{\text{bg}} \frac{c^4}{|\vec{v}|^2} \hat{v}. \quad (7)$$

2. Contact interactions

We now turn to the velocity-independent interactions arising from a contact force. This follows the formalism of Sec. II A 1, except with the appropriate (constant) cross

section. As Ref. [20] shows, for an isotropic interaction, this leads to a drag acceleration of the form

$$\vec{a}_{\text{drag}}^{\text{contact}} = -\frac{1}{4} \left(\frac{\tilde{\sigma}}{m_{\text{DM}}} \right) \rho_{\text{bg}} \vec{v}^2, \quad (8)$$

where $\tilde{\sigma} = \int d\sigma/d\Omega d\Omega$ is the total cross section. Since we are assuming an isotropic cross section here, $\tilde{\sigma}$ is just a constant.

However, this does not take into account the effects of evaporation on the halo. Allowing for evaporation, the drag acceleration is modified to [14,21]

$$\vec{a}_{\text{drag}}^{\text{contact}} = -\frac{\chi_d}{4} \left(\frac{\tilde{\sigma}}{m_{\text{DM}}} \right) \rho_{\text{bg}} \vec{v}^2, \quad (9)$$

where χ_d is the fraction of events that leads to deceleration rather than evaporation. The authors of Ref. [14] find this fraction by considering the momentum change per collision and comparing this to the escape velocity of the halo. This gives

$$\chi_d = 1 - 4 \int_{\sqrt{x^2/(1+x^2)}}^1 dy y^2 \sqrt{y^2 - x^2(1-y^2)}, \quad (10)$$

where $x \equiv |\vec{v}_{\text{esc}}|/|\vec{v}|$ and v_{esc} is the escape velocity for the halo. If we assume a virialized halo, then $v_{\text{esc}} = 2v_{\text{disp}}$, where v_{disp} is the velocity dispersion of the halo.

The velocity dispersion correction does not depend on cross section, and thus has the same form as in the long-range case. Our final equation for the contact drag acceleration is then

$$\vec{a}_{\text{drag}}^{\text{contact}} = -\frac{1}{4} \chi_d \chi_p \left(\frac{\tilde{\sigma}_{\text{DM}}}{m_{\text{DM}}} \right) \rho_{\text{bg}} \vec{v}^2. \quad (11)$$

As described in Sec. I, we cannot reliably constrain contact interactions using the warping of field galaxies. We therefore only implement this drag equation in forecasting possible results for galaxies in clusters, in Sec. IV B.

3. Intermediate-range interactions

As our final case, we consider intermediate-range interactions, where the mediator mass can range from massless to infinitely massive (the contact limit). We do this by interpolating the drag acceleration between the two previous cases,

$$\vec{a}_{\text{drag}}^{\text{inter}} = -\frac{1}{4} \left(\frac{\tilde{\sigma}_{\text{DM}}}{m_{\text{DM}}} \right) \rho_{\text{bg}} \vec{v}^2 \left(\frac{c}{|\vec{v}|} \right)^m \hat{v}, \quad (12)$$

where $0 \leq m \leq 4$. When $m = 0$, this exactly equals the contact case; when $m = 4$, this exactly equals the long-range case. We do not assume a particular differential cross section equation; however, we presume that any actual

differential cross section would map onto this form for the drag force. For example, a common cross section for this type of interaction is (e.g., [21])

$$\frac{d\sigma}{d\Omega} = \frac{\sigma_0 \sin \theta}{2 \left(1 + \frac{(v/c)^2}{w^2} \sin^2 \frac{\theta}{2} \right)^2}, \quad (13)$$

where $w = m_\phi/m_\chi$, ϕ is the mediator, χ is the DM particle, $\sigma_0 = (4\pi\alpha^2 m_\chi^2)/m_\phi^4$, and α is the coupling constant. We find that this gives similar results to our interpolating case (see Sec. IV).

As with the previous cases, we include the effects of both velocity dispersion and evaporation. The velocity dispersion does not depend on the cross section, so this is trivial to add. However, the evaporation effect requires some more thought. The evaporation fraction calculation requires knowing the differential cross section equation [14,21]. We circumvent this by noting that the evaporation rate should be bracketed by the contact and long-range cases, which means it must be a rapidly decreasing function of the variable m that governs the range of the interaction. We adopt

$$\chi_d^{\text{inter}} = 1 - (1 - \chi_d) \exp[-2m]. \quad (14)$$

When $m = 0$, $\chi_d^{\text{inter}} = \chi_d$; however, when $m = 4$, $\chi_d^{\text{inter}} \sim 1$ and there is no evaporation. Unless there is some extra physics that leads to interesting intermediate behavior, the evaporation fraction should smoothly interpolate between the two cases as described approximately by this function.

Thus our final equation for this interaction, including all physics, is

$$\vec{a}_{\text{drag}}^{\text{inter}} = -\frac{1}{4} \chi_d^{\text{inter}} \chi_p \left(\frac{\tilde{\sigma}}{m_{\text{DM}}} \right) \rho_{\text{bg}} \vec{v}^2 \left(\frac{c}{|\vec{v}|} \right)^m. \quad (15)$$

Again we only implement this in forecasting, as the warp predictions for field galaxies are unreliable for small m .

B. Galaxy warping in SIDM

We now know the force on the halo from self-interactions. However, we cannot measure the force directly—we must instead examine its effect on the morphology of the galaxy. In particular, the displacement between the halo and disk induced by dark matter self-interactions sets up a potential gradient across the disk, which warps it into a cup shape. We calculate this warp by considering the difference in acceleration between the disk's center and a general point along the disk, following the methods of Ref. [22].

Let us define the center of the halo to be at the origin of an x - z plane, where \hat{z} points along the disk normal (see Fig. 1). The stars are collisionless, but the halo is subject to the drag force derived above. The total acceleration of the halo is

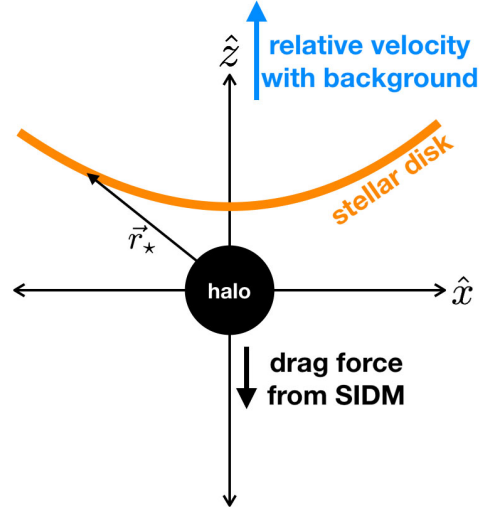


FIG. 1. Cartoon of how a warp is induced by SIDM. In this picture, the Galaxy's stellar disk (orange) and its halo (the COM of the halo is given as the black circle) are falling within an ambient dark matter medium with the relative velocity indicated by the blue arrow. As they fall, the halo experiences a drag force from DM self-interactions, but the stars are collisionless and continue unimpeded. This causes a separation between the centers of the disk and halo, which bends the disk into a U-shaped warp.

