



Structure, Kinematics, and Observability of the Large Magellanic Cloud's Dynamical Friction Wake in Cold versus Fuzzy Dark Matter

Hayden R. Foote¹ , Gurtina Besla¹ , Philip Mocz^{2,3} , Nicolás Garavito-Camargo⁴ , Lachlan Lancaster^{2,4,5} , Martin Sparre^{6,7} , Emily C. Cunningham^{4,5,14} , Mark Vogelsberger⁸ , Facundo A. Gómez^{9,10} , and Chervin F. P. Laporte^{11,12,13}

¹ Steward Observatory, The University of Arizona, 933 North Cherry Avenue, Tucson, AZ 85721, USA; haydenfoote@arizona.edu

² Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA

³ Lawrence Livermore National Laboratory, 7000 East Avenue, Livermore, CA 94550, USA

⁴ Center for Computational Astrophysics, Flatiron Institute, Simons Foundation, 162 Fifth Avenue, New York, NY 10010, USA

⁵ Department of Astronomy, Columbia University, 550 West 120th Street, New York, NY 10027, USA

⁶ Institut für Physik und Astronomie, Universität Potsdam, Karl-Liebknecht-Str 24/25, D-14476 Golm, Germany

⁷ Leibniz-Institut für Astrophysik Potsdam (AIP), An der Sternwarte 16, D-14482 Potsdam, Germany

⁸ Department of Physics, Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

⁹ Instituto Multidisciplinario de Investigación y Postgrado, Universidad de La Serena, La Serena, Chile

¹⁰ Departamento de Astronomía, Universidad de La Serena, Av. Juan Cisternas 1200 Norte, La Serena, Chile

¹¹ Departament de Física Quàntica i Astrofísica (FQA), Universitat de Barcelona (UB), c. Martí i Franquès, 1, E-08028 Barcelona, Spain

¹² Institut de Ciències del Cosmos (ICCUB), Universitat de Barcelona (UB), c. Martí i Franquès, 1, E-08028 Barcelona, Spain

¹³ Institut d'Estudis Espacials de Catalunya (IEEC), c. Gran Capità, 2-4, E-08034 Barcelona, Spain

Received 2023 May 6; revised 2023 June 29; accepted 2023 July 5; published 2023 September 4

Abstract

The Large Magellanic Cloud (LMC) will induce a dynamical friction (DF) wake on infall to the Milky Way (MW). The MW's stellar halo will respond to the gravity of the LMC and the dark matter (DM) wake, forming a stellar counterpart to the DM wake. This provides a novel opportunity to constrain the properties of the DM particle. We present a suite of high-resolution, windtunnel-style simulations of the LMC's DF wake that compare the structure, kinematics, and stellar tracer response of the DM wake in cold DM (CDM), with and without self-gravity, versus fuzzy DM (FDM) with $m_a = 10^{-23}$ eV. We conclude that the self-gravity of the DM wake cannot be ignored. Its inclusion raises the wake's density by $\sim 10\%$, and holds the wake together over larger distances (~ 50 kpc) than if self-gravity is ignored. The DM wake's mass is comparable to the LMC's infall mass, meaning the DM wake is a significant perturber to the dynamics of MW halo tracers. An FDM wake is more granular in structure and is $\sim 20\%$ dynamically colder than a CDM wake, but with comparable density. The granularity of an FDM wake increases the stars' kinematic response at the percent level compared to CDM, providing a possible avenue of distinguishing a CDM versus FDM wake. This underscores the need for kinematic measurements of stars in the stellar halo at distances of 70–100 kpc.

Unified Astronomy Thesaurus concepts: Large Magellanic Cloud (903); Milky Way dynamics (1051); Milky Way dark matter halo (1049); Milky Way stellar halo (1060); Dynamical friction (422); Cold dark matter (265)

1. Introduction

The Large Magellanic Cloud (LMC) is the Milky Way's (MW) largest satellite galaxy, possessing an infall mass of $\sim 1-2 \times 10^{11} M_\odot$, roughly 10% that of the MW (Besla et al. 2010; Peebles 2010; Peñarrubia et al. 2016; Patel et al. 2017; Erkal et al. 2019; Erkal & Belokurov 2020; Shipp et al. 2021; Vasiliev et al. 2021; Correa Magnus & Vasiliev 2022; Koposov et al. 2023). The LMC is currently on its first infall to the MW (Besla et al. 2007; Kallivayalil et al. 2013), and is inducing significant perturbations in the MW's dark matter (DM) halo, including the collective response, MW reflex motion about the MW/LMC barycenter, and a dynamical friction (DF) wake (Weinberg 1998a; Gómez et al. 2015; Garavito-Camargo et al. 2019, hereafter GC19; Garavito-Camargo et al. 2021;

Tamfal et al. 2021; Rozier et al. 2022; see also Vasiliev 2023 for a recent review).

The perturbations to the MW halo potential induced by the LMC have important, widespread effects on the kinematics of halo tracers, including stellar streams (e.g., Vera-Ciro & Helmi 2013; Gómez et al. 2015; Erkal et al. 2018, 2019; Shipp et al. 2019; Vasiliev et al. 2021; Koposov et al. 2023; Lilleengen et al. 2023), globular clusters and satellite galaxies (e.g., Erkal et al. 2020; Garrow et al. 2020; Patel et al. 2020; Correa Magnus & Vasiliev 2022; Telles et al. 2022), and the halo stars in general (e.g., GC19; Cunningham et al. 2020; Erkal et al. 2020; Petersen & Peñarrubia 2020, 2021). The LMC's infall also affects mass measurements of the MW (e.g., Chamberlain et al. 2023; Correa Magnus & Vasiliev 2022; Erkal et al. 2020) and even the shape and dynamics of the MW's stellar disk (e.g., Weinberg 1998b; Laporte et al. 2018a, 2018b).

If the MW halo's response to the LMC depends on the microphysics of the DM particle, then this scenario presents a unique opportunity to constrain the nature of DM. In particular, the LMC's DF wake offers a promising test-bed, as the strength and density structure of DF wakes depends on the physics of

¹⁴ NASA Hubble Fellow.

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

the medium in which they form (e.g., Ostriker 1999; Furlanetto & Loeb 2002; Lancaster et al. 2020; Vitsos & Gourguliatos 2023). However, our limited ability to disentangle the response of halo tracers due to the wake specifically versus other perturbations induced by the LMC (e.g., the LMC’s tidal field and the MW’s reflex motion) presents a barrier to using the wake as a DM laboratory.

GC19 used a tailored suite of high-resolution N -body simulations of the MW/LMC interaction to show that the LMC creates three major responses in the MW’s DM halo, work that was later expanded upon by Garavito-Camargo et al. (2021). These responses are as follows: (1) the collective response, a large-scale overdensity that leads the LMC and arises primarily due to the shift of the inner halo relative to the outer halo; (2) the global underdensity, which surrounds the LMC’s DF wake; and (3) the DF wake itself.

By “painting” a stellar halo onto the DM particles using weighted sampling, GC19 also explored the response of the stellar halo to the perturbations induced by the LMC. They found that there should be an observable stellar overdensity associated with the DM wake, which has been tentatively detected (Belokurov et al. 2019; Conroy et al. 2021). Further, they found that the velocities of stars in the wake converge near the LMC and diverge behind it, which leads to an enhancement in the component of the stellar velocity dispersion that is orthogonal to the wake.

While this approach is effective at capturing the global behavior of the MW’s DM halo in response to the LMC, it is unable to separate the effect of the DM wake from that of the LMC itself and other halo perturbations. In particular, even in the absence of a DM halo, the passage of a massive perturber such as the LMC would be sufficient to form a wake in the stellar halo (Chandrasekhar 1943). If the LMC’s wake is to be used as a DM laboratory, a more detailed understanding of the role of the DM wake’s self-gravity in forming the stellar wake is required.

A complementary study by Rozier et al. (2022) used a linear response formalism to study the effect of the LMC on the MW’s dark and stellar halos. An advantage of linear response theory is that it permits disabling the self-gravity of the DM, giving insight into the DM wake’s role in shaping the response of the stars. Rozier et al. (2022) reported that the DM wake’s self-gravity enhanced the density of the DM wake by $\sim 10\%$, which hints that the stellar response to the wake is likely sensitive to the density field of the DM wake. This further suggests that the stellar response may also reflect changes in the wake structure owing to the nature of the DM particle.

In particular, the behavior of fuzzy DM (FDM) DF wakes can vary significantly from those in cold DM (CDM; e.g. Bar-Or et al. 2019; Lancaster et al. 2020; Chavanis 2021; Traykova et al. 2021; Buehler & Desjacques 2023; Vitsos & Gourguliatos 2023). FDM is an ultralight bosonic scalar field DM with particle masses of $m_a \sim 10^{-22}$ eV (Hu et al. 2000; see also Hui et al. 2017; Ferreira 2021, and Hui 2021 for reviews), with typical particle de Broglie wavelengths on the order of kiloparsecs. FDM exhibits characteristic density fluctuations on size scales comparable to the de Broglie wavelength of the particles, often called “granules,” which arise due to wave interference between the particles. In the context of DF, Lancaster et al. (2020) and Vitsos & Gourguliatos (2023) showed that FDM granules interact with the perturbing object to produce highly stochastic density fields in the wake, which can result in an oscillatory drag

force if the perturber is moving slowly. To test these predictions using the LMC’s DF wake, we must first understand whether such an FDM wake would affect the motions and distribution of halo tracers differently than a CDM wake.

In this paper, we present a suite of windtunnel-style N -body simulations of the LMC’s DF wake under three different assumptions for the DM model: CDM with self-gravity, CDM without self-gravity, and FDM with self-gravity. We aim to determine the extent to which self-gravity and the assumption of the DM model impact the structure and kinematics of the LMC’s DM wake. Additionally, to quantify the effect of the DM wake on the distribution and velocities of halo tracers (halo stars, globular clusters, or satellite galaxies), we include a separate population of stellar tracer particles.

This paper is organized as follows: In Section 2, we outline the setup of our windtunnel simulations, including the motivation for our setup, how we choose our initial conditions, and the specifics of each DM model we consider. In Section 3, we present our results for the structure and kinematics of the DM wakes. Section 4 discusses the response of the stellar halo to both the LMC and the DM wakes. In Section 5, we introduce a toy model for how the stellar wake might be observed from Earth, and determine the robustness of our results to observational errors. We also explore the effect of the chosen DM model on the LMC’s orbit, and discuss the wake’s influence as a perturbation to the MW’s DM halo. Section 6 examines the consequences of changing major assumptions in our simulation setup. Finally, we summarize our findings in Section 7.

2. Simulations

Here, we describe the simulations we use to study the formation of the LMC’s DF wake and corresponding response of the MW’s stellar halo. In Section 2.1, we explain the motivation for and the design of our windtunnel setup. Sections 2.2 and 2.3 describe the motivation for our choices of initial conditions for the DM and stars, respectively. In Section 2.4, we describe our CDM simulations, which we perform with the `Gadget4` code (Springel et al. 2021). Our FDM simulations are performed with the `BECDM` module (Mocz et al. 2017, 2020) for the `Arepo` code (Springel 2010), and are described in Section 2.5.

2.1. Dark Matter Windtunnels

To study the formation of the DM wake behind an LMC-like perturber, we use windtunnel-style simulations, in which the perturber is stationary while a “wind” of particles moves past the perturber with a common bulk velocity. The box’s boundary conditions are set up such that one boundary acts as an inflow, the opposite boundary acts as an outflow, and the boundaries parallel to the wind’s motion are periodic. In this way, the interaction of the perturber with the background wind can be studied in a maximally controlled environment.

Windtunnel setups are commonly used in hydrodynamic simulations (e.g., Salem et al. 2015; Scannapieco & Brüggen 2015; Schneider & Robertson 2017; Sparre et al. 2019, 2020) and when studying DF in FDM backgrounds (Lancaster et al. 2020; Vitsos & Gourguliatos 2023). In hydrodynamic windtunnels, it is common to use inflow/outflow boundaries, where the wind particles are created at the inflow and removed at the outflow. Such boundaries also allow one to change the wind properties with time to mimic a perturber falling deeper into a

Table 1

Summary of the Galaxy Halo Density Profiles We Use in Our Simulations

Galaxy Model	GC19 Model	$M (M_{\odot})$	a (kpc)
MW	MW1	1.2×10^{12}	40.1
LMC	LMC3	1.8×10^{11}	20.0

Note. Both profiles are taken from GC19’s simulation 3, and are Hernquist profiles (Hernquist 1990). We list: our galaxy model; the corresponding galaxy model in GC19; M , the total mass of the profile (if it were integrated to infinity); and a , the scale radius. For more details on these galaxy models, we refer the reader to GC19. Note that unlike in GC19, neither halo is represented as a live system of N -body particles in our simulations; rather, the LMC profile is used to set the external potential over which the wind flows, and the MW profile is used to set wind properties as explained in Section 2.2.

host galaxy’s halo (Salem et al. 2015). In principle, these time-dependent wind properties would seem ideal for our simulations, but in practice, increasing the wind density and speed with time results in the gravitational collapse of the most dense regions of the wind, creating artificial shockwave-like structures.

This restricts us to using a completely uniform background wind, in which the density, dispersion, and velocity remain constant throughout the simulation. Such a wind is most efficiently created by using fully periodic boundary conditions as in Lancaster et al. (2020) and Vitsos & Gourguiliatos (2023). When the wind is given a bulk velocity, it loops through the box and naturally creates inflow/outflow-like boundaries. Of course, care must be taken to stop the simulation prior to the wake wrapping through the box boundary, and so all of our simulations are run for one-half the box crossing time at the bulk wind speed. All of our boxes are cubic and have side lengths of $L = 600$ kpc, which allows us to simulate wakes longer than the MW’s virial radius.

The LMC in our simulations is represented by an external, stationary Hernquist potential (Hernquist 1990) at the center of the simulation volume. This potential is modeled using the density profile of GC19’s LMC3 (see Table 1). A uniform background of DM (the DM “wind”) with constant mass density $\bar{\rho}$ and isotropic velocity dispersion $\bar{\sigma}$ moves across the LMC potential with a constant bulk velocity v in the $+y$ -direction. In this study we choose two sets of wind properties, described in Section 2.2. The advantages of this setup over simulating the full LMC-MW interaction with live halos are threefold:

1. For FDM in particular, a windtunnel is *far* less computationally expensive. Specifically, live FDM halos require exceptionally high spatial and temporal resolution (see Section 2.5) that makes an FDM simulation analogous to those in GC19 prohibitively expensive. A windtunnel, by contrast, requires only that we resolve the relatively uniform wind instead of the complex structure of a halo.
2. A windtunnel setup allows us to study the role of the wake’s self-gravity by running simulations both with and without gravity between DM particles. N -body simulations with live halos by nature require self-gravity between their DM particles to keep the halos bound, while a uniform DM wind is not subject to this restriction. This allows us to separate how the MW’s stellar halo reacts to *only* the LMC, versus the LMC plus a DM wake.

If this difference is observed, it will provide independent evidence that the LMC is moving through a DM medium.

3. Idealized windtunnels present the best stage for studying DF wakes in the absence of other complicating factors present in live interaction simulations such as tides from the host galaxy, the host’s reflex motion, and orbital resonances. Our setup thus streamlines analysis because we do not have to disentangle DF from any other process.

Naturally, the drawback of the windtunnel is that there is no MW potential. As a result, the LMC “moves” in a straight line (as opposed to a curved orbit), and the wind speed and density are constant (as opposed to varying as the LMC plunges deeper into the MW’s halo). Nevertheless, we use GC19’s fiducial Simulation 3 (their LMC3 and MW1 galaxy models, summarized in Table 1) as a reference simulation to guide the setup of our simulations in an effort to make our wakes as realistic as possible. In Appendix A, we show that the wake in our fiducial CDM windtunnel simulation closely resembles the wake formed in GC19’s Simulation 3.

2.2. Dark Matter Wind Parameters

To select the DM wind parameters $\bar{\rho}$, $\bar{\sigma}$, and v , we choose a point along the LMC’s orbit from our reference simulation, and obtain the Galactocentric position and velocity of the LMC at this point. Then, we calculate $\bar{\rho}$ and σ analytically at the orbital radius of interest, using the MW1 density profile from GC19 (see Table 1). The wind bulk velocity v is then simply the LMC’s orbital speed.

Using this procedure to determine wind parameters, we simulate two different cases along the LMC’s orbit:

1. GC19 determined that the stellar halo’s response to the wake is most easily observed at a Galactocentric distance of 70 kpc to maximize the stellar density while avoiding contamination from the Clouds and the Sagittarius stream. To best reproduce this response with a windtunnel, we want our “Fiducial” CDM wake to match GC19’s wake at 70 kpc, which requires taking the wind parameters from 70 kpc as opposed to the LMC’s present-day location or pericenter passage. Therefore, our Fiducial orbit case represents the MW’s halo at 70 kpc, when the LMC is moving at ~ 313 km s $^{-1}$.
2. To study the behavior of FDM versus CDM wakes and the effect of self-gravity as a function of the LMC’s speed and the MW halo’s density, we also simulate an “Infall” orbit case. This Infall case represents the MW’s halo at a distance of ~ 223 kpc (between our MW model’s R_{200} and R_{vir}), when the LMC is moving at 120 km s $^{-1}$.

Figure 1 illustrates the selection of these parameters. In the left panel, we show the LMC’s orbit since it first crossed the MW’s virial radius, until the present day in the reference simulation. The orbit is projected onto the y - z plane, and we mark the locations from which we take each set of wind parameters. Meanwhile, the other panels show the LMC/MW separation, LMC orbital speed, and MW DM density and dispersion at the LMC’s location as a function of time. We also mark each choice of windtunnel parameters in each panel.

For both orbit cases, we run two CDM simulations and one FDM simulation, described in Sections 2.4 and 2.5, respectively. See Table 2 for a summary of the DM wind parameters in each simulation.

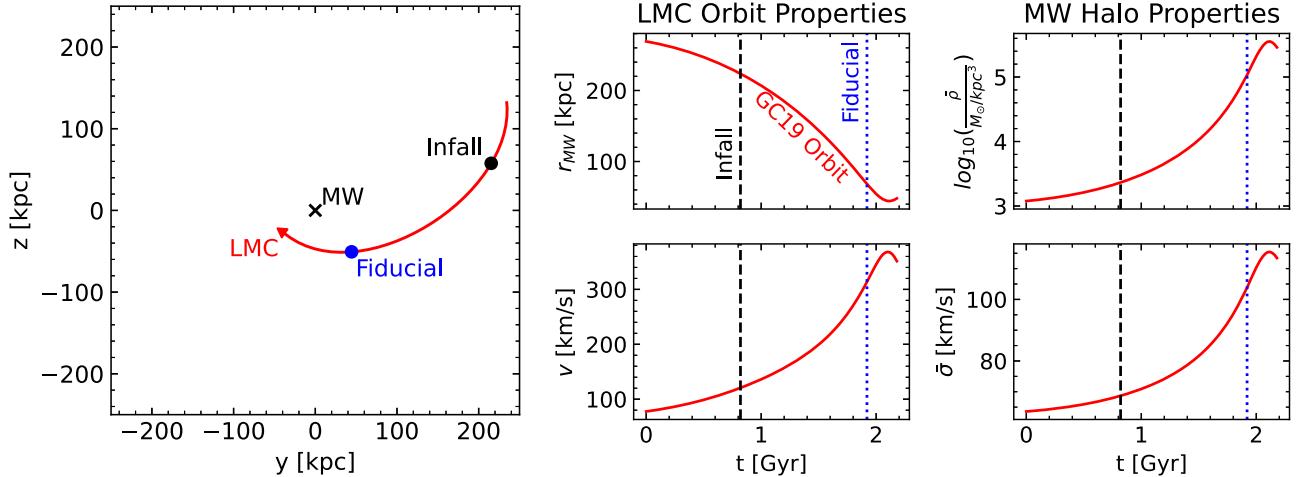


Figure 1. Schematic illustrating how we choose DM windtunnel parameters from the LMC’s orbit. The left panel shows the LMC’s orbit up to the present day from the reference simulation (simulation 3 in GC19). The orbit (red line) is nearly polar, and it is shown projected onto the y – z plane in the rest frame of the MW’s center of mass (denoted by the black x). The blue and black points show the locations on the orbit that we use to select the wind parameters for the Fiducial and Infall orbit cases, respectively. The center column of panels shows the LMC’s distance from the MW center (r_{MW} ; top) and orbital speed (v ; bottom) as a function of time during the reference simulation, while the right column shows the MW’s DM halo density ($\bar{\rho}$; top), and velocity dispersion ($\bar{\sigma}$; bottom) at the LMC’s location vs. time. Blue dotted and black dashed lines mark where we select the Fiducial and Infall wind parameters, respectively.

2.3. Stellar Wind Parameters

In addition to the DM, all of our simulations include a uniform wind of star particles to test the response of the stellar halo to both the LMC and the DM wake. In all simulations, regardless of orbit case or DM model, the stellar wind is composed of test particles at $1 M_{\odot}$ resolution. Their density is calculated from the K-giant stellar halo density profile of Xue et al. (2015) at a Galactocentric distance of 70 kpc, assuming the stellar halo has a total mass of $10^9 M_{\odot}$ inside the MW’s virial radius and is composed entirely of K-giants. The stars’ velocity dispersion is 90 km s^{-1} , again motivated by measurements at 70 kpc (Deason et al. 2012; Cohen et al. 2017; Bird et al. 2021), and they move at the same bulk speed as the DM wind. See Table 3 for a summary of the stellar wind properties. We reiterate that while the DM and stellar winds of the Fiducial case are both calibrated for a Galactocentric distance of 70 kpc, we use the same stellar wind for the Infall case (at 223 kpc) as there are few observational constraints on the stellar halo’s properties at large distances.

2.4. Cold Dark Matter

Our CDM simulations are performed with the Gadget-4 N -body and smoothed particle hydrodynamics (SPH) code (Springel et al. 2021). We use 10^8 DM particles in all CDM simulations, which results in a mass resolution of $5.0 \times 10^3 M_{\odot}$ for the Infall wind, and $2.3 \times 10^5 M_{\odot}$ for the Fiducial wind. All simulations use a softening length of 0.16 kpc, from Equation (15) of Power et al. (2003) with GC19’s MW1 model.

For our CDM initial conditions, we begin by determining the particle mass based on the box volume, number of particles, and the desired wind density $\bar{\rho}$. Particle positions are set randomly throughout the box to create a wind of uniform density. All three velocity components are sampled from a Gaussian according to the isotropic velocity dispersion $\bar{\sigma}$. Finally, every particle is boosted by the bulk wind velocity v in the $+y$ -direction.

An identical procedure is used to create the star initial conditions for all simulations in our suite, though we note again

that the stellar wind uses a different density and velocity dispersion than the dark matter wind (see Table 3).

For each orbit case, we run two CDM simulations: one without self-gravity between the DM particles (i.e., the ONLY forces on simulation particles are from the LMC), and one with self-gravity between the DM particles but NOT the star particles (i.e., all particles feel gravity from the LMC and DM particles, but not from the stars). Comparisons between these simulations allow us to isolate the effects of the DM wake’s self-gravity from the influence of the LMC.

2.5. Fuzzy Dark Matter

Our FDM simulations are performed using the BECDM module (Mocz et al. 2017, 2020) for the Arepo code (Springel 2010). BECDM uses a second-order pseudo-spectral method to solve the FDM equations of motion on a discretized, fixed grid, similar to the AxIREPO module introduced by May & Springel (2021) and May & Springel (2023).

For more detailed background on FDM as a DM candidate, we refer the reader to reviews by Hui et al. (2017), Ferreira (2021), and Hui (2021), and references therein. For detailed descriptions of the numerical methods used here, we refer the reader to Mocz et al. (2017, 2020) and May & Springel (2021, 2023), and references therein. However, we provide an abridged description and information specific to our windtunnel simulations here for completeness.