$$\vec{a}_h = \vec{a}_{\text{bg}} - \vec{a}_{\text{drag}}, \quad (16)$$

where \vec{a}_{bg} is the gravitational acceleration due to surrounding matter and \vec{a}_{drag} is the drag acceleration due to SIDM. The total acceleration of a point on the stellar disk is

$$\vec{a}_* = \vec{a}_{\text{bg}} - \frac{GM_h}{r_*^2} \hat{r}, \quad (17)$$

where r_* is the equilibrium distance from the point to the center of the halo and M_h is the halo mass enclosed within r_* . The second term is the restoring force caused by the offset of the disk from the halo center.

Since we are looking for the equilibrium positions of the stars, we require that the stars and DM halo move together. This sets $\vec{a}_* = \vec{a}_h$, which gives

$$\vec{a}_{\text{drag}} = \frac{GM_h}{r_*^2} \hat{z}. \quad (18)$$

If we assume a spherically symmetric halo, then the points along the stellar disk experience different accelerations,

$$\vec{a}_{\text{drag}} = \frac{GM_h}{r_*^2} \hat{z} \cos \theta = \frac{GM_h}{r_*^2} \hat{z} \left(\frac{z}{x} \right). \quad (19)$$

We assume that the warp is slight and thus $x \approx r_*$. This now allows us to write an equation for the z positions of the

stars in terms of the drag and the mass of the background halo,

$$z = a_{\text{drag}} \frac{|x|^3}{GM_h}. \quad (20)$$

To go further, we must assume a mass profile for the halo. We use a power-law density profile

$$\rho(r) = \rho_s \left(\frac{r_s}{r} \right)^n \quad (21)$$

within the extent of the disk, with scale radius r_s , $\rho(r_s) \equiv \rho_s$, and a free index n [e.g., $n = 1$ for a Navarro-Frenk-White (NFW) profile]. This gives an enclosed mass $M_h = 4\pi\rho_s/(3-n)r_s^{3-n}$. Substituting this into our equation for the warp curve above, we find

$$z = a_{\text{drag}} \frac{3-n}{4\pi\rho_s} |x|^n. \quad (22)$$

In order to compare to observations, we would like a summary statistic that can quantitatively describe the warp. We employ the w_1 statistic used by Refs. [23] and [22]—this is essentially a measure of the average z position across the disk,

$$w_1 = \frac{2}{L^3} \int_0^L z' x dx, \quad (23)$$

where $z' = z - \langle z \rangle$, $\langle z \rangle$ is the average $z(x)$ value across the disk, and we implicitly assume a stellar disk that is symmetric about the z axis. Substituting in Eq. [(22)] and integrating, we find

$$w_1 = \frac{n(3-n)}{(n+1)(n+2)} \frac{a_{\text{drag}}}{4\pi G \rho_s} \left(\frac{L}{r_s} \right)^n \frac{1}{L}. \quad (24)$$

III. METHODS

In this section, we describe the construction of our galaxy sample and explain how we measure the warp curve. Then we describe our model for the estimated warp produced by SIDM.

A. Candidate selection and warp measurement

We use the NASA Sloan Atlas (NSA) [24] v.1.0.1 catalog, [25] a catalog based mainly on Sloan Digital Sky Survey (SDSS) photometry, to select our galaxies. This catalog contains 641,409 galaxies. General quality cuts (positive mass, radius, flux, and redshift measurements) reduce this number to 640,566. We select only those galaxies that have stellar mass greater than $10^9 M_\odot$ and an axis ratio of $b/a = 0.15$, which leaves us with 22,414 galaxies. This mass cut allows us to use abundance

matching to set the dark matter halo masses for our galaxies: the galaxy-halo connection for lower-mass galaxies is considerably more uncertain. The axis ratio cut selects galaxies that are both thin and viewed edge on. There is some degeneracy between inclination and warp—an inclined galaxy always has a smaller warp measurement if we do not properly account for the inclination. Selecting only edge-on galaxies therefore makes the warp curve measurement more robust. We select galaxies within 250 Mpc, which allows us to use the BORG algorithm to estimate the background density at their positions (see below). Finally, we cut five galaxies with defects in their images (cosmic ray streaks across the disk or no galaxy in the r -band image at the NSA catalog position or corrupted image file). This leaves a final sample of 3,213 galaxies.

To measure the warp curves we employ the methods of Ref. [22]. We give a short summary of the procedure here. First, we rotate the r -band image of a galaxy such that the major axis is aligned with the x -axis. The warp curve is given by the intensity-weighted z value at each x slice. We then measure the warp using the w_1 statistic introduced in Ref. [23],

$$w_{1,\text{obs}} = \frac{\int_{-L}^L \frac{x z dx}{L L L}}{\int_{-L}^L \frac{x dx}{L L}} = \frac{1}{L^3} \int_{-L}^L x z dx, \quad (25)$$

where integration from $-L$ to L allows for asymmetry across the z -axis (perpendicular to x on the plane of the sky). In practice, we set $L = 3R_{\text{eff}}$, where R_{eff} is the stellar effective radius. We also mask the images for $|z| > 3b/aR_{\text{eff}} = 0.45R_{\text{eff}}$.

B. Parameters for estimating the warp

To calculate the expected warp due to SIDM, we require several pieces of information for each galaxy: the effective radius of the stellar disk (R_{eff}), the density of the background at the position of the galaxy (ρ_{bg}), the relative velocity between the galaxy and the background overdensity (v), the angle between this relative velocity and the disk normal (θ), the scale radius of the DM halo (r_s), the density of the DM halo at the scale radius (ρ_s), the power-law index for the DM density profile (n), and the velocity dispersion of the halo (v_{disp}).

We estimate R_{eff} by multiplying the measured Sérsic half-light radius from the NSA catalog, `SERSIC_TH50`, by the angular diameter distance to the galaxy,¹ with the redshift given by the NSA parameter `ZDIST`.

We find the halo parameters (r_s , ρ_s , and v_{disp}) using halo abundance matching and N-body simulations. Abundance matching (AM) assigns dark matter halos to galaxies by assuming a positive, monotonic relationship between the

¹Assuming a flat Λ CDM cosmology with $h = 0.7$, $\Omega_\Lambda = 0.7$, and $\Omega_m = 0.3$.

luminosity or stellar mass of the galaxy and the proxy, a function of the halo mass and concentration [26]. Specifically, we use the AM model of Ref. [27], which maps the r -band absolute magnitude, M_r , to a halo proxy given by $v_{\text{vir}}(v_{\text{max}}/v_{\text{vir}})^\alpha$, with a Gaussian scatter σ_{AM} . We take the values $\alpha = 0.6$ and $\sigma_{\text{AM}} = 0.16$ dex, which best reproduce clustering statistics. We use the DARKSKY-400 simulation [28] postprocessed with the ROCKSTAR halo finder [29] for the halo properties. For each matched galaxy-halo pair we calculate r_s and ρ_s from the ROCKSTAR output, assuming a NFW profile [30]. Velocity dispersions, v_{disp} , are calculated by applying the virial theorem to the halos.

The density of the background, ρ_{bg} , is estimated from the Bayesian origin reconstruction from galaxies (BORG) algorithm [31–37]. This algorithm reconstructs the dark matter density field with a resolution of ~ 2.3 Mpc/ h out to ~ 250 Mpc by forward-modeling primordial density perturbations with a particle-mesh code and comparing this to the number density field of galaxies in the 2M++ survey [38]. To fill in the smaller-scale power, we also include the mass associated with the 2M++ galaxies themselves, which are linked to halos using the same AM routine as above [39].

We use one of two models for galaxy velocities. First, we set v to the same constant for all of our galaxies, where we consider velocities from 50–10 000 km/s. This is clearly an idealized case, but it gives us a basic idea of the constraining power of our dataset. Second, we use the CosmicFlows-3 (CF3) catalog [40] of peculiar velocities. We first assign each galaxy a peculiar velocity, v_{pec} , belonging to the CF3 galaxy closest to it in three-dimensional space. We then assume that the galaxy is falling towards the nearest 2M++ galaxy. The SIDM prediction for the warp we see on the sky is proportional to the relative velocity projected onto the sky. We assign the galaxy velocity in the plane of the sky to be equal to the peculiar velocity.

We must then subtract the velocity of the ambient dark matter medium. We use the public large-scale velocity maps of Ref. [41] for this purpose, [42] evaluated at the positions of our galaxies. These maps are estimated using linear perturbation theory and a reconstruction of the large-scale density in the nearby Universe from the 2M++ catalog. They have resolution 4 Mpc/ h ,² and do not provide uncertainty information. We take each of these background velocities and project them onto the sky. We then subtract this velocity from the total galaxy velocity on the sky. The magnitude of this projected velocity is what we call v . We then assign the on-sky angle between this velocity and the disk normal, θ , again assuming that our Galaxy is falling towards the nearest 2M++ galaxy. The CF3 peculiar velocities and 2M++ galaxy directions should give us a better idea of the order of magnitude of these relative

velocities. However, we also consider fractions f of the relative velocity when we use the CF3 velocities—from 1% to 500%. Note that all of our velocities are in the cosmic microwave background rest frame.

We note that the relative velocities we find here are similar to those seen in simulations. In particular, we find that the distribution of fractional velocity differences $[(v_{\text{halo}} - v_{\text{LS}})/v_{\text{halo}}]$, with v_{halo} being the average velocity of DM particles within R_{vir} and v_{LS} the average velocity of DM particles out to $10 R_{\text{vir}}$ in the direction of halo velocity] in the horizon-active galactic nuclei simulation [43] is comparable to that of the galaxies in our model, with v_{halo} approximated by v_{CF3} and v_{LS} from the large-scale velocity reconstruction described above.

With all of these parameters, we can calculate the predicted w_1 statistic for each galaxy using Eq. (24), for any given $\tilde{\sigma}/m_{\text{DM}}$. However, this equation is for a single set of parameter values. We instead want a likelihood function for w_1 that takes into account the uncertainties on these parameters. For each parameter, we either set it directly (v , θ , R_{eff}) or we sample over some prior distribution (all the rest). For the halo parameters, we perform the AM step independently 200 times, in each case producing a slightly different galaxy-halo connection due to the stochasticity introduced by σ_{AM} . This generates distributions for ρ_s and r_s , separately for each galaxy. We then build our prior for the background density, ρ_{bg} , by finding the density within BORG at the position of the galaxy, \vec{x} , at ten independent steps of the BORG Markov chain. Finally, we use a uniform prior for n from 0.5 to 1.5 independently for each galaxy. This range is chosen to include the NFW value ($n = 1$) as well as profiles that are slightly shallower or steeper.

We then perform Monte Carlo sampling for each galaxy independently to determine the w_1 likelihood function. Since $w_1 \propto \tilde{\sigma}/m_{\text{DM}}$, we can generate the likelihood function at $\tilde{\sigma}/m_{\text{DM}} = 1$ cm²/g and then simply scale it up or down when sampling $\tilde{\sigma}/m_{\text{DM}}$,

$$\begin{aligned} \mathcal{L}\left(w_1 \middle| \frac{\tilde{\sigma}}{m_{\text{DM}}} = 1 \text{ cm}^2/\text{g}, m\right) &= \int \mathcal{L}\left(w_1 \middle| \frac{\tilde{\sigma}}{m_{\text{DM}}} = 1 \text{ cm}^2/\text{g}, m, \rho_s, r_s, n, \rho_{\text{bg}}, v, \theta\right) \\ &\times \mathcal{L}(\rho_s, r_s | M_r; \alpha, \sigma_{\text{AM}}) \\ &\times \mathcal{L}(\rho_{\text{bg}}, v | \vec{x}) P(\theta) P(n) d\rho_s dr_s d\rho_{\text{bg}} dv d\theta dn, \end{aligned} \quad (26)$$

where the probability distributions for each of these priors are given in Table I. We test for convergence of the likelihood function for each galaxy by requiring that the mean, variance, and skew of $\mathcal{L}(w_1 | \tilde{\sigma}/m_{\text{DM}} = 1 \text{ cm}^2/\text{g}, m)$ does not change by more than 1% in the last 10% of the samples, which we find requires at least 100 000 Monte Carlo draws from the prior distributions. Note that by building these distributions directly into the likelihood

²This distance is large enough that it corresponds to long-range (i.e., effectively infinite) interactions between the halo and its surroundings.

TABLE I. Priors used to find the likelihood of the warp statistic for given $\tilde{\sigma}/m_{\text{DM}}$ and m .