The FDM is described by a single wave function, which takes the form of a complex-valued scalar field

$$\psi = \sqrt{\rho} e^{i\theta}, \quad (1)$$

where $\rho = |\psi|^2$ is the mass density of the FDM, and $\theta \in [0, 2\pi)$ is the phase. ψ obeys the Schrödinger–Poisson (SP) equations of motion in the nonrelativistic limit:

$$i\hbar \frac{\partial \psi}{\partial t} = \left[-\frac{\hbar^2}{2m_a} \nabla^2 + m_a V \right] \psi \quad (2)$$

$$\nabla^2 V = 4\pi G(\rho - \bar{\rho}), \quad (3)$$

Table 2
Summary of Our Simulation Suite, Listing the DM Wind Properties for Each Simulation

Orbit Case	DM Model	DM Self-gravity	r_{MW} (kpc)	v (km s $^{-1}$)	Particles or Cells	$\bar{\rho}$ (M_{\odot} kpc $^{-3}$)	$\bar{\sigma}$ (km s $^{-1}$)
Fiducial	CDM	No	70	313.6	10^8	1.083×10^5	103.9
	CDM	Yes	70	313.6	10^8	1.083×10^5	103.9
	FDM, $m_a = 10^{-23}$ eV	Yes	70	313.6	1024^3	1.083×10^5	103.9
Infall	CDM	No	223	120.5	10^8	2.315×10^3	68.68
	CDM	Yes	223	120.5	10^8	2.315×10^3	68.68
	FDM, $m_a = 10^{-23}$ eV	Yes	223	120.5	1024^3	2.315×10^3	68.68

Note. We list the following: the orbit case; the dark matter model with the FDM particle mass if applicable; whether DM self-gravity is enabled; r_{MW} , the separation between the MW and LMC at the point in the GC19 reference simulation that the wind parameters are drawn from; v , the wind bulk velocity; the number of DM resolution elements (N -body particles for CDM, grid cells for FDM); $\bar{\rho}$, the DM wind density; and $\bar{\sigma}$, the DM wind velocity dispersion. All simulations have a box side length of $L = 600$ kpc. The LMC is represented as a static Hernquist potential and is static at the center of the box throughout each simulation.

Table 3
Summary of the Stellar Wind Parameters, Which Are Identical in Every Simulation

Quantity	Value
N_*	1.257×10^6
$\bar{\rho}_*$ (M_{\odot} kpc $^{-3}$)	5.818×10^{-3}
$\bar{\sigma}_*$ (km s $^{-1}$)	90.00

Note. We list the following: N_* , the number of star test particles; $\bar{\rho}_*$, the stellar wind density; and $\bar{\sigma}_*$, the stellar wind velocity dispersion. The stellar wind bulk velocity is matched to that of the DM wind in each simulation.

where m_a is the FDM particle mass, and V is the gravitational potential. Additionally, the velocity field of the FDM is encoded by the phase θ via

$$\mathbf{u} = \frac{\hbar}{m_a} \nabla \theta, \quad (4)$$

where \mathbf{u} is the velocity of the FDM.

The FDM wave function (Equation (1)) is discretized onto a grid of N^3 cells of size $dx = L/N$, where L is the side length of the simulation box, and evolved using a kick-drift-kick algorithm. During one time step dt , the potential is first calculated as

$$V = \text{ifft}(-\text{fft}(4\pi G(\rho - \bar{\rho})/k^2) + V_{\text{LMC}}), \quad (5)$$

where “fft” and “ifft” indicate fast Fourier and inverse fast Fourier transforms, respectively, k represents the wavenumbers associated with the grid cells, and V_{LMC} is the external LMC potential. Then, the first “kick” is performed using half the time step:

$$\psi \leftarrow \exp \left[-i \frac{m_a}{\hbar} \frac{dt}{2} V \right] \psi. \quad (6)$$

Next is the “drift,” performed in Fourier space as

$$\hat{\psi} = \text{fft}(\psi) \quad (7)$$

$$\hat{\psi} \leftarrow \exp \left[-i \frac{\hbar k^2}{2m_a} dt \right] \hat{\psi} \quad (8)$$

$$\psi = \text{ifft}(\hat{\psi}), \quad (9)$$

and, finally, the time step is completed by an additional half-step “kick” via Equation (6).

Directly solving the SP equations as we do here has the advantage that it self-consistently describes the full wave dynamics of the FDM, including interference patterns (sometimes

called “granules” or “fringes”) that arise from the velocity dispersion of the FDM and interactions with the LMC potential. Capturing the full wave behavior of the FDM is especially important in studies of DF, as the interference patterns that arise in FDM DF wakes can cause significant deviations from CDM, including stochastic oscillation of the drag force (Bar-Or et al. 2019; Lancaster et al. 2020; Traykova et al. 2021; Buehler et al. 2023; Vitsos & Gourgouliatos 2023). Other numerical descriptions of FDM such as SPH methods or fluid dynamics approaches via the Madelung transformation (Madelung 1927) either approximate or ignore the detailed wave behavior.

The disadvantage of directly solving the SP equations is the enormous spatial and temporal resolution required for numerical convergence. The resolution criteria arise from the wave function phase θ , which cannot vary by more than 2π in a grid cell during one time step (which gives the temporal resolution requirement), or between adjacent grid cells in the same time step (which gives the spatial resolution requirement).

To satisfy the temporal resolution requirement, BECDM uses the time step criterion

$$dt \leq \max \left[\frac{m_a dx^2}{6\hbar}, \frac{h}{m_a |V|_{\max}} \right] \quad (10)$$

where $|V|_{\max}$ is the maximum of the absolute value of the potential (Schwabe et al. 2016; Mocz et al. 2017).

The spatial condition may equivalently be thought of as the requirement that all velocities are resolved, i.e., that the largest velocity in the simulation does not exceed $2\pi\hbar/m_a dx$ (see Equation (4)), or that the smallest de Broglie wavelengths in the problem are resolved:

$$dx = \frac{L}{N} \leq \frac{h}{m_a u_{\max}}. \quad (11)$$

In practice, to ensure that the largest velocities in our simulations are well below u_{\max} , we set the limit on dx according to the bulk wind velocity (the largest velocity scale in the simulation) and then divide by a further factor of 2π , such that our grid cell sizes follow

$$dx \leq \frac{\hbar}{m_a v}. \quad (12)$$

For $m_a = 10^{-23}$ eV and our highest (Fiducial) wind speed of 313.6 km s $^{-1}$, the right-hand side evaluates to 0.611 kpc, which satisfies Equation (12) when $dx = 600$ kpc/1024 = 0.586 kpc.

To generate our FDM initial conditions, we take advantage of the property that ψ can be constructed according to a desired

distribution function f as

$$\psi(\mathbf{x}) \propto \sum_{j=0}^{N^3} \sqrt{f(\mathbf{x}, \mathbf{u}_j)} \exp \left[i \frac{m_a}{\hbar} \mathbf{x} \cdot \mathbf{u}_j + i 2\pi(\phi_{\text{rand},j}) \right], \quad (13)$$

where the sum is over all grid cells in 3D, and $\phi_{\text{rand},j} \in [0, 2\pi)$ is a random number that ensures the phases of each mode are random and uncorrelated, i.e., the FDM has some isotropic velocity dispersion (Widrow & Kaiser 1993). In practice, we desire an FDM wind that is equivalent to our CDM wind, such that it is uniform on the scale of the box and follows an isotropic, Maxwellian velocity distribution. To do this, we take the equivalent approach of constructing the initial conditions in frequency space before taking the inverse Fourier transform and then normalizing such that the mean FDM density is the desired wind density $\bar{\rho}$:

$$\hat{\psi} \propto \sqrt{\exp \left[-\left(\frac{\hbar}{m_a} \right)^2 \frac{(2\pi/L)^2 k^2}{2\sigma^2} \right]} e^{i 2\pi \phi_{\text{rand},j}} \quad (14)$$

$$\psi = \text{ifft}(\hat{\psi}) \quad (15)$$

$$\psi \leftarrow \psi \sqrt{\bar{\rho} / \left(\frac{1}{N^3} \sum_{j=0}^{N^3} |\psi_j|^2 \right)}. \quad (16)$$

Finally, we apply the bulk wind velocity boost by calculating the wavenumber associated with the desired wind velocity

$$k_{\text{boost}} = \frac{v m_a L}{2\pi\hbar} \quad (17)$$

and then applying the boost via

$$\psi \leftarrow \psi \exp \left[i k_{\text{boost}} \frac{2\pi y}{L} \right]. \quad (18)$$

For each orbit case (Fiducial, Infall; see Table 2 and Figure 1), our primary choice for the FDM particle mass is 10^{-23} eV. This is the largest particle mass that is feasible to simulate with $N = 1024$ and $L = 600$ kpc.¹⁵

Lastly, we justify our choice to use the same wind parameters for both our FDM and CDM simulations, as FDM halos differ fundamentally from CDM halos. Instead of being constructed from individual DM particles that obey a particular distribution function (as in CDM), FDM halos are better described as a superposition of eigenmodes that combine to produce a ground-state soliton core surrounded by a “skirt” of excited states that follow a Navarro–Frenk–White (NFW)-like (Navarro et al. 1997) density profile (e.g., Schive et al. 2014; Mocz et al. 2017; May & Springel 2021; Chan et al. 2022; Yavetz et al. 2022; Zagorac et al. 2022). Thus, it is important to verify that our choice of DM wind parameters v , $\bar{\rho}$, and $\bar{\sigma}$ is reasonable in FDM given that we motivate them from a CDM simulation.

As described in Section 2.2, v is given by the LMC’s orbital speed, while $\bar{\rho}$ and $\bar{\sigma}$ come from the MW’s halo. We discuss each parameter in turn:

¹⁵ Our FDM simulations each take $\sim 370,000$ CPU hours at this resolution. We are restricted to $L \geq 600$ kpc to simulate a sufficiently long wake, so increasing the particle mass by a factor of just 2 requires 2048^3 cells. Using the characteristic $\mathcal{O}(N \log(N))$ scaling of the FFT calculations that BECDM relies on, such simulations would take $\sim 800,000$ CPU hours each, which we consider prohibitively expensive.

1. In Section 5.2, we argue that the LMC’s orbit is the same in both a CDM and FDM universe, so our choices of v are valid in both DM models.
2. The MW halo’s density profile is expected to match in CDM and FDM provided we are interested in a regime well outside the soliton core such that the FDM halo follows an NFW-like density profile similar to a CDM halo. Schive et al. (2014) showed that the MW’s soliton would have a radius of ≈ 0.18 kpc, so at the orbital distances of the LMC (≥ 40 kpc), we expect our choices of $\bar{\rho}$ to be valid in both DM models.
3. Yavetz et al. (2022) showed in their Appendix A that far from the soliton core, there is a direct correspondence between the classical particle distribution function of a CDM halo and the eigenmodes that comprise an FDM halo. As such, we expect that for the region of interest in our windtunnel (i.e., a volume many times larger than the de Broglie wavelength and far from the core), using a CDM distribution function to set the FDM eigenmodes (Equation (14)) is a reasonable approach (T. Yavetz 2023, personal communication).

Ultimately, we expect our choice of wind initial conditions to be equally valid in CDM and FDM. It is also worth noting that the inner density profile of the LMC would likely be different in FDM due to the presence of a core. However, in this work we use the same LMC model in both our CDM and FDM simulations to ensure that any differences in our wakes are due purely to our choice of DM model and not the density profile of the perturber. We leave an investigation of the wake’s dependence on the perturber’s density profile to future work.

3. Dark Matter Wakes

In this section, we compare the structure and kinematics of the DM wakes in (1) CDM without self-gravity, (2) CDM with self-gravity, and (3) FDM with $m_a = 10^{-23}$ eV.

3.1. Density

Figure 2 shows the density structure of the simulations with the Fiducial wind (see Table 2) for our three primary DM models/scenarios. In this figure and throughout this work, when we discuss the density of simulation particles, we will use the overdensity

$$\delta\rho = \rho/\bar{\rho} - 1 = \Delta\rho/\bar{\rho}, \quad (19)$$

which measures the relative change of the density compared to the input wind density, i.e., an overdensity of 0.1 corresponds to a 10% increase in density over the background. Figure 2 shows the projected overdensity of each simulation after they have been evolved for 0.7 Gyr,¹⁶ which is the latest time at which there is no evidence that the wake has begun wrapping through the box’s periodic boundaries.

In Figure 2, we begin by taking a 120 kpc-wide slice about the box’s midplane in z ; that is, we select particles/cells with $z \in [-60, 60]$. For the CDM simulations, we then calculate the projected (column) density of DM particles in a grid of 2 kpc-wide bins in x and y , before calculating the overdensity according to Equation (19). For the FDM simulation, we calculate and display the column overdensity in each z -column

¹⁶ The wind travels ≈ 225 kpc during this time.

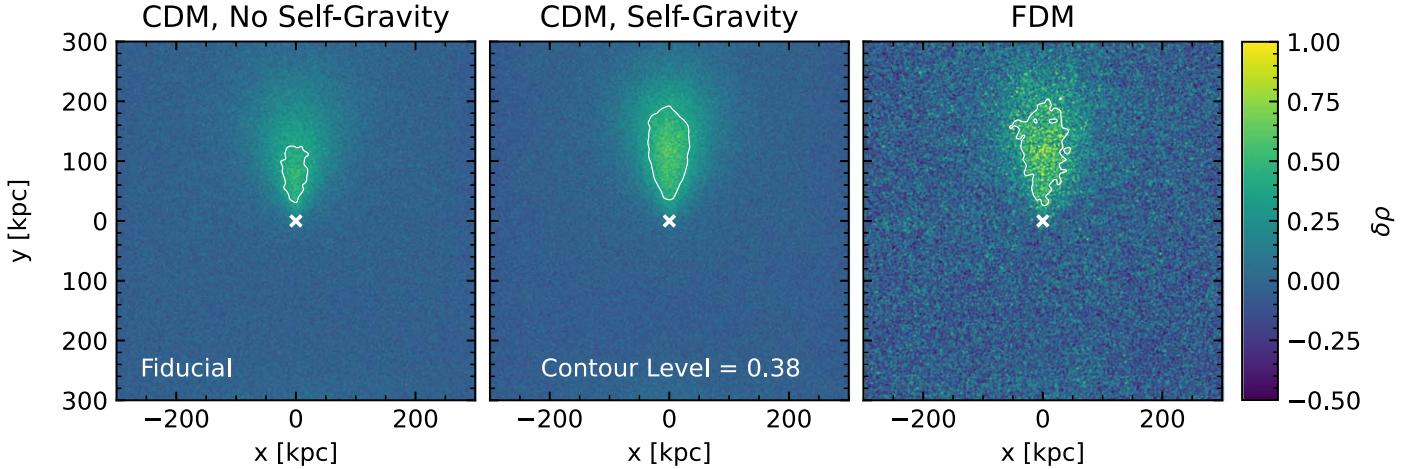


Figure 2. DM overdensity of simulations from the Fiducial case after $t = 0.7$ Gyr of integration. Each panel shows the projected overdensity (Equation (19)) in a 120 kpc-thick slice of the box centered on the LMC potential (i.e., we select particles with $z \in [-60, 60]$) for different DM models. A white cross in each panel denotes the center of the LMC potential. In each panel, the wake manifests as a net overdensity trailing the LMC. The left panel shows the CDM wake without self-gravity, while the central panel shows the CDM wake with self-gravity. To compare the strength and extent of the wake across the panels, the contour in each panel encloses the region with an overdensity greater than the half-max of the CDM wake with self-gravity, which is 0.38. Comparing the two CDM simulations, the addition of self-gravity increases the size and strength (see also Figure 4) of the wake. Specifically, while the wake without self-gravity decays to $\delta\rho < 0.38 \sim 130$ kpc behind the LMC, the wake with self-gravity takes ~ 200 kpc to do the same. This indicates self-gravity has a significant effect on the wake structure, keeping the wake coherent over larger distances. The right panel shows the FDM wake. The signature “granular” structure of the FDM is apparent both in the background and in the wake. Some granules within the wake reach much higher overdensities than either CDM wake achieves, though the FDM wake is still qualitatively similar to the CDM wake with self-gravity.

of cells in the 120 kpc slab with the same x - y coordinates. The white cross in each panel marks the location of the center of the LMC potential. In each simulation, the DM wake is apparent as an overdensity extending from the center of the box in the $+y$ -direction. To ease comparison between the DM models, we calculate the half-max of the overdensity in the CDM simulation with self-gravity ($\delta\rho = 0.38$), and enclose the region with $\delta\rho$ higher than this with a contour in each panel. When placing the contours, we smooth the density with a Gaussian kernel of $\sigma = 4$ kpc, which reduces the noise associated with the FDM granules.

The two leftmost panels show the two CDM simulations. Comparing these two panels, the DM wake becomes larger when adding self-gravity: in the left panel (without self-gravity), the region enclosed by the contour reaches a maximum width of ~ 50 kpc and extends ~ 130 kpc behind the LMC. Adding self-gravity (middle panel) increases the width of the contour to ~ 80 kpc, and the length to ~ 200 kpc. Importantly, the augmentation in wake length demonstrates that the DM wake’s self-gravity plays a significant role in the wake’s structure, acting to hold the wake together at larger distances behind the LMC.

The right panel shows the FDM simulation (with $m_a = 10^{-23}$ eV; see Table 2). We stress that the relative fuzziness of the FDM wakes is *not* a resolution effect (in fact, the FDM simulation is at higher resolution than the CDM). Rather, this granularity is a characteristic property of the FDM that arises due to wave interference between the FDM particles in a velocity-dispersed medium. The FDM wake looks qualitatively similar to the CDM wake with self-gravity aside from the granularity. While some granules near the center of the wake reach much higher overdensities than are seen in CDM, these granules are small and the overall density structure is qualitatively similar to the CDM wake. In Section 6.1, we will discuss the impact of FDM particle mass on these results.

We quantify the wake overdensity by plotting a time-averaged, cross-sectional profile of the wake along the

x -direction (perpendicular to the wind motion). Figure 3 describes this process in a schematic. We begin by taking the same z -slice as we do for the projection plots ($z \in [-60, 60]$). Then, we select particles/cells in a 100 kpc-thick slice in $y \in [50, 150]$ just behind the LMC, before binning the particles/cells along the x -direction in 10 kpc-wide bins. Within each x -bin, we calculate the overdensity. To reduce noise and limit errors related to our choice of a specific snapshot, we repeat this process for five time-adjacent snapshots, spanning 100 Myr. The density in each x -bin is then averaged over the five snapshots, giving us a time-averaged profile of the density as a function of x across the wake.

Figure 4 shows the resulting profiles of the overdensity across the wakes generated in our Fiducial simulations (Figure 2) from $t = 0.6 - 0.7$ Gyr. The upper panel shows the overdensity of each DM wake as a function of x , i.e., across the wake, and the lower panel shows the residuals with respect to CDM with self-gravity. The wakes show up as strong density peaks at the center of the box. The addition of self-gravity to the CDM wake raises the peak overdensity by roughly 10%, in agreement with the results of Rozier et al. (2022). The granularity of the FDM wake shows up as oscillations with an amplitude of $\delta\rho \sim 0.05$, though the average profile of the FDM wake matches the CDM wake with self-gravity.

Figure 5 is the same as Figure 2 but for the Infall orbit case simulations after 2 Gyr (again the last time step at which there is no evidence for the wake wrapping through the box). As a reminder, the wind in this case is roughly 100 times less dense and moving at one-third the speed of the wind in the Fiducial case (see Table 2). The lower wind speed means particles spend a longer time near the LMC, creating a wider wake when compared to the Fiducial case: in the CDM simulation with self-gravity, the contour is ~ 20 kpc wider in the Infall case than the Fiducial case (compare to Figure 2).

The slower speed greatly reduces the effect of the wake’s self-gravity, as the relative importance of the LMC’s influence on the particles’ motions increases. Comparing the CDM

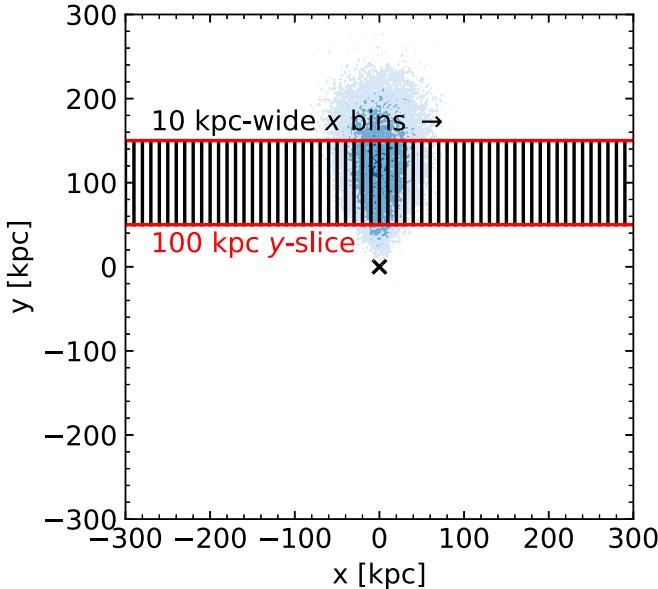


Figure 3. Schematic of the binning procedure to make wake profile plots, in which we compute a quantity of interest across the width of the wake. The filled contours show the typical location of a wake behind the LMC potential, whose center is shown by the x. We begin by selecting particles/cells from $z \in [-60, 60]$, as in the projection plots (e.g., Figure 2). Then, as pictured, we take a 100 kpc-wide slice of the box of $y \in [50, 150]$ (red), and select particles within 10 kpc-wide bins according to their x-position (black). In each x-bin, we calculate the quantity of interest, either density (as in Figures 4 and 6), or z-velocity dispersion (as in Figures 8 and 10). For each figure, this process is repeated for five snapshots spanning 100 Myr. The corresponding profile plots illustrate the average of the five snapshots.

simulations (two leftmost panels) shows they are now almost indistinguishable in projection. Just as in the Fiducial case, the FDM wake appears similar to the CDM wake with self-gravity but is more granular.

The density profiles in Figure 6 reinforce this result, as we see the density profiles across the wakes of the two CDM simulations are very close, only showing a $\sim 3\%$ difference at the peak. Meanwhile, the FDM wake’s density oscillates about the CDM simulation with self-gravity with an amplitude of $\delta\rho \sim 0.05$ as in the Fiducial wind case. An additional effect of the slower wind speed is that the wakes in the Infal case reach much higher overdensity peaks ($\delta\rho \sim 1.6$, compared to ~ 0.56 in the Fiducial case; see Figure 4).

Overall, these results imply that the wake’s self-gravity is only expected to become relevant at higher orbital speeds, i.e., once the LMC reaches a Galactocentric distance of ~ 100 kpc. Therefore, observable effects of the wake’s self-gravity (i.e., halo tracers’ reaction to the wake) will likely not be present outside of ~ 100 kpc. Meanwhile, the SMC is at a distance of 60 kpc and extends $\sim 30^\circ$ on the sky from the LMC (Grady et al. 2021). The LMC’s orbit extends past the SMC on the sky at a distance of ~ 70 kpc (GC19). Together, the decreased effect of DM self-gravity outside of 100 kpc and the need for avoiding SMC contamination suggest the effects of the wake’s self-gravity are best searched for at distances of 70–100 kpc.