Parameter	Source of uncertainty	Model used
$P(n)$	Inner DM halo density slope	Uniform prior $n \in [0.5, 1.5]$
$P(\rho_s, r_s M_r; \alpha, \sigma_{\text{AM}})$	Stochasticity in galaxy-halo connection	200 mock AM catalogs at fixed α and σ_{AM}
$P(\rho_{\text{bg}} \vec{x})$	Background DM density	Ten draws from BORG posterior
$P(v)$	Galaxy relative velocity	Delta function at set velocity (see Sec. III B)
$P(\theta)$	Unknown relative velocity direction	Delta function at set angle (see Sec. III B)

we are effectively sampling from the priors in these quantities rather than the posteriors, which would be computationally too expensive. While this likelihood is written for general m , we remind the reader that the inference is most reliable for larger m , corresponding to a longer-range interaction.

C. Parameter inference

We now have a measured warp statistic for each galaxy, $w_{1,\text{obs}}$, and the likelihood of a given warp statistic under a SIDM model with $\tilde{\sigma}/m_{\text{DM}} = 1 \text{ cm}^2/\text{g}$. This enables us to derive constraints on $\tilde{\sigma}/m_{\text{DM}}$ and m using Bayes' theorem and a Markov chain Monte Carlo (MCMC) algorithm. Note that Eq. (24) is linear in $\tilde{\sigma}/m_{\text{DM}}$ and in the other factors that affect the particular physics of the interactions, outside of m . Thus, for our parameter estimation of $\tilde{\sigma}/m_{\text{DM}}$, we simply sample from $\mathcal{L}(w_1 | \tilde{\sigma}/m_{\text{DM}} = 1 \text{ cm}^2/\text{g}, m)$ and then scale by the particular $\tilde{\sigma}/m_{\text{DM}}$ value the Markov chain is sampling. We then compare this to the measured $w_{1,\text{obs}}$ value for each galaxy, as described below. In the long-range and contact cases, we fix m at the appropriate values and do not sample over it.

For the most part, the measured warp values are many orders of magnitude larger than the estimated warp parameters, given a reasonable cross section. In other words, noise dominates the warp signal. Given that we have no reasonable model for how other processes may produce U-shaped warps, we assume that the noise is normally distributed and marginalize over its variance, $\sigma_{w_1}^2$. This modifies the w_1 likelihood to

$$\begin{aligned} \mathcal{L}\left(w_{1,\text{obs}} \left| \frac{\tilde{\sigma}}{m_{\text{DM}}}, m, \sigma_{w_1}\right.\right) \\ = \int \frac{dw_1}{\sqrt{2\pi\sigma_{w_1}^2}} \times \mathcal{L}\left(w_1 \left| \frac{\tilde{\sigma}}{m_{\text{DM}}}, m\right.\right) \exp\left[\frac{-(w_{1,\text{obs}} - w_1)^2}{2\sigma_{w_1}^2}\right], \end{aligned} \quad (27)$$

In practice, we evaluate this integral by discretizing w_1 into 50 bins between its minimum and maximum values, separately for each galaxy.

We sample this likelihood using the *emcee* affine-invariant Markov sampler [44]. We set the flat prior $\tilde{\sigma}/m_{\text{DM}} \in (0, 10^4)$ and check that varying this prior does not significantly

change the results. For the intermediate-range results, we sample $\log_{10}(\tilde{\sigma}/m_{\text{DM}}/[\text{cm}^2 \text{ g}^{-1}]) \in (-20, 2)$ with a uniform prior in log. The power-law index, m , has a flat linear prior over $m \in [0, 4]$. Finally, we sample in $\log_{10} \sigma_{w_1}$, with no restrictions on its range.

For the long-range and contact interactions, we use ten walkers and take 20 000 samples, after burn-in. This gives a Gelman-Rubin convergence parameter $R < 0.01$. For the intermediate-range case we require $\gtrsim 25$ 000 samples after burn-in to give the same level of convergence.

IV. RESULTS

A. Long-range interactions

Our main results are

- (1) There is no preference for SIDM ($\tilde{\sigma}/m_{\text{DM}} > 0$) over the null hypothesis that warps are generated purely by astrophysical or measurement noise. This indicates no net correlation between the direction of the warps and the galaxy velocities on the plane of the sky, or between the warp magnitude and the expectation of Eq. (24).³
- (2) For long-range interactions, we place a limit of $\tilde{\sigma}/m_{\text{DM}}(v=300\text{km/s}) < 3 \times 10^{-13} \text{ cm}^2/\text{g}$, scaling as $\sim v^{1.0}$ assuming a constant velocity $v < 1000 \text{ km/s}$. This scaling goes as $\sim v^{-0.028}$ for the velocities set using fractions of the CF3 velocities. Including the possibility of variation in the galaxy velocities, we find a range of 68% upper bounds on the cross section from $\tilde{\sigma}/m_{\text{DM}} \lesssim 2 \times 10^{-13} - 10^{-10} \text{ cm}^2/\text{g}$, with assumed median galaxy velocity from $v \sim 50 \text{ km/s} - 10^4 \text{ km/s}$.
- (3) For intermediate-range interactions we find a slight preference for smaller m , with a corresponding constraint on the cross section.

In the rest of this section, we use different assumptions about the relative velocities of our galaxies to give more detailed results. These results are summarized in Fig. 2, and Table II. Note that in all cases we marginalize over the variance of the noise term, σ_{w_1} . We find that σ_{w_1} is not

³This is predictable from the results of Ref. [22], which show that there is a positive correlation between the warp direction and the orientation of the fifth-force field in thin-shell-screened modified gravity theories, which is largely antialigned with galaxies' velocities.

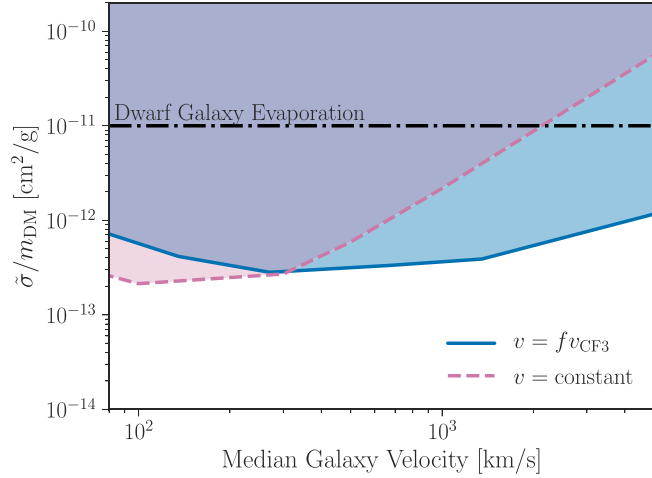


FIG. 2. 68% upper limits on the SIDM cross section assuming a long-range interaction versus the median assumed velocity. We show limits assuming either that all galaxies have the same relative velocities (dashed pink), or that they have velocities proportional to their CF3 velocities (solid blue). The black dotted-dashed line gives an upper limit from dwarf galaxy evaporation rates [11].

degenerate with any other model parameter and its posterior is invariant for all of the models we consider. It is peaked at the measured variance of $w_{1,\text{obs}}$, indicating that it picks up the overall magnitude of the measured warps. The constraints on SIDM parameters instead depend on the correlation of \hat{w}_1 with environment and galaxy/halo properties.

Our long-range limits are given in Fig. 2. As described in Sec. III B, we use either of two assumptions for the velocities: (1) we set all velocities to the same value; or (2) we set the velocities to some fraction of the measured velocities from the CF3 data. The limit on the cross section differs by at most a factor ~ 2 between these models. The results in Fig. 2 include the effects of velocity dispersion but not evaporation (Sec. II). We find our limits to be considerably stronger than those from dwarf galaxy evaporation [11] for all but very highest galaxy velocities.⁴

B. Forecasted constraints for contact and intermediate-range interactions

As discussed in Sec. I (and described further in Sec. V), we cannot place reliable limits on contact interactions with our current sample of galaxies. Because of the low background densities, there are few interactions between the halo and background DM particles before they are able to relax to the same average velocity. In other words, there is no difference in velocity between the halo and the

⁴Note that the evaporation considered in Ref. [11] is different to what we discuss in Sec. II: they consider evaporation due to long-range interactions over a very long timescale (“cumulative” evaporation). In addition, evaporation is more pronounced for dwarf galaxies than the larger-mass galaxies we consider.

TABLE II. Limits on the self-interaction cross section for long-range interactions.

Assumed velocity km/s	Evaporation?	Dispersion?	68% upper limit cm ² /g	95% upper limit cm ² /g
$v = 300$	N/A	...	2.0×10^{-13}	4.4×10^{-13}
	N/A	✓	2.7×10^{-13}	6.1×10^{-13}
$v = v_{\text{CF3}}$	N/A	...	4.7×10^{-14}	1.2×10^{-13}
	N/A	✓	3.9×10^{-13}	9.4×10^{-13}

background particles in the contact interaction case, eliminating the expected drag. In addition, we find that our intermediate results slightly prefer the contact interactions case, so we also cannot give reliable limits in this case with our current sample.

To give an idea of the contact and intermediate range constraints that *could* be achieved with galaxies in clusters, we repeat our analysis using the corresponding drag forces. We use the same galaxy parameters as in the long-range case, but change the cross section and its prior. This is a conservative forecast in that the background densities (and relative velocities) of our galaxies are significantly lower than in clusters, leading to underestimation of the drag force. The cross section constraints should scale as ρ_{bg}^{-1} . The average background density for our galaxies is $\rho_{\text{bg}} = 330 M_{\odot}/\text{kpc}^3 \sim 2.4\rho_{\text{crit}}$. For galaxies in clusters we would expect $\rho_{\text{bg}} > 200\rho_{\text{crit}}$, which would strengthen constraints by a factor $\gtrsim 100$. However, these limits may underpredict the amount of evaporation. Note that the degree of evaporation depends only on the ratio of the relative velocity to the escape velocity of particles within the halo. Thus evaporation effects would only become overwhelming if the relative velocities were, on average, many times larger than the escape velocities. Given that the escape velocity in a typical disk galaxy is >500 km/s and the dispersion of galaxies in clusters is ~ 1000 km/s, we do not expect this to be the case. Our forecasts should therefore provide conservative upper limits for our sample size, although we caution that it will be observationally challenging to find this many thin disks in cluster environments.

We give our main results in Fig. 3. As with the long-range case, we report our limits as a function of the assumed velocity and show both the constant and CF3 velocity models. These limits for average velocities greater than ~ 500 km/s are tighter than the Bullet Cluster constraints [11,14,15]. In Fig. 4, we show how the evaporation and velocity dispersion effects change our limits. Adding both of these effects (as is done in Figs. 5 and 3) weakens the limits by about 1 order of magnitude, regardless of the velocity scale.

Finally, we consider the intermediate case in Fig. 5. This shows the posterior distributions for $\bar{\sigma}/m_{\text{DM}}$ and m , the power-law index for the velocity dependence of the

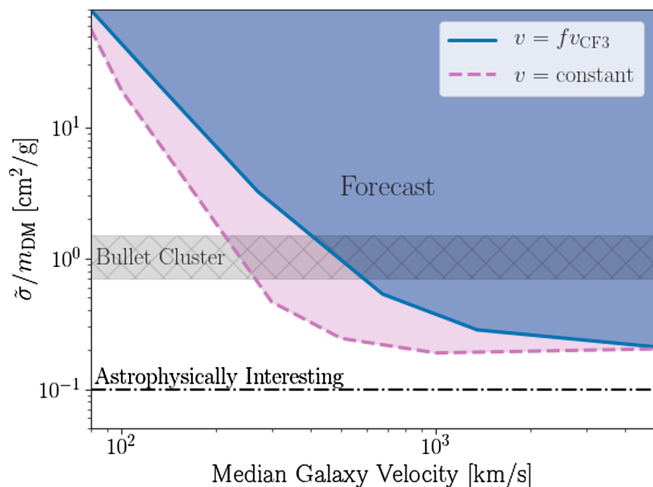


FIG. 3. Forecasted 68% upper limits on the SIDM cross section assuming a contact interaction versus the median assumed velocity, assuming a sample similar to ours but in environments where multistreaming and the fluid approximation obtain. We show limits assuming either that all galaxies have the same relative velocities (pink), or that they have velocities proportional to their CF3 velocities (blue). The grey hatched region gives the range of constraints on the cross section from the Bullet Cluster [11,14,15]. The dotted-dashed line gives the minimum SIDM cross section needed to provide astrophysically interesting effects (i.e., suppression of small-scale structure and DM halo cores) [8].

interaction. Contact interactions (low m values) are slightly preferred, although this may be solely because they allow a larger volume of the $\tilde{\sigma}/m_{\text{DM}}$ prior. In addition, the very low m values are unreliable—they suffer the same issues as the

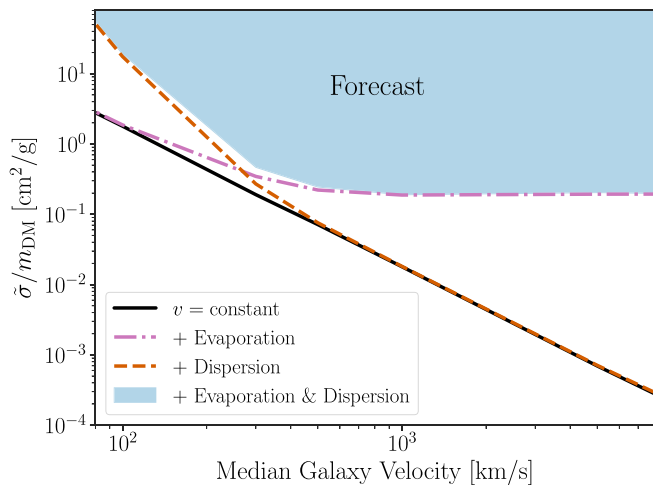


FIG. 4. Forecasted 68% upper limits on the SIDM cross section assuming a contact interaction versus the median assumed velocity, as Fig. 3. Here we show the effects of velocity dispersion and evaporation on the results. The black line shows the limits if we do not consider either of these physical effects. The pink, dotted-dashed line includes evaporation and the orange, dashed line includes velocity dispersion. The blue region shows the same limits as Fig. 3, which includes both effects.