3.2. Velocity Dispersion

Figure 7 shows the z-velocity dispersion of the DM in each simulation of the Fiducial case, analogous to Figure 2. We follow the same binning procedure as in the previous section, with a few small differences. The z-slice is still from

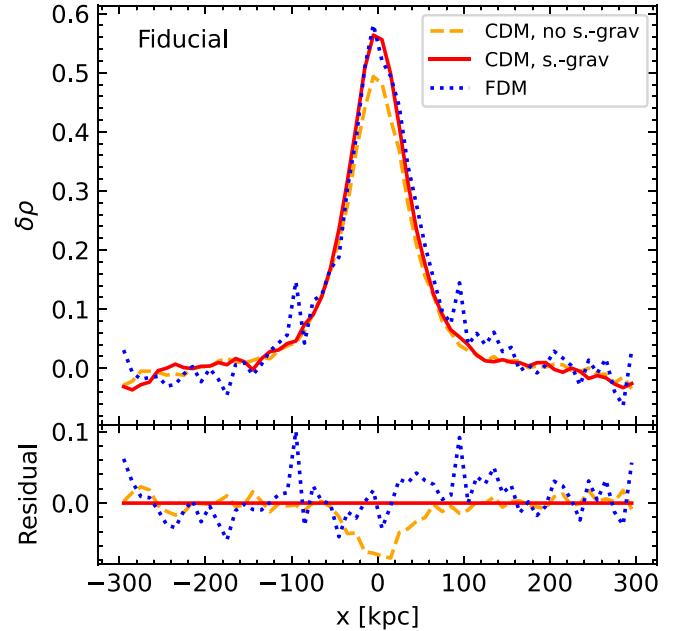


Figure 4. Time-averaged density profile plots for the simulations in Figure 2, showing how the overdensity varies across the wake in the x-direction for each DM model over the past 100 Myr. The procedure for making these plots is explained in Figure 3 and in the text. The upper panel shows the overdensity as a function of x, while the lower panel shows the residuals of the upper panel with respect to CDM with self-gravity. The CDM simulation without self-gravity is shown as the orange dashed line, the CDM simulation with self-gravity is the red solid line, and the FDM simulation is the blue dotted line. The peak overdensity of the CDM wake without self-gravity is $\sim 10\%$ lower than in the CDM simulation with self-gravity included, demonstrating the addition of self-gravity increases the peak wake density. The density across the FDM wake oscillates about that of the CDM wake with self-gravity included, with an amplitude of $\delta\rho \sim 0.05$.

$z \in [-60, 60]$; however, we use an x–y grid of 3 kpc bins, and calculate the z-velocity dispersion in each bin. Finally, we apply a 6 kpc-wide Gaussian smoothing kernel. Similar to the overdensity, we report the dispersion as its relative difference from the mean dispersion, which we refer to as the velocity dispersion enhancement:

$$\delta\sigma_z = \sigma_z/\bar{\sigma} - 1 = \Delta\sigma_z/\bar{\sigma}. \quad (20)$$

We also include a single contour placed at the half-max of the CDM simulation with self-gravity.

The wake signature is an increase of the dispersion resulting from particles being deflected as they move past the LMC. Comparing the two CDM simulations in Figure 7 (left and center panels), the effects of the wake’s self-gravity on the velocity dispersion are similar to the density: when self-gravity is turned on, the wake becomes larger. Specifically, the region enclosed by the contour becomes ~ 40 kpc longer and ~ 20 kpc wider.

For the FDM simulation (right), the granularity is still present in the velocity dispersion, causing an oscillatory behavior that washes out the smooth wake. The contour is much more irregular in shape, and encloses a ~ 40 kpc narrower region than in CDM with self-gravity.

In Figure 8, we compute the dispersion profile across the simulated wakes. We calculate these profiles identically to their density versions (Figures 4 and 6), where, instead of overdensity, we calculate the z-velocity dispersion enhancement in each bin.

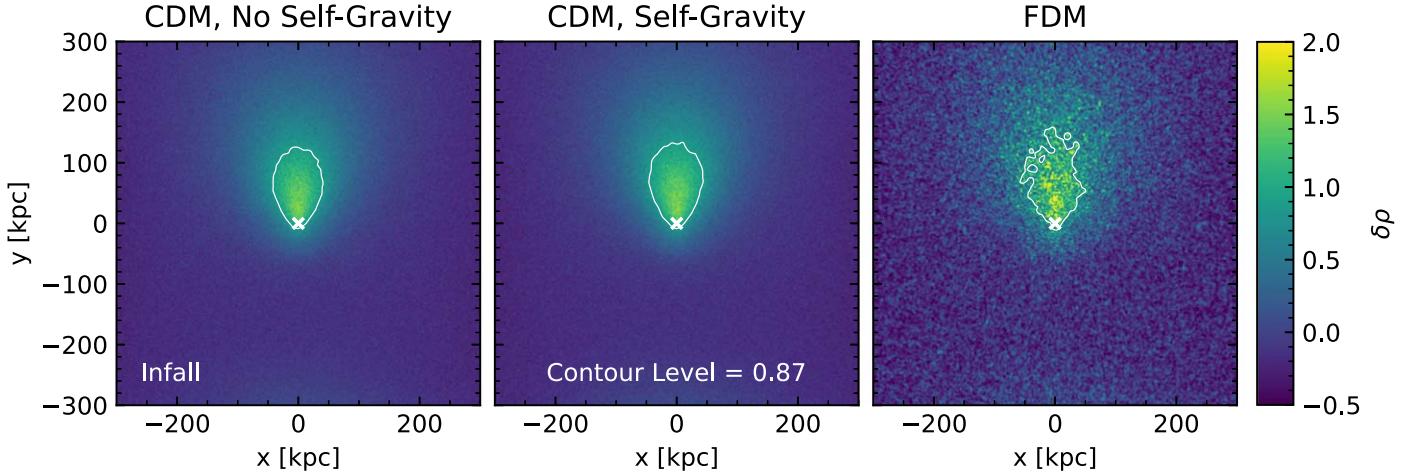


Figure 5. Overdensity of the Infall orbit case simulations after $t = 2$ Gyr of evolution (similar to as Figure 2). Comparing the CDM simulations in the left and central panels reveals that the effect of the wake’s self-gravity is almost negligible at this speed. The CDM with self-gravity half-max contour is placed at $\delta\rho = 0.87$, and extends to ~ 130 kpc behind the LMC in both CDM simulations. The FDM wake (right panel) is once again more granular than the CDM wakes, though overall it remains qualitatively similar to CDM. At this slower wind speed, the wakes reach higher overdensities and are wider when compared to the Fiducial case (i.e., the max. width of the contour is ~ 80 kpc wide in the Fiducial case, compared to ~ 100 kpc in the Infall case here).

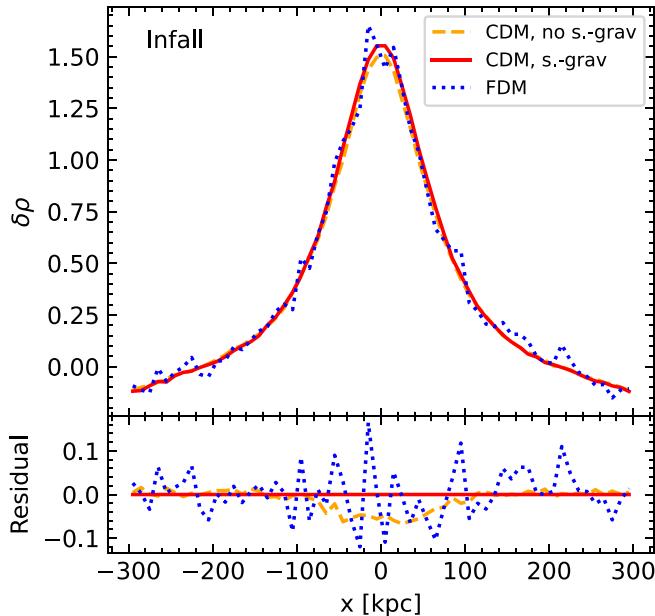


Figure 6. Density profile (same as Figure 4), but for the Infall orbit case after ~ 2 Gyr of evolution, averaged over five snapshots spanning 100 Myr. The CDM wake with self-gravity reaches a peak overdensity of ~ 1.6 , much higher than in the Fiducial case (~ 0.56 ; see Figure 4). Removing self-gravity lowers the peak overdensity only by $\sim 3\%$ in the Infall case, showing that the wake’s self-gravity is relatively unimportant in this regime. The FDM density continues to oscillate about the CDM density with an amplitude of $\delta\rho \sim 0.05$.

Here, we again see a stronger ($\delta\sigma_z \sim 0.03$) peak in the CDM wake when self-gravity is on versus when it is not included. Interestingly, unlike the density, the mean of the FDM wake’s oscillations does *not* trace the CDM wake with self-gravity. Instead, the FDM profile is consistently similar to the CDM profile *without* self-gravity, showing that a self-gravitating FDM wake is colder than a self-gravitating CDM wake.

Figure 9 shows the z -velocity dispersion within the simulated wakes, but for the Infall orbit case. Overall, we see that the slower wind speed results in a stronger but less extended (in the y -direction) response in velocity dispersion compared to the Fiducial case.

Like the density, the CDM wakes show a much smaller difference in the Infall case, as the LMC has more time to influence the particle velocities. The contours in both CDM simulations extend ~ 75 kpc behind the LMC, and are ~ 90 kpc wide. The FDM wake retains its characteristic stochasticity, though in the Infall case, the FDM response in velocity dispersion is significantly weaker than even the no-self-gravity CDM simulation, as the FDM contour is ~ 30 kpc thinner and shorter than in CDM.

The profile plots in Figure 10 illustrate the dispersion profile across the wakes (in the x -direction). The peak dispersion is *slightly* ($\delta\sigma_z \sim 0.01$) higher in the self-gravity-on case. The peak of the FDM wake is now much weaker than either CDM simulation, reaching $\delta\sigma_z \sim 0.23$ as opposed to ~ 0.29 in CDM.

Taken together, Figures 8 and 10 show that FDM wakes are dynamically colder overall than CDM wakes. This can be explained by considering how FDM granules react to a gravitational potential. FDM particles collect into the characteristic granules that have a size of approximately the de Broglie wavelength. When the gravitational potential changes significantly on a scale comparable to or smaller than a granule, gravity becomes less effective at doing work on the granule (Khlopov et al. 1985). This reduces the effectiveness of the LMC at heating the wake, and produces an FDM wake with lower dispersion than a CDM wake.

The $\sim 20\%$ reduction in the velocity dispersion response of the FDM wake compared to CDM is consistent across both the Infall and Fiducial orbit cases. This result suggests that DF wakes in FDM will be $\sim 20\%$ colder than in CDM independent of the density of the medium or speed of the perturber.

3.3. Velocity Divergence

To help explain our results for the wake density and velocity dispersion, we also plot the divergence of the bulk velocity field to study how the particles are deflected by the LMC and the self-gravity of the wake. We again begin with the same 120 kpc-wide slice about the z -midplane, and bin the particles/cells into an x – y grid, this time with 4 kpc bins. In each bin, we calculate the mean x - and y -velocity components, leaving us with a 150×150 grid of

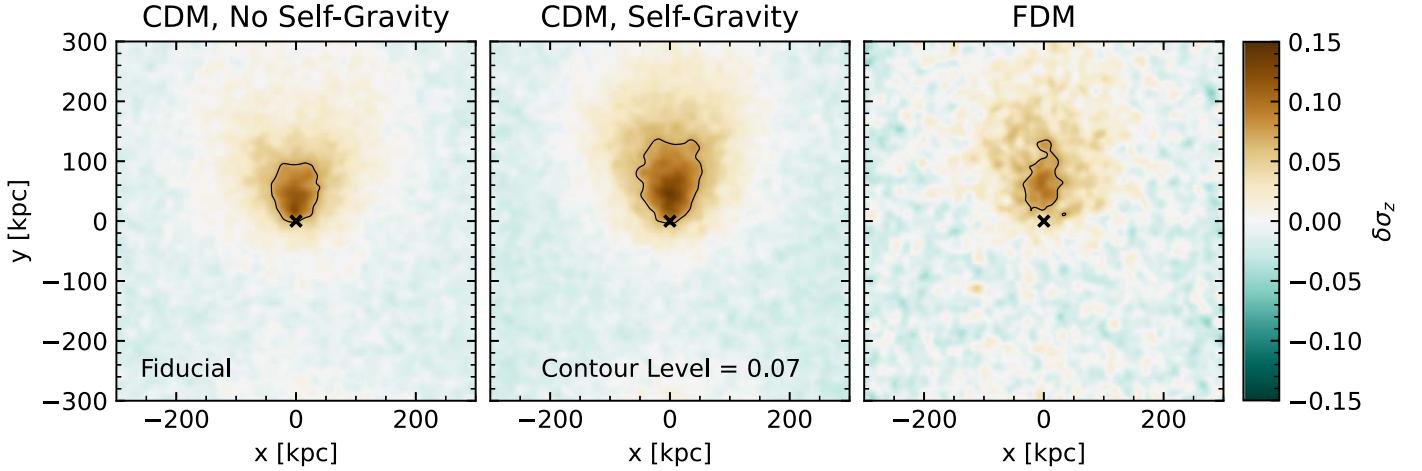


Figure 7. Same as Figure 2, but showing the z (into/out of the page) component of the velocity dispersion for the Fiducial case. The wake shows up as an enhancement of the velocity dispersion behind the LMC as DM particles converge. The dispersion is computed in a similar fashion to the overdensity: the simulated dispersion is divided by the mean dispersion, then subtract 1 (see Equation (20)). The effect of self-gravity in the CDM simulations is evident, as the region enclosed by the contour is larger (~ 20 kpc wider and ~ 40 kpc longer) with self-gravity. The FDM granularity is apparent in the velocity dispersion signature as well. Overall, the granularity results in a weaker dispersion signature in FDM compared to CDM: while the contour in the FDM simulation reaches the same length as in the CDM simulations with self-gravity, it is ~ 40 kpc thinner and tapers more quickly than CDM without self-gravity. We discuss this effect further in Figure 8 and in the text.

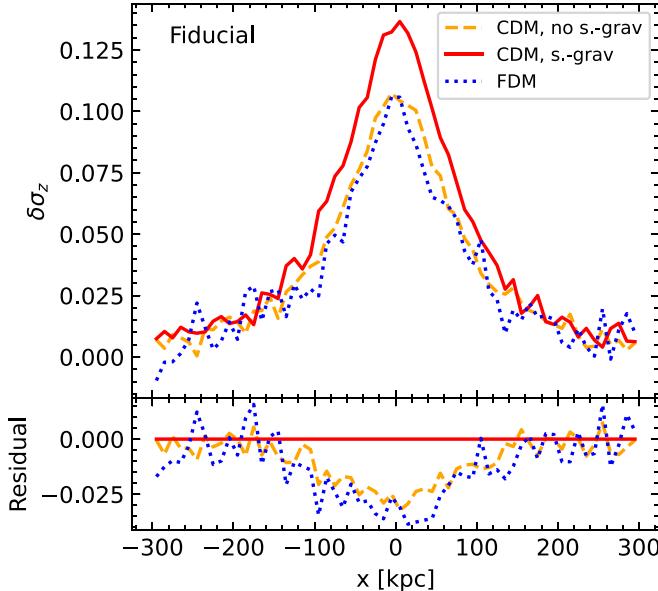


Figure 8. The profile of the z -velocity dispersion computed across the wakes in the Fiducial simulations, shown in the same fashion as the overdensity (Figures 4 and 6). Results are averaged over five snapshots spanning $t = 0.6 - 0.7$ Gyr of evolution. The addition of self-gravity to the CDM simulations raises the peak of the dispersion by $\delta\sigma_z \sim 0.03$. The FDM dispersion profile across the wake is consistently weaker than either CDM case, though the FDM peak is very close to the peak in CDM without self-gravity. Overall, the FDM wake is colder than the CDM wake.

2D velocity vectors. We then calculate the divergence of this 2D velocity field. Finally, we apply a Gaussian kernel of $\sigma = 12$ kpc to the result to reduce noise.

Figure 11 shows the resulting divergence maps for the Fiducial simulations. The wake signature shows up as regions of negative divergence (blue) tracing where the bulk flow of wind particles is converging. In all DM models, the region of strongest convergence is directly behind the LMC, where its gravity most strongly deflects particles. After being deflected, the particles cross the undeflected wind at larger impact

parameters and create a region of converging flow that effectively traces the boundary of the wake. The crossing streams of particles behind the LMC produce the enhancement in the velocity dispersion seen in Figure 7.

Comparing the CDM simulations, we can now pinpoint the effect that the wake's self-gravity has on the particle kinematics and wake structure. In the simulation without self-gravity, the particles deflected by the LMC simply continue on straight paths, creating a region of diverging velocities immediately downstream of the LMC. When self-gravity is turned on, the pull of the wake continues to deflect particles toward the center of the box, eliminating the diverging region, narrowing the wake boundaries, and enhancing the wake's density and velocity dispersion.

As in Figure 7, the FDM reacts less coherently to the LMC in velocity space, and the granularity persists in this kinematic signature. Despite the FDM simulations having self-gravity, the FDM wake shows regions of diverging velocity within the wake of a similar size scale to the granules in velocity space.

Figure 12 illustrates the velocity divergence of the wake produced in the Infall case for all three DM models. At this lower wind speed, the particles can be deflected significantly *before* they reach the LMC center. Overall, particles are deflected more strongly when the wind speed is reduced, leading to wider wake boundaries where these strongly deflected particles cross over the undisturbed wind. The larger deflection angles make it more difficult for the wake's gravity to keep the deflected streams together, and while the self-gravity results in a slight narrowing (by ~ 20 kpc at $y = 150$ kpc) of the downstream diverging region, it is not sufficient to eliminate the diverging flow behind the LMC. Larger deflections also cause a stronger velocity dispersion signature in the Infall case (see Figure 9) compared to the Fiducial case (see Figure 7).

In the Infall case, the wake boundaries are much less clear in FDM, just as they are in the Fiducial case (see Figure 11). While the converging region in front of the LMC is clear, the granularity almost entirely washes out the wake boundaries.

Overall, the velocity divergence illuminates several results from the previous two sections. In the Fiducial case, the wake's self-gravity eliminates the diverging flow in the center of the wake, raising the wake's density by 10% and increasing the

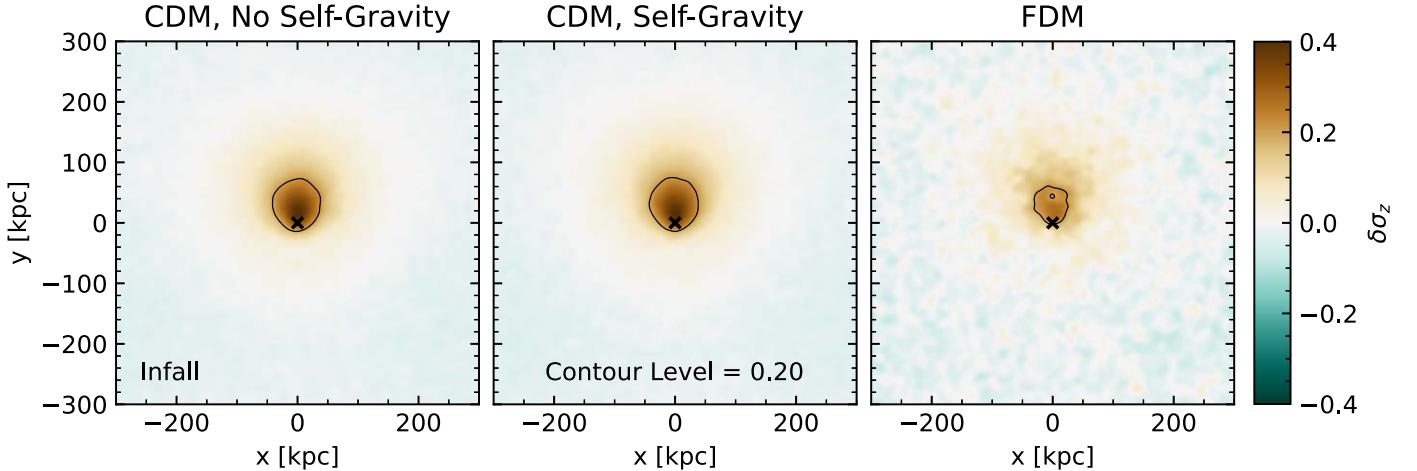


Figure 9. z -velocity dispersion for the Infall orbit case simulations, where the simulations are evolved for 2 Gyr. The two CDM simulations agree closely due to the reduced importance of the wake’s self-gravity at this lower wind speed and density. The velocity dispersion within the FDM wake is lower than in all CDM wakes, as the region enclosed by the contour is ~ 30 kpc shorter and thinner in the FDM case. Overall, the velocity dispersion enhancement is larger in the Infall case than the Fiducial case, as particles spend more time near the LMC, owing to the lower orbital speed (120.5 km s^{-1} vs. 313.6 km s^{-1}). Consequently, the DM is more deflected by the LMC.

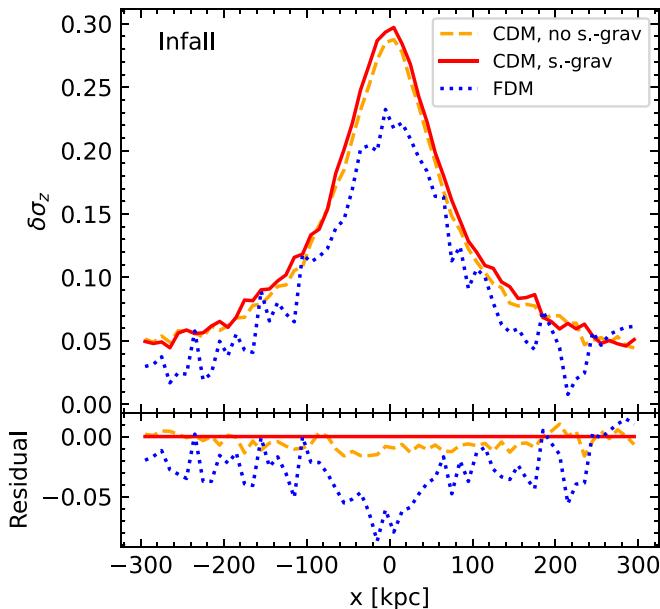


Figure 10. Profile plots of the z -velocity dispersion across the wake in the Infall simulations. The CDM simulations agree closely, with the self-gravity-on version having a slightly larger ($\delta\sigma_z \sim 0.01$) peak. The FDM wake is now much colder than either CDM wake, with a peak $\delta\sigma_z \sim 0.07$ lower than either CDM simulation.

distance the wake takes to decay by $\sim 35\%$. In the Infall case, the diverging region remains regardless of the wake’s self-gravity due to the increased deflection angles of the particles, which explains why the CDM Infall wakes look very similar regardless of self-gravity. FDM’s granularity in the velocity divergence persists across both wind speeds and densities, showing that FDM does not react as coherently to a perturber as CDM. In turn, FDM wakes have lower velocity dispersions than their CDM counterparts.

4. Stellar Wakes

Now, we turn our attention to the observable stellar counterpart of the LMC’s wake. As a reminder, the stellar wind input

parameters (see Table 3) are meant to mimic the MW’s stellar halo at 70 kpc, just as the DM initial conditions in the Fiducial orbit case match the MW’s DM halo at 70 kpc (see Table 3).

In Section 3, we argued that the influence of DM self-gravity on the wake properties should be most observable at distances of 70–100 kpc. Extending this argument to the stellar wake, the best observational signatures of the stellar wake and the DM wake’s influence on it should be between 70 and 100 kpc.

For this reason, we focus only on the Fiducial simulations in our discussion of stellar wakes. As with the DM wakes, we examine the density and velocity structure of the stellar wakes and identify signatures with which to confirm: (1) the presence of a stellar wake; (2) the presence of a DM wake; and (3) distinguishing features between a CDM or FDM wake. The observability of these signatures will be discussed further in Section 5.1.