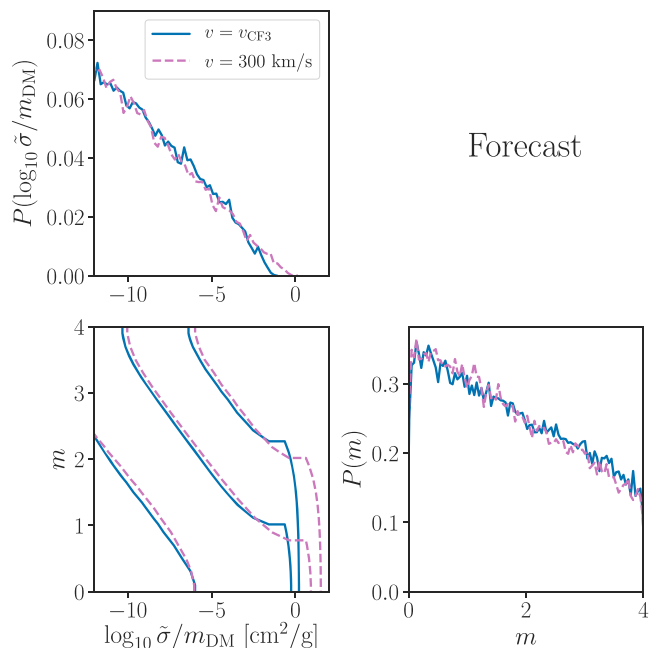


FIG. 5. Forecasted corner plots for an intermediate-range DM self-interaction, assuming a sample similar to ours but in environments where multistreaming and the fluid approximation obtain. We show our limits assuming all galaxies have $v = 300$ km/s (pink) and assuming they have velocities set by their CF3 velocities (blue). m determines the dependence of a_{drag} on the relative velocity of the halo and background [Eq. (12)]. Note that because we use a Jeffrey's prior here for $\tilde{\sigma}/m_{\text{DM}}$ and the posterior peaks at $\tilde{\sigma}/m_{\text{DM}} = 0$ cm²/g, the confidence levels depend sensitively on the arbitrary lower limit of the prior and should not be used: the contour lines in the off-diagonal panel are meant merely to show the degeneracy direction.

contact interactions case (see Sec. V). Note that we use a logarithmic prior on the cross section in this case due to the enormous width of the posterior as m varies. However, since the posterior peaks at $\tilde{\sigma}/m_{\text{DM}} = 0$ cm²/g, confidence limits depend on the arbitrary lower limit of the prior and are therefore not reliable. The shapes of the posteriors and their dependence on velocity are nevertheless robust.

V. DISCUSSION

Our results in the previous section show that we can place new constraints on SIDM cross sections by measuring the warps of stellar disks. In this section, we discuss possible systematics and our attempts to mitigate them. We also discuss the prospects for improving the constraints with next-generation surveys.

As shown in Figs. 2 and 3, the bounds on the cross section are dictated by the magnitudes of the galaxies' relative velocities. We have provided a range of constraints based on different reasonable assumptions, but more robust limits require more precise velocity measurements. The CF3 velocities have very large errors, in excess of 100% at

times. In addition, the CF3 catalog does not include most of the galaxies in our sample, forcing us to assign velocities by means of a nearest neighbor algorithm. Most of our sample is within ~ 10 Mpc of a CF3 galaxy. We find that, within the CF3 catalog, the velocities are well-correlated on these scales. We therefore expect this to be an adequate estimator of the true velocity, but caution that it must introduce some uncertainty. Note also that we do not include uncertainties on the peculiar velocities in our likelihood function.

Another possible systematic is the effect of baryonic physics on galaxy morphology. Most warps caused by tidal or baryonic effects are S shaped [45], and are therefore effectively filtered out by our choice of warp statistic. Any non-SIDM contribution to w_1 is captured to leading order by our noise model (marginalization over σ_{w_1}), but only under the assumption that this contribution is Gaussian and independent of environment and galaxy/halo properties. Baryonic and tidal effects are likely to break this assumption to some degree. In addition, gas in the galaxy experiences hydrodynamical drag from interaction with gas in the intergalactic medium (IGM), which leads to a U-shaped warp in the same direction as SIDM. Thus, including this IGM contribution would tighten our limits, making our results again conservative. The location of the gas as well as the dependence of the measured warp on gas mass would help break the degeneracy between these two types of physics in the context of future, more precise constraints.

We also neglect the effects of tidal interactions, which could contribute to anisotropy in halo and galaxy profiles. These would be largest within clusters while our galaxies are mainly in the field, so we do not expect it to significantly bias our results. However, future constraints on the contact interaction would require a sample of galaxies within clusters, where tidal interactions may need to be considered more carefully.

On the theory side, we use the fluid approximation to derive the SIDM prediction for the warp. As we have mentioned, given the low background densities of our field galaxies, the fluid approximation is not valid for contact interactions. The average background densities near our galaxies are $\rho_{\text{bg}} = 330 M_{\odot}/\text{kpc}^3 \sim 2.4\rho_{\text{crit}}$, where ρ_{crit} is the critical density today. For $v = 300$ km/s and $\tilde{\sigma}/m_{\text{DM}} = 1$ cm²/g this gives an interaction time larger than $1/H_0$. In this time, most of the background particles would easily have time to relax into the potential of the halo, giving them the same velocity as the halo particles. There would then be no relative velocity between the background and halo particles, and no drag effect. If our galaxies were located in clusters, the background density would be high enough to ensure interactions could occur before relaxation. In other words, a warping effect from short-range interactions requires multistreaming. The long-range results also depend on the fluid approximation, but there are many more interactions because of the nature of the force: the

interaction times in this case are closer ~ 50 Myr, which is less than the typical dynamical times for these galaxies.

Finally, we neglected to include the self-gravity of the disk in our calculations. Reference [22] found this to be a small effect. However, as imaging and analysis techniques improve, disk self-gravity may become a relevant systematic.