Figure 13 shows the density structure of the stellar wakes in the Fiducial simulations. To make these plots, we use a procedure identical to Figure 2, with a single additional step of smoothing the resulting density fields with a Gaussian kernel with $\sigma = 4 \text{ kpc}$. This additional smoothing is done to reduce the noise that results from sampling ~ 100 times fewer stars than DM particles. We again include a contour that encloses the region with overdensities higher than the half-max of the CDM simulation with self-gravity.

The left panel shows the stellar wake in the absence of DM self-gravity, i.e., the stellar wake that would form due to *only* the passage of the LMC. The contour extends for roughly 150 kpc behind the LMC. In contrast, the central panel shows the stellar wake that forms when the stars feel the gravity from the CDM wake.

The more striking difference is that the contour extends ~ 50 kpc farther behind the LMC than in the wake without self-gravity. This demonstrates that the DM wake’s self-gravity holds the stellar wake together. Observationally confirming the existence of a stellar wake with $\delta\rho_* \gtrsim 0.6$ more than 150 kpc behind the LMC would provide strong evidence for the existence of a DM DF wake behind the LMC. Comparing the CDM simulation with self-gravity to the FDM simulation (right), however, reveals little difference in the density of the stellar wakes formed under the gravity of different DM particles.

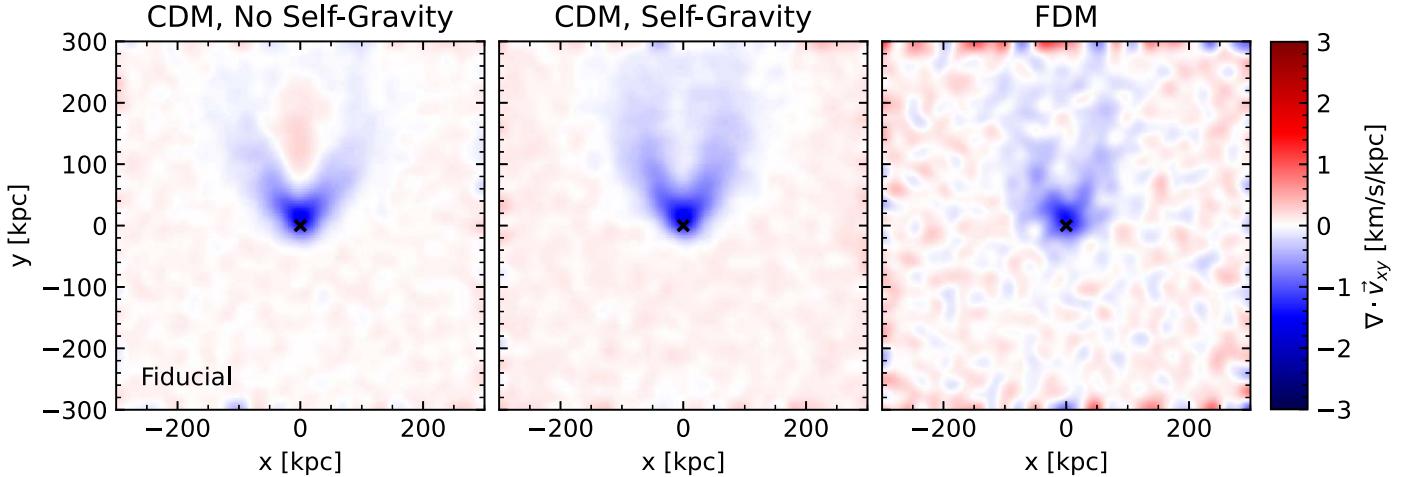


Figure 11. Divergence of the projected x - y velocity field for the Fiducial case. See Section 3.2 for how this is calculated. Regions of converging flows (blue, negative $\nabla \cdot \vec{v}$) clearly identify the wake boundaries and velocity structure in each simulation. The self-gravity-off CDM simulation in the left panel shows only the effect of the LMC on the particle velocities: particles that pass close to the center of the LMC potential are deflected most strongly, producing a region of highly converging flow directly behind the LMC. After they are deflected, particles continue on relatively straight paths in the absence of wake self-gravity, producing a region of diverging flow at the center of the wake. Meanwhile, strongly deflected particles continue to move toward the box edges, passing near particles that have been deflected less strongly, creating a converging flow at the wake boundary. Turning on wake self-gravity (center panel) means that strongly deflected particles continue to feel a force toward the wake after they pass the LMC. This eliminates the diverging region behind the LMC that is seen in the no-self-gravity case, and keeps the wake boundaries narrower. The net result is the effect seen in Figures 2 and 7: a denser, more coherent wake. The FDM simulation on the right shows a similar overall structure to the self-gravity-on CDM case; however, the granularity of the FDM results in a less smooth signature in the velocity divergence. The wake boundaries are less smooth, and some regions of converging flow can be seen inside the wake. Together with Figure 7, it is clear that the FDM velocity field is not as coherently affected by the LMC as the CDM velocity field.

The profile-style plots (e.g., Figures 4 and 8) become very noisy when made with star particles due to the ~ 100 times smaller sample sizes in each bin (with respect to DM). Instead, to compare the overall strength of the stellar response in each DM model, we compute an estimate of the overall wake density (Figure 14), time-averaged over five snapshots spanning 100 Myr of evolution. In detail, for each of the five snapshots, we compute a 2D histogram of the quantity of interest (exactly as in Figure 13 for the density, or in Figure 15 for the dispersion). For each histogram, we select bins with values that are over half that of the maximum bin, then take the median of these. The values reported in Figures 14 and 16 are the time-average and standard deviation of the medians.

Figure 14 shows the time-averaged median density of the stellar wakes in each Fiducial simulation. The stellar wake reaches an overdensity of ~ 0.48 when the DM wake’s gravity is not included, compared to ~ 0.58 when including the gravity of a CDM wake. The stellar wake in the FDM simulation reaches $\delta\rho_* \sim 0.56$, similar to the CDM with self-gravity case.

In short, the gravity of a DM wake raises the density of the stellar wake by $\delta\rho_* \sim 0.1$, and extends the density response by ~ 50 kpc. CDM and FDM wakes do not leave significantly different signatures in the density of the stellar wake.

Figure 15 shows the z -velocity dispersion of the stars in the Fiducial simulations, exactly as Figure 7 but for the star particles. The smoothing length is also increased to 9 kpc in order to mitigate the increased noise associated with the relatively low number of star particles. The velocity dispersion signature in the CDM simulation without self-gravity (left) is ~ 20 kpc narrower than when DM self-gravity is included (center and right). Additionally, when compared to the CDM simulation without self-gravity, the contour tapers more slowly in the CDM simulation with self-gravity and more slowly still in the FDM simulation.

Figure 16 shows the time-averaged median enhancement in the z -velocity dispersion in the same fashion as Figure 14. A

CDM wake’s gravity raises the dispersion of the stellar wake by $\delta\sigma_{z*} \sim 0.01$ compared to when the DM wake’s gravity is not present. Importantly, an FDM wake heats the stars more than a CDM wake: the velocity dispersion of the stars is $\delta\sigma_{z*} \sim 0.01$ higher in the FDM simulation than the CDM simulation with self-gravity.

We plot the divergence of the x - y velocity field to illuminate the density and kinematic structure of the stellar wakes (Figures 13 and 15) in Figure 17. In the absence of the DM wake’s gravity (left panel), we again see a region of converging flows immediately behind the LMC (blue), followed by a region of diverging flows (red) farther downstream as deflected stars pass by each other. Adding the gravity of the DM wake eliminates the diverging region just as it does for the DM particles, enhancing the density and velocity dispersion of the stellar wakes with DM self-gravity. The divergence of the stellar velocities looks similar between CDM with self-gravity and FDM.

Altogether, the velocity dispersion enhancement of the stellar wake is slightly ($\sim 5\%$) higher in response to an FDM wake compared to a CDM wake. The only difference in the forces on the stars in both cases is caused by the differences in the density fields of the DM wakes. In Section 3.1, we showed that FDM granules persist and are even strengthened inside of a DF wake. Therefore, we expect that the additional heating of the stellar wake in the FDM simulation is due to the scattering of stars by FDM granules. This so-called “granule heating” has been well studied in other contexts (e.g., Hui et al. 2017; Bar-Or et al. 2019, 2021; Church et al. 2019; Chavanis 2021; Dalal et al. 2021; Dalal & Kravtsov 2022) and is a known property of FDM. In Section 6.1, we will discuss the role of granule heating and its dependence on FDM particle mass further.

Ultimately, we have demonstrated that the gravity of the DM wake plays an important role in shaping the response of the stars. Specifically, the gravity of the DM wake raises the

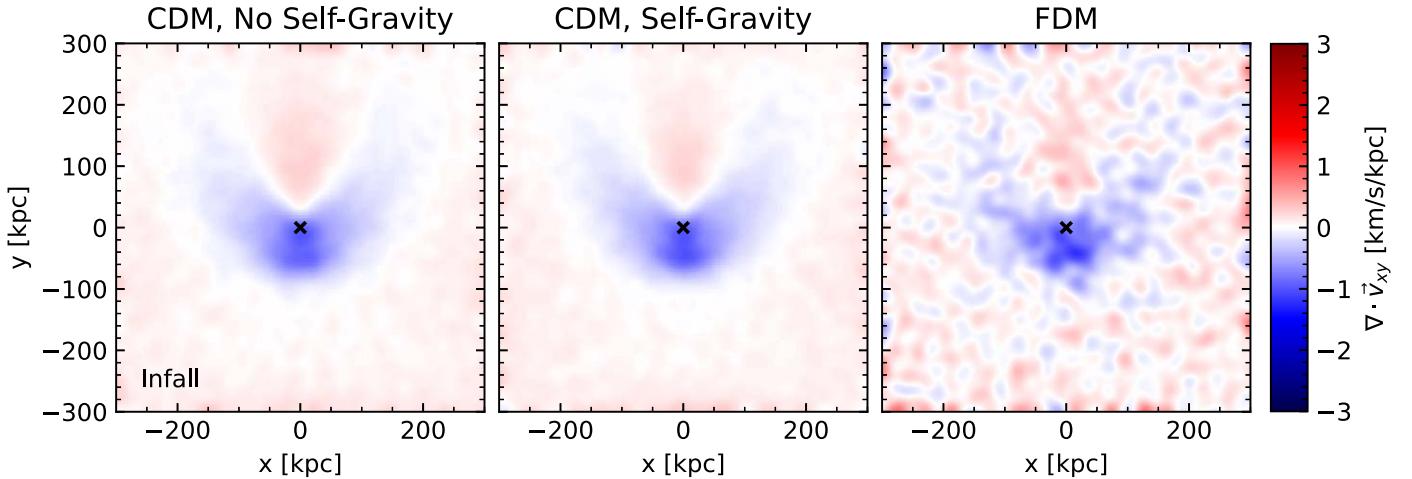


Figure 12. Same as Figure 11, but for the Infall case after 2 Gyr of evolution. As in Figure 11, the divergence traces the boundary of the wake. However, in all cases, the converging regions behind the LMC illustrate a wider opening angle than in the Fiducial case. This results directly from the lower wind speed for the Infall case. In the Infall case, the wake’s self-gravity is not sufficient to erase the diverging flow at the center of the wake, although it is slightly narrower with self-gravity on, i.e., at $y = 150$ kpc, the diverging region is ~ 160 kpc wide with self-gravity, compared to ~ 180 kpc wide without self-gravity. Meanwhile, the differences between the FDM simulation and its CDM counterparts are similar to those in Figure 11: the resistance of FDM granules to strong accelerations washes out the coherent velocity signature seen in CDM.

overdensity of the stellar response by $\sim 10\%$ and extends the stellar wake’s density response by ~ 50 kpc. The enhancement in the velocity dispersion within the stellar wake is $\sim 5\%$ higher when CDM self-gravity is turned on, and $\sim 5\%$ higher in FDM compared to CDM.

5. Discussion

In this section, we discuss the implications of our results in a wider context. We introduce a toy model for the observables of the stellar wake in Section 5.1, assess the sensitivity of the LMC’s orbit to the choice of DM particle in Section 5.2, and discuss the DM wake’s mass and its impact as a perturber of the MW’s dark halo in Section 5.3.

5.1. Observational Predictions

In Section 4, we presented three key predictions for the stellar wake. The gravity of a DM wake will (1) enhance the overdensity of the stellar wake by roughly 10%; (2) extend the length of the stellar overdensity and kinematic response by a few tens of kiloparsecs; and (3) the velocity dispersion enhancement of the stellar wake will be mildly ($\sim 5\%$) higher in response to an FDM wake than a CDM wake. In this section, we assess the extent to which these results could be observable by introducing a toy model to approximate how our windtunnel wakes would be viewed from Earth. Using this toy model, we study the density and radial velocity dispersion of the stellar wake with the addition of simulated distance and radial velocity errors.

To study how the stellar wake will appear when observed from Earth, we transform our “windtunnel” or simulation-box coordinate system to Galactic (l, b, r) coordinates, in which the origin is the solar system barycenter, the x -axis points toward the Galactic center, and the z -axis is normal to the Galactic plane. The LMC’s path in the windtunnel is straight as opposed to a curved orbit, so we cannot exactly reproduce the appearance of the GC19 wake on the sky, nor can we reproduce the effect of the collective response.

However, we can carefully choose the transformation to ensure we are best-reproducing the orientation and location of

the GC19 wake in the region of sky where we want to make our observations. In this case, following GC19 and our argument in Section 3.1, we want to focus our observations where the wake is at a Galactocentric distance of 70–100 kpc. Therefore, our goal is to transform from windtunnel coordinates such that the straight windtunnel path is tangent to the LMC’s orbit at a Galactocentric distance of 70 kpc, while the LMC itself is as close to its present-day location on the sky as possible. Our coordinate transformations are performed with Astropy version 4.2.1 (Astropy Collaboration et al. 2013, 2018, 2022) and are described in Appendix B.

The result of the coordinate transformation is shown in Figure 18. We plot the LMC’s orbit in the GC19 reference simulation in Galactocentric coordinates in red as in Figure 1. In our toy observational model, the path of the LMC in the windtunnel is tangent to the LMC orbit from GC19 at 70 kpc from the Galactic center, which is the distance that our Fiducial wind parameters are taken from.

To estimate how observational uncertainties affect our results, we also include Gaussian distance and radial velocity errors in our model. We choose two levels of errors, motivated by the performance of contemporary surveys. The distance errors are 10% and 20%, typical for spectrophotometric distance measurements from DESI (Cooper et al. 2023) and the H3 survey (Conroy et al. 2019). For the radial velocity errors, we choose 1 and 10 km s^{-1} . Note that 1 km s^{-1} velocity errors reflect the performance of spectroscopic radial velocity measurements from DESI (Cooper et al. 2023), Gaia (Katz et al. 2019; Seabroke et al. 2021), and H3 (Conroy et al. 2019), while 10 km s^{-1} provides a reasonable worst-case scenario.

With our toy model in-hand, we can now use it to study how the stellar wake might be observed. Figure 19 shows all-sky Mollweide projections (made with Healpy; Gorski et al. 2005; Zonca et al. 2019)¹⁷ in Galactic coordinates of the overdensity of stars with distances of 70–100 kpc in the Fiducial simulations after 0.7 Gyr of evolution. The bin size is 1.16° , and the resulting density map is smoothed by a Gaussian kernel with $\sigma = 15^\circ$. Each panel corresponds to a DM model,

¹⁷ <https://healpix.sourceforge.io/>

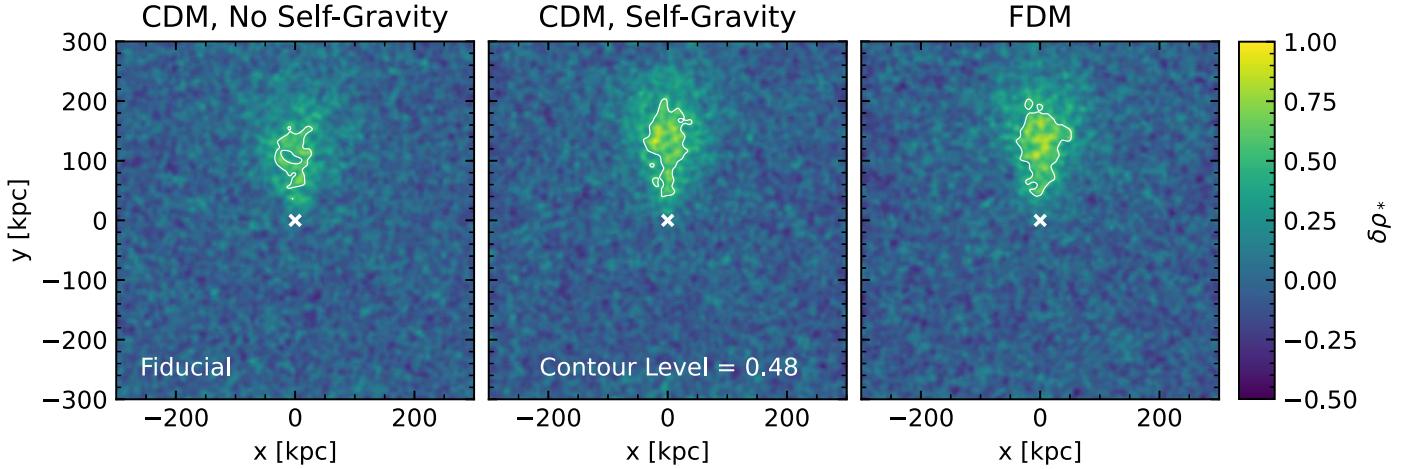


Figure 13. Projected overdensity (similar to Figures 2 and 5) of stars in the Fiducial simulations after $t = 0.7$ Gyr of integration. As in Figures 2 and 5, a white contour in each panel traces the region with a density greater than the half-max of the CDM simulation with self-gravity (center). The CDM simulation without self-gravity (left) shows the stellar wake that is present due to only the passage of the LMC, i.e., we would expect this wake to form in the stellar halo regardless of the behavior (or existence) of the DM wake. When compared to the stellar wake with CDM self-gravity (center), we see that the entire wake is not enclosed by the contour in the absence of self-gravity (left), i.e., there is a region in the center of the wake without CDM self-gravity that drops below $\delta\rho_* = 0.48$. This demonstrates that adding the gravity of the DM wake enhances the density of the stellar wake, which we quantify in Figure 14. Additionally, the contoured region extends ~ 50 kpc farther behind the LMC when the DM wake’s gravity is included. Comparing the CDM simulation with self-gravity to the FDM simulation (right), the stellar wake’s density looks similar in both of these DM models.

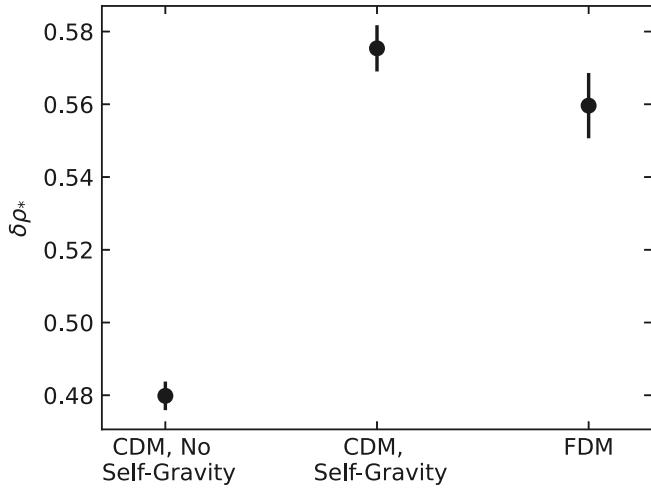


Figure 14. Median densities of the stellar wakes in the Fiducial simulations, averaged over 100 Myr of evolution from $t = 0.6$ – 0.7 Gyr. Error bars show the standard deviation of the median density during this time period. The stellar wake’s density is increased by $\delta\rho_* \sim 0.1$ when the gravity of the DM wake is added. The stellar wake is less dense in response to an FDM wake than a CDM wake, but this difference is small, only $\delta\rho_* \sim 0.02$.

with CDM without self-gravity on the left, CDM with self-gravity in the center, and FDM on the right. Each panel shows the path of the LMC in the windtunnel as the solid white line, and the LMC orbit from the reference simulation as the dashed white line; we see good agreement between the position of both paths on the sky. As in Figure 13, we enclose the region with an overdensity higher than 0.34 with a contour; this level is the half-maximum density of the CDM simulation with self-gravity.

In agreement with GC19, the stellar wake appears as an overdense region in the Galactic southeast, ranging from $l \sim 0$ –120, and $b \sim -80$ –0. Notably, the extension of the stellar wake owing to the DM wake’s gravity is readily observable: while the stellar wakes in the CDM simulation with self-gravity

and FDM simulation do not decay below $\delta\rho_*$ of 0.38 until $b \approx 0$, the stellar wake decays to this level by $b \approx -20$ in the simulation without self-gravity.

To quantify the differences in the strength of the response in this observed frame, we use the same procedure as in Section 4: we calculate the median wake density or velocity dispersion in bins that are higher than half of the maximum bin. For each simulation, we repeat this for five snapshots spanning 100 Myr, and then report the average and standard deviation of the medians from the five snapshots. In this section, we calculate the quantity of interest in on-sky bins as in Figures 19 and 21.

To estimate the number of stars that need to be observed to distinguish between the simulations, we also downsample the number of star particles, i.e., after adding simulated errors, we sample a fixed number of stars with distances between 70 and 100 kpc from the entire sky. Without downsampling, there are approximately 10^5 stars with distances in this range based on the stellar wind density and the volume of the shell. For our plots, we choose three different levels of downsampling, selecting 1.5×10^4 , 10^4 , and 10^3 stars. These sampling levels correspond to selecting approximately 1300, 900, and 70 stars within the wake (i.e., inside the contours in Figure 19), respectively.

Figure 20 shows the time-averaged median overdensity of the stellar wake between 70 and 100 kpc with different observational errors and sampling rates. The black circles show the mean and standard deviation of the median overdensity without any observational errors. The error bars on each point are computed via bootstrapping, i.e., for each of the five snapshots, we randomly sample errors and star particles 50 separate times such that the final reported median overdensity is over 250 samples.

With no errors and 1.5×10^4 stars, when we compare the CDM simulations with and without self-gravity, the stellar wake’s overdensity increases by ~ 0.05 with self-gravity. In the observational frame, we now also see a further increase in density in the FDM simulation, with the FDM simulation reaching $\delta\rho_* \sim 0.07$ higher than the CDM simulation with self-

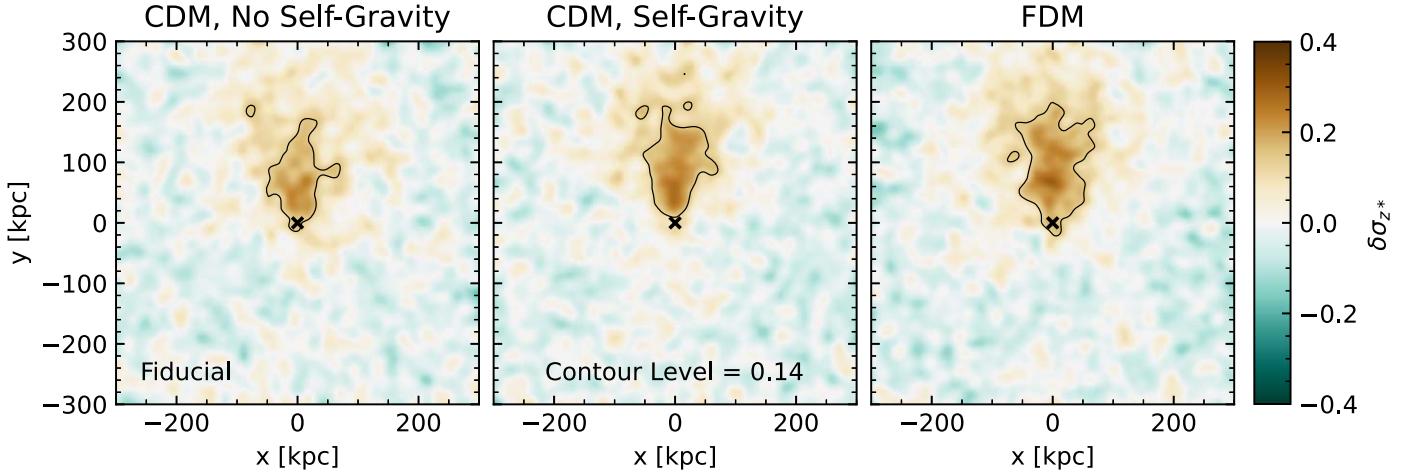


Figure 15. z -velocity dispersion of the stellar wakes in the Fiducial simulations (similar to Figure 7 but for the stars). Comparing the CDM simulations, we see that the addition of the DM wake’s gravity widens the contour by ~ 20 kpc. While the contour extends roughly 180 kpc behind the LMC in both CDM simulations, the wake tapers more quickly without the DM self-gravity. Comparing the FDM simulation to the CDM simulation with self-gravity, the FDM wake tapers more slowly than either CDM simulation, though it reaches a similar maximum width to the CDM simulation with self-gravity.

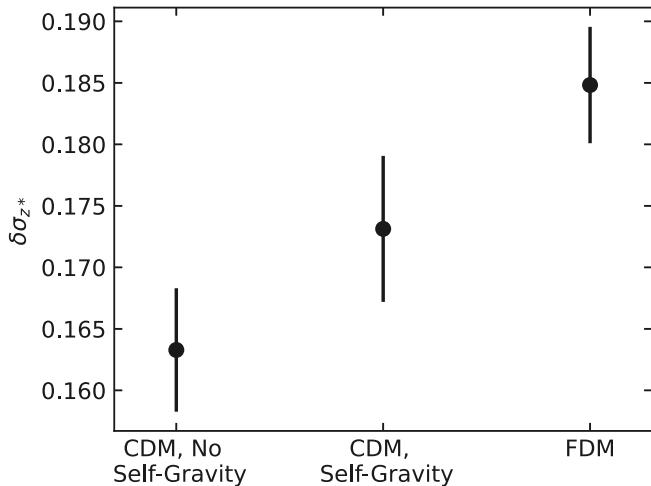


Figure 16. Time-averaged median enhancement of the z -velocity dispersion in the stellar wakes. The stellar wake’s velocity dispersion is raised by $\delta\sigma_{z*} \sim 0.01$ when adding the gravity of the DM wake. The stellar wake is heated $\delta\sigma_{z*} \sim 0.01$ more by an FDM wake than a CDM wake.

gravity. Note that this is opposite to the trend we saw in Section 4, where the stellar wake was slightly less dense in FDM compared to CDM. In the observational model, we are now looking at a 30 kpc-thick slice of the wake, as opposed to 120 kpc in Section 4, so this is most likely an effect of the viewing angle and distance selection of stars. When adding observational errors and reducing the number of stars to 10^4 , the differences between the simulations remain visible. Sampling only 10^3 stars, however, is not sufficient to see the differences between the simulations.

Figure 21 shows all-sky maps of the enhancement in the radial velocity dispersion of the stars in the same fashion as Figure 19. In this plot, we report the velocity dispersion as the difference from the shell average $\Delta\sigma_{r*} = \sigma_{r*} - \bar{\sigma}_{r*}$. In all panels, the velocity response traces the location of the density response well. The addition of the DM wake’s gravity extends the length of the velocity response, as it decays to below 5.26 km s^{-1} by $b \approx -30$ in the CDM simulation without self-gravity, compared to $b \approx 0$ in both simulations with the DM self-gravity.

It is also worth mentioning that we expect an increase in both the longitudinal and latitudinal velocity dispersion in the wake. At these distances (70–100 kpc), we measure this increase to be approximately 0.03 mas yr^{-1} . For our purposes of distinguishing between DM models, the qualitative differences between the simulations are the same as for the radial velocity dispersion, so we do not elaborate on the proper motions here for brevity.

Figure 22 shows the median velocity response averaged over 100 Myr with observational errors and different numbers of stars, similar to Figure 20. Here, we see the same trend in the observational frame that we did in the simulation-box frame in Section 4: with 1.5×10^4 stars, the velocity dispersion enhancement in the stellar wake is lowest ($\sim 5.0 \text{ km s}^{-1}$) without a DM wake’s gravity, higher in response to a CDM wake ($\sim 5.6 \text{ km s}^{-1}$), and highest in response to an FDM wake ($\sim 6.5 \text{ km s}^{-1}$). The addition of observational errors does not affect this trend, i.e., the simulations are still distinguishable with the largest errors we consider. Note that 10^4 stars is also enough to distinguish the simulations, though the differences between CDM with self-gravity and FDM become close to 1σ with 10 km s^{-1} radial velocity errors and 20% distance errors. The differences between the simulations are not visible while sampling only 10^3 stars.

With the caveat that our observational framework is only a toy model, we find that the general results reported in Section 4 still hold. In particular, we have demonstrated several important qualitative results as follows: Distinguishing the strength of the density and kinematic response of the stellar wake between DM models should be possible with $\geq 10^4$ stars across the entire sky ($\gtrsim 900$ stars within the wake) with distances between 70 and 100 kpc. This sampling rate corresponds to a number density of $3.6 \times 10^{-3} \text{ kpc}^{-3}$, which agrees with the number density of stars that GC19 reported is required to confidently detect the wake. In other words, if we observe enough stars to detect the wake, we have enough stars to distinguish between the DM models considered here.

Provided this sampling rate is achieved, we find that the telltale sign of the presence of a DM wake is the length of the response, as both the density and velocity dispersion responses are lengthened by over 20° on the sky when the self-gravity of

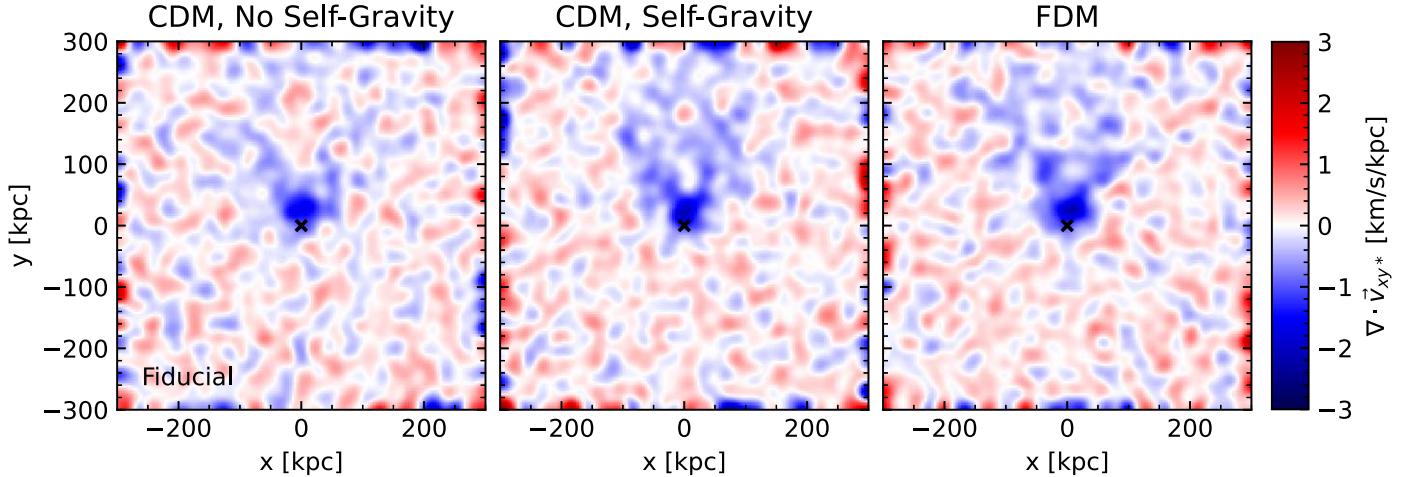


Figure 17. Divergence of the x - y velocity field for the stars in the Fiducial simulations. As in Figure 11, the star particles in the no-self-gravity CDM simulation (left) converge just behind the LMC and then continue on relatively straight paths, forming a divergent region farther downstream. Turning on DM gravity in the other panels largely eliminates this diverging region just as it does for the DM. This demonstrates how the DM wake enhances the density of the stellar wake and keeps it coherent for larger distances behind the LMC (see Figure 13), but there are no discernible differences between CDM and FDM here.

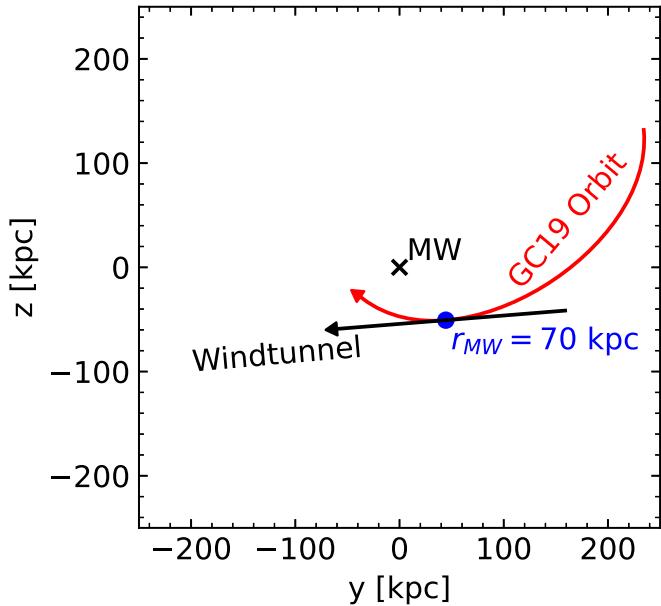


Figure 18. Schematic showing the location of the Fiducial windtunnel wake in our toy observational model in comparison to the LMC's orbit from the [GC19](#) reference simulation. As in Figure 1, we show both orbits in Galactocentric coordinates, projected onto the y - z plane. The MW center of mass is denoted by the cross, the [GC19](#) orbit is shown by the red line, and the black line shows the path of the LMC in our windtunnel simulations (equivalent to the y -axis in our simulation box). Arrows at the head of each path show the present-day location of the LMC. A blue dot shows the location at which the LMC is at a Galactocentric distance of 70 kpc, which is where we draw our Fiducial wind parameters. The path of the LMC in the windtunnel is tangent to the [GC19](#) orbit at this location.

the DM wake is included. Additionally, we find that differences in the kinematics of the stellar wake between a CDM and FDM universe are still visible when accounting for the viewing perspective and observational errors. As also reported by [GC19](#), we find that the increased velocity dispersion is a characteristic signature of the wake that differentiates it from cold substructure such as stellar streams. Ultimately, these results demonstrate that kinematic information is crucial when making observations of DF wakes, both for detecting the wake and inferring the nature of its DM component.

5.2. Dynamical Friction Drag Forces and the LMC's Orbit

In this section, we compare the behavior of the DF drag force felt by the LMC due to the DM wakes in our simulations and discuss the impact of DM microphysics on the LMC's orbit.

To determine the acceleration due to DF in our simulations, we calculate the y -component of the gravitational acceleration that would be felt by a constant-density sphere 5 kpc in radius at the center of the box due to all DM particles in the simulation. When done at each time step, this gives us an approximation of the DF acceleration felt by the LMC as a function of time.

Additionally, we calculate the expected DF acceleration using the classic formula from Chandrasekhar (1943):

$$a_{\text{DF}} = \frac{4\pi^2 MG^2 \bar{\rho} \ln(\Lambda)}{\sigma^2} \frac{1}{2X^2} \left[\text{erf}(X) - \frac{2X}{\sqrt{\pi}} e^{-X^2} \right], \quad (21)$$

where ‘‘erf’’ is the error function and

$$X = v/\sqrt{2}\sigma. \quad (22)$$

In these equations, we use the input wind parameters and LMC mass, i.e., M from Table 1, and $\bar{\rho}$, v , and σ from Table 2.

For the Coulomb logarithm, we follow van der Marel et al. (2012), Patel et al. (2017), and [GC19](#), using

$$\ln(\Lambda) = \max [L, \ln(r/Ca)^\alpha], \quad (23)$$

where r is the distance between the satellite and its host, a is the satellite's scale radius, and $L=0$ and $\alpha=1$, and C are constants. Here, r is the separation between the LMC and MW at the point in the reference simulation that we base our wind parameters on (70 kpc for the Fiducial case and 223 kpc for the Infall case), and a is the LMC's scale radius from Table 1. For C , we pick values such that the analytic DF acceleration roughly agrees with the measured acceleration when the wake reaches the end of the box. For the Fiducial wind, this is $C=0.52$, and for the Infall wind, $C=2.90$.

Figure 23 shows the measured and analytic DF accelerations for our simulations, with the Fiducial wind case in the left panel and the Infall wind case on the right. In each simulation, the strength of the drag increases with time as the wake forms, before plateauing once the faster-moving wake particles begin to wrap though the box. Overall, the drag from the Fiducial

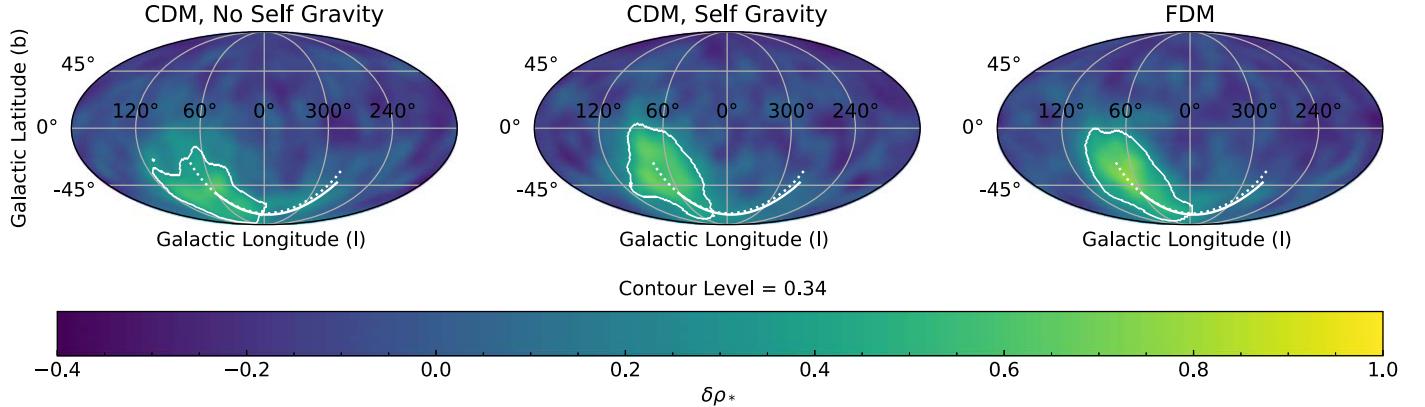


Figure 19. All-sky maps of the overdensity of stars in the Fiducial simulations after 0.7 Gyr of integration. Each panel shows a Mollweide projection in Galactic coordinates as defined by the *Astropy* convention, and includes stars with distances between 70 and 100 kpc. White lines show the paths of the LMC from our GC19 reference simulation (dashed) and in our windtunnel (solid). As in Figure 13, the contour in each panel encloses the region with a response greater than the half-maximum of the CDM simulation with self-gravity, which is $\delta\rho_* > 0.34$ here. Panels show different DM models, with CDM without self-gravity on the left, CDM with self-gravity shown in the central panel, and FDM on the right. In all three panels, the wake appears as an overdensity covering nearly an entire octant of the sky, from $l \in [0, 120]$ and $b \in [-80, 0]$. The addition of self-gravity extends the contoured region to $b \approx 0$, while without self-gravity the wake decays to below the contour by $b \approx -20$. The extent of the stellar wake is similar in CDM and FDM when self-gravity is included.

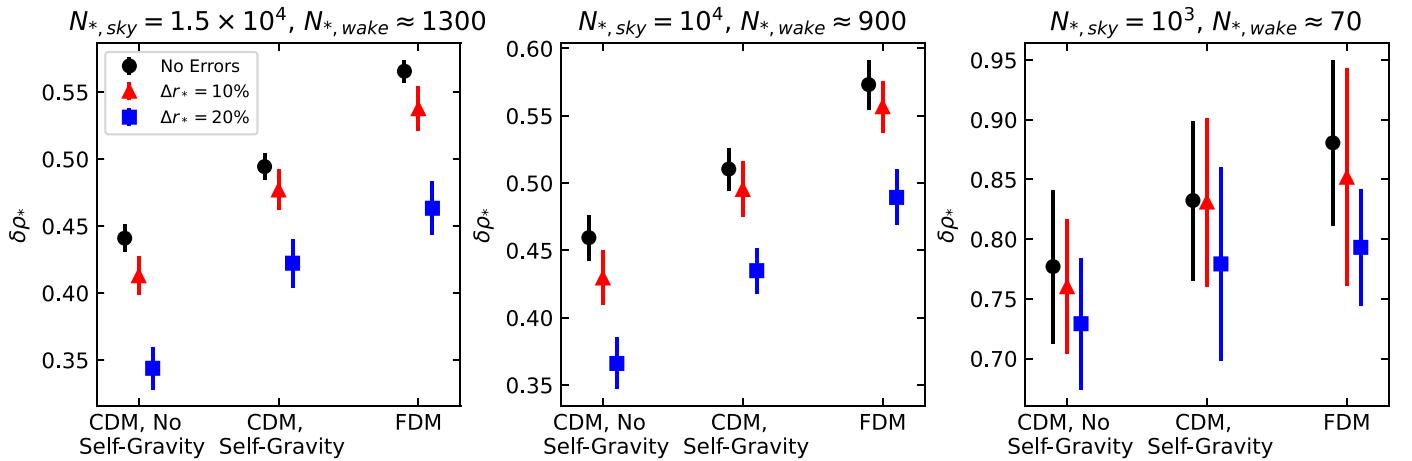


Figure 20. Time-averaged median wake densities (as in Figure 14) in the observable frame with observational errors. In each panel, we sample a different number of stars from the entire sky with distances between 70 and 100 kpc to compute the median wake density. Each point shows the average and standard deviation (computed via bootstrapping) of the median density across all sampled snapshots. Black circles show the simulation data with no errors. The red triangles (blue squares) have simulated distance errors of 10% (20%). With no observational errors at our highest sampling rate (1.5×10^4 stars; left panel), the stellar wake is $\delta\rho_* \sim 0.05$ higher when the CDM wake’s gravity is included, and the stellar wake in the FDM simulation is $\delta\rho_* \sim 0.07$ higher than in the CDM simulation with self-gravity. As observational errors increase, the measured overdensity of the wake decreases but the same trend between all three simulations remains. With 10^4 stars (middle panel), the differences between DM models are lower but still distinguishable with 20% distance errors. With only 10^3 stars (right panel), the DM models are indistinguishable.

wake is slightly more than an order of magnitude stronger than the drag from the Infall wake, which aligns with the ρ/v^2 scaling expected from Equation (21).

In the Fiducial case, we see that the reduction in wake size and density when self-gravity is removed translates to a weaker drag force—the acceleration is $\sim 10\%$ weaker in the CDM simulation without self-gravity versus with self-gravity. Meanwhile, the behavior of the FDM drag is consistent with the predictions of Lancaster et al. (2020), who calculated that the time-averaged drag force on the LMC should be well-approximated by classical DF (i.e., with noninteracting background particles). In the Infall case, we see closer agreement between the two CDM simulations, as the effect of the wake’s self-gravity is diminished at this lower wind speed.

Ultimately, our result that both DM models produce a similar drag force regardless of the wind speed and density (when DM

self-gravity is included) implies that the LMC’s orbit would not be impacted by the assumption of a CDM versus FDM universe.

5.3. The Mass of the Wake

In this section, we calculate the mass of the DM wakes in our simulations, and develop a basic framework to understand the DM wake as a perturbation to the MW’s DM halo.

To calculate the wake mass in each of our simulations, we begin by defining a rectangular region that roughly contains the wake (i.e., that contains where $\delta\rho \geq 0.1$; $x \in [-100, 100]$, $y \in [-50, 300]$, $z \in [-100, 100]$ for the Fiducial wake; $x \in [-150, 150]$, $y \in [-50, 300]$, $z \in [-150, 150]$ for the Infall case). At a particular time step, the wake mass is estimated by taking the difference between the total DM mass within the region at that time step and the mass within the region at the start of the simulation, i.e., the region’s volume multiplied by $\bar{\rho}$

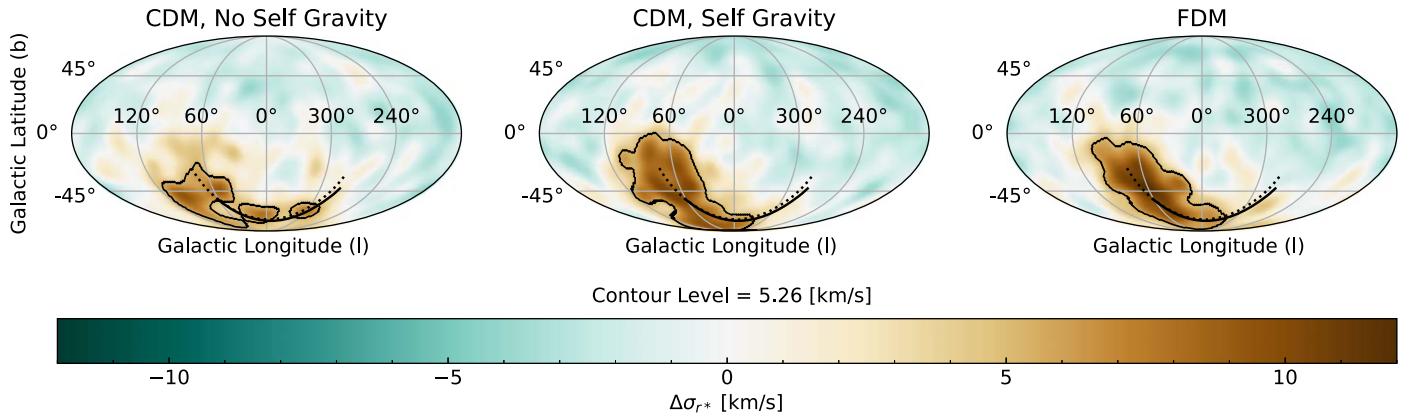


Figure 21. All-sky maps of the change in the radial velocity dispersion with respect to the shell average, $\Delta\sigma_{r*} = \sigma_{r*} - \bar{\sigma}_{r*}$, for the fiducial simulations after 0.7 Gyr. Similar to the simulation-box projection (Figure 15), the wake signature is an enhancement of the stars' radial velocity dispersion. The location of the velocity dispersion response is similar to the location of the density response in Figure 19. Without self-gravity, the wake decays below the contour by $b \approx -30$. Both FDM and CDM with self-gravity extend the stellar wake to $b \approx 0$.