With the coming era of large and deep photometric surveys (e.g., LSST[46] [47], WFIRST[48] [49], and Euclid[50] [51]), we can expect to have a much larger sample of edge-on galaxies to test in the future. Assuming that we can continue to measure the properties of the DM background (density and velocity) in these survey volumes and the galaxies' peculiar velocities, we can expect these samples to yield considerably tighter constraints. To quantify this, we repeat our analysis for the long-range interaction case (for $v = 300$ km/s and without the velocity dispersion effect) using random subsets of size N of our galaxy sample. This produces a range of results depending on the subset of galaxies chosen. For each subset size, we run 1000 separate MCMC chains and record the 68% upper limits on $\tilde{\sigma}/m_{\text{DM}}$ for each chain. We find that the upper 16% of these limits is well fit by $\sigma/m_{\text{DM}}|_{1\sigma} \propto N^{-0.9}$. In other words, in the worst-case scenario that all of the future galaxies have the same constraining power as our least-constraining few hundred galaxies, with 10 000 galaxies the limits would be tighter than those of Sec. IV B by a factor of ~ 2.8 . The median limits show that with this same number we can more likely expect at least an order of magnitude better constraints. This is even without accounting for any improvements in the velocity determination and other modeling. We can further improve these constraints by finding more thin, edge-on galaxies in high density environments—these would be expected to have the largest warp signature and thus the greatest constraining power. We would also want to choose galaxies at relatively low redshifts and with high stellar masses, which would reduce uncertainties both in measuring the warps and assigning halo properties to the galaxies.

In summary, we calculate the expected stellar disk warp due to DM self-interactions for a variety of interaction types and additional physical processes. We then compare these to the measured warps of edge-on disk galaxies in the SDSS to place constraints on long-range interactions that are stronger than those from dwarf galaxy evaporation. These results are conservative given our treatment of the interstellar medium and velocity uncertainties, although there remain modeling challenges (e.g., the use of the fluid approximation and the precise values of galaxies' peculiar velocities). We also show that a similar sample of galaxies in cluster environments would place highly competitive constraints on contact and intermediate-range interactions. Given the power of this probe, we believe this to be a fruitful avenue for future work. With more galaxies, better photometry, and more accurate velocities, we can hope to use galaxy structure either to detect SIDM, or to rule it out as an astrophysically interesting possibility.

ACKNOWLEDGMENTS

We are very grateful to Michael Blanton for sharing the NSA v1.0.1 images necessary for this work. We thank James Binney, Hume Feldman, Michael Hudson, Bhuvnesh Jain, Lucas Secco, David Spergel, and Richard Watkins for useful discussions, and an anonymous referee for comments that greatly improved the manuscript. K. P. received support

from the Balzan foundation via New College, Oxford, and from the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1656466. H. D. is supported by St John’s College, Oxford, and acknowledges financial support from ERC Grant No. 693024 and the Beecroft Trust. P. G. F. is supported by the ERC, STFC, and the Beecroft Trust.

-
- [1] M. Tegmark, M. A. Strauss, M. R. Blanton, K. Abazajian, S. Dodelson, H. Sandvik, X. Wang, D. H. Weinberg, I. Zehavi, and N. A. Bahcall, Cosmological parameters from SDSS and WMAP, *Phys. Rev. D* **69**, 103501 (2004).
- [2] V. C. Rubin, W. K. Ford, Jr., and N. Thonnard, Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 ($R = 4$ kpc) to UGC 2885 ($R = 122$ kpc), *Astrophys. J.* **238**, 471 (1980).
- [3] D. S. Akerib *et al.* (LUX Collaboration), Results from a Search for Dark Matter in the Complete Lux Exposure, *Phys. Rev. Lett.* **118**, 021303 (2017).
- [4] X. Cui *et al.* (PandaX-II Collaboration), Dark Matter Results from 54-Ton-Day Exposure of Pandax-ii Experiment, *Phys. Rev. Lett.* **119**, 181302 (2017).
- [5] E. Aprile *et al.* (Xenon Collaboration), Dark Matter Search Results from a One Ton-Year Exposure of XENON1T, *Phys. Rev. Lett.* **121**, 111302 (2018).
- [6] J. S. Bullock and M. Boylan-Kolchin, Small-scale challenges to the Λ CDM paradigm, *Annu. Rev. Astron. Astrophys.* **55**, 343 (2017).
- [7] D. N. Spergel and P. J. Steinhardt, Observational Evidence for Self-Interacting Cold Dark Matter, *Phys. Rev. Lett.* **84**, 3760 (2000).
- [8] M. Rocha, A. H. G. Peter, J. S. Bullock, M. Kaplinghat, S. Garrison-Kimmel, J. Oñorbe, and L. A. Moustakas, Cosmological simulations with self-interacting dark matter—I. Constant-density cores and substructure, *Mon. Not. R. Astron. Soc.* **430**, 81 (2013).
- [9] S. Tulin and H. B. Yu, Dark matter self-interactions and small scale structure, *Phys. Rep.* **730**, 1 (2018).
- [10] O. Y. Gnedin and J. P. Ostriker, Limits on collisional dark matter from elliptical Galaxies in clusters, *Astrophys. J.* **561**, 61 (2001).
- [11] F. Kahlhoefer, K. Schmidt-Hoberg, M. T. Frandsen, and S. Sarkar, Colliding clusters and dark matter self-interactions, *Mon. Not. R. Astron. Soc.* **437**, 2865 (2014).
- [12] J. Miralda-Escudé, A test of the collisional dark matter hypothesis from cluster lensing, *Astrophys. J.* **564**, 60 (2002).
- [13] A. H. G. Peter, M. Rocha, J. S. Bullock, and M. Kaplinghat, Cosmological simulations with self-interacting dark matter—II. Halo shapes versus observations, *Mon. Not. R. Astron. Soc.* **430**, 105 (2013).
- [14] M. Markevitch, A. H. Gonzalez, D. Clowe, A. Vikhlinin, W. Forman, C. Jones, S. Murray, and W. Tucker, Direct constraints on the dark matter self-interaction cross section from the merging Galaxy cluster 1E 0657–56, *Astrophys. J.* **606**, 819 (2004).
- [15] S. W. Randall, M. Markevitch, D. Clowe, A. H. Gonzalez, and M. Bradač, Constraints on the self interaction cross section of dark matter from numerical simulations of the merging Galaxy cluster 1E 0657–56, *Astrophys. J.* **679**, 1173 (2008).
- [16] A. Robertson, R. Massey, and V. Eke, What does the bullet cluster tell us about self-interacting dark matter?, *Mon. Not. R. Astron. Soc.* **465**, 569 (2017).
- [17] R. Massey, T. Kitching, and D. Nagai, Cluster bulleticity, *Mon. Not. R. Astron. Soc.* **413**, 1709 (2011).
- [18] D. Harvey, R. Massey, T. Kitching, A. Taylor, E. Jullo, J. P. Kneib, E. Tittley, and P. J. Marshall, Dark matter astrometry: Accuracy of subhalo positions for the measurement of self-interaction cross-sections, *Mon. Not. R. Astron. Soc.* **433**, 1517 (2013).
- [19] A. Banerjee, S. Adhikari, N. Dalal, S. More, and A. Kravtsov, Signatures of self-interacting dark matter on cluster density profile and subhalo distributions, [arXiv:1906.12026](https://arxiv.org/abs/1906.12026).
- [20] L. F. Secco, A. Farah, B. Jain, S. Adhikari, A. Banerjee, and N. Dalal, Probing self-interacting dark matter with disk Galaxies in cluster environments, *Astrophys. J.* **860**, 32 (2018).
- [21] J. Kummer, F. Kahlhoefer, and K. Schmidt-Hoberg, Effective description of dark matter self-interactions in small dark matter haloes, *Mon. Not. R. Astron. Soc.* **474**, 388 (2018).
- [22] H. Desmond, P. G. Ferreira, G. Lavaux, and J. Jasche, Fifth force constraints from galaxy warps, *Phys. Rev. D* **98**, 083010 (2018).
- [23] V. Vikram, A. Cabré, B. Jain, and J. T. VanderPlas, Astrophysical tests of modified gravity: the morphology and kinematics of dwarf galaxies, *J. Cosmol. Astropart. Phys.* **08** (2013) 020.
- [24] M. R. Blanton, E. Kazin, D. Muna, B. A. Weaver, and A. Price-Whelan, Improved background subtraction for the sloan digital sky survey images, *Astron. J.* **142**, 31 (2011).
- [25] <https://www.sdss.org/dr13/manga/manga-target-selection/nsa/>.
- [26] A. V. Kravtsov, A. A. Berlind, R. H. Wechsler, A. A. Klypin, S. Gottlöber, B. o. Allgood, and J. R. Primack, The dark side of the halo occupation distribution, *Astrophys. J.* **609**, 35 (2004).
- [27] B. V. Lehmann, Y.-Y. Mao, M. R. Becker, S. W. Skillman, and R. H. Wechsler, The concentration dependence of the