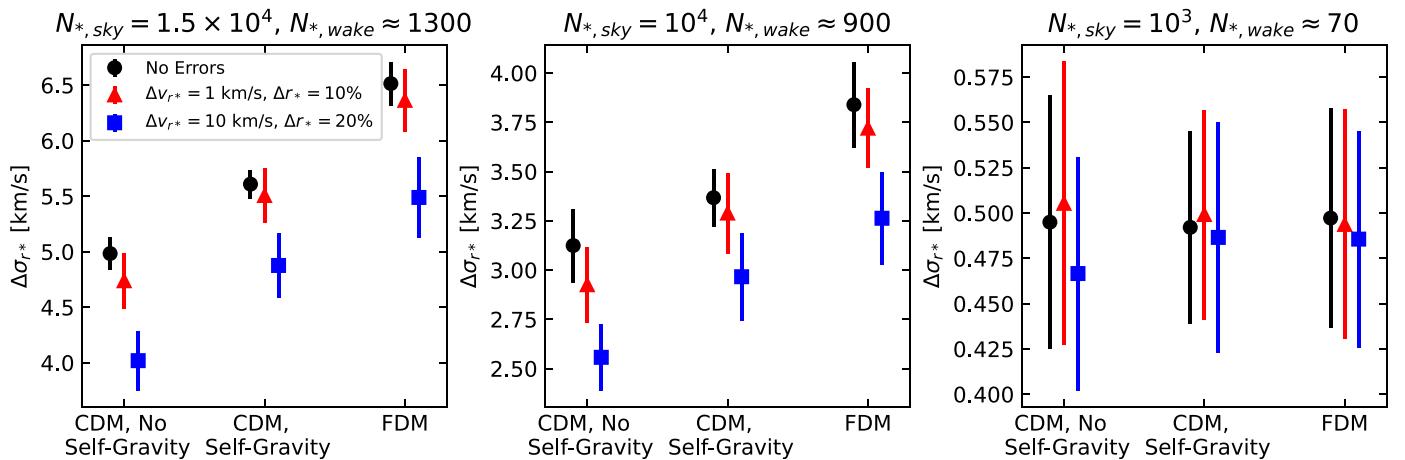


Figure 22. Time-averaged median radial velocity dispersion enhancement (same as Figure 20 but showing $\Delta\sigma_{r*}$) of the stars in our toy observational model. In addition to distance errors, the red triangles (blue squares) have simulated radial velocity errors of 1 (10) km s^{-1} . With 1.5×10^4 stars and no observational errors, the velocity dispersion of the stellar wake is $\sim 0.6 \text{ km s}^{-1}$ higher than the simulation without self-gravity with the CDM wake's gravity, and $\sim 1.5 \text{ km s}^{-1}$ higher with the FDM wake's gravity. Adding observational errors reduces the measured velocity dispersion enhancement, though the differences between the simulations remain above 1σ . With 10^4 stars, the simulations are still distinguishable, but the differences reach just above 1σ with 10 km s^{-1} and 20% distance errors. Sampling 10^3 stars is insufficient to distinguish the simulations at any error level.

from Table 2. As we have done throughout this work, when estimating the wake mass, we average over five snapshots spanning 100 Myr of evolution.

The top row of Figure 24 shows the masses of all DM wakes in our simulations. The left panel shows the Fiducial wind after 0.7 Gyr of evolution, and the right panel shows the Infall wind after 2 Gyr of evolution. The mass of the Fiducial wake is roughly comparable to the LMC, while the mass of the Infall wake is roughly an order of magnitude lower. In both the Infall and Fiducial cases, the FDM wake and CDM wake with self-gravity have similar masses, while the CDM wake without self-gravity is of the order of 10% less massive than either wake with self-gravity.

To get a rough approximation of the impact of the DM wake as a perturbation to the MW's DM halo, we also calculate the distance at which an object with the wake's mass would need to be behind the LMC to produce a similar drag force as the wake. The middle row in Figure 24 lists the DF acceleration during the same time frames as the top panel, taken from Figure 23.

The distances at which an object of the wake mass would produce a gravitational acceleration equivalent to DF are shown in the bottom row of panels in Figure 24. In the Fiducial case, the distances all agree, and are approximately 100 kpc. The Infall distances are roughly 135 kpc, and also show agreement between each DM model.

In summary, we see that the Fiducial wake acts like an additional LMC-mass object that trails the LMC at a distance of 100 kpc, while the Infall wake is equivalent to an object with roughly 10% the mass of the LMC trailing at a distance of 135 kpc. Additionally, this behavior is insensitive to the assumption of CDM or FDM.

6. Discussion: Simulation Parameters

In this section, we explore how our results are affected by changing certain assumptions in our simulation setup. We assess the importance of the FDM particle mass to our results in Section 6.1. Section 6.2 discusses the impact of the uncertainty in the LMC's mass on our observational predictions. We

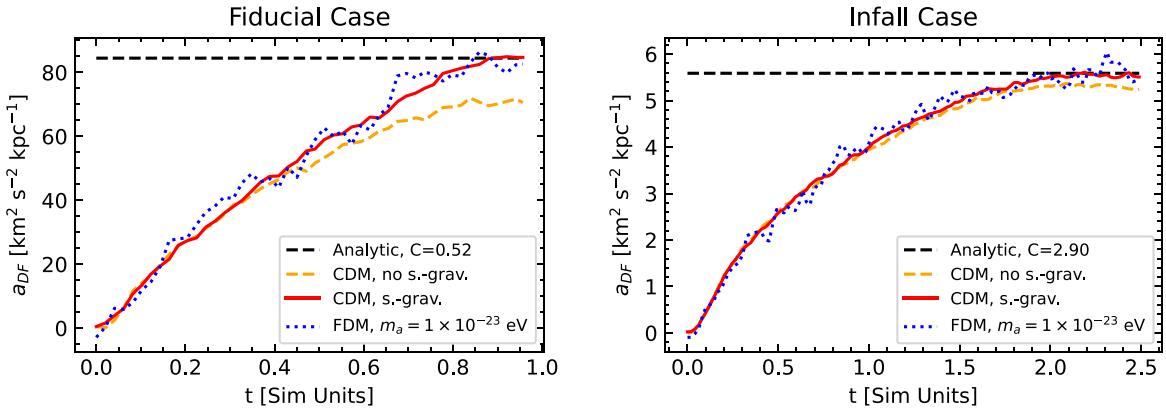


Figure 23. DF accelerations vs. time. The CDM simulations with and without self-gravity are shown by the solid red and dashed orange lines, respectively. The primary FDM mass ($m_a = 10^{-23}$ eV) is the blue dotted line. The black dashed line in each panel shows the analytic acceleration calculated from the classical Chandrasekhar formula as described in the text. In both orbit cases, the measured acceleration increases in strength with time as the wake forms. Left panel: the Fiducial orbit case. CDM produces a $\sim 10\%$ stronger acceleration when self-gravity is turned on, consistent with the enhancement in the wake’s density and longevity when self-gravity is included. FDM produces a similar acceleration to the self-gravity-on CDM wake. The steady-state drag force is similar to the analytic approximation when $\ln(\Lambda)$ is calculated using $C = 0.52$. Right panel: the Infall orbit case. The reduced effect of DM self-gravity in the Infall scenario is apparent in the similar accelerations between the two CDM simulations. Again, the FDM traces the CDM acceleration well. The steady-state drag force matches the analytic approximation when $\ln(\Lambda)$ is calculated using $C = 2.90$. The drag force in the Fiducial case is much higher than the drag force in the Infall case, as the background density is roughly two orders of magnitude larger, while the wind velocity is only a factor of ~ 3 larger. Overall, the FDM drag force matches the CDM drag force well, consistent with the predictions of Lancaster et al. (2020) with respect to DF on the LMC from an FDM wake.

quantify the effect of the stellar halo’s velocity dispersion in Section 6.3 and discuss implications for the wake’s impact on cold stellar substructures. Finally, we discuss the prospects for using the wake to constrain alternative DM models beyond FDM in Section 6.4. To study each of these effects, we run additional simulations, which are summarized in Tables 4 and 5.

6.1. The Effect of FDM Particle Mass

As the behavior of FDM is strongly dependent on the particle mass m_a , it is important to place our choice of $m_a = 10^{-23}$ eV into context within the literature and test the extent to which a different choice would affect our results. Table 6 compiles a list of recent papers that report a constraint on m_a through an astrophysical technique (see also Ferreira 2021 for a recent review, and Figure 1 of Dome et al. (2023) for a graphical approach). We do not guarantee that this list is exhaustive, nor do we include constraints from laboratory or direct-detection experiments. Nevertheless, we hope to demonstrate that FDM particle mass constraints are abundant and may be derived with a very wide range of methods. Notably, the constraints we list here span the entire range of FDM masses (10^{-26} – 10^{-16} eV), though almost all come with caveats.

One common method of constraining m_a relies on trying to detect soliton density cores in dwarf galaxies (e.g., Bar et al. 2019; Desjacques & Nusser 2019; Safarzadeh & Spergel 2020). The widest constraint comes from Safarzadeh & Spergel (2020), who reported that a single-component FDM is incompatible with the observed differences between Fornax and Segue 1’s central density profiles. This result relies heavily on the measurement of the ultra-faint dwarf (UFD) density profile slopes, and relaxing the core profile slope constraint from Walker & Peñarrubia (2011) results in a lower bound of $m_a > 6 \times 10^{-22}$ eV. Moreover, Chan et al. (2022) reported that there is significant scatter in the FDM soliton core–halo mass relation, which may weaken constraints derived by examining DM density profiles.

Meanwhile, there is a growing tension between the requirements that FDM is light enough that it produces

sufficiently large cores in dwarf galaxies (Marsh & Pop 2015; González-Morales et al. 2017) and heavy enough that it is consistent with the small-scale matter power spectrum as inferred from the Ly α forest (Armengaud et al. 2017; Iršič et al. 2017; Kobayashi et al. 2017; Rogers & Peiris 2021) and other cosmological probes (e.g., Dentler et al. 2022). Such cosmological probes of the FDM mass typically rely on comparisons to simulations performed with traditional N -body codes, which modify the linear power spectrum of the initial conditions (and sometimes the transfer function) to match that expected of FDM using axionCAMB (Hložek et al. 2017). Schive et al. (2016) argue that this approach is valid for power spectrum modeling, though such methods do not consider nonlinear effects. Large-scale (box sizes of $L \gtrsim 10$ Mpc/ h) cosmological simulations with full SP solvers are becoming available (May & Springel 2021, 2023), which could be used to test the Ly α results with higher-fidelity simulations.

Other methods (such as this work) rely on examining the gravitational effect of FDM granules and/or subhalos on luminous matter. There is a growing class of papers that examine dynamical heating of stars by FDM substructures as a method of placing upper limits on m_a (e.g., Amorisco & Loeb 2018; Church et al. 2019; Benito et al. 2020; Schutz 2020; Banik et al. 2021). Again, none of these studies utilizes a fully self-consistent numerical treatment of the FDM, instead approximating granules as massive extended particles or utilizing only the subhalo mass function.

A very recent study by Dalal & Kravtsov (2022) provides one of the more stringent kinematic constraints of $m_a > 3 \times 10^{-19}$ by examining the heating of stars by granules in the Segue 1 and 2 UFDs. Their simulation technique, outlined in Dalal et al. (2021), approximates FDM granules as linear perturbations to a static potential. This results in a computationally inexpensive, relatively accurate treatment of the wave behavior of FDM in the idealized case of a spherically symmetric, equilibrium halo.

While very few of these existing constraints have been confirmed with self-consistent, nonlinear SP simulations, our

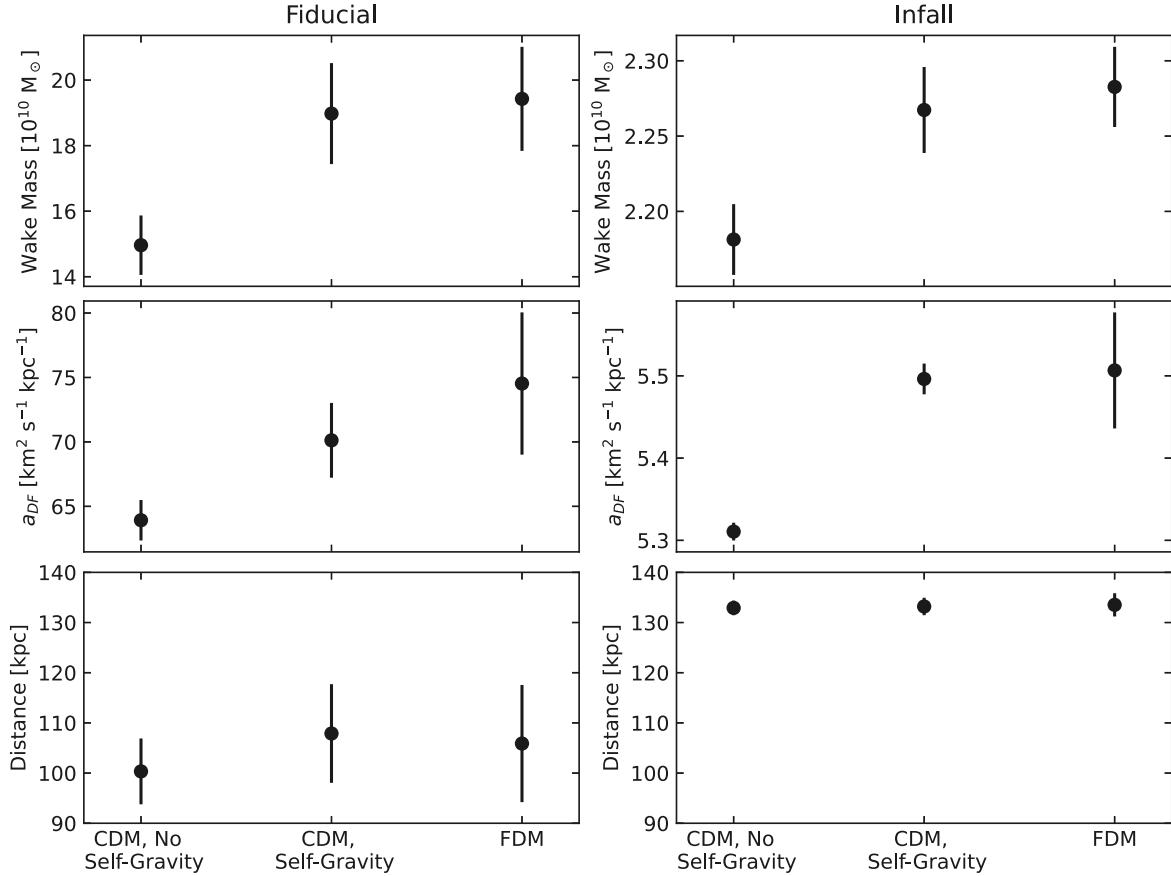


Figure 24. Wake masses in our simulations, calculated as described in the text. In all panels, quantities are averaged over five snapshots spanning 100 Myr, and the points show the mean and standard deviation of the quantity of interest. The left column of panels shows the Fiducial wind simulations from 0.6–0.7 Gyr of evolution, and the right column shows the Infall wind simulations from 1.9–2 Gyr of evolution. The top row shows the wake masses. In the Fiducial case, the wake masses are roughly that of the LMC’s virial mass at infall, while in the Infall case, they are roughly an order of magnitude less. In both wind cases, the CDM wake without self-gravity is of the order of 10% less massive than the CDM wake with self-gravity. Both wakes with self-gravity have similar masses. The middle row of panels shows the DF acceleration on the LMC due to the wake during at the same times as we measure the wake mass. See Section 5.2 and Figure 23 for further discussion of the drag forces. The bottom row of panels shows the distance at which a point object of the mass in the upper row would produce the gravitational acceleration in the middle row, for each simulation. Thus, the Fiducial wake acts like an LMC-mass perturber trailing the LMC at a distance of ∼100 kpc, while the Infall wake acts as a perturber with ∼1/10 the mass of the LMC at a distance of ∼135 kpc.

choice of $m_a = 10^{-23}$ eV is clearly inconsistent with a wide range of observational probes. As justified in Section 2.5, this is the largest mass we can feasibly simulate, so it is important to explore how our results are affected by another choice of m_a . To determine this, in addition to our $m_a = 10^{-23}$ eV simulations, we have performed another set of simulations with $m_a = 2.5 \times 10^{-24}$ eV (see Table 4). We discuss each of our FDM-specific results and their dependence on m_a in turn.

Dark matter wake structure. Figure 25 shows the over-density projections of both FDM simulations with the Fiducial wind parameters (similar to Figure 2). Reducing the mass by a factor of 4 correspondingly increases the de Broglie wavelength of the FDM particles by a factor of 4. As expected, this increases both the size and relative strength of the granule density fluctuations, with peak granule densities within the wake reaching overdensities of ∼2.4 (1.7) for the lower (higher) particle mass. In the low-mass case, some of the background granules (those outside the wake) reach higher overdensities than the half-max of the CDM wake with self-gravity. At higher masses than we are able to simulate, the granules would decrease in size and strength and the density field of the wake would approach the behavior of CDM.

Dark matter wake velocity dispersion. In Figure 26, we reproduce Figure 8 with the inclusion of our $m_a = 2.5 \times 10^{-24}$ eV

FDM simulation as a dashed-dotted, purple line. The CDM simulation without self-gravity is removed for clarity. The increased de Broglie wavelength of the low-mass simulation causes larger velocity granules, which can be seen as the increased oscillation amplitude in the profile of the low-mass FDM wake. This roughly four-fold increase in the oscillation strength is inversely proportional to the decrease in mass compared to the primary (higher) FDM mass. Despite the oscillations, the two particle masses we consider here show very similar overall/averaged behavior, i.e., when comparing the two masses tested here, our result that the dispersion enhancement of an FDM wake is ∼80% that of CDM is unchanged. We caution that this result may not hold for higher particle masses, especially as FDM phenomenology approaches CDM when m_a increases. It is, however, suggestive that the kinematic signatures of FDM wakes are less sensitive to m_a than their density field signatures.

Kinematics of the stellar response. In Section 4, we argued that FDM granule heating is responsible for raising the velocity dispersion of the stellar wake in an FDM universe compared to a CDM universe. Following an argument similar to that of Dalal & Kravtsov (2022), we can roughly estimate the extent to which granule heating is expected to operate within our windtunnel simulations: FDM granules are approximated as objects of mass $\delta M \approx \bar{\rho} r^3$, where we assume that the granule

Table 4
Summary of Our Additional Simulations for Section 6, Listing the Properties for Each Simulation

Category	DM Model	DM Self-gravity	Particles or Cells	DM Wind	Notes
Low-mass FDM, Section 6.1	FDM, $m_a = 2.5 \times 10^{-24}$ eV	Yes	256^3	Fiducial	^a
	FDM, $m_a = 2.5 \times 10^{-24}$ eV	Yes	256^3	Infall	^a
Alternative LMC Masses, Section 6.2	CDM	No	10^8	Fiducial	^b
	CDM	No	10^8	Fiducial	^c
Low Stellar Dispersion, Section 6.3	CDM	Yes	10^8	Fiducial	^d
	FDM, $m_a = 10^{-23}$ eV	Yes	1024^3	Fiducial	^d

Notes. We list the following: the simulation category along with the subsection where the simulations are discussed; the dark matter model with the FDM particle mass if applicable; whether DM self-gravity is enabled; the number of DM resolution elements (N -body particles for CDM, grid cells for FDM); the orbit/DM wind case (see Table 2); and how the simulation differs from those described in Table 2.

^a FDM particle mass is reduced by a factor of 4.

^b Uses the “Light LMC,” see Table 5.

^c Uses the “Heavy LMC,” see Table 5.

^d The stars in this simulation have a velocity dispersion of $\bar{\sigma}_* = 30.0$ km s⁻¹.

overdensity fluctuation is of order unity, and the granule radius $r \approx \hbar/m_a \bar{\sigma}$ is set by the de Broglie wavelength associated with the FDM velocity dispersion. Thus, the FDM granules will cause a perturbation in the gravitational potential $\delta\Phi \approx G\delta M/r = G\bar{\rho}r^2$. Stars that encounter granules at a relative velocity of $\sim\bar{\sigma}_*$ will have their velocities perturbed by $\delta v \approx \delta\Phi/\bar{\sigma}_* = G\bar{\rho}r^2/\bar{\sigma}_*$. Repeated encounters would increase the velocity dispersion of the stars by $\Delta\sigma_* \approx \sqrt{N\delta v^2}$, where $N \approx \bar{\sigma}_* t/r$ is the number of star–granule encounters during a time t . Putting all of this together gives

$$\Delta\sigma_* \approx \sqrt{\frac{G^2 \bar{\rho}^2 t}{\bar{\sigma}_* \bar{\sigma}^3} \left(\frac{\hbar}{m_a}\right)^3}. \quad (24)$$

Notably, Equation (24) is derived assuming a uniform density and velocity dispersion of both DM and stars, i.e., similar to our initial conditions. In addition to the increase in granule density within the wake, Lancaster et al. (2020) and Vitos & Gourgouliatos (2023) demonstrated that FDM wakes grow additional interference fringes during the interaction with length scales set by the de Broglie wavelength associated with the wind velocity. Therefore, we do not necessarily expect Equation (24) to hold for our simulations, but it illustrates that we may expect granule heating to become stronger for lower values of the FDM particle mass.

Figure 27 reproduces Figure 16 but includes the low-mass FDM simulation in place of the CDM simulation without self-gravity. The leftmost two points are the same as the rightmost two points in Figure 16. The stellar wake’s velocity dispersion is increased more by the lower-mass FDM wake when compared to CDM, confirming that granule heating becomes stronger when the FDM particle mass decreases. This demonstrates that future observations of the stellar wake’s velocity dispersion may be used to place an independent constraint on m_a .

Overall, we find that our choices of m_a do not affect our result that an FDM wake is $\sim 20\%$ colder than a comparable CDM wake. We cautiously suggest that these results may hold at higher values of m_a , but emphasize the need for higher-resolution simulations conducted with values of m_a that are permitted by other astrophysical constraints to verify this conclusion. Additionally, we should expect granule heating of the stellar wake to decrease as m_a increases, and vice versa.

Table 5
Summary of Our Additional LMC Mass Models, Also from [GC19](#)

Galaxy Model	GC19 Model	$M (M_\odot)$	a (kpc)
Light LMC	LMC2	1.0×10^{11}	12.7
Heavy LMC	LMC4	2.5×10^{11}	25.2

Note. We list the following: our galaxy model; the corresponding galaxy model in [GC19](#); M , the total mass of the profile (if it were integrated to infinity); and a , the scale radius.

6.2. The Effect of the LMC’s Mass

In Section 4, we argued that the length of the stellar wake could be used to reveal the presence of the DM wake, as the DM wake’s self-gravity enables the stellar wake to persist for longer than without self-gravity. However, the LMC’s mass will affect the strength and length of the stellar wake in a manner that could be degenerate with the presence of the DM wake. To investigate this possibility, we ran two additional simulations with alternative LMC models (see Tables 4 and 5) of different masses. Both of these additional simulations are run in CDM without DM self-gravity to assess whether a more or less massive LMC could cause a density enhancement in the stellar wake similar to that caused by the addition of the DM wake’s gravity.

Figure 28 compares the density of the stellar wakes (similar to Figure 13) in all three simulations without self-gravity. The LMC mass differs between each column, and increases from left to right. To compare the density response in these simulations to that expected with CDM self-gravity, the contours are set at the half-maximum of the wake’s density *with* DM self-gravity, i.e., at the same level as in Figure 13. Increasing the LMC mass increases the strength and length of the response. The wake produced by the Light LMC (left) barely reaches an overdensity of 0.48, while the wake produced by the Heavy LMC is ~ 25 kpc longer than that produced by the Fiducial LMC.

Therefore, we see that the LMC mass is mildly degenerate with the presence of the DM wake for increasing the length of the stellar wake. However, in Section 4 we showed that the DM wake’s gravity lengthens the stellar wake by ~ 50 kpc, roughly twice the increase resulting from raising the LMC mass to $2.5 \times 10^{11} M_\odot$. Additional constraints on the LMC’s mass will

Table 6
Noncomprehensive List of Recent FDM Particle Mass Constraints in the Literature

Allowed Values of m_a (eV)	Technique	References
$<1.1 \times 10^{-22}$	Cusp-core problem solution	Marsh & Pop (2015)
$>2.9 \times 10^{-21}$	$\text{Ly}\alpha$ forest	Armengaud et al. (2017)
$<4 \times 10^{-23}$	Internal kinematics of dSph MW satellites	González-Morales et al. (2017)
$>3.75 \times 10^{-21}$	$\text{Ly}\alpha$ forest	Iršič et al. (2017)
$>10^{-21}$	$\text{Ly}\alpha$ forest	Kobayashi et al. (2017)
$>1.5 \times 10^{-22}$	Stellar stream heating	Amorisco & Loeb (2018)
$<2 \times 10^{-20}$ OR $>8 \times 10^{-19}$	Soliton gravity measurements using M87*	Bar et al. (2019)
$>4 \times 10^{-22}$	Soliton gravity measurements using Sag A*	Bar et al. (2019)
$>6 \times 10^{-23}$	MW disk star heating	Church et al. (2019)
$>10^{-18}$	Core–halo and BH–halo mass relations	Desjacques & Nusser (2019)
$>5.2 \times 10^{-21}$	Subhalo mass function via lensing and stellar streams	Benito et al. (2020)
$>6 \times 10^{-22}$	UFD density profiles	Safarzadeh & Spergel (2020)
$>2.1 \times 10^{-21}$	Subhalo mass function via lensing and stellar streams	Schutz (2020)
$>2.2 \times 10^{-21}$	Stellar stream heating and MW satellite counts	Banik et al. (2021)
$>2 \times 10^{-20}$	$\text{Ly}\alpha$ forest	Rogers & Peiris (2021)
$>3 \times 10^{-19}$	Internal kinematics of UFDs	Dalal & Kravtsov (2022)
$>10^{-23}$	Planck and Dark Energy Survey Year-1 shear measurements	Dentler et al. (2022)

Note. Mass constraints are listed as the region of *allowed* parameter space, in electronvolts. We include a brief note as to the technique used to derive the constraint as well as the reference. Confidence levels for the constraints vary, but are typically at least 2σ .

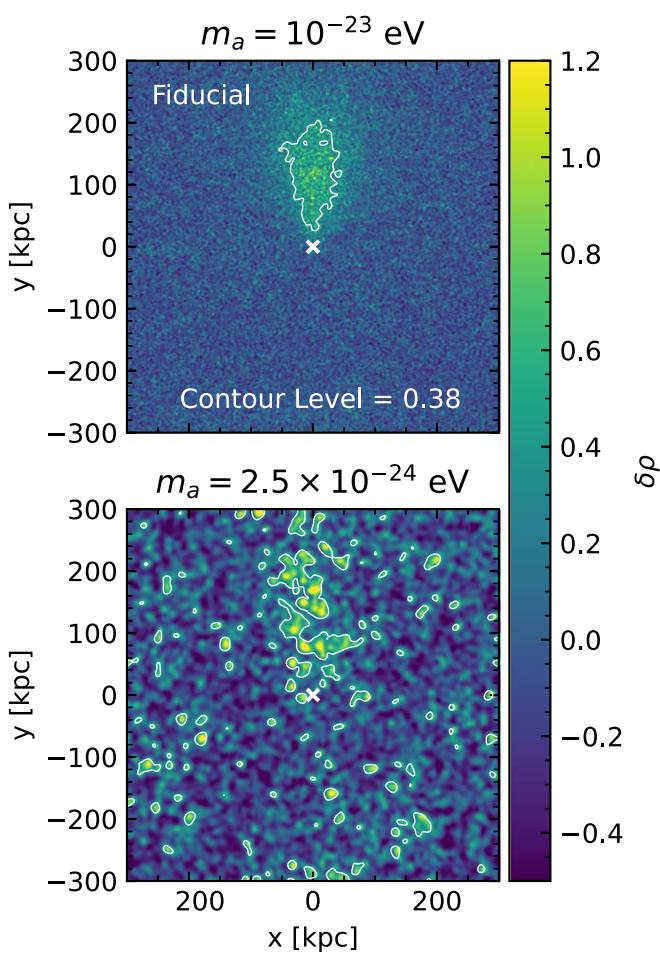


Figure 25. Density of the wakes in both FDM simulations with the Fiducial DM wind parameters (similar to Figure 2). The top panel shows the simulation with $m_a = 10^{-23}$ and is identical to the right panel of Figure 2. The bottom panel shows the simulation with $m_a = 2.5 \times 10^{-24}$. Decreasing the FDM particle mass increases the de Broglie wavelength, granule size and strength, and acts to wash out the wake slightly.

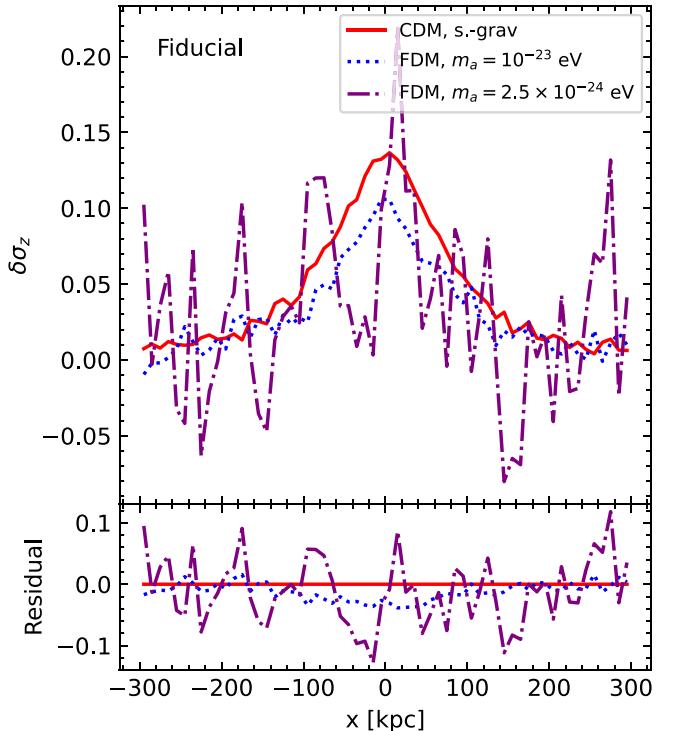


Figure 26. Profile of the z -velocity dispersion across the wake (same as Figure 8) with the $m_a = 2.5 \times 10^{-24}$ eV FDM simulation included as the purple dashed-dotted line. For clarity, we omit the CDM simulation without self-gravity. The increased granule strength at lower particle mass is also seen in the velocity dispersion. The four-fold decrease in particle mass manifests as a roughly four-fold increase in the amplitude of the oscillations in the low-mass FDM profile. On average, the two FDM profiles show a similar $\delta\sigma_z \sim 0.03$ drop in the peak dispersion compared to CDM.

also help break this degeneracy. For example, a different assumption for the mass of the LMC will change its orbit (e.g., Gómez et al. 2015; Patel et al. 2017). Characterizing the location and kinematics of the wake will constrain the LMC’s orbit and in turn its mass (GC19). Thus, both the length and location of the

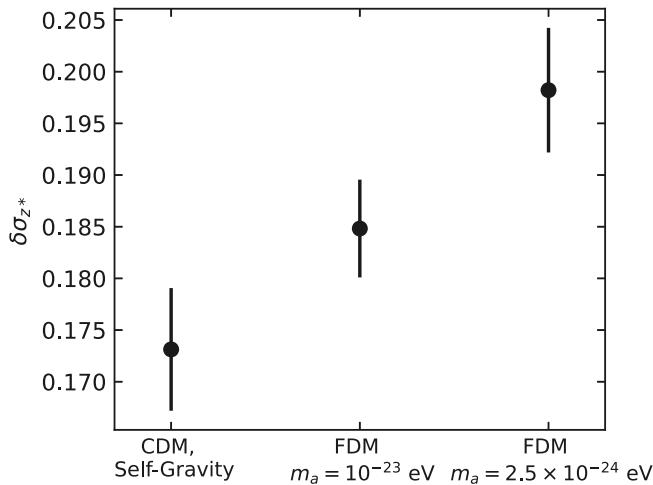


Figure 27. Time-averaged median z -velocity dispersion enhancement of the stellar wake (same as Figure 16) including the CDM simulation with self-gravity and both FDM simulations. The velocity dispersion of the stellar wake is increased more compared to CDM when the FDM particle mass is lower. This demonstrates that the strength of granule heating becomes increases as the FDM particle mass decreases.

wake could be used to break the degeneracy between DM gravity and the LMC’s mass for a given MW model.

6.3. The Effect of the Stellar Velocity Dispersion

While we have so far assumed that the MW stellar halo is smooth and isotropic beyond 70 kpc from the Galactic center, the stellar halo at these distances is likely highly substructured, consisting of streams, shells, and other partially phase-mixed debris from the MW’s past accretion events (e.g., see the review by Helmi 2020, and references therein). These substructures have lower local velocity dispersions than the phase-mixed component of the stellar halo and will complicate measurements of the wake’s influence on halo stars (Cunningham et al. 2020).

While we leave a detailed study of the LMC and wake’s impact on cold substructure for future work, we can ask whether the differences between CDM and FDM due to granule heating become more pronounced if the stars have an initially low velocity dispersion. To this end, we re-run our Fiducial simulations in CDM and FDM with DM self-gravity and a stellar velocity dispersion of $\bar{\sigma}_* = 30 \text{ km s}^{-1}$ (see Table 4), a factor of 3 lower than in our Fiducial simulations.

Figure 29 shows the z -velocity dispersions of the stellar wakes in these new simulations. The black contour is set at the half-max of the CDM simulation with low stellar dispersion. To compare the size of the wake with the Fiducial simulations (in which $\bar{\sigma}_* = 90 \text{ km s}^{-1}$), we reproduce the wake boundary contours from Figure 15 in blue. The impact of lowering the initial stellar velocity dispersion is to narrow the wake (by ~ 60 kpc) in both CDM and FDM.

In Figure 30 we compare the time-averaged median z -velocity dispersion within the stellar wakes in our Fiducial simulations to those with a lowered initial stellar dispersion. The strength of the response increases roughly five-fold in the low-stellar-dispersion simulations compared to the Fiducial simulations. While granule heating is still present in the low-stellar-dispersion simulations, i.e., the stellar wake is hotter in FDM than CDM, the difference between CDM and FDM is $\delta\sigma_{z*} \sim 0.01$, identical to the Fiducial simulations.

These results suggest that the LMC and its wake will leave much stronger kinematic signatures in cold stellar substructures compared to phase-mixed populations of stars. These results warrant further testing with simulations of the LMC’s infall through a substructured stellar halo.

6.4. Self-interacting Dark Matter

Given the differences we find in the LMC’s DM and stellar wakes between FDM and CDM, it is also worth asking whether other DM particle candidates might also impact the DM and/or stellar wake in unique ways. In particular, self-interacting DM (SIDM; Carlson et al. 1992; Spergel & Steinhardt 2000; see also Tulin & Yu 2018 and Adhikari et al. 2022 for recent reviews), in which the DM particle has some nonnegligible cross section for self-scattering, has emerged as another promising alternative to CDM. In the context of DF wakes, self-scattering between DM particles could potentially alter both the density and velocity of DM particles within the wake, as well as induce a bow shock or mach cone in the DM (Furlanetto & Loeb 2002).

For a constant cross section of $\sigma_\chi = 1 \text{ cm}^2 \text{ g}^{-1}$, the mean free path of an SIDM particle within the wake (i.e., at \sim twice the Fiducial wind density of $1.083 \times 10^5 M_\odot \text{ kpc}^{-3}$) is ~ 45 Mpc, so we do not expect the wake itself to differ from CDM in this case. However, in SIDM, the LMC’s DM halo would be subject to ram pressure from the MW’s halo. At the central density of our LMC model’s halo ($\sim 7 \times 10^9 M_\odot \text{ kpc}^{-3}$), the mean free path for an SIDM particle with $\sigma_\chi = 1 \text{ cm}^2 \text{ g}^{-1}$ is reduced to 0.67 kpc, where the scattering may have a nonnegligible effect.

Complicating matters further, velocity-dependent cross sections are compatible with a much wider range of astrophysical observations than constant cross sections (e.g., Yoshida et al. 2000; Vogelsberger et al. 2012, 2016; Correa 2021). Such velocity-dependent cross sections are typically smaller for larger relative velocities, which would alter the efficacy of SIDM ram pressure as the LMC’s orbital velocity changes during its infall. A detailed study of the effects of different SIDM cross sections on the LMC’s wake and DM halo along its orbit would require representing the LMC with a live halo of N -body particles, which we leave to future work.

7. Conclusions

In this paper, we have presented a suite of windtunnel-style simulations of the LMC’s DF wake. Our simulation suite compares the wake at two different points in the LMC’s orbit (223 and 70 kpc from the MW), and with three different assumptions for the DM model (CDM with and without self-gravity, and FDM). We also explored the impact of the LMC alone and the LMC plus the DM wake on the MW’s stellar halo using the three different DM models.

Our goals were to quantify the impacts of self-gravity and the DM particle assumption on the DM wake’s structure and kinematics. We also sought to determine the response of the stellar halo both with and without the gravity of a DM wake, whether different DM particles leave different signatures in the stellar wake, and if these differences are observable when considering typical observational errors.

We summarize our conclusions about the DM wakes as follows:

1. *The FDM and CDM (with self-gravity) wakes both reach comparable peak densities of ~ 1.6 times higher than the background.*

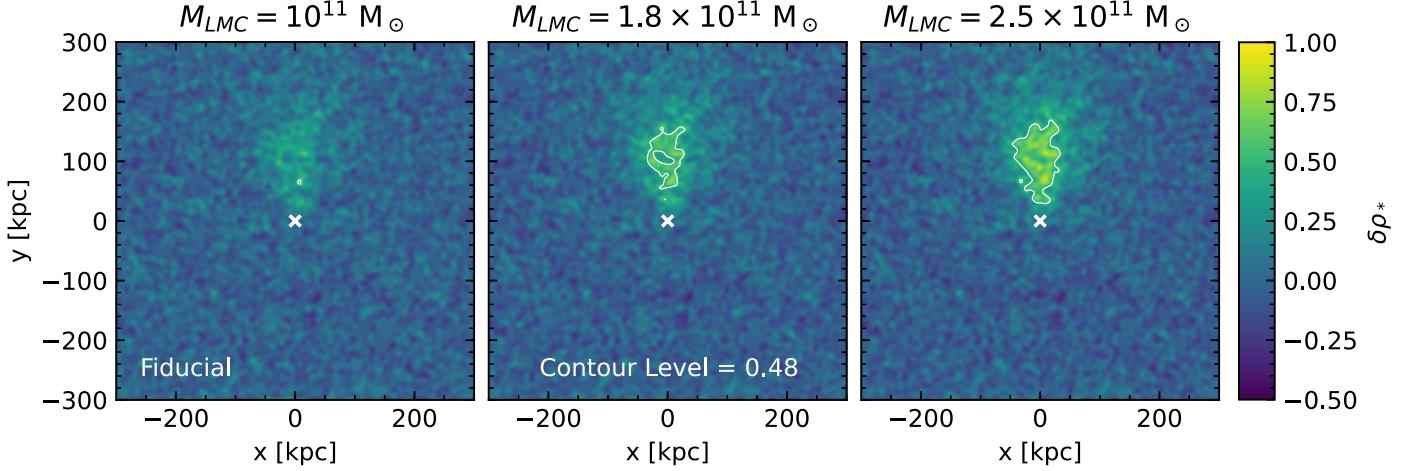


Figure 28. Density of the stellar wake in the three simulations without CDM self-gravity with different LMC mass models (see Table 5), computed identically to Figure 13. From left to right, the panels show the simulation with the Light LMC, Fiducial LMC (i.e., this panel is identical to the left panel of Figure 13), and the Heavy LMC. Increasing the LMC mass increases the length of the response. However, raising the LMC’s mass from $1.8 \rightarrow 2.5 \times 10^{11} M_{\odot}$ increases the wake’s length by only ~ 25 kpc, compared to the ~ 50 kpc increase caused by the addition of the DM wake’s gravity (see Figure 13).

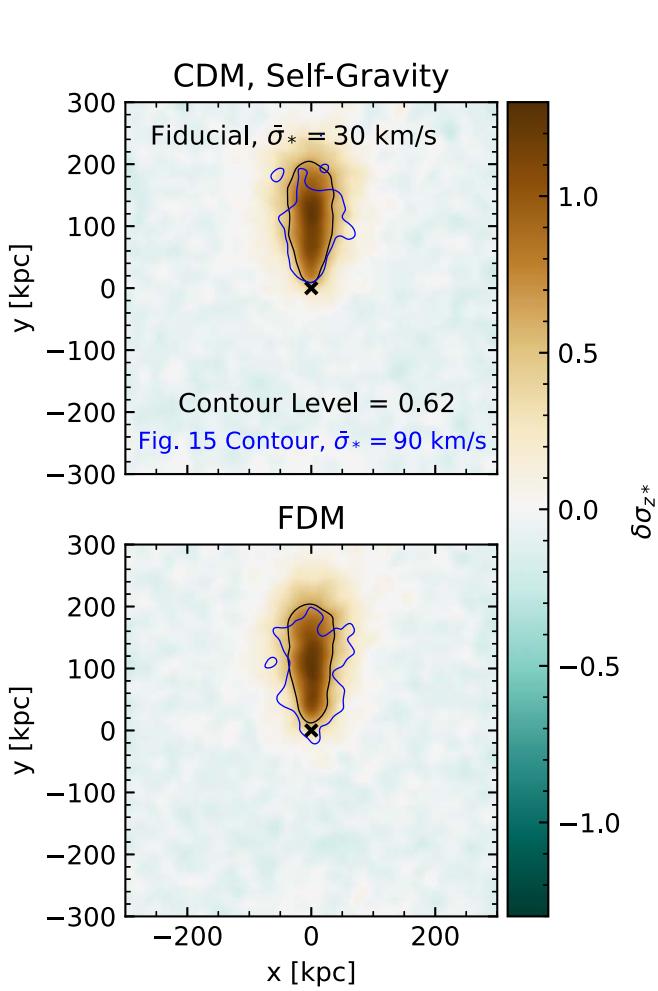


Figure 29. z -velocity dispersion of the stellar wake for the simulations with low stellar velocity dispersion ($\bar{\sigma}_* = 30 \text{ km s}^{-1}$), similar to Figure 15. The response is similar in both DM models. The black contours enclose the region in which the velocity response is greater than half its maximum in CDM, while the blue contours show the wake boundaries from Figure 15 for comparison (in which the initial stellar velocity dispersion is 90 km s^{-1}). The wake is ~ 60 kpc narrower in x when the stellar velocity dispersion is reduced, though the length of the wake does not change significantly.

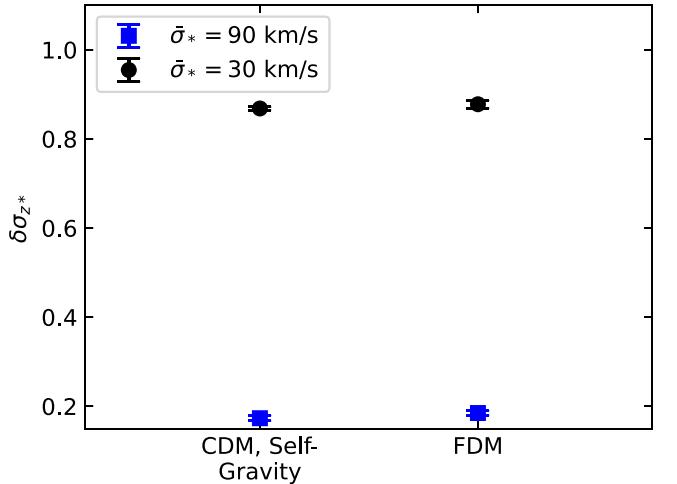


Figure 30. Time-averaged median z -velocity dispersion enhancement of the stellar wake as a function of DM model and stellar wind dispersion (similar to Figure 16). The blue squares show the Fiducial simulations, while the black circles show the simulations with a lowered stellar velocity dispersion. The strength of the response increases by roughly a factor of 5 when the initial stellar velocity dispersion is lowered by a factor of 3. Meanwhile, the strength of granule heating in the FDM simulations is unaffected by the initial velocity dispersion of the stars, i.e., $\delta\sigma_{z*}$ is ~ 0.01 higher in FDM compared to CDM in all cases.

2. *The inclusion of self-gravity increases the density of the CDM wake, and extends its length. The self-gravity of the DM wake cannot be ignored.* The inclusion of self-gravity increases the peak overdensity of the wake by $\sim 10\%$, in agreement with Rozier et al. (2022). In addition, the LMC DM wake sustains a density that is a factor of 1.38 times larger than the background over a distance ~ 50 kpc larger than if self-gravity is ignored. The impact of self-gravity on the properties of the wake is dependent on the LMC’s orbital properties and is maximized after the LMC falls within 100 kpc of the Galactic center. At larger distances, the LMC is moving at lower speeds. As such, particles spend more time under the influence of the LMC’s gravity, which reduces the relative contribution of the wake’s gravity to its structure.

This suggests that the best possible region to search for the influence of the DM wake’s gravity observationally is at Galactocentric distances of 70–100 kpc (which also avoids contamination from the Clouds themselves).

3. *In FDM, the DM wake is more granular as the DM background grows stochastic interference patterns that interact with the LMC.* While individual granules can reach much higher overdensities than are seen in the CDM wake, the overall density of the FDM wake is similar to CDM when self-gravity is included.
4. *The dispersion of the CDM wake with self-gravity is ~ 1.13 times higher than the mean dispersion.* The inclusion of self-gravity increases the velocity dispersion of the CDM wake by $\sim 20\%$.
5. *FDM wakes are 20% colder than CDM wakes regardless of the wind speed and density.* This is due to the reduced response of FDM granules to a steep gravitational potential. Consequently, FDM wakes have a granular structure in kinematic signatures as well (see, e.g., Figure 7), compared to the smooth signatures of CDM. This result is insensitive to the FDM particle mass within the range tested here ($m_a = 2.5 \times 10^{-24}$ – 10^{-23} eV), suggesting this result may hold for higher FDM masses.
6. *The DF drag forces felt by the LMC are similar in FDM and CDM when self-gravity is included.* This result holds across all simulation parameters that were varied (e.g., wind speed, density: Infall versus Fiducial case). As such, we do not expect the LMC’s orbit to change in an FDM universe compared to a CDM universe. When self-gravity is turned off, the drag force is reduced by $\sim 10\%$, consistent with the $\sim 10\%$ reduction in wake density when self-gravity is removed.
7. *The LMC’s DM wake reaches a mass comparable to the LMC’s infall mass in the Fiducial wind case, regardless of the DM model.* To a first approximation, the wake acts like an additional subhalo with a mass of $\sim 1.9(1.5) \times 10^{11} M_\odot$ when self-gravity is on (off)—comparable to the LMC’s infall mass—that trails the LMC by ~ 100 kpc. This implies that the wake is a nonnegligible perturber to the dynamics of MW halo tracers.

We summarize our conclusions about the stellar wakes as follows:

1. *The stellar counterparts to the FDM and CDM (with self-gravity) wakes both reach comparable peak densities of ~ 1.6 times higher than the background.* This is similar to the behavior of the DM wakes alone.
2. *The self-gravity of the DM wake causes the stellar wake to peak at higher densities (by 10%) and persist over larger distances behind the LMC than if there were no DM wake.* The LMC’s gravity will cause the formation of a stellar wake in the absence of DM. However, the stellar wake persists over a larger distance (by ~ 50 kpc) if the DM wake self-gravity is included.
3. *In the CDM simulation with self-gravity, the stellar wake velocity dispersion is ~ 1.173 times higher than the mean stellar dispersion.* The self-gravity of the DM wake causes the stellar wake’s velocity dispersion relative to the mean to increase by $\sim 5\%$.
4. *In the FDM simulations, scattering of stars by FDM granules increases the stellar velocity dispersion relative to the mean by $\sim 5\%$.* Interestingly, this behavior is opposite to

that of the DM wake: while the stellar wake is dynamically hotter in FDM, the DM wake is colder in FDM when compared to CDM. The effect of granule heating in the stellar wake will decrease for higher values of m_a .