- Galaxy-Halo connection: Modeling assembly bias with abundance matching, *Astrophys. J.* **834**, 37 (2017).
- [28] S. W. Skillman, M. S. Warren, M. J. Turk, R. H. Wechsler, D. E. Holz, and P. M. Sutter, Dark sky simulations: Early data release, [arXiv:1407.2600](https://arxiv.org/abs/1407.2600).
- [29] P. S. Behroozi, R. H. Wechsler, and H.-Y. Wu, The ROCKSTAR phase-space temporal halo finder and the velocity offsets of cluster cores, *Astrophys. J.* **762**, 109 (2013).
- [30] J. F. Navarro, C. S. Frenk, and S. D. M. White, A Universal density profile from hierarchical clustering, *Astrophys. J.* **490**, 493 (1997).
- [31] J. Jasche, F. S. Kitaura, B. D. Wandelt, and T. A. EnBlin, Bayesian power-spectrum inference for large-scale structure data, *Mon. Not. R. Astron. Soc.* **406**, 60 (2010).
- [32] J. Jasche and B. D. Wandelt, Bayesian inference from photometric redshift surveys, *Mon. Not. R. Astron. Soc.* **425**, 1042 (2012).
- [33] J. Jasche and B. D. Wandelt, Bayesian physical reconstruction of initial conditions from large-scale structure surveys, *Mon. Not. R. Astron. Soc.* **432**, 894 (2013).
- [34] J. Jasche, F. Leclercq, and B. D. Wandelt, Past and present cosmic structure in the SDSS DR7 main sample, *J. Cosmol. Astropart. Phys.* **01** (2015) 036.
- [35] J. Jasche and G. Lavaux, Matrix-free large-scale Bayesian inference in cosmology, *Mon. Not. R. Astron. Soc.* **447**, 1204 (2015).
- [36] G. Lavaux and J. Jasche, Unmasking the masked Universe: The 2M++ catalogue through Bayesian eyes, *Mon. Not. R. Astron. Soc.* **455**, 3169 (2016).
- [37] J. Jasche and G. Lavaux, Physical Bayesian modelling of the nonlinear matter distribution: new insights into the nearby Universe, *Astron. Astrophys.* **625**, A64 (2019).
- [38] G. Lavaux and M. J. Hudson, The 2M++ galaxy redshift catalogue, *Mon. Not. R. Astron. Soc.* **416**, 2840 (2011).
- [39] H. Desmond, P. G. Ferreira, G. Lavaux, and J. Jasche, Reconstructing the gravitational field of the local Universe, *Mon. Not. R. Astron. Soc.* **474**, 3152 (2018).
- [40] R. B. Tully, H. M. Courtois, and J. G. Sorce, Cosmicflows-3, *Astron. J.* **152**, 50 (2016).
- [41] J. Carrick, S. J. Turnbull, G. Lavaux, and M. J. Hudson, Cosmological parameters from the comparison of peculiar velocities with predictions from the 2M++ density field, *Mon. Not. R. Astron. Soc.* **450**, 317 (2015).
- [42] <https://cosmicflows.iap.fr/>.
- [43] Y. Dubois *et al.*, Dancing in the dark: Galactic properties trace spin swings along the cosmic web, *Mon. Not. R. Astron. Soc.* **444**, 1453 (2014).
- [44] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, emcee: The MCMC Hammer, *Publ. Astron. Soc. Pac.* **125**, 306 (2013).
- [45] J. Binney, Warps, *Annu. Rev. Astron. Astrophys.* **30**, 51 (1992).
- [46] <https://www.lsst.org/>.
- [47] P. A. Abell, J. Allison, S. F. Anderson *et al.* (LSST Science Collaboration), LSST Science Book, Version 2.0, [arXiv:0912.0201](https://arxiv.org/abs/0912.0201).
- [48] <https://wfirst.ipac.caltech.edu/>.
- [49] D. Spergel, N. Gehrels, J. Breckinridge, M. Donahue *et al.*, WFIRST-2.4: What every astronomer should know, [arXiv:1305.5425](https://arxiv.org/abs/1305.5425).
- [50] <https://www.euclid-ec.org/>.
- [51] R. Laureijs *et al.*, Euclid: ESA's mission to map the geometry of the dark universe, *Proc. SPIE Int. Soc. Opt. Eng.* **8442**, 84420T (2012).