5. *Reducing the initial velocity dispersion of the stellar halo by a factor of 3 (from 90 – 30 km s^{-1}) results in an increase in the stellar wake dispersion relative to the mean by a factor of 5 in both the CDM (with self-gravity) and FDM simulations.* This implies that the LMC’s wake will have a stronger imprint on the motions of cold substructures in the stellar halo than on phase-mixed halo stars. Meanwhile, the effect of FDM granule heating remains the same when the initial stellar velocity dispersion is lowered.
6. *The angular extent of the stellar wake on the sky can indicate the existence of a DM wake.* When viewed in Galactic coordinates between distances of 70 and 100 kpc, the stellar wake appears as an enhancement in the density and radial velocity dispersion of the stellar halo. The response appears in the Galactic southeast, traces the past orbit of the LMC, and extends up to the Galactic plane in b when the DM wake’s gravity is included, in agreement with the results of GC19. Without the gravity of the DM wake, the stellar wake decays below an overdensity of 0.34 by $b \approx -20$, and the velocity dispersion enhancement decays below 5.26 km s^{-1} by $b \approx -30$. Thus, the length of the stellar wake is an observational sign of the presence of a DM wake, though this is partially degenerate with the LMC’s mass. Independent constraints on the LMC’s mass or orbit (such as by determining the wake’s location) will help break this degeneracy.
7. *The differences in the density and velocity dispersion of the stellar wake found across the three models considered (CDM with or without self-gravity, and FDM) persist when the wake is viewed in Galactic coordinates with simulated observational errors, provided at least 10^4 (900) stars are observed across the sky (in the wake).* Distinguishing FDM from CDM through measurements of the stellar wake’s velocity dispersion will be difficult, as the effect of granule heating is only at the percent level. However, our results demonstrate granule heating does play a role in DF wakes and merits further study. Additionally, we find in general that the increased velocity dispersion and extent of the stellar wake are telltale features of a DM wake that would distinguish it from cold stellar streams, confirming the findings of GC19. These results underscore the importance of making kinematic measurements when designing observations of the stellar wake.

In this work, we have demonstrated that there are marked differences in the density structure and kinematics of the LMC’s DF wake in a CDM versus an FDM universe, but these differences may be challenging to distinguish in observations of the stellar halo. Significantly, we have also illustrated that the self-gravity of the DM wake plays a crucial role in strengthening and extending the stellar halo’s response to the DM wake—providing a new avenue to test for the existence of DM. Next-generation spectroscopic surveys like DESI (DESI Collaboration et al. 2016a, 2016b), LSST/Vera Rubin Observatory (Ivezic 2019), and the Nancy Grace Roman Space Telescope are poised to provide precision radial velocity and distance measurements of increasing numbers of stars in the stellar halo. These measurements will provide an unprecedented window into the underlying DM

distribution of our Local Group (Dey et al. 2023), and potentially the nature of DM itself.

Acknowledgments

H.R.F. would like to thank Peter Behroozi, Arjun Dey, and Dennis Zaritsky for productive discussions that greatly improved this manuscript, Tomer Yavetz for a particularly helpful discussion on the correspondence between CDM and FDM halos, and the anonymous referee for providing insightful comments that improved the clarity of the paper and presentation of our results.

This work is based upon High Performance Computing (HPC) resources supported by the University of Arizona TRIF, UITS, and Research, Innovation, and Impact (RII) and maintained by the UArizona Research Technologies department. H.R.F. thanks Derrick Zwickl and Chris Reidy for their assistance with MPI troubleshooting, which was made possible through University of Arizona Research Technologies Collaborative Support program.

We respectfully acknowledge the University of Arizona is on the land and territories of Indigenous peoples. Today, Arizona is home to 22 federally recognized tribes, with Tucson being home to the O’odham and the Yaqui. Committed to diversity and inclusion, the University strives to build sustainable relationships with sovereign Native Nations and Indigenous communities through education offerings, partnerships, and community service.

H.R.F. and G.B. are supported by NSF CAREER AST-1941096. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. E.C.C. acknowledges support for this work provided by NASA through the NASA Hubble Fellowship Program grant No. HST-HF2-51502.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS5-26555. F.A.G. acknowledges support from ANID FONDECYT Regular 1211370 and by the ANID BASAL project FB210003. F.A.G.

acknowledges funding from the Max Planck Society through a “Partner Group” grant. C.F.P.L. acknowledges funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 852839).

Software: Arepo (Springel 2010); Astropy (Astropy Collaboration et al. 2013, 2018, 2022); BECDM (Mocz et al. 2017, 2020); Gadget4 (Springel et al. 2021); Healpy (Górski et al. 2005; Zonca et al. 2019); H5py (Collette 2013); Matplotlib (Hunter 2007); Numpy (Harris et al. 2020); Scipy (Virtanen et al. 2020).

Appendix A

Comparison of the Windtunnel Wake to a Wake in a Live Halo

In this appendix, we compare our Fiducial CDM (with self-gravity) wake to our reference simulation (simulation 3 from GC19). Our windtunnel simulations have several important inherent differences from the full-interaction scenarios presented in GC19: in our simulations, the LMC is represented by a fixed potential instead of as a live halo, there are no Galactic tides from the MW, the LMC “travels” in a straight line instead of a curved orbit, and the DM wind parameters are time-independent instead of varying along the orbit as the LMC plunges into the MW’s halo. Therefore, it is vital to compare the CDM wake simulated in this paper to the wake in our reference GC19 Simulation 3 to ensure that our windtunnel is a reasonable laboratory to study differences between CDM and FDM and the impact of self-gravity on the wake.

Figure 31 compares the strength and size of the DM wake in GC19’s Simulation 3 (left) to our Fiducial CDM simulation with self-gravity (right). Similar to Figure 2, we draw a contour enclosing the region with an overdensity greater than half the maximum overdensity in the right panel. Both wakes are very similar in strength and size, demonstrating that our windtunnel simulation framework can well reproduce the wake formed in a full-interaction simulation.

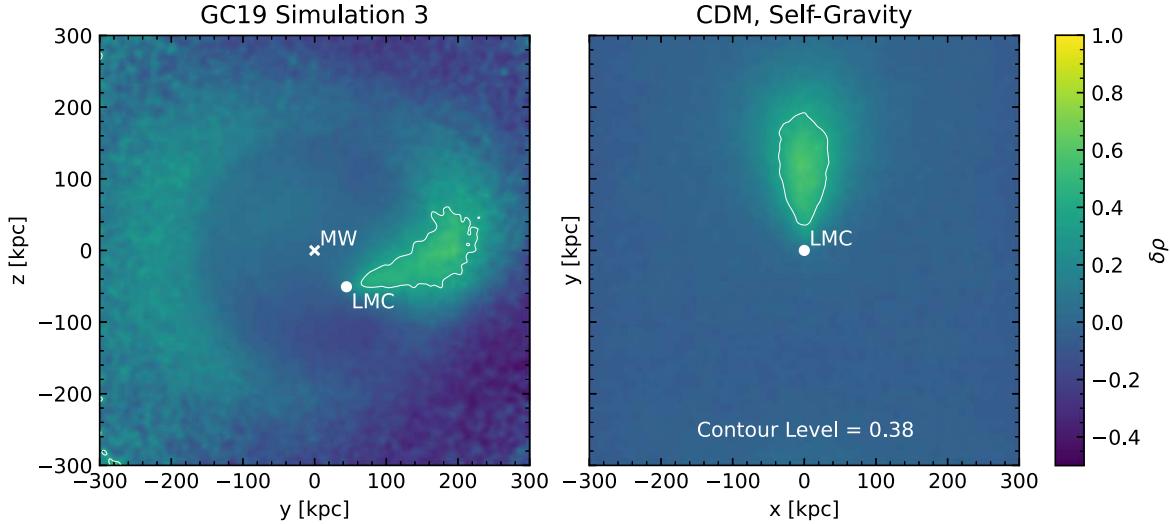


Figure 31. Comparison between the DM overdensity of GC19’s Simulation 3 and our Fiducial CDM simulation with self-gravity. The right panel is a reproduction of the center panel of Figure 2. On the left, we compare this to the overdensity of MW DM particles in GC19’s Simulation 3 when the LMC is at a Galactocentric distance of 70 kpc. We compute the density in the GC19 simulation within a 120 kpc-thick slab centered on the Galactic $y-z$ plane such that the volume of the slab is identical to our Figure 2. We also draw a contour around the region with an overdensity greater than 0.38 as in the right panel. The size and strength of the wake agree between both panels, showing that the wakes formed in our windtunnels are a good approximation of the wake formed in a live-halo simulation.

Appendix B Transformation from Windtunnel to Galactic Coordinates

To perform transformations between coordinate frames (windtunnel simulation box, Galactocentric, and Galactic), we make use of Astropy version 4.2.1 (Astropy Collaboration et al. 2013, 2018, 2022), and adopt the definitions and conventions of Galactocentric and Galactic coordinates as given in this version of Astropy. Our steps for transforming between simulation-box coordinates and Galactocentric coordinates are as follows:

1. For reference, we use the LMC’s present-day location and velocity vector from Kallivayalil et al. (2013). The LMC’s orbit, as usual, is taken from GC19’s Simulation 3.
2. We rotate the simulation box such that the LMC’s unit velocity vector ($-\hat{y}$ in the windtunnel frame) points in the same direction as the LMC’s velocity vector at $r_{\text{MW}} = 70$ kpc from the reference simulation.
3. Next, we translate the simulation box such that the center of the box (where the LMC potential is located) matches the present-day location of the LMC. This ensures that the LMC is as close as possible to the correct location on the sky after the next step.
4. A further translation matches the location of the straight windtunnel orbit and the curved orbit from the reference simulation at a Galactocentric distance of 70 kpc. Together with the rotation, this ensures the LMC’s path in the windtunnel is tangent to the LMC’s orbit at 70 kpc, which is the location our Fiducial wind parameters are drawn from.
5. Finally, the particles are given a velocity boost to remove the bulk wind velocity, i.e., to ensure the wind particles have no net motion in a Galactocentric frame.

Figures 19–22 in Section 5.1 are in Galactic coordinates, and we use Astropy’s built-in functionality for transforming between Galactocentric and Galactic coordinates. When

plotting the velocities of stars in Galactic coordinates, we also remove the Sun’s motion about the Galactic center.

ORCID iDs

Hayden R. Foote <https://orcid.org/0000-0003-1183-701X>
 Gurtina Besla <https://orcid.org/0000-0003-0715-2173>
 Philip Mocz <https://orcid.org/0000-0001-6631-2566>
 Nicolás Garavito-Camargo <https://orcid.org/0000-0001-7107-1744>
 Lachlan Lancaster <https://orcid.org/0000-0002-0041-4356>
 Martin Sparre <https://orcid.org/0000-0002-9735-3851>
 Emily C. Cunningham <https://orcid.org/0000-0002-6993-0826>
 Mark Vogelsberger <https://orcid.org/0000-0001-8593-7692>
 Facundo A. Gómez <https://orcid.org/0000-0002-1947-333X>
 Chervin F. P. Laporte <https://orcid.org/0000-0003-3922-7336>

References

Adhikari, S., Banerjee, A., Boddy, K. K., et al. 2022, arXiv:2207.10638
 Amorisco, N. C., & Loeb, A. 2018, arXiv:1808.00464
 Armengaud, E., Palanque-Delabrouille, N., Yèche, C., Marsh, D. J. E., & Baur, J. 2017, *MNRAS*, **471**, 4606
 Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, *ApJ*, **935**, 167
 Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
 Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
 Banik, N., Bovy, J., Bertone, G., Erkal, D., & de Boer, T. J. L. 2021, *JCAP*, **2021**, 043
 Bar, N., Blum, K., Lacroix, T., & Panci, P. 2019, *JCAP*, **2019**, 045
 Bar-Or, B., Fouvry, J.-B., & Tremaine, S. 2019, *ApJ*, **871**, 28
 Bar-Or, B., Fouvry, J.-B., & Tremaine, S. 2021, *ApJ*, **915**, 27
 Belokurov, V., Deason, A. J., Erkal, D., et al. 2019, *MNRAS: Letters*, **488**, L47
 Benito, M., Criado, J. C., Hützi, G., Raidal, M., & Veermäe, H. 2020, *PhRvD*, **101**, 103023
 Besla, G., Kallivayalil, N., Hernquist, L., et al. 2007, *ApJ*, **668**, 949
 Besla, G., Kallivayalil, N., Hernquist, L., et al. 2010, *ApJL*, **721**, L97
 Bird, S. A., Xue, X.-X., Liu, C., et al. 2021, *ApJ*, **919**, 66
 Buehler, R., & Desjacques, V. 2023, *PhRvD*, **107**, 023516

Carlson, E. D., Machacek, M. E., & Hall, L. J. 1992, *ApJ*, **398**, 43

Chamberlain, K., Price-Whelan, A. M., Besla, G., et al. 2023, *ApJ*, **942**, 18

Chan, H. Y. J., Ferreira, E. G. M., May, S., Hayashi, K., & Chiba, M. 2022, *MNRAS*, **511**, 943

Chandrasekhar, S. 1943, *ApJ*, **97**, 255

Chavanis, P.-H. 2021, *EJPP*, **136**, 703

Church, B. V., Mocz, P., & Ostriker, J. P. 2019, *MNRAS*, **485**, 2861

Cohen, J. G., Sesar, B., Bahnlözer, S., et al. 2017, *ApJ*, **849**, 150

Collette, A. 2013, Python and HDF5 (Sebastopol, CA: O'Reilly & Associates)

Conroy, C., Bonaca, A., Cargile, P., et al. 2019, *ApJ*, **883**, 107

Conroy, C., Naidu, R. P., Garavito-Camargo, N., et al. 2021, *Natur*, **592**, 534

Cooper, A. P., Koposov, S. E., Allende Prieto, C., et al. 2023, *ApJ*, **947**, 37

Correa, C. A. 2021, *MNRAS*, **503**, 920

Correa Magnus, L., & Vasiliev, E. 2022, *MNRAS*, **511**, 2610

Cunningham, E. C., Garavito-Camargo, N., Deason, A. J., et al. 2020, *ApJ*, **898**, 4

Dalal, N., Bovy, J., Hui, L., & Li, X. 2021, *JCAP*, **2021**, 076

Dalal, N., & Kravtsov, A. 2022, arXiv:2203.05750

Deason, A. J., Belokurov, V., Evans, N. W., et al. 2012, *MNRAS*, **425**, 2840

Dentler, M., Marsh, D. J. E., Hložek, R., et al. 2022, *MNRAS*, **515**, 5646

DESI Collaboration, Aghamousa, A., Aguilar, J., et al. 2016a, arXiv:1611.00036

DESI Collaboration, Aghamousa, A., Aguilar, J., et al. 2016b, arXiv:1611.00037

Desjacques, V., & Nusser, A. 2019, *MNRAS*, **488**, 4497

Dey, A., Najita, J. R., Koposov, S. E., et al. 2023, *ApJ*, **944**, 1

Dome, T., Fialkov, A., Mocz, P., et al. 2023, *MNRAS*, **519**, 4183

Erkal, D., Belokurov, V., Laporte, C. F. P., et al. 2019, *MNRAS*, **487**, 2685

Erkal, D., & Belokurov, V. A. 2020, *MNRAS*, **495**, 2554

Erkal, D., Belokurov, V. A., & Parkin, D. L. 2020, *MNRAS*, **498**, 5574

Erkal, D., Li, T. S., Koposov, S. E., et al. 2018, *MNRAS*, **481**, 3148

Ferreira, E. G. M. 2021, *A&ARv*, **29**, 7

Furlanetto, S. R., & Loeb, A. 2002, *ApJ*, **565**, 854

Garavito-Camargo, N., Besla, G., Laporte, C. F. P., et al. 2019, *ApJ*, **884**, 51

Garavito-Camargo, N., Besla, G., Laporte, C. F. P., et al. 2021, *ApJ*, **919**, 109

Garrow, T., Webb, J. J., & Bovy, J. 2020, *MNRAS*, **499**, 804

Gómez, F. A., Besla, G., Carpintero, D. D., et al. 2015, *ApJ*, **802**, 128

González-Morales, A. X., Marsh, D. J. E., Peñarrubia, J., & Ureña-López, L. A. 2017, *MNRAS*, **472**, 1346

Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759

Grady, J., Belokurov, V., & Evans, N. W. 2021, *ApJ*, **909**, 150

Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357

Helmi, A. 2020, *ARA&A*, **58**, 205

Hernquist, L. 1990, *ApJ*, **356**, 359

Hložek, R., Marsh, D. J., Grin, D., et al. 2017, *PhRvD*, **95**, 123511

Hu, W., Barkana, R., & Gruzinov, A. 2000, *PhRvL*, **85**, 1158

Hui, L. 2021, *ARA&A*, **59**, 247

Hui, L., Ostriker, J. P., Tremaine, S., & Witten, E. 2017, *PhRvD*, **95**, 043541

Hunter, J. D. 2007, *CSE*, **9**, 90

Iršič, V., Viel, M., Haehnelt, M. G., Bolton, J. S., & Becker, G. D. 2017, *PhRvL*, **119**, 031302

Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, **873**, 111

Kallivayalil, N., van der Marel, R. P., Besla, G., Anderson, J., & Alcock, C. 2013, *ApJ*, **764**, 161

Katz, D., Sartoretti, P., Cropper, M., et al. 2019, *A&A*, **622**, A205

Khlopov, M. Y., Malomed, B. A., & Zeldovich, Y. B. 1985, *MNRAS*, **215**, 575

Kobayashi, T., Murgia, R., De Simone, A., Iršič, V., & Viel, M. 2017, *PhRvD*, **96**, 123514

Koposov, S. E., Erkal, D., Li, T. S., et al. 2023, *MNRAS*, **521**, 4936

Lancaster, L., Giovanetti, C., Mocz, P., et al. 2020, *JCAP*, **2020**, 001

Laporte, C. F. P., Gómez, F. A., Besla, G., Johnston, K. V., & Garavito-Camargo, N. 2018a, *MNRAS*, **473**, 1218

Laporte, C. F. P., Johnston, K. V., Gómez, F. A., Garavito-Camargo, N., & Besla, G. 2018b, *MNRAS*, **481**, 286

Lilleengen, S., Petersen, M. S., Erkal, D., et al. 2023, *MNRAS*, **518**, 774

Madelung, E. 1927, *ZPhy*, **40**, 322

Marsh, D. J. E., & Pop, A.-R. 2015, *MNRAS*, **451**, 2479

May, S., & Springel, V. 2021, *MNRAS*, **506**, 2603

May, S., & Springel, V. 2023, *MNRAS*, **524**, 4256

Mocz, P., Fialkov, A., Vogelsberger, M., et al. 2020, *MNRAS*, **494**, 2027

Mocz, P., Vogelsberger, M., Robles, V., et al. 2017, *MNRAS*, **471**, 4559

Navarro, J. F., Frenk, C. S., & White, S. D. M. 1997, *ApJ*, **490**, 493

Ostriker, E. C. 1999, *ApJ*, **513**, 252

Patel, E., Besla, G., & Sohn, S. T. 2017, *MNRAS*, **464**, 3825

Patel, E., Kallivayalil, N., Garavito-Camargo, N., et al. 2020, *ApJ*, **893**, 121

Peebles, P. J. E. 2010, arXiv:1009.0496

Peñarrubia, J., Gómez, F. A., Besla, G., Erkal, D., & Ma, Y.-Z. 2016, *MNRAS*, **456**, L54

Petersen, M. S., & Peñarrubia, J. 2020, *MNRAS*, **494**, L11

Petersen, M. S., & Peñarrubia, J. 2021, *NatAs*, **5**, 251

Power, C., Navarro, J. F., Jenkins, A., et al. 2003, *MNRAS*, **338**, 14

Rogers, K. K., & Peiris, H. V. 2021, *PhRvL*, **126**, 071302

Rozier, S., Famaey, B., Siebert, A., et al. 2022, *ApJ*, **933**, 113

Safarzadeh, M., & Spergel, D. N. 2020, *ApJ*, **893**, 21

Salem, M., Besla, G., Bryan, G., et al. 2015, *ApJ*, **815**, 77

Scannapieco, E., & Brüggen, M. 2015, *ApJ*, **805**, 158

Schive, H.-Y., Chiueh, T., & Broadhurst, T. 2014, *NatPh*, **10**, 496

Schive, H.-Y., Chiueh, T., Broadhurst, T., & Huang, K.-W. 2016, *ApJ*, **818**, 89

Schneider, E. E., & Robertson, B. E. 2017, *ApJ*, **834**, 144

Schutz, K. 2020, *PhRvD*, **101**, 123026

Schwabe, B., Niemeyer, J. C., & Engels, J. F. 2016, *PhRvD*, **94**, 043513

Seabroke, G. M., Fabricius, C., Teyssier, D., et al. 2021, *A&A*, **653**, A160

Shipp, N., Erkal, D., Drlica-Wagner, A., et al. 2021, *ApJ*, **923**, 149

Shipp, N., Li, T. S., Pace, A. B., et al. 2019, *ApJ*, **885**, 3

Sparre, M., Pfrommer, C., & Ehlers, K. 2020, *MNRAS*, **499**, 4261

Sparre, M., Pfrommer, C., & Vogelsberger, M. 2019, *MNRAS*, **482**, 5401

Spergel, D. N., & Steinhardt, P. J. 2000, *PhRvL*, **84**, 3760

Springel, V. 2010, *MNRAS*, **401**, 791

Springel, V., Pakmor, R., Zier, O., & Reinecke, M. 2021, *MNRAS*, **506**, 2871

Tamfai, T., Mayer, L., Quinn, T. R., et al. 2021, *ApJ*, **916**, 55

Traykova, D., Clough, K., Helfer, T., et al. 2021, *PhRvD*, **104**, 103014

Trelles, A., Valenzuela, O., Roca-Fábregas, S., & Velázquez, H. 2022, *A&A*, **668**, A20

Tulin, S., & Yu, H.-B. 2018, *PhR*, **730**, 1

van der Marel, R. P., Fardal, M., Besla, G., et al. 2012, *ApJ*, **753**, 8

Vasiliev, E. 2023, *Galax*, **11**, 59

Vasiliev, E., Belokurov, V., & Erkal, D. 2021, *MNRAS*, **501**, 2279

Vera-Ciro, C., & Helmi, A. 2013, *ApJL*, **773**, L4

Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261

Vitsos, A., & Gourgouliatos, K. N. 2023, *JCAP*, **2023**, 071

Vogelsberger, M., Zavala, J., Cyr-Racine, F.-Y., et al. 2016, *MNRAS*, **460**, 1399

Vogelsberger, M., Zavala, J., & Loeb, A. 2012, *MNRAS*, **423**, 3740

Walker, M. G., & Peñarrubia, J. 2011, *ApJ*, **742**, 20

Weinberg, M. D. 1998a, *MNRAS*, **297**, 101

Weinberg, M. D. 1998b, *MNRAS*, **299**, 499

Widrow, L. M., & Kaiser, N. 1993, *ApJL*, **416**, L71

Xue, X.-X., Rix, H.-W., Ma, Z., et al. 2015, *ApJ*, **809**, 144

Yavetz, T. D., Li, X., & Hui, L. 2022, *PhRvD*, **105**, 023512

Yoshida, N., Springel, V., White, S. D. M., & Tormen, G. 2000, *ApJL*, **544**, L87

Zagorac, J. L., Sands, I., Padmanabhan, N., & Easter, R. 2022, *PhRvD*, **105**, 103506

Zonca, A., Singer, L., Lenz, D., et al. 2019, *JOSS*, **4**, 1298